Table of Contents

International Productivity Monitor — Number 41, Fall 2021

PUBLICATIONS MAIL AGREEMENT NO 40049476
RETURN UNDELIVERABLE CANADIAN ADDRESSES TO:
CSLS
604-170 Laurier Avenue West
Ottawa ON K1P 5V5
email: csls@csls.ca

Editors' Overview

We are pleased to feature in the 41^{st} issue the *International Productivity Monitor* a symposium on the relationship between productivity and pay, an important topic the journal has addressed on many previous occasions. The contributions in this symposium provide up-to-date estimates on this relationship for the United States, the United Kingdom and Canada. In addition, they offer new and original interpretations of what is driving the gap between productivity and pay, and the different ways it has manifested itself across those three economies.

The first article of the symposium by **Jacob Greenspon, Anna Stansbury** and **Lawrence H. Summers** provides a comparative perspective between Canada and the United States. The authors make a distinction between two sources of the gap between productivity and pay: *divergence* which measures the degree to which productivity has grown faster than pay, and *delinkage* which refers to the degree to which incremental increases in the rate of productivity growth translate into incremental increases in the rate of growth of pay. The authors find that while *divergence* has occurred in both countries, there has been little *delinkage* as periods of faster productivity growth also saw an increase in pay. This implies that pro-productivity policies tend to raise middle class incomes.

The second article by **Andreas Teichgräber** and **John Van Reenen** finds much less of a decoupling between productivity and median wages in the United Kingdom than other studies have found for the United States. The divergence they do find is largely explained by rising wage inequality and to a lesser extent by a rise in non-wage compensation costs. The authors also address the relatively large role of lower increases in compensation of self-employed workers for their activities because of the relatively large share of "solo self-employed" and a large fall in hours worked by the self-employed.

The third article by **Lawrence Mishel** and **Josh Bivens** provides complementary arguments to the observation in the first article that factors which are independent to productivity growth have been driving productivity and typical pay in the United States further apart. Excessive unemployment, eroded collective bargaining, and corporate-driven globalization explain more than half of the divergence according to the authors, where a diminished overtime salary threshold, employee misclassification, employer-imposed noncompete agreements, and corporate fissuring-subcontracting and major-buyer dominance also explain a fair component.

The final article in the symposium by **Andrew Sharpe** and **James Ashwell** shows that the gap between productivity and real median wage growth in Canada has fallen quite considerably since 2000. They argue that the bargaining power of workers fell dramatically in the last quarter of the 20^{th} century due to high unemployment, falling unionization rates and a rising import share, but that since 2000 trends in these factors have reversed or stabilized since.

Together the four articles in this symposium provide good and bad news for the majority of wage earners in the three countries under consideration. On the positive side, productivity remains a key driver of earnings. However, on the negative side, the link between productivity and pay is often affected by other factors, not directly related to pay, but due to institutions, labour market imperfections and political preferences regarding income policies, taxation, etc. Restoring the link could also be an important incentive for reverse causality, namely that higher wages could provide an incentive to productivity.

For decades the Penn World Table (PWT) has been a widely used data resource on comparative measures of prices and income levels. In the most recent versions of PWT, starting with 8.0, measures of output and productivity have also been introduced, and the last version 10.0 now includes such time series for 183 economies from 1950 to 2019. Surprisingly, in the latest version, several less developed countries have a total factor productivity (TFP) level well above that of the United States. The article by **Robert Inklaar** and **Pieter Woltjer** discuss the case of Egypt, which in 2017 had a TFP level 123 per cent that of the United States. They trace this anomalous outcome to the underlying measurement and modelling issues on comparative inputs. The authors argue that the development accounting framework in PWT is a useful guide to distinguishing outliers from regular patterns in the data.

The dispersion of productivity within industries has been a key topic for productivity researchers, as the issue is pervasive across countries, industries and time. The article by **Cindy Cunningham, Sabrina Wulff Pabilonia, Jay Stewart, Lucia Foster, Cheryl Grim, John Haltiwanger** and **Zoltan Wolf** uses new dispersion measures on productivity in US manufacturing industries, describing how periods of innovation are initially associated with a surge in business start-ups, followed by increased experimentation that leads to rising dispersion potentially with declining aggregate productivity growth, and then a shakeout process that results in higher productivity growth and declining productivity dispersion.

In the Spring 2021 issue of the *International Productivity Monitor*, we published a review article by Bert Balk of the volume *Measurement of Productivity and Efficiency: Theory and Practice* by **Robin Sickles** and **Valentin Zelenyuk**. In a response to this review article, the authors explain how the material they cover in the first seven chapters of their book builds on and expands in important ways Balk's own book, *Industrial Price, Quantity, and Productivity Indices: The Micro-Economic Theory and Applications on Productivity*, which was published more than twenty years ago.

Productivity and Pay in the United States and Canada

Jacob Greenspon
Harvard Kennedy School of Government
Anna Stansbury
MIT Sloan School of Management
Lawrence H. Summers
Harvard University[1]

Abstract

We study the productivity-pay relationship in the United States and Canada along two dimensions. The first is *divergence*: the degree to which productivity has grown faster than pay. The second is *delinkage*: the degree to which incremental increases in the rate of productivity growth translate into incremental increases in the rate of growth of pay, holding all else equal. In both countries there has been divergence: the pay of typical workers has grown substantially more slowly than average labour productivity over recent decades, driven by both rising labour income inequality and a declining labour share of income. Even as the levels of productivity and pay have grown further apart, however, we find evidence for a substantial degree of linkage between productivity growth and pay growth: in both countries, periods with faster productivity growth rates have been periods with faster rates of growth of the pay of average and typical workers, holding all else equal. This linkage appears somewhat stronger in the US than in Canada. Overall, our findings lead us to tentatively conclude that policies or trends which lead to incremental increases in productivity growth, particularly in large relatively closed economies like the USA, will tend to raise middle class incomes. At the same time, other factors orthogonal (i.e. statistically independent) to productivity growth have been driving productivity and typical pay further apart, emphasizing that much of the evolution in middle class living standards will depend on measures bearing on relative incomes.

1 Jacob Greenspon is a student in the Master in Public Policy program at the Harvard Kennedy School of Government. Anna Stansbury is an Assistant Professor of Work and Organization Studies at the MIT Sloan School of Management. Lawrence H. Summers is Charles W. Eliot University Professor at Harvard University. The authors thank Andrew Sharpe, Bart van Ark, and two anonymous referees for comments and suggestions. Corresponding author: Anna Stansbury, amms@mit.edu.

Introduction and Conceptual Framework: Divergence and Delinkage

Economic theory predicts that workers in aggregate should receive higher compensation as the value of what they are able to produce increases (assuming their supply is not perfectly elastic). The growing gap between labour productivity and the pay of typical workers in the United States since the early 1970s has therefore been the subject of much attention. In Canada, there has been similar focus on a growing gap between productivity and pay in recent decades. If productivity has been growing fast, yet pay of typical workers has been growing slowly, this raises the question: does productivity growth benefit typical workers by raising their pay?

Stansbury and Summers (2019) approached this question for the United States, finding that despite the growing gap between productivity and typical pay over 1948-2016, incremental increases in the rate of productivity growth appeared to translate close to one-for-one into incremental increases in the rate of growth of pay of typical workers. This result suggests that there is still transmission from productivity growth to pay growth – implying that faster productivity growth would all else equal benefit typical workers by raising their pay – but that other factors not related to productivity growth have been suppressing pay over recent decades even as productivity has been acting to raise it.

Has this relationship held for Canada as well? How does Canada's productivity-pay relationship compare to that of the United States? And what can this comparison tell us about the productivity-pay link in general? We tackle these questions in this article, emphasizing two dimensions of the relationship between productivity and pay, which we call divergence and delinkage.

We use divergence to refer to a growing gap between productivity and pay in levels: to the extent that productivity has grown faster than pay over time, the two series have diverged over the period.[2] As documented by Williams (2021) and Sharpe and Ashwell (2021) in Canada, and Bivens and Mishel (2015, 2021) among others in the United States, there has been divergence in both countries between labour productivity and the pay of both the average worker and typical workers. In this article we compare the degree of divergence across the two countries along different metrics and time periods.

We use *delinkage* to refer to the relationship between the growth rates of productivity and pay, holding all else equal. At one extreme, if a one percentage point higher growth rate in productivity is associated with a one percentage point higher growth rate in pay, all else equal, we consider the two series to be linked: it suggests that an incrementally higher rate of productivity growth will lead to an incrementally higher rate of pay growth. At the other extreme, if a one percentage point lower growth rate in

2 Labour productivity will always be higher than pay if returns to other factors of production are non-zero. By divergence, we refer to an increase in the gap between productivity and pay over time. The degree of divergence can be visualized by indexing both productivity and pay to 100 in a given year, and charting the extent to which the two series diverge over time.

productivity is associated with no change in pay all else equal, we consider the two series to be delinked: it suggests that an incrementally higher or lower rate of productivity growth will have no effect on the rate of growth of pay.

In Stansbury and Summers (2019), we showed that despite divergence in levels between productivity and median pay in the United States, there was still substantial linkage in the growth rates of the series. This suggested that incremental increases in aggregate productivity growth could be expected to translate close to one-for-one into increases in typical pay, holding all else equal. In this article we update our analysis for the United States and examine the degree of linkage or delinkage in Canada.

Why does the degree of linkage or delinkage, as measured through annual changes, matter? If two series – like productivity and pay – have diverged, studying the degree of linkage or delinkage helps us diagnose why this has happened. On one extreme is complete delinkage: something may be blocking the transmission mechanism from productivity to pay, so that an incremental increase in productivity growth does not translate into an incremental increase in pay growth. Indeed, the factor causing increased productivity growth may itself be acting to suppress workers' pay – for example, a technological change which increases productivity but leads to the substitution of labour for capital.[3] On the other extreme is *complete linkage*: an incremental increase in productivity growth translates one-for-one into a boost to workers' pay growth, all else equal – suggesting that the transmission mechanism from productivity to pay is functioning, but that at the same time other factors orthogonal (i.e. statistically independent) to productivity growth are suppressing pay and therefore responsible for the rising productivity-pay gap. Understanding the degree of linkage in the productivity-pay relationship can therefore shed light on the degree to which incremental productivity growth helps boost workers' pay.

The concept can be illustrated by a simple metaphor: water in a bucket. Think of pay as the level of water in a bucket. Think of water running into the bucket from a hose as productivity growth. Over the last forty years, the faucet has been running – productivity has been growing – but the level of water in the bucket has barely risen. Why might this be? It is possible that the hose is broken and is leaking water somewhere between the faucet to the bucket: the transmission mechanism from productivity to pay is broken (*delinkage*). Or, it is possible that there is a hole in the bucket: even as the water flowing from the hose is increasing the water level in the bucket, water is draining from the hole at the bottom of the bucket, meaning that on net the water level does not rise (*linkage*). This has implications for what we expect to happen if we increase the water pressure at the faucet. If there is a hole in the hose, this may not make any difference: more water

3 For example, capital-augmenting technological change with an elasticity of substitution between labour and capital of greater than 1, or labour-augmenting technological change with an elasticity of substitution between labour and capital of less than 1 (Lawrence, 2015).

flowing into the hose will not affect the water level in the bucket. However, if there is instead a hole in the bottom of the bucket, increasing the rate of inflow of water into the bucket will indeed make a difference – it will slow the rate of draining of the bucket, and may even increase the water level on net.

As this metaphor illustrates, understanding the degree of linkage or delinkage between productivity and the pay of typical workers is therefore of considerable policy significance. Much of traditional economic policy discussion is directed at accelerating productivity growth, whether through promoting investment, technological progress or better functioning markets. If the delinkage hypothesis holds, then success or failure on these dimensions will have little impact on middle class well-being: an incrementally higher rate of productivity growth will not be expected to feed through into pay growth for typical workers, and the benefits instead will be reaped by others. Success in raising middle class incomes will depend *entirely* on distributional measures. In contrast if, despite divergence between productivity and middle-class pay, there is still linkage between the two, then it is still the case that increased productivity growth should be expected to increase the pay of typical workers – an incrementally higher rate of productivity growth will exert upward pressure on typical workers' pay, and so a rising tide can be expected to lifts all boats to some extent – even as there may at the same time be other variables reducing the relative growth in incomes of the lower or middle parts of the distribution (like the declining bargaining power of workers, or the increase in globalization).

Similarly, these conclusions can help us understand the dynamics of the past. Counterfactually, if there had been slower growth in productivity over recent decades, what would have happened to typical workers' pay? Under the delinkage hypothesis, it would have made little difference to the rate of growth of typical pay (which was relatively slow in real terms over recent decades in both Canada and the United States). Under the linkage hypothesis, typical workers' wage growth may have already been slow for other reasons (like declining worker bargaining power), and so slower productivity growth would have meant even slower (or perhaps even negative) real wage growth for typical workers over the period.

We begin our analysis in section 2, examining the degree of divergence between productivity and compensation in the United States and Canada since 1961. We ask the question "To what extent has productivity grown faster than average pay, or than the pay of typical workers?". We use two measures to proxy for the pay of typical workers in each country. For the United States, we use the average compensation of production and nonsupervisory workers, as well as median hourly compensation (available only since 1973). For Canada, we use a new measure we have constructed reflecting average hourly compensation of hourly-paid workers in five sectors for which historical time series data is available: manufacturing, mining, construction, laundries, and hotels and restaurants; since 1976, we also use an estimate of median hourly compensation.[4]

In both countries, the last four to five decades have seen substantial divergence in the growth paths of labour productivity and the pay of typical workers: labour productivity has grown much faster than real compensation for typical workers. This divergence can be thought of in terms of three wedges (as illustrated by Bivens and Mishel 2015): (i) a decline in the labour share of income (faster growth in labour productivity than average compensation, deflated by the same price deflator), (ii) a rise in inequality in (pre-tax and transfer) labour incomes caused by faster growth in average compensation than the compensation of typical workers, and (iii) a decline in labour's terms of trade (faster growth in consumer than producer price deflators).

These three trends have played out somewhat differently in the two countries. In both countries, a declining labour share of income has generated increasing divergence between labour productivity and average compensation. In both countries, market labour income inequality has also risen over the period we study, generating increasing divergence between average labour productivity and the real compensation of typical workers – but this trend has been much more pronounced in the United States. Trends in relative price deflators have differed across the two countries and time periods. On net, since labour productivity growth has been much faster in the United States than in Canada – particularly since 1976 – average compensation has also grown much faster in the United States than in Canada even though the US labour share has fallen. But since labour income inequality as measured by the mean-median ratio has risen so much faster in the United States than in Canada, the growth in real median compensation has been about the same in both countries despite the much faster growth in productivity in the United States.

In section 3, we examine the degree of linkage or delinkage between hourly labour productivity and measures of average and typical worker compensation in the United States and Canada since 1961 – asking the question "To what extent does an incremental increase in productivity growth translate into an incremental increase in compensation growth, all else equal?".

In the United States, we find relatively strong linkage between productivity and pay, both for average compensation and our measures of typical workers' compensation (updating the evidence in Stansbury and Summers 2019). In regressions of the three-year moving average of the change in log compensation on the change in log productivity, controlling for unemployment, we find that over recent decades a one percentage point increase in the rate of productivity growth in the United States has been associated with 0.6-0.8 percentage points faster average compensation growth, 0.5-0.7 percentage points faster median compensation growth, and 0.3-0.9

4 1976 is the first year for which median pay data for Canada are available.

5 Coefficients depend on the time period considered (1948-2019 or 1973-2019), the price deflator used for the compensation series, and the moving average length. All coefficients were strongly significantly different from

percentage points faster growth in the compensation of production and nonsupervisory workers.[5]

In Canada, we find moderate linkage between productivity and average compensation: a one percentage point increase in productivity growth has been associated all else equal with 0.3-0.7 percentage points higher average compensation growth (with coefficients in later decades not statistically significantly different from zero). We find strong evidence of linkage between productivity and typical compensation as proxied by hourly-rated workers in five sectors over 1961-2019, but estimates for more recent periods are too noisy to rule out either of our extreme cases of strong linkage (one-for-one translation from productivity to pay) or strong delinkage (no translation of productivity to pay).[6] We find no evidence of linkage between productivity and median compensation, but large standard errors and measurement error concerns in the median hourly compensation series suggest that this result should be interpreted with substantial caution.

Why does there seem to be a weaker link between pay and productivity in Canada than in the United States? In section 4, we explore possible explanations. One possibility is that Canada is a smaller, more internationally open economy than the United States. The smaller and more open the economy, the greater the degree to which productivity and real compensation may be determined abroad rather than domestically – and the less, therefore, one might expect researchers to be able to detect a process where productivity increases translate into increases in real compensation. We present evidence consistent with this hypothesis: coefficient estimates on regressions of average compensation on productivity in US regions – which are similar in GDP and population to Canada, and could be considered small open economies – are substantially lower than for the United States as a nation, and similar in magnitude to the estimates for Canada. A second possibility is that there is more meaningful high frequency variation in productivity in the United States than in Canada, but we do not find substantial evidence to support this.

Finally, in section 5 we conclude, discussing other possible drivers of the US-Canada differences and implications for policy.

To What Extent Have Productivity and Pay Diverged?

Several studies have explored the divergence between productivity and pay in both Canada and the United States, with conclusions dependent in large part on measurement choices, especially regarding the output measure used for productivity. Williams (2021) provides an overview for Canada of measures and data sources, the recent literature, and an analysis of trends. Bivens and Mishel (2015, 2021) and Stansbury and Summers (2019) are among the

zero and many coefficients were not significantly different from one.

6 In addition, breakpoint tests indicate a structural break in the relationship between productivity and this five-sector measure of typical pay around 1997 or 2000.

papers on the United States which document the divergence between productivity and pay, and examine the role of measurement choices for depreciation, price deflators, and productivity and compensation series. In this section we describe the measures used in our analysis and compare key trends in the US and Canada in light of previous research.

Data

Data for the measures used in this analysis are carefully constructed from a variety of sources, which are provided in the appendix (available online at [http://www.csls.ca/IPM])

The main measure of labour productivity is total economy output per hour worked, where output is net of depreciation. This better reflects the production output available for labour compensation than a productivity measure based on gross domestic product, especially when the capital depreciation rate is increasing (as it has in both Canada and the United States since the mid 1970s).[7] In both countries, this total economy productivity measure is a conservative estimate of productivity growth: changes in output in the government and non-profit sectors are generally calculated from changes in inputs, due to challenges quantifying public sector output that is not priced, which implicitly assumes zero productivity growth (BEA, 2018). If we assumed instead that there was positive productivity growth in the government and non-profit sectors, we would find an even larger divergence between productivity and pay (due to a faster growth rate for the former).

We explore the divergence of productivity from compensation by looking at four measures of compensation for each country:

Average hourly compensation of all workers

This measure, used in our baseline analyses, reflects average hourly compensation of all employed persons in the total economy.[8] For Canada we construct this measure from total compensation divided by total hours worked (obtained from Statistics Canada). For the United States we use the average hourly compensation of all employed persons (obtained from the BLS total economy productivity data set). Note that since it is difficult to determine whether the income of self-employed workers is considered to be compensating labour or capital, our measure of average hourly compensation for both the United States and Canada includes an imputed mea-

7 Unlike in Williams (2021), our productivity measure does not exclude taxes less subsidies in order to be consistent with US data. In addition, a decrease in tax rates leads to an apparent increase in productivity (which can significantly affect the year-to-year changes which are the focus of our study).

8 The long-term total compensation and hours worked series were calculated from combining 1961-1997 data with 1997-2019 data (see on-line Appendix for details). One caveat is that the publicly-available 1961-1997 data (from Statistics Canada Table 36-10-0303-01, the same table that Uguccioni (2016) use to construct 1976-1997 data) was last revised in 2007 and is no longer in use. Williams (2021) instead uses historical data provided by Statistics Canada on special request that differs slightly (comparison available upon request).

9 The imputation is carried out by the respective statistical agencies of the United States and Canada. Williams (2021) reports that in Canada self-employment as a share of total labour compensation peaked in the 1990s and is now approximately 4 per cent. Average compensation including the self-employed has grown faster than

sure of compensation for the self-employed based on wages of similar employee workers.[9]

Average hourly compensation of employees only

This measure excludes the self-employed and therefore avoids challenges with imputing their labour income. For both Canada and the United States this is constructed as total compensation of employees divided by total hours worked by employees (in the total economy).

Median hourly compensation

Median compensation is a better reflection than average compensation of the experience of typical middle-income workers, as it is not skewed by large changes at the top or bottom of the income distribution. Median compensation measures are, however, not available over the entire period of our analysis for either country. For the United States, we use the Economic Policy Institute's estimates of median hourly compensation for each year since 1973. This is calculated by estimating median hourly wage and salary income from the Current Population Survey, and then multiplying this by the average ratio of total compensation to wage and salary income obtained from the Bureau of Economic Analysis. For Canada, we estimate a measure of median hourly compensation using Statistics Canada's data on median annual wage and salary income, divided by our estimate of median annual hours worked as estimated from the Labour Force Survey, and multiplied by the ratio of total compensation of employees to total wage/salary income of employees (obtained from Statistics Canada). Unfortunately, our measure of median hourly compensation in Canada must be approached with caution given substantial measurement concerns in our estimates of the number of hours worked. (A direct measure of median hourly compensation, which would be preferable to use, is only available from the Labour Force Survey from 1997 onwards).[10]

Typical compensation for hourly paid workers

Since median compensation measures are not available for the whole period, and given measurement error concerns for the Canadian median compensation series, we also provide an additional measure of 'Typical compensation' for each country.

Production/nonsupervisory compensation (US)

In the United States, we use average hourly compensation of production/nonsupervisory employees as another

---

average compensation of employees in Canada over 1961-2019, as illustrated in Appendix Chart 4. There is much less difference between the growth rates of compensation including or excluding the self-employed in the United States.

10 Median compensation in Canada is calculated by dividing median annual wage/salary income by a median annual hours worked measure constructed from the Labour Force Survey. However, both the numerator and denominator contain possible measurement error. The income measure includes wages/salaries of anyone that has worked at all in a year (even for one week), so the median may be skewed by, for example, changing numbers of seasonal workers. The denominator calculates annual hours worked as 52 multiplied by median weekly hours worked, which ignores changes over time in weeks worked per year (for example, due to parental leave).

proxy for the average compensation of typical workers. This is estimated from the average hourly earnings of production/nonsupervisory employees multiplied by the national average compensation-wage ratio (as described also in Stansbury and Summers (2019) and Bivens and Mishel (2015)). Production/nonsupervisory employees represent about 80-84 per cent of all employees in the US private sector.

Five-sector hourly compensation (Canada)

Unfortunately, for Canada an analogous measure of the compensation of production and nonsupervisory employees is only available from 1997 on (from the Statistics Canada Labour Force Survey). Instead, for our alternative proxy for typical workers' compensation, we use historical data from the Survey of Employment, Payrolls and Hours (and precursors) to construct a measure of the (weighted) average hourly wage of workers paid by the hour across five sectors of the economy – manufacturing; mining (including oil/gas); construction; laundries and cleaners; and hotels and restaurants – and then multiply this wage by the national average compensation-wage ratio.[11] We choose these five sectors since they are the only sectors with hourly wage data for hourly-paid workers for the entire period. They made up 34 per cent of total employment in 1961 and 24 per cent by 2011.[12] While this is not as comprehensive a measure as we would like, it is worth noting that the hourly-rated jobs in these sectors are generally among those with lowest barriers to entry for workers (e.g. credential requirements) and thus represent some of the sectors that are most likely to provide employment opportunities during times of growth.[13] The construction of this measure is detailed further in the appendix.

Note that in both the United States and Canada our typical compensation measures (median compensation, production/nonsupervisory compensation, and five-sector hourly compensation) are estimated from data on wage and salary income, and then adjusted to include an estimate of non-wage compensation using the national average ratio of total compensation to wages. This adjustment is important given the increase in non-wage compensation over recent decades (see for example Williams (2021) for Canada and Bivens and Mishel (2015) for the United States). Indeed, over the 1961-2019 period the ratio of total compensation to average wage/salary compensation increased from

11 Note the Survey of Employment, Payrolls and Hours and precursors only survey businesses with 20+ employees.

12 1961 and 2011 employment shares are from, respectively: the 1961 Census of Canada: Volume III, Part 2: Labour Force, Industries by Sex, Table 1; and 2011 National Household Survey Data tables: Industry - (NAICS) 2007 (425), Class of Worker (5), Age Groups (13B) and Sex (3) for the Employed Labour Force Aged 15 Years and Over, in Private Households of Canada

13 For example, Beach (2016) notes that "The rapid growth of Canada's resource and energy sectors out West and in Atlantic Canada and the strong housing construction boom... have provided distinctly strong employment opportunities for relatively lower-educated, largely male workers." The major omission is retail trade, for which there is no publicly available hourly-rated workers average wage data from the SEPH until 1983.

14 However, this adjustment may introduce three issues of concern to our empirical analysis. First, to the extent that the share of compensation which comes from non-wage benefits is higher at the top of the income distribution, our adjustment will overstate the level of total compensation for typical workers. Second, to the extent that the share of compensation which comes from non-wage benefits may have grown by more for workers

1.13 to 1.25 in Canada and from 1.05 to 1.16 in the United States.[14]

Before examining trends in productivity and compensation, it is helpful to summarize the price deflators used in this analysis to bring these series into real terms by adjusting for inflation. Productivity series are adjusted for changes in producer prices using the GDP deflator in Canada and the NDP deflator in the US.[15] There are then three choices of deflators to adjust the compensation series for inflation. Real 'product wages' use these producer price indexes to adjust for output price inflation. Real 'consumer wages' are calculated using either the Consumer Price Indices or a chain-type price index of goods and services purchased by consumers, called the Personal Consumption Expenditures (PCE) in the USA and the Household final consumption expenditure deflator (HCE) in Canada.[16]

Trends in Productivity and Pay, 1961-2019

Our primary period of analysis is 1961-2019, which is the longest period of analysis possible with Canadian data on average labour productivity and compensation.[17] Panel A in Chart 1 shows the productivity and compensation trends in Canada and the United States since 1961 (with growth rates by period in Table 1). There is a substantial divergence between productivity and typical compensation in both United States and Canada over the long-term period, regardless of choices about the price deflator. However, for both countries the size of the divergence between productivity and average compensation depends strongly on the price deflator used. In both countries, when productivity and compensation are deflated using the same price index – the NDP or GDP deflator respectively – there is a divergence between productivity and average compensation, which first emerges in the 1990s in Canada and around 2000 in the United States. In Canada average compensation outpaces labour productivity when deflated with either the CPI or HCE consumer price index.[18] For the United States the difference between series based on price deflators is even more stark: average compensation deflated by the PCE kept pace with labour productivity growth, while average compensation deflated by CPI grew about

at the higher end of the distribution than those in the middle, our adjustment will also overstate the growth in compensation for typical workers. Third, as the share of non-wage compensation in total compensation grows over time, this imputation becomes more substantial in its impact on our estimated median/typical compensation estimates – and to the extent we are measuring true typical compensation with error as a result of imputing non-wage compensation using the average figure, this measurement error may grow over time.

15 As Williams (2021) notes, Statistics Canada does not produce a price deflator for NDP at basic prices and the best alternative measure is to use the GDP deflator as the price index for NDP.

16 In general these price indexes differ because they are based on different underlying concepts, are constructed differently (a Fisher-type 'chain' price index vs. a Laspeyres-type price index), rely on different weights, cover some different items, and use different seasonal adjustment processes (McCully *et al.* 2007).

17 US data on average labour productivity and compensation extends back to 1948; we present results with this longer time period for the United States in the Appendix.

18 These results are largely in line with the analysis in Williams (2021), which finds over the same period for Canada that average real 'product wages' (deflated by output prices) tracked productivity measure while average real 'consumption wages' (deflated by CPI) outpaced both measures.

Chart 1: Productivity and Compensation, Canada and the United States

Panel A: Indexed 1961=100 for 1961-2019 period

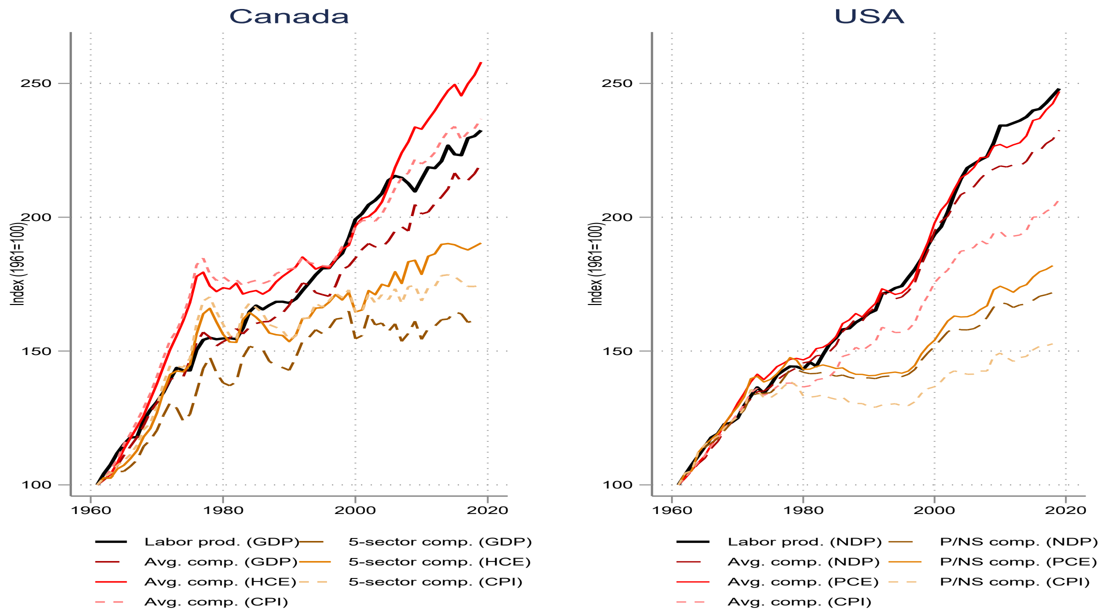

Panel B: Indexed 1976=100 for 1976-2019 period

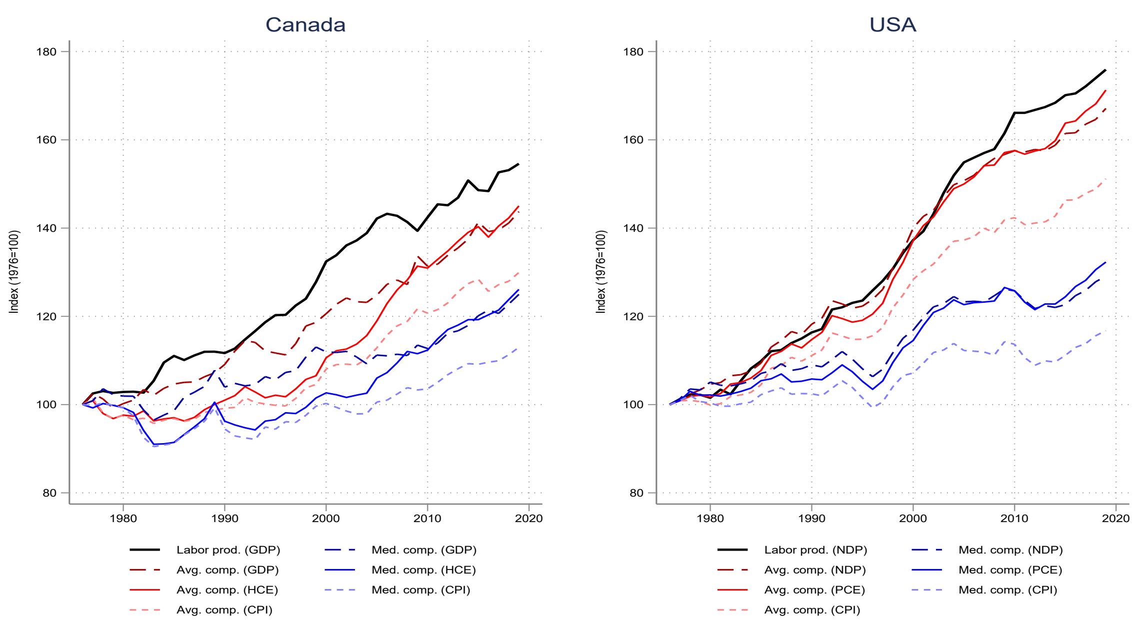

Note: “Avg. comp" refers to average compensation, “Med. comp." to median compensation, “P/NS comp." to the average hourly compensation of production and nonsupervisory workers (US only) and "five-sector comp." to our constructed measure of pay for “typical workers" in Canada, the average hourly compensation of hourly-rated workers in five sectors. Parentheses refer to the deflator used to adjust the series to real terms. Series are indexed to 1961=100 in Panel A and 1976=100 in Panel B.

41 percentage points less than productivity.

The starting point of the analysis is also important for our conclusion on the degree of divergence between productivity and average pay. Panel B of Chart 1 charts these productivity and pay trends over the 1976-2019 period for which there is (imperfect) median hourly compensation data for Canada.[19] Over this slightly shorter period, average compensation grew slower than labour productivity for both countries, although the degree of divergence again depends on the deflator used to adjust compensation.

Bivens and Mishel (2015) discuss the divergence between net labour productivity and median real compensation in terms of three separate 'wedges'. First is the divergence between average productivity and average compensation (measured with the product price deflator) as labour's share of income decreases. In the United States, the decline in the labour share has been the subject of a great deal of research, and is attributed variously to technological changes, globalization and labour offshoring, reductions in worker bargaining power, higher firm concentration, increased markups, and housing market dynamics (see overview in Stansbury and Summers, 2020). In Canada, Sharpe *et al.* (2008) argue for three key drivers of the decline in the labour share: the declining bargaining power of workers, rising commodity prices, and an increasing share of GDP going to capital consumption allowances; Williams (2021) notes that the net labour share (compensation divided by net domestic product, which excludes the impact of changes in the share of GDP going to capital consumption allowances) has declined very little over 1961-2019.

Another wedge is the divergence between average and typical compensation (as proxied by median compensation, or by production/nonsupervisory compensation in the US and five-sector hourly compensation in Canada). This wedge reflects rising labour income inequality (pre-tax and transfer) between the middle and top of the distribution. The increase in labour income inequality in both countries since the 1960s/1970s has been well documented in the United States (see overview in Stansbury and Summers, 2019) and in Canada (Green *et al.* 2017).[20] Similar to debates over the falling labour share, the increase has been attributed variously to purely technological explanations or institutional factors (such as declining unionization rates and increased trade with China) as well as slower growth in educational attainment in the face of skill-biased technological change (Goldin and Katz, 2009).

A final wedge relates to the 'terms of trade' divergence between consumer prices

19 Appendix Chart 1 illustrates the productivity-pay divergence in the United States for the full period for which we have data, 1948-2019.

20 The exact temporal dynamics of labour income inequality depend on the period chosen, particularly for Canada. The ratio between hourly mean and median compensation, or between hourly mean compensation and our "five-sector hourly" measure of typical compensation, has risen over most sub-periods from 1961-2019, as illustrated in Table 1 and Appendix Table 1. However, other measures of labour income inequality show a different picture. For example, the Gini coefficient of annual adjusted market income rose sharply in the 1980s and early 1990s and has gradually and incrementally declined a little over the following three decades (Table 11-10-0134-01 Statistics Canada).

Table 1: Productivity and Compensation Measures by Period (Per Cent Compound Annual Growth Rate

Panel A: Canada													
Measure	**Labour Productivity**	**Average compensation**			**Average compensation of employees**			**Median compensation**			**5-sector hourly compensation**		
Deflator	**GDP**	**GDP**	**HCE**	**CPI**	**GDP**	**HCE**	**CPI**	**GDP**	**HCE**	**CPI**	**GDP**	**HCE**	**CPI**
Sub-periods:													
1961-1976	2.80	2.90	3.90	4.10	2.50	3.60	3.70	-	-	-	2.00	3.00	3.20
1976-1997	1.00	0.60	0.20	0.10	0.70	0.20	0.10	0.40	-0.10	-0.20	0.90	0.40	0.30
1997-2019	1.10	1.10	1.50	1.10	1.00	1.40	1.00	0.70	1.20	0.70	0.00	0.50	0.10
Long periods:													
1961-2019	1.50	1.40	1.60	1.50	1.30	1.50	1.40	-	-	-	0.80	1.10	1.00
1976-2019	1.00	0.80	0.90	0.60	0.80	0.80	0.60	0.50	0.50	0.30	0.40	0.50	0.20
Panel B: United States													
Measure	**Labour Productivity**	**Average compensation**			**Average compensation of employees**			**Median compensation**			**5-sector hourly compensation**		
Deflator	**NDP**	**NDP**	**PCE**	**CPI**	**NDP**	**PCE**	**CPI**	**NDP**	**PCE**	**CPI**	**NDP**	**PCE**	**CPI**
Sub-periods:													
1961-1976	2.30	2.20	2.50	2.10	2.10	2.40	2.00	-	-	-	2.10	2.30	2.00
1976-1997	1.20	1.10	1.00	0.80	1.20	1.10	0.90	0.40	0.20	0.00	0.20	0.10	-0.10
1997-2019	1.50	1.30	1.50	1.20	1.30	1.50	1.10	0.80	1.00	0.70	0.90	1.10	0.70
Long periods:													
1961-2019	1.60	1.50	1.60	1.30	1.50	1.60	1.30	-	-	-	1.00	1.10	0.70
1976-2019	1.30	1.20	1.30	1.00	1.20	1.30	1.00	0.60	0.70	0.40	0.50	0.60	0.30

*Growth rates for production/ non-supervisory worker typical compensation calculated up to 2018 due to data availability
Note: Periods selected to correspond to periods used for regression analysis (which, in turn, correspond to the availability of different data series).

used to deflate compensation and the producer prices used to deflate output. Bivens and Mishel (2015) describe this wedge as "the faster price growth of things workers buy relative to the price of what they produce" and report that output prices have been outpaced by consumer price growth in the United States.

Using a gross output per hour worked measure of labour productivity, Uguccioni (2016) decomposes this gap between productivity and median pay in Canada from 1976 to 2014 and find that increased average-median earnings inequality accounts for 51 per cent of the gap while 30 per cent is accounted for by a decrease in labour's share of income and the final 19 per cent by a deterioration in labour's terms of trade. However, a longer-term analysis by Williams (2021) finds that labour's terms of trade improved by 0.3 per cent per year on average from 1961-2019 in Canada.

These productivity and pay trends in Canada and the United States are interesting to consider alongside each other. In both countries there has been concern over slow productivity growth during the 21st century amidst a global productivity slowdown. In the United States, the labour productivity slowdown has been attributed to a mix of mismeasurement, an industrial shift from high to low productivity sectors, slow TFP growth, population aging, and other factors (Moss *et al.* 2020). Canada's particularly poor productivity performance has confounded policymakers (Drummond, 2011).

The gap between Canadian and American productivity growth has also attracted attention. After Canada's stronger business sector labour productivity growth until the mid-1980s narrowed its levels gap with the United States to 5 percentage points, Canada experienced slower productivity growth for the following quarter-century, with its productivity level dropping to 72 per cent of the US level in 2010 (a 28 point gap), before closing to 26 points in 2016 (Sharpe and Tsang, 2018).[21]

Despite Canada's relatively poor productivity performance and worse average pay growth than in the United States since 1976, remarkably median pay growth has been about the same in these countries (as illustrated in Table 1 and Chart 1b). This is in line with a smaller recorded increase in income inequality over recent decades in Canada than the USA (Green *et al.* 2016).

To What Extent Have Productivity and Pay Delinked?

The analysis in the previous section illustrates that the pay of typical workers has grown more slowly than productivity in both the United States and Canada. Studying the degree of linkage or delinkage helps diagnose why this disconnection has happened, and what this might imply in terms of whether incremental increases (or decreases) in the rate of productivity growth in the future will benefit (or reduce) typical workers' pay. As outlined in the introduction, if a one percentage point growth rate in productivity is associated with a one percentage point growth rate in pay, all else equal, we consider the two series to be linked, and if it is associated with no change in pay all else equal, we consider the two series to be delinked.

Stansbury and Summers (2019) showed that despite substantially faster growth in productivity than in median pay in the United States, there was still linkage between the growth rates of the series. This suggests that incremental increases in US productivity growth still translated close to one-for-one into increases in typical pay, holding all else equal, which implies that the transmission mechanism from productivity to pay is functioning, but that at the same time other factors orthogonal (i.e. statistically independent) to productivity growth (like, perhaps, globalization or the declining bargaining power of workers) are suppressing pay and therefore responsible for the rising productivity-pay gap. This section updates this analysis for the United States and examines the degree of linkage or delinkage between productivity and pay in Canada.

Empirical Estimation

We use a simple linear model as in equation (1) below in order to understand the degree of productivity-pay linkage, following Stansbury and Summers (2019). They describe a spectrum of possible interpretations of the productivity-pay divergence: at one end, 'strong delinkage' where an incremental increase productivity growth does not systematically translate into any incremental growth in workers' compensation, holding all else equal; and at the

21 Gu and Willcox (2018) suggest this recent catch-up by Canada is due to its higher total factor productivity growth, larger capital deepening effect, and more gradual realization of the benefits of ICT investment.

other end of the spectrum 'strong linkage' where an incremental increase in productivity growth can be expected to translate one-for-one into an incremental increase in compensation growth. Strong linkage could be compatible with a situation where the levels of productivity and pay have diverged if factors orthogonal (i.e. statistically independent) to productivity growth have been putting downward pressure on worker compensation at the same time as rising productivity has been putting upward pressure on pay.

Under the 'strong linkage' view $\beta = 1$ and under the 'strong delinkage' view $\beta = 0$. A value of β between 0 and 1 suggests some point on the spectrum between productivity-pay linkage and delinkage. While this is a partial model and other factors may affect compensation growth in addition to productivity, these other factors will not affect estimation of β as long as they are orthogonal (i.e. statistically independent) to productivity growth. The parameter α is a constant.

$$\begin{aligned} &Compensation\ growth_t \\ &\quad = \alpha + \beta productivity\ growth_t \end{aligned} \tag{1}$$

We estimate this model by regressing the year-on-year change in hourly compensation on the change in labour productivity, controlling for the unemployment rate (see below). As described above, we use four separate measures of compensation in our regressions: two measures of average compensation (average compensation for all workers, average compensation for employees only), and two measures of typical compensation for each country (median compensation and production/nonsupervisory compensation for the United States; median compensation and our constructed measure of compensation for hourly-rated workers in five sectors in Canada). Since we run the same analyses on all measures, for brevity all four measures are referred to as 'compensation' below. In our baseline specifications, all compensation measures are deflated with the chain-linked consumer price deflator (PCE for United States, HCE for Canada).

In our baseline regression, in equation (2) below, we regress the three-year moving average of the change in log compensation on the three-year moving average of the change in log labour productivity and the current and lagged three-year moving average of the unemployment rate.[22] We use moving averages rather than annual changes in our baseline specification to account for a potentially longer time horizon for the productivity-pay relationship, for example because firms change pay and benefits infrequently or because it takes firms and workers some time to discern that increased output is due to higher labour productivity. In the appendix, we also show results from a similar regression with five-

22 We account for autocorrelation introduced by the moving average specification by using Newey-West heteroskedasticity and autocorrelation robust standard errors, with a lag length of 6 years.

year moving averages.

$$
\begin{aligned}
&\frac{1}{3}\sum_{0}^{2}\Delta log\ comp_{t-i} \\
&= \alpha + \beta\frac{1}{3}\sum_{0}^{2}\Delta log\ prod_{t-i} \\
&+ \gamma\frac{1}{3}\sum_{0}^{2} unemp_{t-i} \\
&+ \delta\frac{1}{3}\sum_{0}^{2} unemp_{t-i-1} + \epsilon_t
\end{aligned} \tag{2}
$$

In all specifications we control for the unemployment rate, for two reasons. First, as an indicator of labour market slack/tightness, the level of unemployment is likely to affect bargaining dynamics. In the context of a slack labour market with a high unemployment rate, employers would be able to raise compensation by less than they otherwise would have for a given productivity growth rate, as more unemployed workers are searching for jobs at that time. Second, as a proxy measure of general labour market conditions, unemployment is likely to reflect broader cyclical economic fluctuations that could impact compensation in the short term. For example, higher unemployment may reflect a downturn, which could mean lower pay rises for a given rate of productivity growth. If unemployment is also related to changes in productivity growth – for example, if the least productive workers are likely to be laid off first in a downturn– then excluding unemployment would bias the results. By controlling for the current and one-year lagged moving average of the unemployment rate we allow for both the level and the change in unemployment to affect compensation growth.

Regression Results

We run our baseline regression in equation (2) above for both the United States and Canada, over various time periods. Table 2 shows our regression results for the United States and Table 3 for Canada. Chart 2 charts the coefficients estimated for the compensation-pay relationship from these regressions (i.e.in equation (2) above), using the PCE or HCE chain-linked consumer price deflator. Each dot and line in the chart shows the point estimate and 95 per cent confidence interval (respectively) for this coefficient in the three-year moving average regression in equation (2). In addition, Chart 3 shows the cyclically-adjusted by residualizing these variables on the unemployment rate, the lagged unemployment rate (both three-year moving averages), and a constant. These results shed light on the degree of productivity-pay delinkage in the United States and Canada (discussed in turn below).[23]

United States

For the United States, these regression results suggest that average compensation is for the most part still strongly linked to net productivity, as found in Stansbury and Summers (2019). For both average compensation measures (i.e. including and excluding the self-employed), regardless of the price deflator used, the coefficient on

23 In the appendix, we present results with five-year rather than three-year moving averages.

Chart 2: Coefficient on Compensation-Productivity Regressions

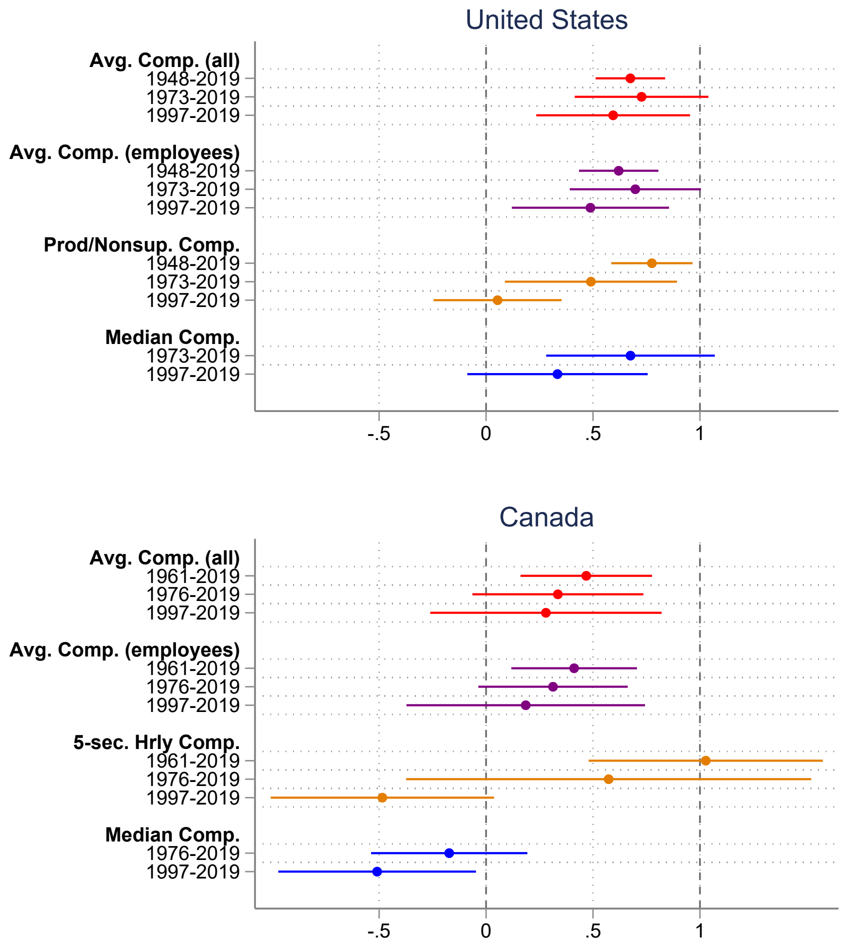

Note: Each dot and line shows the point estimate and 95% confidence interval for the coefficient on productivity, from an annual regression of the change in log of compensation on the change in log of labour productivity, the unemployment rate, and the lagged unemployment rate (with all variables taken as 3-year moving averages). Variables in US regressions are hourly net labour productivity for the entire economy (real NDP/hours worked) and hourly real average compensation for all employed persons including the self-employed (red), hourly real average compensation for employees (purple), hourly real average compensation for production and nonsupervisory employees (orange), and hourly real median compensation for all employees (blue). Compensation is inflation-adjusted using the PCE price index. Variables in Canada regressions are hourly net labour productivity for the entire economy (real NDP/hours worked) and hourly real compensation for all employed persons (red), hourly real average compensation for employees (purple), hourly real average compensation for hourly-paid workers in five large sectors (orange), and hourly real median compensation for all employees (blue). Compensation is inflation-adjusted using the HCE price index. All regressions have Newey-West heteroskedasticity and autocorrelation consistent standard errors with a lag length of six.

Table 2: Coefficients From Regressions of Average Compensation on Productivity, United States, 3 Year Moving Averages

Period	1948-2019			1973-2019			1997-2019		
Specification	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Comp. deflator	NDP	PCE	CPI	NDP	PCE	CPI	NDP	PCE	CPI
Dep. Var.: Average Hourly Compensation, all employed persons, from national accounts, 3yma									
Labour Productivity	0.60***	0.67***	0.79***	0.57***	0.73***	0.73***	0.58**	0.59***	0.53***
	(0.07)	(0.08)	(0.09)	(0.16)	(0.15)	(0.14)	(0.20)	(0.17)	(0.12)
Dep. Var.: Average Hourly Compensation, employees, from national accounts, 3yma									
Labour Productivity	0.55***	0.62***	0.74***	0.54***	0.70***	0.70***	0.47**	0.49**	0.42***
	(0.06)	(0.09)	(0.11)	(0.16)	(0.15)	(0.14)	(0.21)	(0.17)	(0.12)
Dep. Var.: Compensation of Production and Nonsupervisory Workers, hourly, 3yma									
Labour productivity	0.70***	0.78***	0.89***	0.33**	0.49**	0.49**	0.03	0.05	-0.01
	(0.10)	(0.10)	(0.11)	(0.14)	(0.20)	(0.21)	(0.15)	(0.14)	(0.13)
Dep. Var.: Median Hourly Compensation, 3yma									
Labour Productivity	-	-	-	0.52***	0.68***	0.68***	0.32*	0.33	0.27
	-	-	-	(0.18)	(0.20)	(0.18)	(0.17)	(0.20)	(0.23)
Obs.	69	69	69	44	44	44	20	20	20

Note: Each cell contains the coefficient estimate on hourly labour productivity (change in log, 3y trailing moving average) from a regression of the change in log compensation (3yma) on the change in log net hourly labour productivity (3yma), controlling for the current and 1-year lagged unemployment rate (3yma) and a constant, using annual data. Newey-West (HAC) standard errors (6-year lag) are listed below each coefficient estimate in parentheses, **** p<0.01, ** p<0.05, * p<0.1.*

Table 3: Coefficients From Regressions of Average Compensation on Productivity, Canada, 3 Year Moving Averages

Period	1961-2019			1976-2019			1997-2019		
Specification	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Comp. deflator	GDP	HCE	CPI	GDP	HCE	CPI	GDP	HCE	CPI
Dep. Var.: Average Hourly Compensation, all employed persons, from national accounts, change in log, 3yma									
Labour productivity	0.50***	0.47***	0.69***	0.27	0.34*	0.39**	0.32***	0.28	0.24
(ch. log, 3yma)	(0.09)	(0.15)	(0.14)	(0.19)	(0.20)	(0.19)	(0.08)	(0.26)	(0.21)
Dep. Var.: Average Hourly Compensation, employees, from national accounts, change in log, 3yma									
Labour productivity	0.44***	0.41***	0.64***	0.24	0.31*	0.37**	0.23**	0.19	0.15
(ch. log, 3yma)	(0.09)	(0.15)	(0.13)	(0.17)	(0.17)	(0.17)	(0.08)	(0.26)	(0.21)
Dep. Var.: Compensation of Hourly-Paid Workers in Five Sectors, change in log, 3yma									
Labour productivity	1.06***	1.03***	1.25***	0.50	0.57	0.63	-0.44*	-0.48*	-0.53**
(ch. log, 3yma)	(0.22)	(0.27)	(0.27)	(0.40)	(0.47)	(0.49)	(0.24)	(0.25)	(0.22)
Dep. Var.: Median Hourly Compensation, employees, survey-based, change in log, 3yma									
Labour productivity	-	-	-	-0.24*	(0.17)	(0.11)	-0.47***	-0.51**	-0.55***
(ch. log, 3yma)	-	-	-	(0.14)	(0.18)	(0.19)	(0.16)	(0.22)	(0.17)
Obs.	56	56	56	41	41	41	20	20	20

Note: Each cell contains the coefficient estimate on hourly labour productivity (change in log, 3y trailing moving average) from a regression of the change in log compensation (3yma) on the change in log net hourly labour productivity (3yma), controlling for the current and 1-year lagged unemployment rate (3yma) and a constant, using annual data. Newey-West (HAC) standard errors (6-year lag) are listed below each coefficient estimate in parentheses, **** p<0.01, ** p<0.05, * p<0.1.* Observations refer to 3-year trailing moving average periods for annual change in log data. For example, the 56 observations in column (1) are for each year from 1964-2019, where the data point for 1964 incorporates data from 1961-2, 1962-3, and 1063-4.

Chart 3: Cyclically-adjusted Productivity and Compensation Growth, United States and Canada

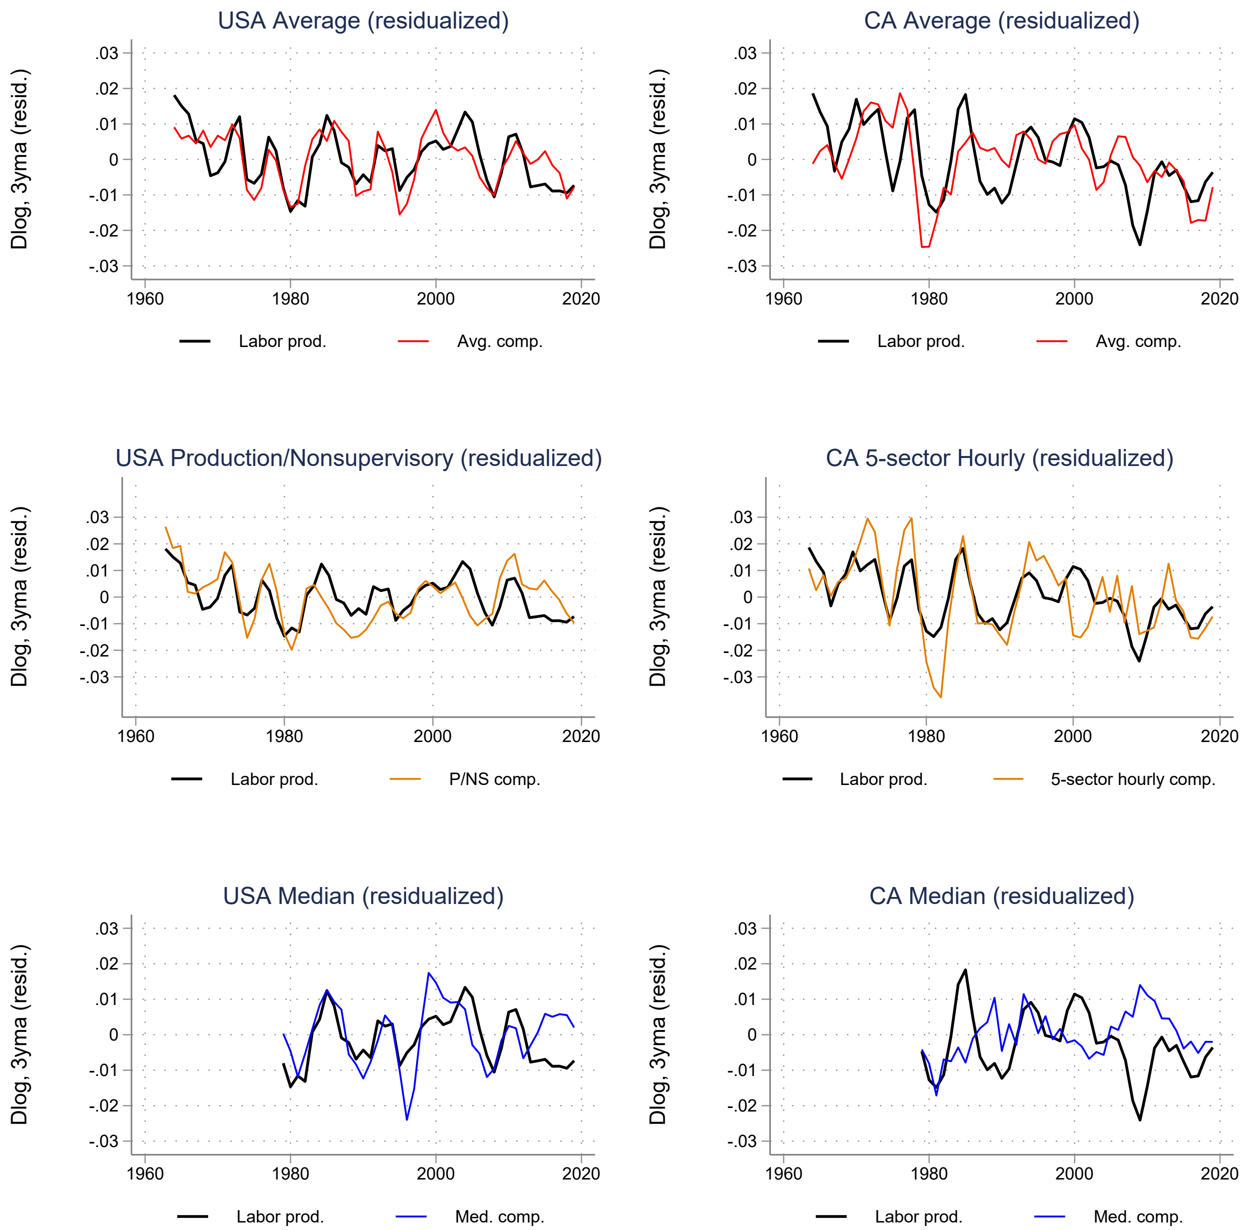

Note: Each chart shows the 3-year trailing moving average of the change in log of labour productivity and of compensation, cyclically-adjusted by residualizing these variables on the unemployment rate, the lagged unemployment rate (both three-year moving averages), and a constant. The first row of charts shows average compensation for all employed workers, the second row shows our two measures of compensation for typical workers (production/nonsupervisory workers in the US, and hourly-paid workers in five large sectors in Canada), and the third row shows median hourly compensation. All US compensation measures are deflated using the PCE; all Canadian compensation measures are deflated using the HCE price index.

productivity is close to 1. The confidence intervals illustrate that these coefficients are strongly significantly different from 0 and, for the 1973-2019 regressions with the largest estimated coefficients, not significantly different from 1. Interestingly, the coefficients are smallest for the more recent 1997-2019 period, suggesting that the linkage may be getting weaker over time (though there is still a somewhat strong link and the confidence interval of the point estimate does not contain zero).

This weaker relationship in recent decades in the United States is even more apparent when looking at the link between productivity and either the average compensation of production/nonsupervisory workers, or median compensation. With production/nonsupervisory compensation as the dependent variable, the coefficient on productivity is very large for the sample over the entire postwar period but very small (and not significantly different from zero) for the subsample over the last two decades. The coefficient on productivity in the regression with median compensation is also very small for the 1997-2019 period, although it is larger for the longer-term 1973-2019 sample for which median compensation data is available. Breakpoint tests in years 1997, 2000, and 2008 (Appendix Table 5) fail to reject the null hypothesis of no structural break for *average* compensation, but suggest structural breaks in the productivity-pay relationship for both of our typical worker measures: production/nonsupervisory compensation and median compensation. It is possible that this reflects a decline in the linkage between productivity and typical workers' pay in recent decades. It is also possible that, to the degree that there may be a growing bias in our estimates of median or production/nonsupervisory compensation over time (due to the growing influence of the non-wage compensation imputation), this may lead to attenuation of the estimated regression coefficients over time.

Chart 3 illustrates the temporal dynamics of the relationship, showing that in the US compensation growth generally co-moved with productivity growth until the 2000s. In the early 2000s, cyclically-adjusted productivity grew faster than compensation, while the decline in productivity growth in the 2010s was not matched by as substantial a decline in real compensation growth. This pattern held for both average compensation and the compensation of production and nonsupervisory workers. For median compensation, the picture is a little more nuanced: the spikes in the blue line from 1995-2005 illustrate that median compensation grew more slowly than would have been predicted by the level of unemployment in the mid 1990s and more quickly in the tight labour market of the late 1990s.

Canada

The regression results for Canada are shown in Table 3. For Canada, average compensation (all measures deflated using the HCE price index) is positively and significantly linked to productivity over the longer-term 1961-2019 sample (whether or not the self-employed are included). The point estimates on these regressions are somewhat smaller than those for the US – between 0.4 and 0.5, with confidence intervals ruling out both 0 (strong delinkage) and 1 (strong linkage). Estimated coefficients are smaller, and not significantly different from zero, for the more recent 1976-2019 period and especially the 1997-2019 period (although structural break tests fail to reject the null hypothesis that there is no change in the productivity-average compensation relationship in 1997, 2000, or 2008 – see Appendix Table 5). Coefficients are relatively similar regardless of which price deflator is used to deflate compensation, as can be seen by comparing across columns (1)-(3) in Table 3.

Results for typical compensation are a little harder to interpret. One of our typical compensation measures, average wages for hourly-rated workers in five sectors, is strongly linked to productivity over the whole period with the coefficient very close to one and strongly significantly different from zero. Yet breakpoint tests indicate a structural break for this five-sector hourly compensation measure in Canada in 1997 (Appendix Table 5), and there is no evidence of this compensation measure being linked to productivity in the past two decades: the point estimate, in fact, is negative, although the standard errors are so large that a coefficient of zero is also within the estimated confidence interval.[24]

Chart 3 illustrates the dynamics behind this: in the early 2000s, cyclically-adjusted productivity growth rose (and so did cyclically-adjusted average compensation growth), but cyclically-adjusted compensation growth in our five sectors of interest moved in the opposite direction, spiking sharply downwards. In contrast, in the Great Recession era there was a large decline in cyclically-adjusted productivity that was not matched by as large of a decline in cyclically-adjusted real compensation for typical workers in these five sectors.

Finally, there is no evidence of linkage for our other typical compensation measure – median compensation. Estimated coefficients are negative, with large standard errors, for both the 1976-2019 period and the more recent 1997-2019 period. Mechanically, this is driven by the 1997-2019 period featuring very few meaningful fluctuations in the growth rate of median compensation, with an exception during the Great Recession period, where the cyclically-adjusted rate of productivity growth fell but the cyclically-adjusted rate of median compensation growth rose.

We treat these results for median compensation with caution, however, given the large potential measurement error for this measure of median hourly compensation (discussed above). To the extent that we are measuring total hours worked per year with noise, and to the extent that year-to-year fluctuations in true median hourly compensation are relatively small, the noise in the hours worked measure could swamp the "signal" of true median hourly compensation. In addition, to the extent that our imputation of non-wage compensation is becoming more significant over time, and to the extent this may introduce measure-

24 Comparisons across the specifications with different deflators for compensation illustrate that a differential growth rate of different price indices is not the primary driving force behind this result: whether compensation is deflated with the GDP deflator, CPI, or HCE, estimated coefficients for the most recent period are negative, with large standard errors.

ment error and/or upward bias in our estimate of total compensation, the bias introduced by this imputation may also be increasing over time.

Exploring US-Canada Differences

In the previous section, we found that the relatively strong linkage found between productivity and pay in the United States in Stansbury and Summers (2019) held with updated data. This affirms the conclusion of that previous analysis that factors orthogonal (i.e. statistically independent) to productivity have been acting to suppress average and typical compensation in the United States even as productivity growth has been acting to raise it.[25]

We have found less conclusive evidence for strong linkage between year-to-year fluctuations in productivity growth and compensation growth in Canada. It is possible that the relatively small size of these regression coefficients on productivity growth, and large standard errors, may be related to explanations specific to the Canadian context: there may, for example, be a weaker productivity-pay link in Canada than the United States because the former is a smaller, more open economy. The Canadian economy is less than one-eleventh the size of the USA economy and has more than twice the share of trade as a percentage of GDP.[26] As a result, it is reasonable that international factors are likely to be a relatively more important determinant of worker compensation in Canada than in the United States.[27]

To evaluate this argument, we analyze the productivity-pay relationship at the level of other similar small open economies: US regions. The economies of some US regions may be better analogs for the Canadian economy than the economy of the United States as a whole. Comparing Canada to the eight regions of the continental US as defined by the Bureau of Economic Analysis, Canada sits somewhere in the middle in terms of GDP and population.[28] Canada's population in 2019 was 37.6 million and its GDP in 2019 (in current USD) was $1.74 trillion: this made it larger in both population and GDP terms than the Rocky Mountain, New England, or Plains regions, but smaller in both population and GDP terms than the South West, Great Lakes, Mideast, Far West, or South East regions of the United States.

It should be noted, of course, that the US regions are not perfect comparators:

25 While the linkage between productivity and pay in the United States appears weaker for the first two decades of the 21st century than for the second half of the 20th century, this appears to be a result of extremely low cyclically-adjusted productivity growth in the post-Great Recession era not being matched by as large a decline in the cyclically-adjusted rate of compensation growth.

26 Economy size in terms of PPP-adjusted GDP. Trade is sum of exports and imports of goods and services as a share of GDP. Data from World Development Indicators database, World Bank. Over 2010-2019, Canada's trade share of GDP ranged between 60 and 67 percent; the US' trade share of GDP ranged between 26 and 30 percent.

27 To the extent that these international factors are correlated with both productivity and compensation growth in Canada, these may be expected to reduce the strength of the linkage we can estimate between productivity growth and compensation growth in Canadian data.

28 The eight regions and the states which comprise them are listed in Appendix Table 8.

Chart 4: Coefficient on Compensation-Productivity Regressions: Comparison to US Regions

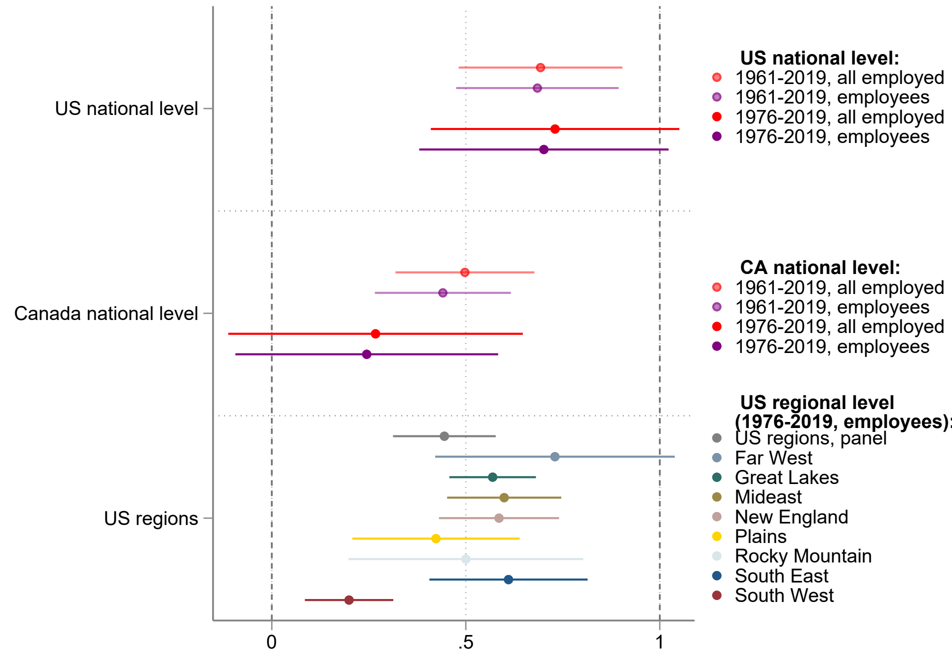

Note: Each dot and line shows the point estimate and 95% confidence interval for the coefficient on productivity, from an annual regression of the change in log of compensation on the change in log of labour productivity, the level of unemployment, and the lagged level of unemployment (with all variables taken as 3-year moving averages). Variables in US national level regressions are hourly net labour productivity for the entire economy (real NDP/hours worked) and hourly real compensation for all employed persons including the self-employed (shown in red), and for employees only (shown in purple). Compensation is inflation-adjusted using the PCE price index. Variables in Canada national level regressions are hourly net labour productivity for the entire economy (feal NDP/hours worked) and hourly real compensation for all employed persons including the self-employed (shown in red), and for employees only (shown in purple). Compensation is inflation-adjusted using the HCE price index. Variables in US regional level regressions are labour productivity per worker (real NDP/employed persons) and annual real compensation for employees (employee compensation/employees), inflation-adjusted using the PCE price index. The "US regions, panel" regression is a panel regression which includes region and year fixed effects and has robust standard errors clustered at region level. All other regressions are time series regressions with Newey-West heteroskedasticity and autocorrelation consistent standard errors (using a six year lag).

their economies are more open (to each other) than Canada's economy is to the rest of the world, since trade frictions between US regions are lower than between Canada and the United States or other countries, and migration is substantially easier between US regions than across an international border into or out of Canada. If migration across US regions fully arbitrages wage differences between US regions, one would have no reason to expect a productivity-pay linkage at the US region level. This comparison therefore relies to some extent on the assumption that US regions represent distinct labour markets, with less-than-perfect arbitrage of wages through migration between regions.

To analyze the productivity-pay relationship for the eight US regions, we construct worker-level measures of annual

29 To estimate annual productivity per worker for each region, we start with region-level GDP, adjust it for

labour productivity and annual employee compensation from the BEA Regional Economic Accounts.[29] We use these to run our baseline regressions at the region level, regressing the change in log of compensation on the change in log of productivity (3-year moving average), the current and lagged regional unemployment rate (3-year moving average), and a constant. Coefficients from these regressions are shown in the bottom panel of Chart 4, where the first coefficient estimate shows the estimate from a panel regression across all regions (including region and year fixed effects), and the other coefficient estimates show the estimates from separate time series regressions for each region. Across most regions, the point estimates are close to and not significantly different from 0.5, but are significantly different from 0 for all US regions and significantly different from 1 for seven of the eight US regions.

Notably, the coefficients from these region-level regressions are much smaller than those for the US national level regressions for average compensation. Indeed, the magnitude of the coefficients for US regions are quite close to those for average compensation for Canada – especially for the US regions most geographically proximate and similar to Canada: the Plains, Great Lakes, and Rocky Mountain regions. This evidence is consistent with the idea that the more open an economy is, the less tightly linked are year-to-year changes in domestic productivity and domestic compensation.

A further point of comparison is whether labour productivity differentials between the United States and Canada translate into differences in pay. To test this, we regressed the difference in the change in log compensation on the difference in the change in log productivity (United States minus Canada), controlling for the difference in unemployment. These regression results (in Appendix Table 4) suggest that a one percentage point faster productivity growth rate in the United States than in Canada translates into a 0.2-0.5 percentage point faster average compensation growth rate, holding all else constant. This fits with the finding noted above that the US experienced faster productivity growth and average compensation growth than Canada between 1976 and 2019. On the other hand, we find no evidence that a one percentage point faster productivity growth rate in the US than in Canada translates into a faster median compensation growth rate in the former. This is consistent with the fact that despite much faster productivity growth in the US than in Canada, real median compensation grew at about the same rate over the period.

Why might the Canadian economy, as a more open economy, have a smaller measured linkage between productivity and pay than the United States? One notable feature of the Canadian economy

---

depreciation by the national GDP-NDP ratio, deflate it with the national NDP price index, and then divide by the number of employed persons in the state (employees and self-employed). To estimate annual real compensation per worker for each region, we divide total employee compensation by the number of employees (from the BEA Regional Economic Accounts), and deflate compensation by the national PCE price index. Note that our compensation measure is for employees, not all employed persons, because of difficulties in obtaining estimates of compensation per worker for the self-employed at the region level.

is its dependence on commodities which are subject to international price fluctuations. For example, international commodity price 'supercycles' are an important influence on Canada's terms of trade, exchange rate, employment, income and inflation Buyuksahin *et al.* (2016), and Green *et al.* (2019) find that the most recent Canadian resource boom in the 2000s substantially boosted real average wages in Canada broadly — that is, not only in the resource sector nor only in resource regions—which may have also hurt firms in non-resource regions by raising wage costs without the increase in demand enjoyed by firms in resource regions. It is possible that the influence of international commodity price fluctuations may explain the weakness of the measured productivity-pay relationship – though, it does not provide a completely obvious account of why the productivity-wage relationship would be affected if both series were deflated by a (correctly measured) product price index. To us this appears to be an avenue worth of further exploration.

Another avenue worthy of further exploration would be the degree to which productivity fluctuations which occur for different reasons may be expected to translate into pay growth. For example, a larger share of the fluctuations in measured productivity growth rates in Canada are driven by natural resource prices, as compared to changes in technology or other drivers of economy-wide productivity. It is possible that the transmission mechanism from measured productivity gains as a result of export price increases differs as compared to the transmission mechanism from productivity gains arising from, for example, new technologies.[30]

Concluding Remarks

The analysis above sheds light on the productivity-pay relationship in the United States and Canada. First, we studied the *divergence* in levels between productivity and pay. In both countries the pay of typical workers has diverged substantially in real terms from average labour productivity. However, these divergences are attributed to slightly different forces. While the labour share of income has declined in both countries, and labour income inequality has also risen, the latter has been much more pronounced in the United States. As a result, despite much faster growth in both labour productivity and average compensation in the United States than in Canada – particularly since 1976 – the growth in real median compensation has been about the same in both countries.

Second, we studied whether this divergence has come alongside a *delinkage* in the growth rates of productivity and pay.

30 The timing of the transmission process may also differ according to the source of productivity growth, and the extent to which it is perceived as transitory or permanent. In a study of the US oil and gas field services industry, Kline (2008) finds that wage increases lag price increases in the oil extraction sector. A further possible explanation for the smaller degree of linkage between productivity and pay in Canada as compared to the US might have been a smaller degree of meaningful high frequency fluctuations in Canadian productivity growth compared to the United States. If this were the case, we would expect attenuated coefficients on the productivity-pay relationship due to classical measurement error. We do not, however, find evidence that productivity fluctuations are smaller in Canada than in the United States: the variance of the change in log productivity is the same or higher in Canada over the 1961-2019 period.

We find evidence for some linkage between productivity and pay in both countries: a one percentage point higher productivity growth rate is associated with significant increases in the rate of compensation growth in both countries. However, our evidence suggests that this linkage between productivity and *average* pay growth may be stronger in the United States than in Canada.

In the United States there is also evidence for linkage between productivity and typical workers' pay, although the magnitude appears to attenuate somewhat over time. In Canada, the evidence on typical workers' pay is more mixed and depends on both the measure chosen and the time period: there is evidence of strong linkage of productivity with our measure of the compensation of hourly-rated workers in five sectors (although estimates in more recent periods become too noisy to rule out either strong linkage (one-for-one translation from productivity to pay) or strong delinkage (no translation of productivity to pay)), and there is no evidence of linkage of productivity with our measure of median compensation. Given the substantial concerns about measurement error in our measures of typical compensation in Canada, particularly median compensation, these estimates should be treated with caution however.

We explore possible explanations for the difference in the degree of estimated linkage between productivity and average compensation in the United States and Canada. In particular we emphasize the possibility that since Canada is a smaller, more internationally open economy than the United States, there may be a greater role for international factors (such as global commodity prices) that affect both productivity and compensation growth to reduce the strength of the linkage between domestic productivity growth and domestic compensation growth in Canada. This is supported by a comparison of our productivity-pay regression results from Canada with estimates of regressions of average compensation on productivity in US regions, which are similar in GDP and population to Canada, and could also be considered small open economies. These estimates suggest that linkage in US regions is substantially lower than for the United States as a nation, and is similar in magnitude to the estimates for Canada.

The argument that the productivity-pay linkage is weaker in small open economies merits further study. Combined analysis of multiple countries may shed light on the importance for the size and trade-dependence of a country on its degree of productivity-pay linkage. Another focus of further research should be the degree to which productivity fluctuations which occur for different reasons may be expected to translate into pay growth. For example, productivity growth arising from a new labour-substituting technology may be expected to have substantially different implications for the pay of typical workers, than would productivity growth arising from different types of technological change – or, indeed, from productivity growth arising from a change in export demand or export prices. Understanding the role of these factors in the productivity-pay relationship will enable policymakers to respond appropriately. In particular, as policymakers in Canada continue to express concern over its slow productivity growth rate, it is important to understand the extent to which slower productivity growth may impact workers at different points in the income distribution.

Further research will illuminate the extent to which increases in productivity growth raise living standards across the board and which sources of increased growth impact most strongly on different parts of the income distribution. We think it is unlikely that the view that growth augments the incomes of most workers will be overturned. In addition to its direct impact on middle class incomes we expect that by augmenting government revenue collections increased growth will lead to at least some extent to augmented public spending. Finally, we emphasize that to suggest that productivity growth influences middle class living standards is not of course to imply that fluctuations in growth are primary determinants of middle class living standards – other distributional factors are also important.

References

Baldwin, J. R. and N. Rylska (2007) "Revisions to the Quarterly Labour Productivity Estimates," *Canadian Productivity Review*, Statistics Canada.

BEA (2018) "NIPA Handbook: Concepts and Methods of the U.S. National Income and Product Accounts," U.S. Bureau of Economic Analysis (BEA), https://www.bea.gov/resources/methodologies/nipa-handbook

Beach, C. M (2016) "Changing Income Inequality: A Distributional Paradigm for Canada," *Canadian Journal of Economics,* Vol. 49, No. 4, pp. 1229–1292.

Bivens, J. and L. Mishel (2015) "Understanding the Historic Divergence between Productivity and a Typical Worker's Pay: Why It Matters and Why It's Real," Economic Policy Institute.

Bivens, J. and L. Mishel (2021) "The Productivity-Median Compensation Gap in the United States: The Contribution of Increased Wage Inequality and the Role of Policy Choices," *International Productivity Monitor,* No. 41, Fall, pp.

Buyuksahin, B., Mo, K. and K. Zmitrowicz (2016) "Commodity Price Supercycles: What are they and What lies ahead?" *Bank of Canada Review,* Autumn.

Drummond, D (2011) "Confessions of a Serial Productivity Researcher," *International Productivity Monitor,* No. 22, Fall, pp. 3-10.

Goldin, C. and L. F. Katz (2009) *The Race between Education and Technology,* Harvard University Press.

Green, D. A., W. Craig Riddell, and F. St-Hilaire (2016) "Income Inequality in Canada: Driving Forces, Outcomes and Policy," *Income Inequality: The Canadian Story,* (Montreal: Institute for Research on Public Policy), pp. 1-73.

Green, D. A., Morissette, R., Sand, B.M. and I. Snoddy (2019) "Economy-wide Spillovers From Booms: Long-distance Commuting and the Spread of Wage Effects," *Journal of Labor Economics,* Vol. 37, No. S2, pp. S643-S687.

Gu, W. and M. Willox (2018) "Productivity Growth in Canada and the United States: Recent Industry Trends and Potential Explanations," *International Productivity Monitor,* No. 35, pp. 73-94.

Kline, P (2008) "Understanding Sectoral Labor Market Dynamics: An Equilibrium Analysis of the Oil and Gas Field Services Industry," Cowles Foundation Discussion Papers 1645, Cowles Foundation for Research in Economics, Yale University.

Lawrence, R. z. (2015) "Recent declines in labor's share in US income: A preliminary neoclassical account," National Bureau of Economic Research, No. w21296.

McCully, C.P., Moyer, B. C. and K. J. Stewart (2007) "Comparing the Consumer Price Index and the Personal Consumption Expenditures Price Index," *Survey of Current Business,* November, pp. 26–33.

Moss, E., Nunn, R. and J. Shambaugh (2020) "The Slowdown in Productivity Growth and Policies That Can Restore It," The Hamilton Project, the Brookings Institution, https://www.brookings.edu/wp-content/uploads/2020/06/Productivity_Framing_LO_6.16_FINAL.pdf.

Sharpe, A., Arsenault, J-F. and P. Harrison (2008) "Why Have Real Wages Lagged Labour Productivity Growth in Canada?" *International Productivity Monitor,* No. 17, Fall, pp. 16-27.

Sharpe, A. and J. Tsang (2018) "The Stylized Facts about Slower Productivity Growth in Canada," *International Productivity Monitor,* No. 35, Fall, PP. 52-72.

Sharpe, A. and J. Ashwell (2021) "The Evolution of the Productivity Median Wage Gap in Canada, 1976-2019," *International Productivity Monitor*, No. 41, Fall.

Stansbury, A. and L. Summers (2019) “Productivity and Pay: Is the Link Broken?” in *Facing Up To Low Productivity Growth,* Peterson Institute for International Economics.

Stansbury, A. and L. Summers (2020) “The Declining Worker Power Hypothesis,” *Brookings Papers on Economic Activity,* Spring.

Uguccioni, J (2016) “Explaining the Gap Between Productivity and Median Wage Growth in Canada, 1976-2014,” *International Productivity Monitor,* No. 31, Fall, pp. 37-56.

Williams, David M. (2021) "Pay and Productivity in Canada: Growing Together, Only Slower than Ever, " *International Productivity Monitor,* No. 40, Spring, pp. 3-26.

Have Productivity and Pay Decoupled in the UK?

Andreas Teichgräber
London School of Economics
John Van Reenen
London School of Economics and Massachusetts Institute of Technology[1]

Abstract

In the long-run at the macro level, the growth in real pay of workers tends to follow that of labour productivity. In recent years, however, there have been concerns that this relationship has broken down and that pay has become "decoupled" from productivity, growing much more slowly, and leading to a fall in the labour share. This has been a well-documented phenomenon in the United States (US) since the early 1980s. By contrast, we show that in the United Kingdom (UK), employee mean hourly compensation has grown at the same rate as labour productivity between 1981 and 2019. However, there has been a divergence between median employee hourly wage growth and productivity growth of about 25 percentage points. About three-fifths of this "overall decoupling" is due to increasing inequality (mean wages growing faster than median wages) and one-third is due to the increased non-wage compensation costs, in particular employer pension contributions. However, this analysis relates to employee compensation. The average self-employed worker has seen their income grow by only 50 per cent, compared to 80 per cent for the average employee. Using micro-data, we show that this gap can essentially be explained by (i) the growth in the numbers of "solo self- employed" (who have relatively low incomes), and (ii) a much greater fall in hours worked by the self-employed than for the employed. Finally, if we "correct" the labour share for self- employment and non-wage labour costs, the UK labour share has fallen by about 3.5 percentage points over the last four decades.

---

1 Andreas Teichgräber is a Research Assistant at the Centre for Economic Performance at the London School of Economics. John Van Reenen is Ronald Coase School Professor at the London School of Economics and the Gordon Y. Billiard Professor of Management and Economics at the Massachusetts Institute of Technology. We thank the Economic and Social Research Council for their financial support through the Programme On Innovation and Diffusion (POID). Useful comments have come from Richard Blundell, Angus Deaton, Steve Machin, Ralf Martin, Andrew Sharpe, Bart Van Ark, three anonymous referees, and seminar participants at the LSE and IFS. Many people have helped tremendously with the data, including Brian Bell, Alex Bryson, Rui Costa, Richard Dickens, Giulia Giupponi, Helen Miller, Larry Mishel, João Pessoa, Felix Ritchie, Matt Whittaker, and several people from the Department for Work and Pensions (in particular Clive Warhurst) and the Office of National Statistics (in particular Joshua Abramsky and Karen Grovell). Emails: a.o.teichgraeber@lse.ac.uk and j.vanreenen@lse.ac.uk

The "decoupling" of wages and labour productivity is a common phenomenon in many rich countries of the world (OECD, 2018; Nolan, Roser, and Thewissen, 2019). In the US for example, labour productivity (GDP per hour) has grown substantially more than worker compensation per hour since the 1980s. This is illustrated in Chart 1 that uses data from Mishel and Bivens (2021) over the same time period that we will use later in the article for the UK. This is closely related to the fall in the labour share of GDP, which has been the subject of a vast literature (e.g. Autor *et al.* 2020). The consensus is that there has been a substantial fall in the US labour share. There is less consensus regarding other countries, but most studies do find a general fall since 1980, albeit with very different speeds and magnitudes (e.g. Karabarbounis and Neiman, 2014).

Pessoa and Van Reenen (2013) showed that UK trends in decoupling looked different from the US in some respects. Using data for the 1972-2010 period, they found that (unlike the US) average employee compensation rose at a similar rate to labour productivity.[2] However, like the US, median wages of employees had risen much more slowly than labour productivity. This article first revisits the question of UK decoupling, using another decade of data through to 2019, the year before the start of the pandemic (in order to avoid confounding longer-term trends with the COVID-19 shock).

In the first part of this article, we still do not find "net decoupling" of labour productivity and employee compensation looking over the period from 1981 to 2019 as a whole. We use "net" to indicate this is the difference from "overall decoupling" when we take into account inequality, non-wage compensation and some statistical factors. However, there has been substantial "overall decoupling" of labour productivity and employee median wages.

In our decomposition analysis we find that most of the divergence between overall and net decoupling (three-fifths) can be explained by an increase in inequality which drove a large wedge between mean wages (whose growth was dominated by the most highly paid) and median wages. A further one third of overall decoupling is accounted by the increase in non-wage benefits (the difference between compensation and wages). Although one might regard non-wage compensation such as employer pension contributions a bona fide element of (deferred) labour compensation, it turns out that in the UK a substantial part of this is because of firms re-financing their past pension commitments (which counts as compensation under Office of National Statistics (ONS) conventions). The other components of overall decoupling, such as the statistical discrepancies between data sources and the consumer versus producer price deflator, are generally small in magnitude and offset each other. The compensation and wage measures in the first part

2 In this article, we follow the standard intuitive definition of defining workers as the sum of employees and the self-employed. This should not be confused with legal definitions. For example, under English law, a "worker" is a person in an employment relationship that confers less rights than an "employee", but has more labour rights than a self-employed person.

of the article only include employees. For a comprehensive analysis of the whole UK economy, it is important to look at the self-employed as well.

The second part of our article examines the self-employed, who have increased from 11.8 per cent of the workforce in 1981 to 15.7 per cent in 2019 - and whose compensation is therefore missing from the employee average compensation and wage series. The self-employed do contribute to GDP, however, so ignoring them is clearly problematic (e.g. Gollin, 2002; Guitérrez and Piton, 2020; Smith *et al*, 2019). This turns out to matter a lot in the UK context. The self-employed as a group appear to have done much worse than the employed in terms of their income trends since 1981 and especially after 2001. In our baseline estimates, we find that the average real compensation of an employee grew by 80 per cent between 1981 and 2019 compared to only 50 per cent for the income of a self-employed person: a 30 percentage point difference.

A difficulty with self-employed income data is to determine what part can be classified as labour income and what part as capital income. The Office for National Statistics divide the overall "mixed income" (income derived from the business they run) of the self-employed into a part which is labour compensation and the residual (capital income). This fraction is a difficult object, as the self-employed have a lot of latitude to determine exactly how they will split their income, and this decision is heavily influenced by tax rules. Nonetheless, we show that however one does this split, the self-employed as a group have been doing very poorly compared to the employed.

Using data from the Family Resources Survey (FRS), we examine employed and self- employed income since the mid-1990s, and show that two factors play a key role in explaining the slower growth of self-employed income. First, there has been a big increase of the share of solo self-employed in total self-employment. This group earns substantially less on average than the non-solo self-employed (i.e. those who employ other people). This compositional shift explains over half of the divergence. The rest of the gap is essentially all explained by hours worked, which have fallen dramatically for the self-employed, but have been stable for the employed.

We combine our findings to trace the impact on the labour share of GDP. As our decoupling analysis implies, if the average income of the self-employed had grown at the same rate to that of employees, the labour share of GDP would have been flat for the period from 1981 to 2019. Incorporating the self-employed and taking out non-wage compensation implies a 3.5 percentage point fall in the UK labour share.

The article is structured as follows. The first section provides a short literature review, Section 2 presents the decoupling analysis. Section 3 includes the analysis of self-employed income and the implications for the labour share, Section 4 provides an analysis of potential mechanisms behind our findings. Section 5 concludes.

Literature Review

As noted above, Bivens and Mishel (2015) and Mishel and Bivens (2021) provide the facts on decoupling in US data

Chart 1: Growth of Average Compensation and Labour Productivity in the United States, 1981-2019

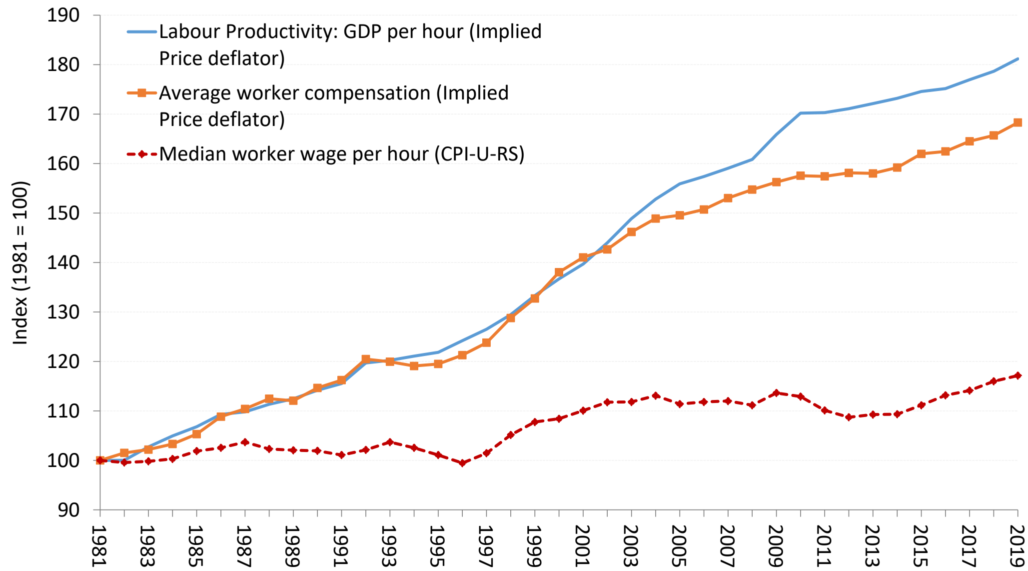

Source: Data from Mishel and Bivens (2021).

Note: GDP data come from Bureau of Economic Analysis (BEA). Compensation is approximated via a wage to compensation ratio based on BEA and BLS data and includes all workers (i.e. including self-employed). Hours worked also come from the BLS. GDP and average compensation are deflated by the implied price deflator (we later refer to the difference between these two series as "net decoupling"), and median wages by the CPI-U-RS. "Average" refers to the mean.

since the 1970s. As shown in Chart 1, US labour productivity grew faster than mean compensation (net decoupling) and much faster than median wages (overall decoupling). Apart from studies of individual countries, there are several cross-country comparisons of labour productivity and wage/compensation growth (e.g. OECD, 2018 and Greenspon, Stansbury and Summers, 2021). For example, Nolan, Roser, and Thewissen (2019) find that most countries have experienced decoupling of productivity and median household income growth, but note that divergence is particularly large in the US.

Although one of Kaldor's (1957) "stylized facts" was the stability of the labour share, much recent work documents a fall in the labour share globally and in individual countries over time. An extensive literature discusses reasons behind the fall (see the survey in Grossman and Oberfield, 2021). Karabarbounis and Neiman (2014) argue that rapid falls in the quality-adjusted price of information and communication technology has led firms to shift from labour to capital (although others have expressed scepticism that the labour- capital elasticity could be large enough to generate this). Autor *et al.* (2020) emphasise the fact that median firm labour shares have been stable, and that the aggregate fall in the labour share is due to "superstar firms" (that have low labour shares and high mark-ups) becoming more dominant in the economy. De Loecker, Eeckhout, and Unger (2020) argue for a rise in aggregate mark-ups and market power. Other reasons identified in the literature are exposure to trade with China and international outsourcing (Elsby, Hobijn, and Şahin, 2013), changing social norms as well as the role of labour market institutions such as unions (Piketty, 2014), and privatisation (Azmat, Manning, and Van Reenen, 2012). Rognlie (2015) looks at the role that housing plays in the fall of the labour share in more detail. More closely related to our focus on the self-employed, Gollin (2002) explains that neglecting the self-employed in labour share calculations can lead to substantial misinterpretations of labour share trends. Guitérrez and Piton (2020) propose different methods to account for the self-employed in the labour share as do ONS researchers (e.g. Dunn, Heys and Sidhu, 2018).

Cribb, Miller, and Pope (2019), Cribb and Xu (2020), and Giupponi and Xu (2020) provide detailed analyses of self-employed income patterns in the UK over the last 20 years. They show that self-employed have experienced particularly large drops in income after the financial crisis. Boeri *et al.* (2020) focus on the self-employed who do not employ other workers ("solo self-employed"). Based on results from large-scale surveys in the UK, US, and Italy, they show that there are substantial differences in working patterns and income between solo self-employed and self-employed who employ other workers. For example, the solo self-employed earn substantially less on average than other self-employed and a higher share of solo self-employed are dissatisfied with the amount of hours and would like to work more.

Decoupling Analysis in the UK

Data Sources

Our data come from multiple sources.[3] Our baseline measure of labour productivity is GDP divided by total hours worked.[4] An alternative output measure would be Gross Value Added (GVA)[5] and we use this as a robustness check in Appendix A1. The core measure of labour compensation is from the ONS national accounts and is defined as wages and salaries plus non-wage benefits. The ONS obtains this information mostly based on tax information from HM Revenue and Customs (HMRC), the UK IRS. Non-wage benefits include employers' contributions to pensions (the main item), national insurance, health insurance (unlike the US, a minor element in the UK due to the NHS) and other benefits.

Our baseline wage data to construct the median wages is from the Labour Force Survey (LFS). Earnings in LFS include basic pay, benefits, tips, and other smaller components and it currently covers about

3 See the online Appendix A for details on data available at http://www.csls.ca/ipm/41/IPM_41_Decoupling_IPM_appendix.pdf

4 The hours estimate of the UK national accounts come from the Labour Force Survey (LFS). We prefer to use hourly measures due to rises in part-time working and longer holidays. These changes would lead to an underestimation of labour productivity when using per worker measures.

5 GDP = GVA + product taxes – product subsidies. Another option would be to take net domestic product (NDP), which equals GDP minus depreciation. However, the ONS only provides such a series from 1987 onwards. When deflating both GDP and NDP by the GDP deflator, growth rates of the two differ by about one percentage point between 1987 and 2019.

6 We follow Pessoa and Van Reenen (2013) and splice LFS with the General Household Survey pre-1992. See Appendix A for details about the sources.

Exhibit 1: Elements of the Overall Decoupling of Productivity and Pay

$$\begin{aligned} OD = {} & \Delta ln(prod_{PD}^{ONS}) - \Delta ln(comp_{PD}^{ONS}) && \textit{(i. Net Decoupling)} \\ & + \Delta ln(comp_{PD}^{ONS}) - \Delta ln(meanwage_{PD}^{ONS}) && \textit{(ii. Non-wage Compensation)} \\ & + \Delta ln(meanwage_{PD}^{ONS}) - \Delta ln(meanwage_{PD}^{LFS}) && \textit{(iii. LFS/ONS divergence)} \\ & + \Delta ln(meanwage_{PD}^{LFS}) - \Delta ln(medwage_{PD}^{LFS}) && \textit{(iv. Inequality)} \\ & + \Delta ln(medwage_{PD}^{LFS}) - \Delta ln(medwage_{CPI}^{LFS}) && \textit{(v. Deflators)} \end{aligned}$$

40,000 households every quarter.[6] There are well-known issues with standard international surveys like the LFS. First, the earnings and wage data are self-reported so may be incorrect. Second, the LFS has a response rate of about 60 per cent and this has been declining over time, like most voluntary surveys. there are sampling weights that try to correct for non-response bias, but these might be inadequate.

To address these potential issues with LFS, we repeat the entire analysis using the Annual Survey of Hours and Earnings (ASHE) in Appendix B. ASHE is a panel of 1 per cent of employees that are selected randomly by the last two digits of their National Insurance (Social Security) number. Because the data come directly from payrolls of employers, it is likely to be very accurate and there is close to 100 per cent compliance. ASHE does have the disadvantage however, that it has only been conducted from 2004 onwards. For years prior to that, data from the New Earnings Survey (NES) are needed to construct a longer time series. Major breaks in wage data can be observed between 2003 and 2004 in NES-ASHE, and also in subsequent years when the methodology of ASHE adjusted further, bringing in more part-time employees. Note that the ONS wages and salaries series are derived from different sources to the LFS (and ASHE). So some of the divergence between mean and median wages could be from the alternative data sources and we examine this explicitly.

We convert employee wages and compensation to hourly values by dividing the series by employee hours. Employee hours are obtained by multiplying the share of employees in total employment (taken from the UK national accounts which base their estimates on the LFS) with total hours worked. This implicitly assumes that employees and self-employed work the same number of hours per week on average.[7] We use the ONS' GDP deflator and a Consumer Price Index (CPI) respectively to convert GDP and the different wage and compensation measures into real series. Most of our other data series are consistently available from 1981 onwards, but we also present an analysis beginning in 1972 as a robustness check in the Appendix. However, this requires more assumptions to produce longer time series.

7 As we will see in Section 3, this assumption is problematic. In particular, FRS data suggest that self- employed have worked more hours on average than employees in earlier years, and the gap has only closed recently. This would suggest that our results might slightly underestimate growth of hourly employee wages/compensation. However, the bias should be minor and the approach is in line with other work dealing with the limitations around employee hours data in the UK (Pessoa and Van Reenen, 2014; Whittaker, 2019).

We define *"overall decoupling"* (OD) as the difference between growth of labour productivity (deflated by the GDP deflator) and median hourly employee wages (deflated by the CPI deflator). We define *"net decoupling"* (ND) as the difference between growth of labour productivity and mean hourly employee compensation (both series deflated by the GDP deflator).

In the following, we decompose overall decoupling (OD) into different parts. Denoting the change from the base year to year t with Δ, we define OD as:

$$OD = \Delta ln(prod_{PD}^{ONS}) - \Delta ln(medwage_{CPI}^{LFS}) \quad (1)$$

All variables are on a per hour measure, with the subscript indicating the deflator (PD for producer/GDP deflator) and the superscript indicating the data source (i.e. LFS and ONS in our baseline analysis). Labour productivity, $prod_{PD}^{ONS}$, is defined as real GDP (using the GDP deflator, PD) divided by total worker hours (both from ONS). Median wages, $medwage_{CPI}^{ONS}$ are LFS employee median real wages (using CPI). We decompose overall decoupling into five elements as shown in Exhibit 1.

In row (i), $comp_{PD}^{ONS}$ is mean employee compensation of employees and we have defined "net decoupling" as the difference between productivity growth and this measure. In row (ii), $meanwage_{PD}^{ONS}$ are ONS mean employee hourly wages, so this reflects the difference between compensation and wages ("Non-Wage Compensation"). In row (iii), $meanwage_{PD}^{LFS}$ is LFS mean employee hourly wages, so this difference reflects any divergence between the ONS and LFS mean wage series ("LFS/ONS divergence"). In row (iv), $medwage_{PD}^{LFS}$ is LFS median employee hourly wages, so this difference reflects the wedge between mean and median wage growth ("Inequality"). Finally, in row (v), $medwage_{CPI}^{LFS}$ deflates median hourly by the CPI deflator instead of the producer price deflator, so this difference reflects a difference in the measures of inflation ("Deflators").

Decoupling Analysis

We start by looking at our two baseline measures of decoupling, overall and net, between 1981 and 2019 in Chart 2.

The solid line shows the growth of labour productivity.[8] The line with triangle markers is employee compensation per hour also deflated by the GDP deflator. The dashed line is LFS median earnings deflated by the CPI deflator. It is clear that labour productivity and hourly compensation have grown at a similar rate over 1981 to 2019 as whole, i.e. there has been no net decoupling (i.e. row (i) of equation (1) is a trivial -0.1 percentage points). Both series grew by 82 per cent in the quarter century 1981-2007, and both series have essentially stagnated since the Financial Crisis that began with the collapse of Northern Rock in 2007. In this sense, Britain's major economic problem over the last 14 years has been the dis-

8 It is important to note that our analysis applies to the UK economy as a whole which means that measurement issues in the non-business sector may affect our results.

9 Interestingly, Williams (2021) makes a very similar argument for Canada. He shows that productivity and

Chart 2: Overall and Net Decoupling in the UK, 1981-2019

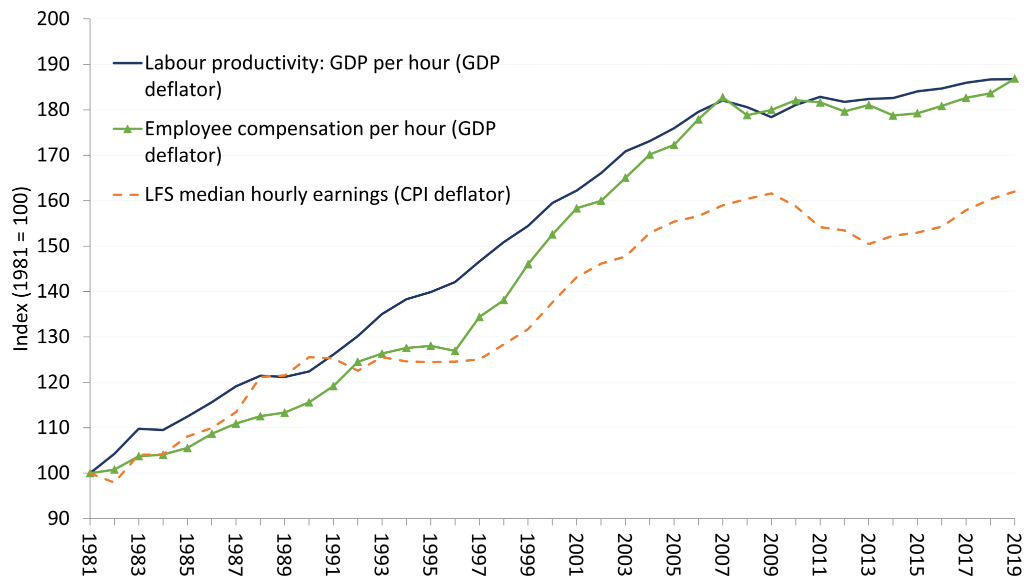

Source: LFS, ONS, and OECD data (see Appendix for details).

Note: Values are shown as an index (1981=100). Labour productivity is total GDP divided by total hours worked deflated by the GDP deflator. Employee compensation is divided by total employee hours and also deflated by the GDP deflator. LFS median hourly earnings are deflated by the CPI deflator. We refer to the difference between the growth rates of labour productivity and average compensation as "net decoupling", and the difference between labour productivity and LFS median earnings as "overall decoupling".

mal record of productivity which grew by a mere 0.21 per cent after 2007, compared to 2.34 per cent in the pre-crisis period.[9]

Looking more closely, it is clear that there has been net decoupling of productivity and employee compensation in certain sub-periods. Under the Thatcher-Major Conservative governments through 1996, labour productivity did grow faster than employee compensation, leading to substantial net decoupling of about 16 percentage points. Under New Labour 1997-2007, compensation grew much faster than productivity, making up all the lost ground in the earlier years. Since then, both series have stagnated alongside each other.

The slow growth of UK productivity has been extensively discussed without any clear resolution of the causes.[10] Some part is due to a general slowdown in productivity across the globe (Bloom *et al.*, 2020), especially after the financial crisis, although the slowdown has been particularly severe in the UK. Explanations include measurement problems (Syverson, 2017), a period of learning about new technologies like Artificial Intelligence (Brynjolfsson, Rock and Syverson, 2021), the overhang of financial market frictions (Besley *et al.*, 2020), the growth of firm market power (Philippon, 2019) and/or too much austerity, especially in the years following the crisis which saw large cuts in public investment (Bagaria, Holland and Van Reenen, 2012).

average compensation have grown at similar rates since 2000, but very slowly.

10 See for example Blundell, Crawford, and Jin, 2014; Patterson *et al.*, 2016; Goodridge, Haskel, and Wallis, 2018; Oulton, 2019; Valero and Van Reenen, 2019.

Chart 3: Detailed Decoupling Analysis in the UK, 1981-2019

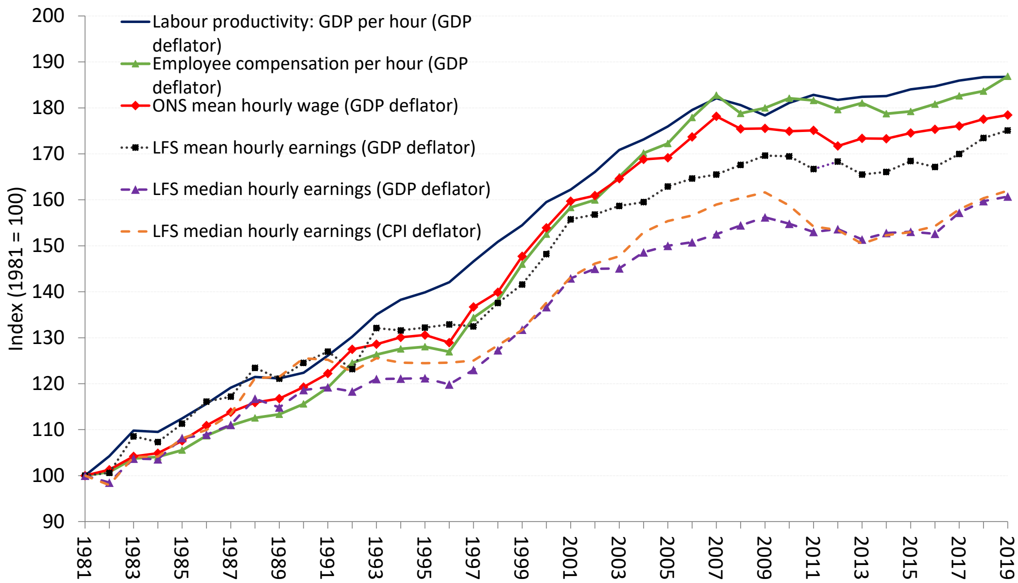

Source: LFS, ONS, and OECD data (see Appendix for details).

Note: Values are shown as an index (1981=100). Labour productivity is total GDP divided by total hours worked (GDP deflator). Employee compensation divided by total employee hours worked (GDP deflator). ONS mean wage is employee total annual wages and salaries earned by total employee hours worked. This series and LFS mean hourly earnings are deflated by the GDP deflator. For median hourly earnings, we provide one series deflated with the GDP deflator and another deflated with the CPI.

Returning to Chart 2, we can also see that there has been substantial overall decoupling. The increase of labour productivity was 87 per cent whereas median wages rose by only 62 per cent, a difference of 25 percentage points. There have been two periods of big divergence. The first was in 1990-1996 when median hourly wages stagnated (average annual growth of -0.12 per cent), and productivity grew consistently (by 2.52 per cent on average per annum). The second was in 2007-2013 when labour productivity stagnated (average annual increase of 0.03 per cent), but median wages actually fell (by -0.91 per cent on average per annum).

Chart 3 extends the previous Chart to present our detailed decoupling analysis. In addition to the growth of labour productivity, employee compensation per hour, and LFS median hourly earnings deflated by the CPI deflator, it displays the growth of ONS mean hourly wages, LFS mean hourly earnings, and LFS median hourly earnings (all deflated by the GDP deflator) since 1981. This allows us to decompose overall decoupling into different parts (following the methodology presented in Section 2.2) and see where differences between growth of labour productivity and median earnings come from.

To simplify the story, Chart 4 has the decomposition over the whole 1981-2019 period. The first bar shows the size of overall decoupling (difference between the growth of labour productivity and LFS median wages deflated by the CPI), which is 24.7 percentage points. The sum of all other five bars on the right hand side equals this overall decoupling. The biggest share

Chart 4: Decoupling Decomposition in the UK (Cumulative Change Between 1981 and 2019, Percentage Points)

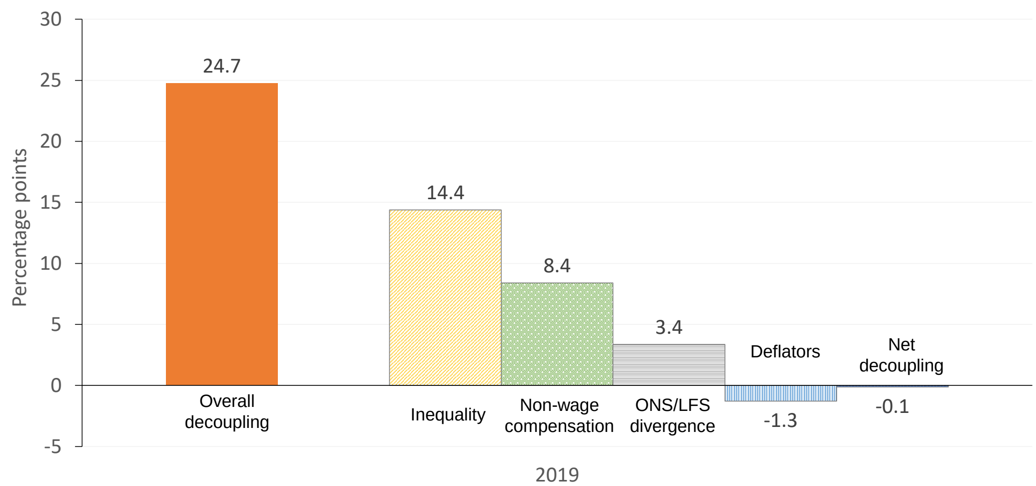

Source: Decomposition of Chart 3 into its components 1981-2019.

Note: Values shown are the percentage point differences between the growth rates. "Overall decoupling" refers to difference between GDP per hour (GDP deflator) and LFS median hourly earnings (CPI deflator); "Inequality" is the difference between LFS mean hourly earnings and LFS median hourly earnings; "Non-wage compensation" is the difference between employee compensation per hour and ONS mean hourly wage; "ONS/LFS divergence" is the difference between ONS mean hourly wage and LFS mean hourly earnings; "Deflators" is the difference between LFS median hourly earnings (GDP deflator) and LFS median hourly earnings (CPI deflator); "Net decoupling" is the difference between GDP per hour and employee compensation per hour.

in overall decoupling comes from inequality, contributing 14.4 percentage points to the overall decoupling number. The second biggest contribution comes from non-wage compensation with 8.4 percentage points. Therefore, between them, inequality and non-wage compensation explain more than 90 per cent of decoupling.

There is a divergence between ONS and LFS mean wages of 3.4 percentage points, with the LFS earnings series lagging behind the ONS wage series.[11] Additionally, the CPI has risen faster than the GDP deflator, resulting in a negative contribution of 1.3 percentage points (offsetting parts of the ONS and LFS divergence). This could reflect increasing price-cost mark-ups (see De Loecker, Obermeier, and Van Reenen, 2021). Putting all this together, net decoupling is essentially zero. As shown in Chart 2, average employee compensation has actually grown trivially faster than labour productivity (0.1 percentage point).

Chart 5 illustrates how the contribution of the different components to overall decoupling has changed over time. Each stacked bar represents a selected year within the period from 1981 until 2019. The individual values within a bar sum up to overall decoupling in that year (note that the values in 2019 correspond to that in Chart 4). We observe that the inequality

11 This seems to occur after 1997 (it was the opposite prior to this). One reason for this is that the LFS may not be picking up some of the very high incomes that HMRC tax data finds, because vary rich individuals are increasingly not participating in voluntary surveys (and due to top-coding in LFS, see Appendix A for details). As we discuss in Appendix B, the divergence is the other way around in the ASHE data which probably better reflects high wage individuals than LFS as it is mandatory.

Chart 5: **Decoupling Decomposition in the UK (1981 until 2019, differences in selected years)**

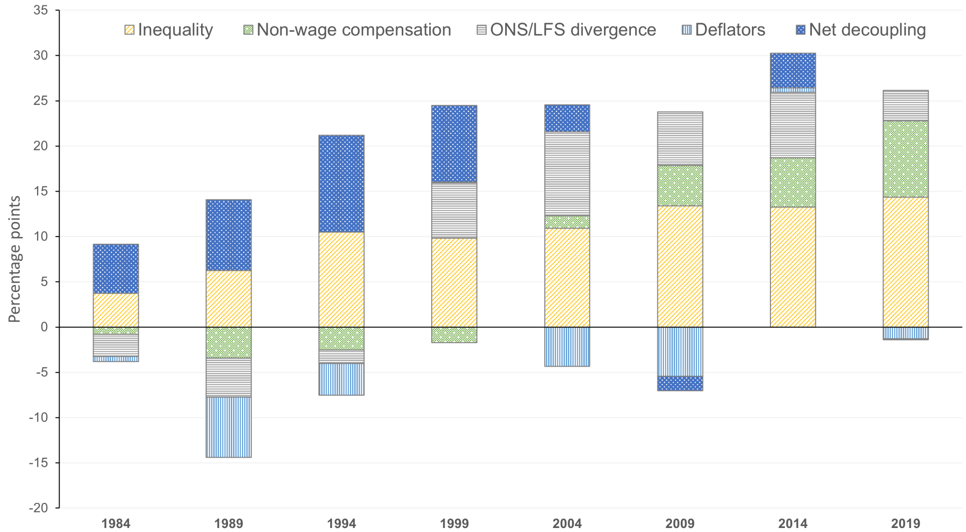

Source: Decomposition of the decoupling analysis in Chart 3 into its single components.

Note: Values shown are the percentage point differences between the growth rates from 1981 until different subsequent years of selected series. Inequality refers to the difference between LFS mean hourly earnings (GDP deflator) and LFS median hourly earnings (GDP deflator); Non-wage compensation to the difference between employee compensation per hour (GDP deflator) and ONS mean hourly wage (GDP deflator); ONS/LFS divergence to the difference between ONS mean hourly wage (GDP deflator) and LFS mean hourly earnings (GDP deflator); Deflators to the difference between LFS median hourly earnings (GDP deflator) and LFS median hourly earnings (CPI deflator); Net decoupling to the difference between GDP per hour (GDP deflator) and employee compensation per hour (GDP deflator).

component has increased consistently with overall decoupling over time. Strikingly, non-wage benefits have not played a major role until the beginning of the 2000s. Whereas wages grew even faster than compensation until the mid-1990s, compensation overtook wage growth in the beginning of the 2000s. As discussed earlier, we also observe major net decoupling until the mid-1990s, and see it vanish afterwards. The components reflecting the ONS/LFS divergence and the deflator difference have also changed over time. The overall growth of LFS mean earnings has been higher than that of ONS mean wages until the mid-1990s. Since then, the overall growth of ONS mean wages from 1981 onwards is higher than the one of LFS mean wages, reaching a difference of almost 10 percentage points in 2004. The overall growth of the GDP deflator from 1981 onwards has been higher than the growth of the CPI in almost all years, with the difference being almost 7 percentage points in 1989 and approximately zero in 2014.

Table 1 additionally shows average annual growth rates of the series depicted in Chart 3 for different time periods. It becomes clear that the main period of overall decoupling has been 1981-1996. Labour productivity has grown by 2.38 per cent on average per annum, whereas median LFS wages deflated by the CPI deflator have only grown by 1.51 per cent per annum on average. In the 1996-2007 period, both productivity and median wages have seen similarly strong growth rates of almost 2.3 per cent per annum. In the most recent 2007-2019 period, both labour productivity and median wages have almost stagnated.

Table 1: Average Productivity and Pay Trends in the UK for Different Time Periods (Average annual per cent change)

	Labour Productivity	Employee Compensation	Mean Wages ONS	Mean Wages LFS	Median Wages LFS	Median Wages LFS CPI
1981-1996	2.38	1.61	1.72	1.96	1.24	1.51
1996-2007	2.28	3.38	3.00	2.03	2.23	2.25
2007-2019	0.21	0.19	0.02	0.48	0.45	0.17
1981-2019	**1.67**	**1.67**	**1.55**	**1.51**	**1.27**	**1.30**

Note: Shown are average annual growth rates (in percent) of the six different lines of Chart 3 for different time periods (1981-1996, 1996- 2007, 2007-2019, and the overall 1981-2019 period). The first 5 columns use the GDP deflator, the last column (Median wages LFS CPI) the CPI deflator.

In summary, we do not observe net decoupling of labour productivity and compensation in the UK, standing in sharp contrast to the US (Bivens and Mishel, 2015; Mishel and Bivens, 2021; Stansbury and Summers, 2018). However, there has been substantial overall decoupling of labour productivity and median wages over 1981-2019.[12] Almost 60 per cent of this divergence can be explained by inequality (mean and median wage difference), and most of the remaining difference by increases in non-wage compensation.

Bell (2015) shows that 85 per cent of the increase in total non-wage compensation between 2003 and 2013 comes from increases in employers' pension contributions. Firms increased these in the beginning of the 2000s to compensate deficits in defined pension systems. Notably, Adrjan and Bell (2018) find that while firms increased pension contributions to close deficit gaps in pension systems, they were able to lower wages of employees to save cost. This implies that employee compensation was potentially only able to keep up with the growth of labour productivity because of increasing employers' pension contributions. However, this increase has potentially not fully benefited large parts of current employees. With a large part of these pension contributions being used to cover deficits in defined pension systems, it is likely that substantial amounts go to retired employees or a rather small share of current employees. This is why we will take a closer look at the non-wage component in compensation when calculating the labour share of income in Section 3.

Extensions and Robustness Checks

We have conducted a large number of robustness checks and extensions to the analysis. We have relegated these to Appendix B and just summarize the main results here. First, we extend the analysis to another decade looking at the trends 1972-2019 instead of 1981-2019 as in our baseline analysis. The data sources become less reliable as we go further back in time, but the qualitative conclusion that large parts of overall decoupling are driven by inequality and non- wage compensation remain the

12 An alternative way to measure decoupling would be to look at the difference between growth of labour productivity and median compensation. Unfortunately, there is no publicly available data tracking median compensation over time in the UK (nor in the US or most other countries). Since non-wage compensation (especially employers' defined pension contributions) also tends to be very unequally distributed overall decoupling based on median compensation may even be higher.

same. The most notable difference is that the non-wage compensation component is much larger over this longer period and we observe slight net decoupling. Second, we switch from using LFS to using ASHE as our main micro-data source to calculate median wages. We note that the divergence between the ONS and ASHE series is larger than with the LFS and in the opposite direction. Third, we switch from using GDP to Gross Value Added (GVA). Fourth, we look at the sensitivity of the results to changing the ordering of the decoupling analysis. Fifth, we use data from the ONS' latest GDP revision in June 2021 which (implements double deflation amongst other changes).[13] Sixth, we present a more detailed analysis of the differences in price deflators. Finally, we look more closely at the role of non-wage compensation.

The bottom line from these extensions is that although the precise magnitude of the contributions to decoupling change, the qualitative results are robust that (i) there is little or no net decoupling of productivity from average employee compensation; (ii) there has been significant overall decoupling between productivity and median wages, and (iii) growing inequality is the main factor and non-wage compensation the second most important factor accounting for overall decoupling.

The Self-Employed

Section 2 showed that the UK has not seen "net decoupling" between labour productivity and average employee compensation over the last four decades. In this respect, our analysis suggests stability of the labour share of GDP. However, it is very important to note that we so far focused on employees only when considering trends in wages and compensation. This is in line with many other comparable decoupling analyses in the literature (OECD, 2018; Whittaker, 2019). Nevertheless, since our productivity growth measure uses estimated hours worked and output from all workers, including the self-employed is potentially important. A criticism of our decoupling analysis thus far is not comparing like with like, as we have implicitly assumed that productivity growth for the employed is the same as it is for the self-employed. There is no simple fix for this issue, as accurately measured the contribution of the self-employed to GDP is very challenging (as well as accurately measuring their income and hours worked).

The self-employed are a very heterogeneous group with major differences in income characteristics and working patterns (Datta, Giupponi, and Machin, 2019; Cribb and Xu, 2020; Giupponi and Xu, 2020). In this section we first analyse differences in income and job trends between employees and the self-employed and then trace out their impact on the aggregate labour share and growth of average worker compensation. We then analyse why self-employed income has grown more slowly, distinguishing between solo and non-solo self- employed using micro-data from the

13 This series has not been produced for the pre-1998 period, so we prefer not to use this for our baseline analysis. It essentially makes no difference to the results.

Chart 6: Growth of Average Employee Compensation and Average Mixed Income in the UK, 1981-2019)

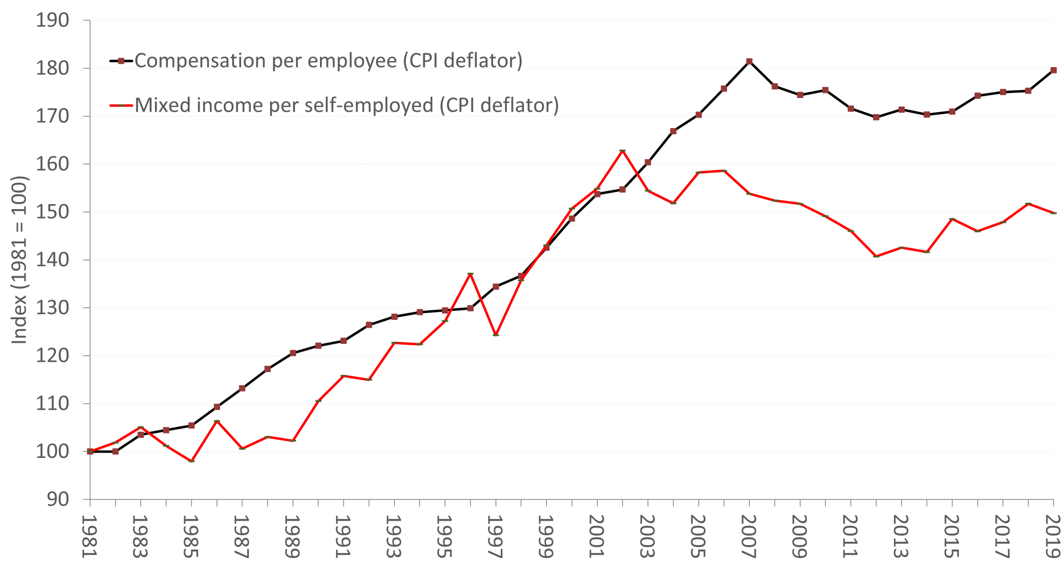

Source: Data from ONS and OECD (see appendix for details).

Note: All values are shown as an index, with the base year 1981 equalling 100. Average compensation is employee compensation divided by number of employees, and average mixed income is total mixed income divided by the number of self-employed. Both series are deflated by the CPI. Mixed income is defined as "the aggregate of a variety of flows of value and rewards accrued by unincorporated businesses owned by households, namely sole proprietors. It contains an element of remuneration for work done by the owner or other members of the household that cannot be disassociated from their profit as an entrepreneur. Mixed income excludes imputed rentals from owner-occupied housing, as this is captured elsewhere in the national accounts.[a]

a https://www.ons.gov.uk/economy/nationalaccounts/uksectoraccounts/articles/nationalaccountsarticles/ areviewofhouseholdsmixedincomeestimatesandplansforupcomingimprovements, last accessed on 29 June 2021

Family Resources Survey (FRS).

Self-Employment, Aggregate Labour Share and Average Worker Compensation, 1981-2019

Trends in Compensation and Employment by Worker Type

Chart 6 compares the growth of average employee compensation and average self-employed income from 1981 to 2019. Note that this is similar to the compensation measure we used in the previous section, but for now, we switch from hourly measures to per worker measures. As the measure of self-employed income, we take average "mixed income" per self- employed person. Mixed income is defined by the ONS as "the aggregate of a variety of flows of value and rewards accrued by unincorporated businesses owned by households, namely sole proprietors" and is sourced from the UK's national accounts. Note that this includes both labour and capital income of the self-employed and we will examine below different ways to divide

14 See Smith *et al.* (2019) for an extensive discussion of differences between capital and labour income. For example, the self-employed can decide what share of profits from their business to take as wages compared to capital income such as dividends. This decision will usually be heavily influenced by tax incentives, which makes it difficult to distinguish the "true" amount of labour compensation accruing to a self-employed person.

Chart 7: Share of the Self-Employed in Total UK Employment, 1981-2019

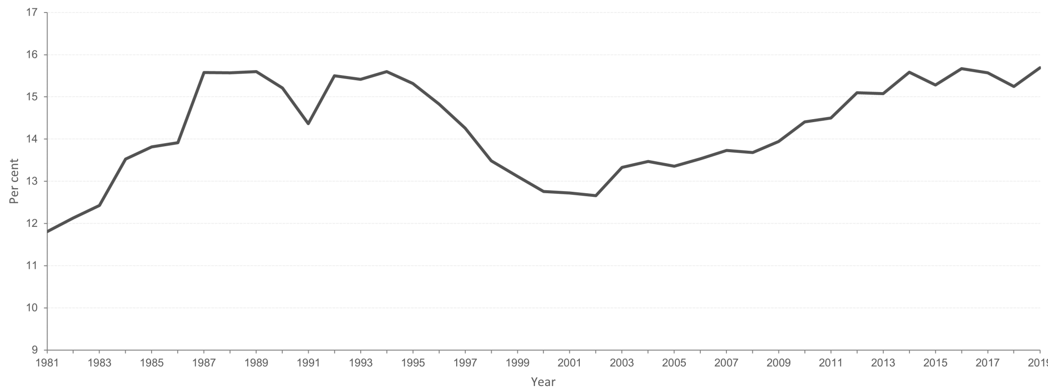

Source: Data from ONS (see appendix for details).
Note: Note that the share of employees in a year equals 100 minus the share of self-employed.

mixed income into labour and capital components.[14]

Chart 6 shows substantial differences between the growth of employee compensation and self-employed income over time. Average compensation increased by about 80 per cent and average mixed income by about 50 per cent (1.55 per cent versus 1.16 per cent average annual increase). This amounts to a 30 percentage point difference between the income growth rates of the two groups. Thus, the average self-employed person has done much worse than the average employee over this period. One caveat is that self-employed average income is estimated by dividing mixed income (from HMRC) by the number of the self-employed (from the LFS). Using administrative data and household survey data certainly creates potential measurement error, although our hope is that this is reasonably stable over time. In any case, it is unclear whether correcting for this would lead to an improvement or a further deterioration in the relative position of the self-employed.

There are two periods of big divergences. First, average mixed income growth stagnated in the 1981-1989 period (average annual increase of 0.37 per cent), whereas employee compensation grew by a substantial 2.37 per cent on average annually. Second, between 2002 and 2007, employees' compensation again grew much faster than mixed income. Post financial crisis, all groups suffered, with average mixed income actually falling.

These results become even more important when looking at the changing share of the self- employed in total employment as shown in Chart 7.

The self-employed share increased from 11.8 per cent in 1981 to 15.7 per cent in 2019. Interestingly, the periods in which the share of self-employed has increased (e.g. 1981 until early 1990s) coincide with slow growth of self-employed income in Chart 6, whereas periods that have seen a decrease in the share of self-employed coincide with fast growth of self-employed income (e.g. mid-1990s until early 2000s). This suggests some selection forces – the people entering self-employment may be more marginal individuals, rather than tal-

Chart 8: Labour Share of GDP in the UK, Estimated via Different Methods, 1981-2019

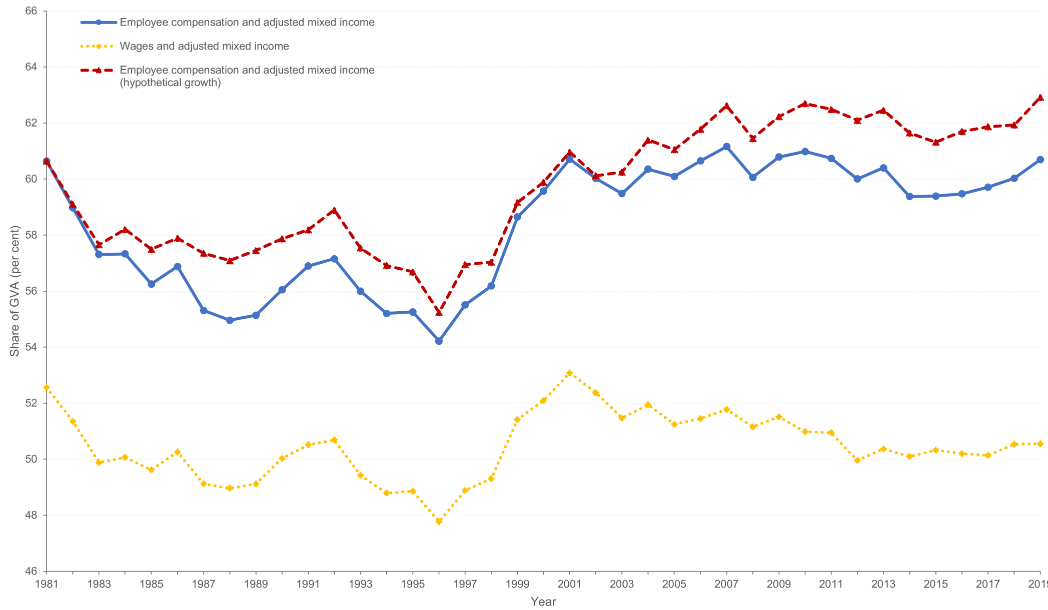

Source: Data from ONS (see Appendix for details).

Note: The solid line with circles (blue) shows compensation and adjusted mixed income (an estimate for self-employed income that can be classified as labour income) over GDP. The dotted (yellow) line shows wages and adjusted mixed income over GDP, i.e. it excludes non-wage benefits of employees (such as employers' pension contributions, employers' national insurance payments etc.). The red (dashed) line takes the value of the blue series in 1981, and then applies a hypothetical growth rate for the years after. The hypothetical growth rate stems from the decoupling analysis in section 2, and equals the growth of employee compensation per hour over growth of GDP per hour. This is to approximate how the labour share could have evolved if all workers (including self-employed) had experienced growth of income equal to that of employees

ented "entrepreneurs".[15]

Together, Chart 6 and Chart 7 clearly show that not only have the self-employed performed much worse than employees since 1981, but at the same time their share in total employment has increased substantially. Unemployed people often select into self-employment if they are unable to find jobs.

Impact of Self-Employment on the Aggregate Labour Share

What does this mean for the UK's labour share of income? Chart 8 shows different estimates of the labour share of GDP.

The blue line corresponds to the ONS' headline measure.[16] It uses employee compensation and self-employed mixed income that can be classified as labour income (the latter being labelled "adjusted mixed

15 This is consistent with the modern empirical entrepreneurship literature, showing that most self-employed have characteristics more similar to the unemployed than high wage employees. Levine and Rubinstein (2017), for example, emphasise that it is important to distinguish between incorporated and unincorporated businesses. Incorporated businesses generally employ workers, whereas unincorporated businesses are the solo self-employed. Unemployed people often select into self-employment if they are unable to find jobs.

16 The ONS uses GVA instead of GDP to calculate the labour share in official publications, e.g. Dunn, Heys and Sidhu (2018). To be consistent with our previous analysis, we use GDP in this Section and repeat the analysis with GVA in Appendix B as a robustness check. Additionally, we show the labour share series using net domestic product (NDP), defined as GDP less capital depreciation.

income") in the numerator and GDP in the denominator. Following national accounting conventions, call CoE_t = employee compensation in year t, MI_t = self-employed mixed income, and the share of mixed income attributed to labour income α_t . Then, the labour share (the blue line in Chart 8) in year t, Ls_t, is:

$$LS_t = \frac{(CoE_t + (\alpha_t \cdot MI_t))}{GDP_t} \tag{2}$$

where

$$\alpha_t = \frac{CoE_t}{(CoE_t + GOS_t)} \tag{3}$$

with GOS_t being the gross operating surplus of corporations. This assumption follows international practice and assumes that in relative terms, the returns to capital and labour of the self-employed are the same as those in the corporate sector. [17] An alternative would be to use the values the self-employed declare as labour income to the tax authorities, but this is likely to be biased as it is heavily influenced by the taxation of the self-employed.

Looking at the "ONS official" labour share (blue line) in Chart 8, we observe a fall of about 2 percentage points between 1981 and 2019 from 56.2 per cent to 54.2 per cent. To examine the extent to which the slow growth of self-employed income contributes to this fall, we construct a hypothetical labour share measure (see Appendix C for details). The red line shows how the labour share would have evolved if self-employed labour income had grown at the same rate as employee compensation per hour.[18] Here, we observe no fall of the hypothetical labour share from 1981 to 2019 (a minimal increase of 0.04 percentage points) as in the net decoupling analysis in section 2. This shows that trends in self-employment were - in an accounting sense - solely responsible for the decline of the labour share over this period.[19]

Next, consider the role of non-wage compensation. The yellow line shows a labour share measure incorporating adjusted mixed income, but just using ONS wages and salaries (thus excluding employers' social contributions). One reason for doing this is that much of non-wage compensation is refinancing of company pension schemes for already retired employees. On this measure, as discussed earlier in this articles, we observe a more substantial fall in the labour share of 3.5 percentage points.[20] We are not arguing that this is the sole "correct" number for the labour share, but rather to highlight the quantitative importance of different assumptions.

As before, we are using adjusted mixed income following Dunn, Heys and Sidhu

17 Dunn, Heys and Sidhu (2018) have an extensive discussion of this in the UK context.

18 When constructing this hypothetical measure, it is important to bear in mind that the share of self-employed in total employment has increased over time. It is not sufficient to multiply the growth of aggregate employee compensation with aggregate mixed income in the base year. This would ignore that the share of self-employed has increased and would lead to an underestimation of the potential labour share.

19 Note that we are using 1981-2019 and this will not be true for all sub-periods or earlier years as discussed above.

20 Note that this assumes no change in the self-employed' share of nonwage compensation. If we assume this grew at the same rate as employees, this would cause the labour share to be another 0.3 percentage points lower (i.e. a 3.8 percentage point fall).

Chart 9: Weighted Average Worker Income, 1981-2019

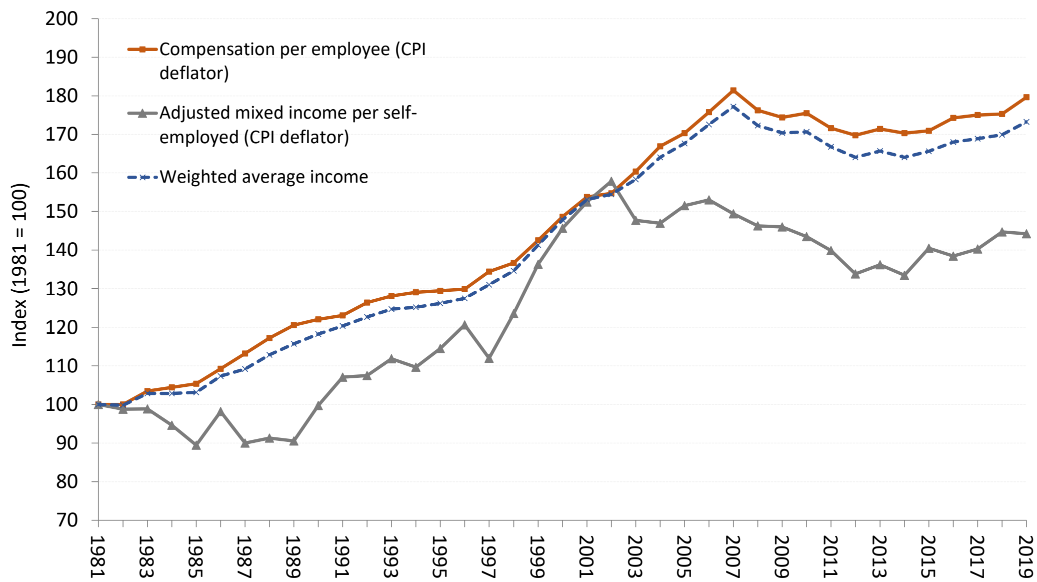

Source: Data from ONS and OECD (see appendix for details).

Note All values are shown as an index, with the base year 1981 equalling 100. The dashed (blue) line, weighted average income, is average compensation plus average adjusted mixed income, weighted by the share of employees and self-employed respectively. It can be interpreted as the income of the average worker.

(2018) to obtain the labour component of mixed income.[21] Average employee compensation Y_t^E (solid line with square markers) has grown by 80 per cent much faster than average self- employed compensation Y_t^M (solid line with triangle markers) at 44 per cent generating a 36 percentage point difference. The weighted average worker compensation, Y_c (dashed line), has grown by 73 per cent. The slower growth of self-employed income drags the average worker compensation line below the employee compensation line, but not by a large amount because the self-employed only make up a relatively small part of the total workforce (15.7 per cent in 2019, as shown in Chart 7).

Comparing columns (ii) and (iii) in Table 2, we see that employee compensation is substantially higher than that of the self-employed. In absolute terms, the difference between employee vs. self-employed compensation increases from £8,001 in 1981 (a 58 per cent employee premium) to £19,234 in 2019 (a 97 per cent premium, see column (iv)).[22] Given the fact that the self-employed earn less than the employed on average, some of the slower growth in average worker compensation comes simply from the compositional shift towards the self-employed.

21 Let S_t^E be the share of workers who are employees and $S_t^M = (1 - S_t^E)$ be the share of workers who are on mixed income (self-employed). Then, Y_c is the average income per worker, with Y_c= S^E*Y_t^E + S^M*Y_t^M

22 Part of this difference is explained by the increase employers' non-wage compensation. We analyse this in Appendix C that shows that the difference in growth rates is still evident if we exclude employers' social contributions

The Role of the Solo Self-employed, 1997-2019

What has caused the slow growth of self-employed income? An important distinction is between people who do not employ any workers – the "solo self-employed" and people who employ workers – "employer firms" (Cribb and Xu, 2020). In what follows, we call these two groups solo SE and non-solo SE. The distinction is close to that in the entrepreneurship literature between incorporated and non-incorporated self-employed.[23]

The ONS and LFS data that we used in the previous section do not allow us to distinguish clearly between solo SE and non-solo SE outcomes. The LFS only provides the numbers of self-employed in total employment, but no self-employed income data. To tackle this we therefore turn to the Family Resources Survey (FRS). The FRS is an annual household survey that covers information such as income, wages, savings, investment, and self-employment.[24] It was first conducted in 1993/1994 and in the last year available to us (2019/2020), about 19,000 households were interviewed.[25] Since this is a much smaller sample size (especially for the self-employed with less than 3,000 respondents in 2019) than the LFS, ASHE or ONS data, we use three-year moving averages to reduce sampling variation.[26] In addition, since the data are known to be less reliable in the earlier years, we present results from 1997/1998 (labelled "1997") through 2019/2020 (labelled "2019"). Chart 6 showed that the largest sustained divergence between employed and self-employed was from 2001 onwards, so it makes sense to focus on this sample period.

As noted above, there are many caveats with self-employed data. First, total income may well be under-reported for tax purposes and although FRS is anonymous and individuals cannot be identified for tax purposes (and respondents are told this), this could still be an issue. In particular, if underreporting has increased over time (although it is unclear why this should be the case) this might help explain slower growth in income trends. Second, hours data are particularly hard to verify. For the employed, hours reporting can be from the employer payroll (e.g. ASHE) or from the worker (e.g. LFS) so the aggregate numbers can be cross-checked. But since there is no administrative series for the hours of the self-employed we have to rely on household surveys. Of course, in the FRS everything is self-reported, but the self-employed may find it more difficult to accurately judge their working hours. Third, business owners may be taking less income as compensation and more as "Gross Oper-

23 A similar distinction is sole traders vs. non-sole traders (e.g. Cribb, Miller, and Pope, 2019). Both coincide in the FRS we will use. For more information about different types of self-employed in the UK, see Blundell (2019).

24 https://www.gov.uk/government/collections/family-resources-survey–2, last accessed on 12 June 2021

25 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment$_{d}$ata/file/874507/family−resources − survey − 2018 − 19.pdf, lastaccessedon12June2021

26 A presentation of the corresponding unadjusted data can be found in Appendix C.

Chart 10: Average Weekly Hours Worked, Weekly, and Hourly Income by order type, 1997-2019

Panel A: Mean Real Weekly Income (2018£)

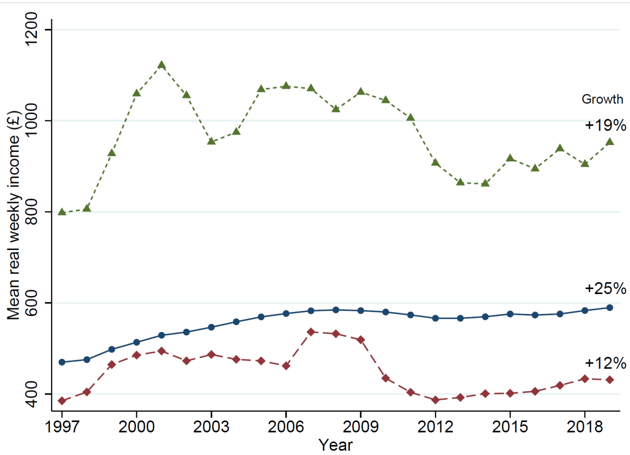

Panel B: Average Weekly Hours Worked

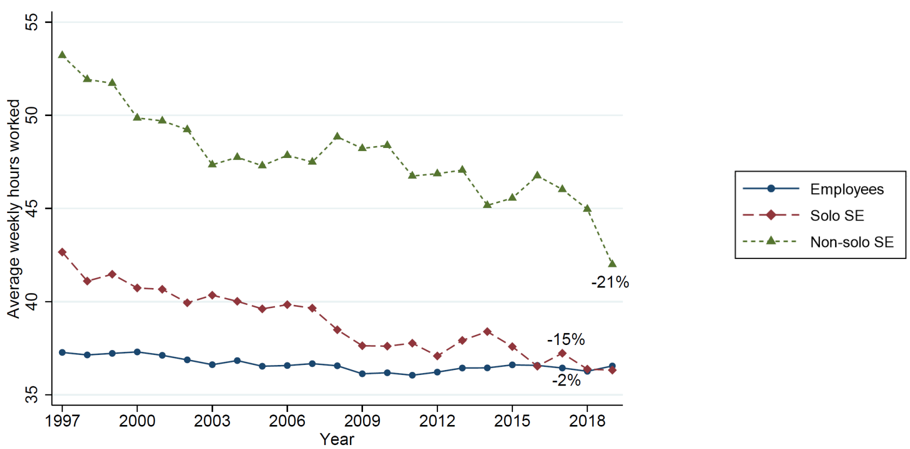

Panel C: Mean Real Hourly Income

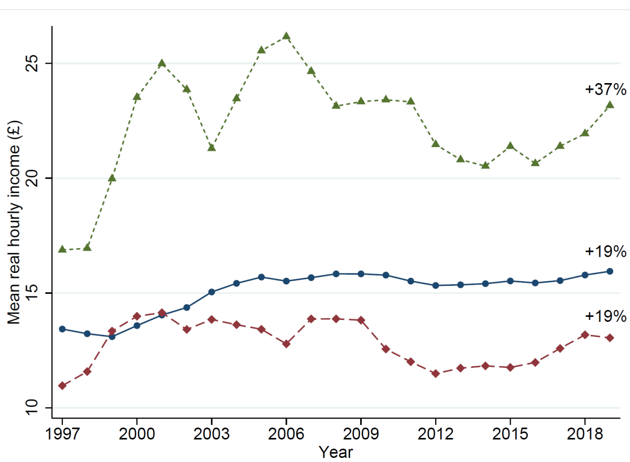

Source: Data from FRS (see Appendix for details).
Note: The number at the end of each line is the growth rate 1997-2019. Panel A: Employee income is gross wages or salaries as shown on their payslip. Income of self-employed is defined as "the total amount of income received from self-employment GROSS of tax and national insurance payments, based on profits where individual considers themselves as running a business, on estimated earnings/drawings otherwise". All data are shown as three year moving averages (except 1997 and 2019, where we use a two- year average). Panel B: Average usual hours worked by a worker on all jobs held excluding unpaid overtime. Panel C: Hourly income divides income (Panel A) over hours (Panel B). All data are shown as three year moving averages (except 1997 and 2019, where we use a two-year average). Income is deflated by the CPI (same CPI deflator as in the decoupling analysis).

Table 2: Share of Self-Employed in Total Employment and Income Statistics in Levels for selected years

	Share of self-employed in total employment, %	Compensation per employee (CPI), £	Adjusted mixed income per self-employed (CPI), £	Income premium of being an employee, %	Average income per worker (CPI), £
1981	11.8	21,750	13,750	58.2	20,806
1997	14.3	29,235	15,401	89.8	27,262
2001	12.7	33,440	20,967	59.5	31,853
2019	15.7	39,065	19,831	97.0	36,047

Source: Data from ONS and OECD (see appendix for details).

Note: Average income per worker (column (v)) is calculated as the average of employee compensation (column (ii)) and adjusted mixed income (column (iii)), weighted by the shares in total employment of the respective groups (column (i)). Column (iv) has the income premium of being an employee compared with the average self-employed person (mark-up of column (ii) over column (iii)).

Chart 11: Share of Solo Self-employed in Total Self-Employment, 1997-2019

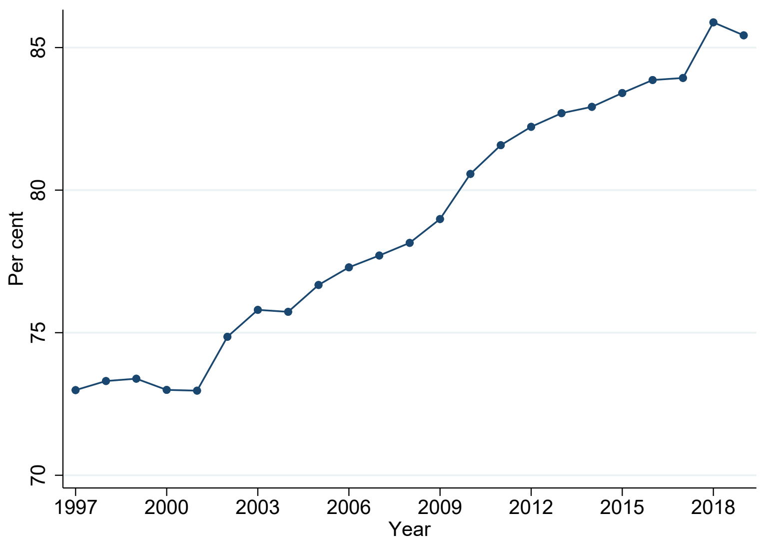

Source: Data from LFS (see appendix for details).

Note: Shown is the share of solo SE in total self-employment from 1997 until 2019. The yearly value is calculated as the average of the four quarters in a year. A corresponding graph with FRS data can be found in the Appendix.

ating Surplus" (Smith *et al.*, 2021). This is why we have focused on all business income (mixed income) for the self-employed so it includes both dividend income and salary. Although these are all concerns, it is not obvious why these measurement issues should have changed over time in such a way to generate the patterns in the data.[27]

Chart 10 shows the FRS information split into three panels. Each panel shows the changes for three workers groups: (i) employees, (ii) solo SE and (iii) non-solo SE for weekly income (Panel A), hours (Panel B) and hourly income (Panel C). In Panel A, and as noted above, employee income has grown by about 25 per cent (from £470

27 In the Appendix we compare trends using the FRS with ONS administrative data. The broad trends are comparable. Employee income growth is nearly identical. Self-employed income has grown more slowly in the FRS than in ONS, however. It is unclear whether this is a problem with the ONS or the FRS, but this caveat should be borne in mind. In what follows, all our comparisons are within the FRS data.

to £590). By contrast, the solo SE have only seen a growth of 12 per cent (£385 to £431) and the non-solo SE of 19 per cent (£798 to £952). Note that most of the growth in weekly income for both self-employed groups occurred pre-2002, consistent with ONS numbers in Chart 6. In terms of income levels, the non-solo SE earn by far the most compared to the other groups. The income of employees is above the solo SE, a gap that grew considerably during and after the financial crisis. In 2019 employees earn about 37 per cent (£160) more than the solo SE per week.

Panel B of Chart 10 shows that in 1997, employees worked the least - about 37 hours per week compared to the solo SE on 42.5 hours and the "Stakhanovite" non-solo SE an enormous 53 hours a week. Whereas there has been little change in hours worked for the employed, there has been a substantial fall for the self-employed, from 42.5 to 36 hours per week (-15 per cent) for solo SE and from 53 to 42 hours per week (-21 per cent) for the non-solo SE. Today, the solo SE now work about the same number of hours per week as employees.[28]

Panel C of Chart 10 shows that in percentage terms, employees and solo SE have seen comparable growth in hourly income of about 19 per cent (albeit from different bases: £11 for solo SE vs. £13 for the employed). Strikingly, non-solo SE have seen by far the highest growth in hourly income of around 37 per cent.

Chart 11 shows that the share of solo SE in total self-employment has increased by more than 12 percentage points over time, from 73 per cent in 1997 to 85 per cent in 2019 (with most of the increase post-2001). Since the solo SE have much lower hourly incomes than the non-solo SE (Panel A of Chart 10), this fundamentally explains most of the slower growth of the self-employed income compared to employee income.[29]

Summary on the Slower Income Growth of the Self-Employed

We summarize our analysis of self-employment in Table 3 based on trends in FRS data. The first three rows show that employee income grew 23.4 (25.5 – 2.1) percentage points more than the self-employed income from 1997 to 2019.

The poor performance of the self-employed may seem surprising as weekly income growth of solo SE was 12 per cent (Row 3) and for non-solo SE was 19 per cent (Row 4). This averages out to a mere 2.1 per cent (Row 5) growth for the SE as a whole through two mechanisms. First, solo SE income is substantially less than non-solo SE income (e.g. in 2019 solo SE earned less than half that of the non-solo

28 Chart 10 shows average hours worked and average income from all jobs that a person has. A person is classified as an employee or self-employed if she works the majority of hours in that job. The fraction of people who perform both employee and self-employee jobs is very small and has not changed much over time (0.95 per cent of all workers in 1997 and 1.02 per cent in 2019). The analysis looks almost identical if we only consider hours and income from the main job type. Corresponding graphs can be found in the Appendix.

29 Note that the data come from LFS, not FRS. We decided to use LFS data to calculate the solo share in total self-employment because of the LFS' larger sample size. We suspect that the LFS estimates should be more accurate. The corresponding charts and results of our analysis with FRS data can be found in the Appendix.

SE: £952 vs. £431 a week). Second, the share of Solo SE in total SE has increased by 12 percentage points as show in Chart 11. If we fixed the fraction of solo-SE at its 1997 level, average SE income would have grown by 15.2 per cent instead of 2.1 per cent and the income growth gap with the employed would fall from 23.4 to 10.3 percentage points. In this sense, the rise of the solo SE explains over half of the slower income growth of the self- employed compared to employees.

A second issue is different trends in hours worked. Although the weekly income change between employed and self-employed was 23.4 percentage points, rows (4)-(6) of Table 2 show that the hourly income difference was only 3.1 percentage points (18.7 per cent for employed – 15.6 per cent for SE). This implies that in hourly terms, the self-employed have not done so badly. 87 per cent ((23.4 - 3.1)/23.4) of the difference in income was due to the big fall in hours worked by the SE. Part of this is related to the compositional shift towards the solo SE who work less hours than non-solo SE, and part of this is the reduction of hours for both types of self-employment (Panel B in Chart 10).

Should this make us more relaxed about the position of the self-employed? It depends whether we think the reduction in hours worked by the self-employed was a voluntary shift to more leisure, or whether it is because the self-employed have been constrained to work fewer hours than they want due to lack of demand. As we will discuss below, it is likely that the solo SE are being constrained to work fewer hours than they would like, so some of the lower hours may be a form of disguised under-employment.[30]

In summary, the declining relative position of the average self-employed worker's weekly income can be explained by these two factors. The majority of the difference is a compositional shift due to the rise in the solo SE. Just about all of the residual difference is explained by the rapidly falling hours worked of the self-employed. Of course, this is just statistical accounting. We now turn to what forces could more fundamentally explain the changing patterns we observe.

Decoupling analysis with self-employed income

The decoupling analysis in section 2 excluded income from the self-employed. This is because estimating the share of labour compensation of the self-employed in mixed income is a difficult task (as discussed above). Bearing this caveat in mind, we now combine results from sections 2 and 3 to include self-employed income in the decoupling analysis. this is reflected by thef black line in Chart 12: it shows average employee and self-employed compensation (the latter being the fraction of mixed income that goes to labour estimated via approach by Dunn, Heys and Sidhu, 2018) per

30 For a general overview of under-employment in the UK, see the ONS' under- and overemployment statistics (sourced from the LFS): url-https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/datasets/underemploymentandoveremploymentemp16 (last accessed on 21 October 2021).

Table 3: Weekly and Hourly Income by Worker Type, 1997-2019 (per cent change)

Growth of weekly income (%)		Growth of hourly income (%)	
(1) Employees	25.5	(5) Employees	18.7
(2) Average SE	2.1	(6) Average SE	15.6
(1) - (2) Difference	23.4	(5) - (6) Difference	3.1
(3) Solo SE	12.0	(7) Solo SE	19.0
(4) Non-solo SE	19.3	(8) Non-solo SE	37.3
(3) - (4) Difference	-7.4	(7) - (8) Difference	-18.3

Source: Data from FRS (see appendix for details).

Note: Shown are growth rates of weekly and hourly income for employees and the average self- employed person (calculated as a weighted average income growth of solo and non-solo self-employed (SE) using their respective shares in employment as weights), growth rates for solo and non-solo SE, and respective differences in percentage points. Growth rates are calculated from two-year averages.

Chart 12: Decoupling Analysis with Self-Employed Income, 1981-2019

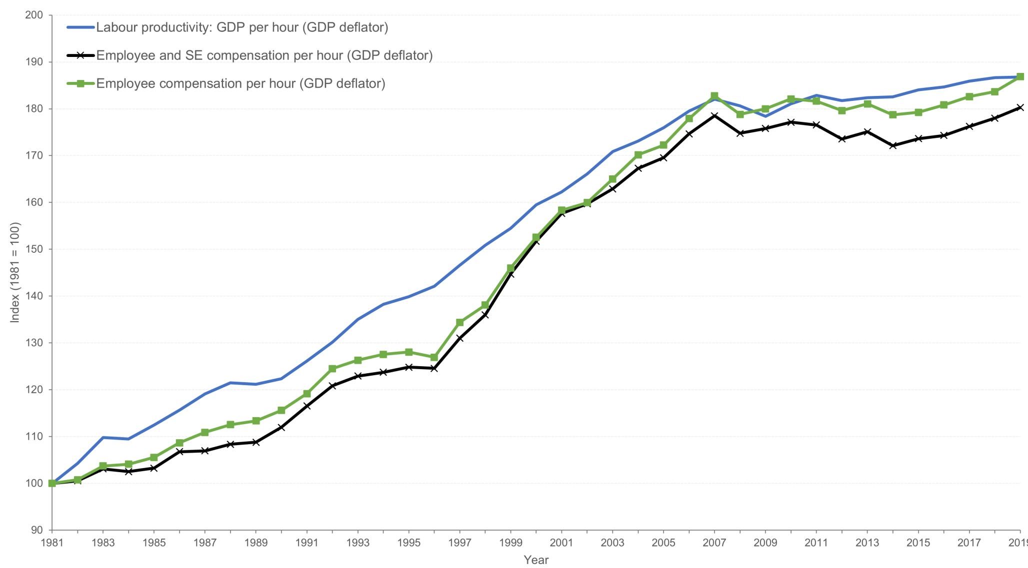

Source: Data from ONS (see appendix for details).

Note: Employee and self-employed (SE) compensation per hour is employee compensation plus mixed income that can be accrued to labour income (estimated via Dunn, Heys and Sidhu, 2018) divided by total hours worked in the economy. All series are deflated with the GDP deflator.

hour worked. Note that this is a per-hour average, so this differs from the per-worker averages we used elsewhere in this section.

Whereas both labour productivity and average employee compensation have grown by about 87 per cent between 1981 and 2019, average employee and self-employed compensation has only grown by 80 per cent. If we re-define net decoupling as the difference between labour productivity and the average employee and self-employed compensation, we obtain net decoupling of about seven percentage points. Thus, in an accounting sense, net decoupling for overall workers is entirely driven by the slower growth of self-employed compensation compared with employee compensation.

Decoupling Mechanisms

Many of the phenomenon discussed in this article are the subject of vast literatures. Increased inequality has been found to be the main reason for overall decoupling between productivity and median wages. The causes of increasing wage disparities

Chart 13: Share of Self-Employed in Total Employment for Selected OECD Countries, 2000-2019

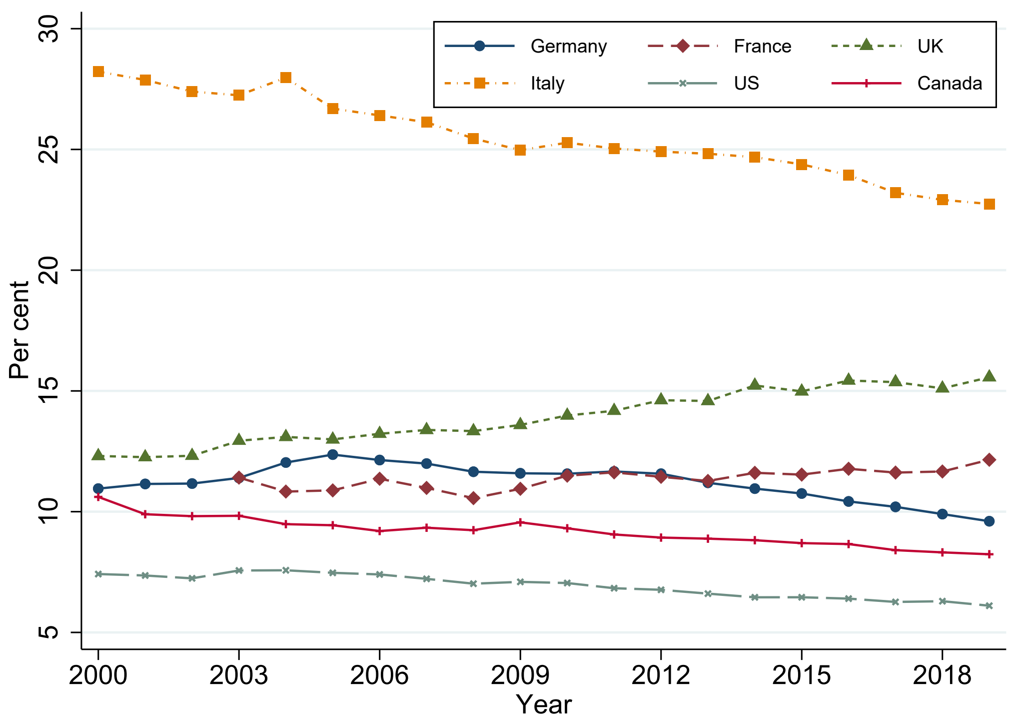

Source: Data from OECD.
Note: Shown are the shares of self-employed in total employment for Canada, Germany, France, Italy, the UK, and the US between 2000 and 2019. The series of France starts in 2003 due to limited data availability.

has been a major topic of economic research in recent decades. Technical change is one major factor (Van Reenen, 2011; Michaels, Natraj and Van Reenen, 2014) which has pushed demand ahead of the supply of skills. Trade may also play a role in reduced demand for the less skilled workers (Autor, 2019). Labour market institutions such as the decline of union power is another major factor (Machin, 2016). The fall of the labour share has also been the subject of a quickly growing literature in the last decade (e.g. Autor *et al.*, 2020; De Loecker, Obermeier and Van Reenen, 2021). Our finding that there has been some fall even in the UK puts it more in line with other countries.

Explaining the Growth of the Self-Employed in the UK

Much less is known about the causes of the changes in the trends for the self-employed. Chart 13 shows trends in self-employment rates for selected OECD countries between 2000 and 2019. Interestingly, the large increase in the share of self-employed seems unique to the UK. In Canada, Germany, Italy, and the US, self-employed shares have fallen since 2000. France has seen a slight increase in the share of self-employed, but not as much and as consistently as the UK. The fraction of solo SE in total self-employment does seem to be increasing across most countries, although it does seem particularly high in the UK (Boeri *et al.*, 2020).

Chart 13 suggests that some UK-specific factors must help explain the increase in self-employment. One factor could be changes in taxation. Evidence by Parker and Robson (2004), Smith *et al.* (2021), and Garin, Jackson, and Koustas (2021) suggests that tax incentives have a major impact on various decisions of the self-employed. The increased tax burden on employees and employers since the mid-1990s could have been a reason for the increase in self-employment. For example, employers' National Insurance contributions were increased substantially in the 2000s.[31] Adam and Miller (2021) argue that in the UK, lower tax rates for self-employed compared with employees (especially through lower national insurance contributions) incentivise people to become self-employed.[32]

A second reason could be related to self-employment as an alternative to unemployment. Giupponi and Xu (2020) call solo self-employment a "fall-back option" for many people and argue that the rise in solo self-employment puts downward pressure on employee wages. UK welfare benefits have become less generous in real terms since 1981 and the strictness of receiving working age benefits such as Job Seekers' Allowance and disability benefits has toughened (e.g. Blundell *et al.*, 2004; Koenig *et al.*, 2019). This may have pushed more non-workers into self-employment, helping deliver the very high employment rates in the UK, even after the Great Recession. Giupponi and Xu (2020) show that solo SE are the group with the highest share of people wishing to work more hours, suggesting that this group is "underemployed".[33] Additionally, Henley (2021) shows that becoming self-employed in the UK is positively associated with performing bad quality jobs (e.g. long hours, low pay, temporary contract) in prior years.

A third factor could be related to regulation. On the one hand, there has been increased regulation of labour contracts often related to EU rules, which could have reduced demand in the formal sector relative to the self-employed. The UK has a relatively liberal labour market compared to other European countries, a large outsourcing industry and thriving "gig economy" with flexible work arrangements. These push-and-pull factors may have helped the growth of self-employment.

It is important to note though that the gig economy only makes up a small share of self- employed workers (e.g. Boeri *et al.*, 2020 estimate that gig workers only make up 7 per cent of total UK self-employment).

---

31 The increase amounted to 36 per cent between 2002 and 2006.

32 They look at the example of a person on gross earnings of £40,000. According to their calculations, the tax of such an employee is £3,300 higher than that of a self-employed person on an equivalent amount. Large parts of the divergence arise due to differences in national insurance (NI) contributions. Including employers' contributions, employees made substantially higher NI contributions.

33 Among non-solo SE, a substantial share (about 17 per cent) wish to work less hours for less pay. Thus, the substantial decrease in average hours worked by both self-employed groups could be in the interest of many non- solo self-employed, but not that much for solo self-employed.

With rising demand for food and grocery delivery services as well as transport providers like Uber, we expect the importance of the gig economy to rise in the future though.

Overall, rather little is known or understood for the pattern of employment and income trends of the self-employed and why they are so different in the UK. We see this as an important avenue for future research.

Conclusions

We have analysed the "decoupling" of aggregate productivity and pay growth in the UK between 1981 and 2019. Real GDP per hour rose by 87 per cent over this period and employee hourly compensation increased by almost exactly the same amount. Consequently, there was no "net decoupling" in the UK, a result that stands in stark contrast to the US where average compensation grew much more slowly than productivity (see Chart 1, Bivens and Mishel, 2015; Mishel and Bivens, 2021; Stansbury and Summers, 2018).

This abstracts from two important factors. First, median employee wages have grown much more slowly than productivity, so in this sense there has been an overall decoupling. About 60 per cent of this decoupling is due to the growth of wage inequality and about 30 per cent is due to an increase in the share of non-wage benefits (in particular employer pension contributions) in overall compensation.

The second important factor is the big divergence in the fortunes of employees compared to the self-employed. Income growth of the self-employed has been substantially lower than that of employees. Using micro-data from the Family Resources Survey over the last two decades, our analysis suggested that the growth of the solo self-employed has been a major factor. The solo self-employed earn substantially less on average than non-solo self-employed and their hourly income growth has been slower. Since their share in total self-employment increased by 12 percentage points, this compositional shift drags down self-employed income growth. A second factor is the sharp reduction in average hours worked by both solo and non-solo self-employed. Some of this may be a welcome choice to take more leisure, but there is also evidence that many solo self-employed would like to work more hours, so it is a less welcome sign of under-employment.

Since the fraction of workers who are self-employed has risen by about six percentage points over the last four decades this has macro-economic consequences. If the compensation of the self-employed had grown at the same rates as that of employees, there would be no fall in the labour share of GDP. Including the estimated labour compensation of the self-employed and dropping non-wage compensation (as a big fraction of this is going to re-finance the pensions of already retired workers) implies a reduction in the labour share in GDP by 3.5 percentage points between 1981 and 2019. The UK may be less dissimilar to other countries like the US than it would initially seem.

Stepping back, the most striking feature of the UK economy is the dismal productivity performance since the Global Financial Crisis. Productivity has stagnated and worker pay has followed suit. Returning to sustainable income growth requires generating much better productivity growth (Van Reenen, 2021).

References

Adam, Stuart and Helen Miller (2021) "Taxing Work and Investment across Legal Forms: Pathways to Well-Designed Taxes," Institute for Fiscal Studies.

Adrjan, Pawel and Brian Bell (2018) "Pension Shocks and Wages," CEP Discussion Paper No 1536.

Autor, David (2019) "Work of the Past, Work of the Future," *AEA Papers and Proceedings*, Vol 109, pp. 1-32

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen (2020) "The Fall of the Labor Share and the Rise of Superstar Firms," pp. *Quarterly Journal of Economics*, Vol 135, No 3, pp. 645–709.

Azmat, Ghazala, Alan Manning, and John Van Reenen (2012) "Privatization and the Decline of Labour's Share: International Evidence from Network Industries," *Economica*, Vol 79, No 315, pp. 470–492.

Bagaria, Nitika, Dawn Holland, and John Van Reenen (2012) "Fiscal Consolidation during a Depression," *National Institute Economic Review*, No 221, F42–F54.

Bell, Brian (2015) "Wage Stagnation and The legacy Costs of Employment", Centre for Economic Performance, *CentrePiece Magazine* Autumn.https://cep.lse.ac.uk/pubs/download/cp458.pdf

Besley, Timothy, Peter Lambert, Isabelle Roland and John Van Reenen (2020) "The Aggregate Effects of Credit Market Frictions: Evidence from Firm-level Default Assessments", CEP Discussion Paper 1672.

Bivens, Josh and Lawrence Mishel (2015) "Understanding the Historic Divergence between Productivity and a Typical Worker's Pay: Why it Matters and Why it's Real," Economic Policy Institute.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb (2020) "Are Ideas Getting Harder to Find?" *American Economic Review*, Vol 110, No 4, pp. 1104–44.

Blundell, Jack (2019) "Clusters in UK Self-Employment," CEP Occasional Paper No 48.

Blundell, Richard, Claire Crawford, and Wenchao Jin (2014) "What Can Wages and Employment Tell us About the UK's Productivity Puzzle?" *Economic Journal*, Vol 124, No 576, pp. 377–407.

Blundell, Richard, Monica Costa Dias, Costas Meghir, and John Van Reenen (2004) "Evaluating the Employment Impact of a Mandatory Job Search Program," *Journal of the European Economic Association*, Vol 2, No 4, pp. 569–606.

Boeri, Tito, Giulia Giupponi, Alan B Krueger, and Stephen Machin (2020) "Solo Self- Employment and Alternative Work Arrangements: A Cross-Country Perspective on the Changing Composition of Jobs," *Journal of Economic Perspectives*, Vol 34, No 1, pp. 170–95.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2021) "The Productivity J- curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, Vol 13, No 1, pp. 333–72.

Cribb, Jonathan, Helen Miller, and Thomas Pope (2019) "Who Are Business Owners and What Are They Doing?" Technical report, The Institute for Fiscal Studies.

Cribb, Jonathan and Xiaowei Xu (2020) "Going Solo: How Starting Solo Self- employment Affects Incomes and Well-Being,"Technical report, IFS Working Papers.

Datta, Nikhil, Giulia Giupponi, and Stephen Machin (2019) "Zero-Hours Contracts and Labour Market Policy," *Economic Policy*, Vol 34, No 99, pp. 369–427.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (2020) "The Rise of Market Power and the Macroeconomic Implications," *Quarterly Journal of Economics*, Vol 135, No 2, pp. 561–644.

De Loecker, Jan, Tim Obermeier, and John Van Reenen (2021) "Firms and Inequality" Technical report, IFS Deaton Inequality Review, LSE mimeo

Dunn, Marianthi, Richard Heys and Sunny Sidhu (2018) "Estimating the Impact of the Self-Employed in the Labour Share", London: Office for National Statistics (ONS). https://www.ons.gov.uk/economy/economicoutputandproductivity/productivitymeasures/methodologies/estimatingtheimpactoftheselfemployedinthelabourshare.

Elsby, Michael WL, Bart Hobijn, and Ayşegül Sahin (2013) "The Decline of the US Labor Share," *Brookings Papers on Economic Activity*, Fall, No 2, 1–63.

Garin, Andrew, Emilie Jackson, and Dimitri Koustas (2021) "New Work or Changes in Reporting? A Framework for Measuring Self-Employment Trends in Administrative Data", Working paper, London School of Economics.

Giupponi, Giulia and Xiaowei Xu (2020) "What Does the Rise of Self-Employment Tell Us about the UK Labour Market?" London: Institute of Fiscal Studies.

Gollin, Douglas (2002) “Getting Income Shares Right,” *Journal of Political Economy*, Vol 110, No 2, pp.458–474.

Goodridge, Peter, Jonathan Haskel, and Gavin Wallis (2018) “Accounting for the UK Productivity Puzzle: A Decomposition and Predictions,” *Economica*, Vol 85, No 339, pp. 581–605.

Greenspon Jacob, Anna Stansburry and Lawrence Summers (2021)"Productivity and Pay in the United States and Canada", *International Productivity Monitor*, No 41, Fall, pp. 3-30

Grossman, Gene and Ezra Oberfield (2021) “The Elusive Explanation for the Declining Labor Share” NBER Working Paper 29165.

Gutiérrez, Germán and Sophie Piton (2020) “Revisiting the Global Decline of the (Non-Housing) Labor Share,” *American Economic Review: Insights*, Vol 2, No 3, pp. 321–38.

Henley, Andrew (2021) “The Rise of Self-Employment in the UK: Entrepreneurial Transmission or Declining Job Quality?” *Cambridge Journal of Economics*, Vol 45, No 3, pp. 457–486.

Kaldor, Nicholas (1957) “A Model of Economic Growth,” *The Economic Journal*, Vol 67, No 268, pp. 591–624.

Karabarbounis, Loukas and Brent Neiman (2014) “The Global Decline of the Labor Share,” T*he Quarterly Journal of Economics*, Vol 129, No 1, pp. 61–103.

Koenig, Felix, Barbara Petrongolo, John Van Reenen, and Nitika Bagaria (2019) “Can Helping the Sick Hurt the Able? Incentives, Information and Disruption in a Welfare Reform,” *Economic Journal*, Vol 129, No 624, pp. 3189–3218.

Levine, Ross and Yona Rubinstein (2017) “Smart and Illicit: Who Becomes an Entrepreneur and Do They Earn More?” *Quarterly Journal of Economics*, Vol 132, No 2, pp. 963–1018.

Machin, Stephen (2016) "Rising Wage Inequality, Real Wage Stagnation and Unions," in Lorenzo Cappellari, Solomon W. Polachek and Konstantinos Tatsiramos (eds.), *Inequality: Causes and Consequences*, Emerald Publishing Limited, Vol 43, pp. 329-354

Michaels, Guy, Ashwini Natraj, and John Van Reenen (2014) “Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years,” *Review of Economics and Statistics*, Vol 96, No 1, pp. 60–77.

Mishel, Lawrence and Josh Bivens (2021) "The Productivity-Median Compensation Gap in the United States: The Contribution of Increased Wage Inequality and the Role of Policy Choices", *International Productivity Monitor* Number 41, Fall.

Nolan, Brian, Max Roser, and Stefan Thewissen (2019) “GDP per Capita versus Median Household Income: What Gives Rise to the Divergence over Time and How Does This Vary Across OECD Countries?” *Review of Income and Wealth*, Vol 65, No 3, pp. 465–494.

OECD (2018) “Decoupling of Wages and Productivity: What Implications for Public Policy?” *OECD Economic Outlook*, Chapter 2, No 2, pp. 51–65.

Oulton, Nicholas (2019) “The UK and Western Productivity Puzzle: Does Arthur Lewis Hold the Key?” *International Productivity Monitor*, No 36, pp. 110-141

Parker, Simon C and Martin T Robson (2004) “Explaining International Variations in Self-Employment: Evidence from a Panel of OECD Countries,” *Southern Economic Journal*, pp. 287–301.

Patterson, Christina, Ayşegül Sahin, Giorgio Topa, and Giovanni L Violante (2016) “Working Hard in the Wrong Place: A Mismatch-Based Explanation to the UK Productivity Puzzle,” *European Economic Review*, Vol 84, pp. 42–56.

Pessoa, João Paulo and John Van Reenen (2013) “Decoupling of Wage Growth and Productivity Growth? Myth and Reality,” *CEP Discussion Paper* No 1246.

Pessoa, João Paulo and John Van Reenen (2014) “The UK Productivity and Jobs Puzzle: Does the Answer Lie in Wage Flexibility?” *Economic Journal*, Vol 124, No 576, 433–452.

Philippon, Thomas (2019) *The Great Reversal*, (Cambridge, MA: Harvard University Press.)

Piketty, Thomas (2014) *Capital in the Twenty-First Century*, (Cambridge, MA: Harvard University Press).

Rognlie, Matthew (2015) “Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?” *Brookings Papers on Economic Activity*, Spring No 1, pp. 1– 69.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick (2019), “Capitalists in the Twenty-First Century,” *Quarterly Journal of Economics*, Vol 134, No 4, pp. 1675–1745.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick (2021), "The Rise of Pass-Throughs and the Decline of the Labor Share," University of Chicago, mimeo.

Stansbury, Anna and Lawrence Summers (2018) "Productivity and Pay: Is the Link Broken?" Technical report, Peterson Institute for International Economics Working Paper 18-5.

Syverson, Chad (2017) "Challenges to Mismeasurement Explanations for the US Productivity Slowdown" *Journal of Economic Perspectives*, Vol 31, No 2, pp. 165-86.

Valero, Anna and Van Reenen, John (2019) "The UK Economy: Policies for Investment and Growth" CEP Election Analysis No. 49.

Van Reenen, John (2011) "Wage Inequality, Technology and Trade: 21st Century Evidence," *Labour Economics*, Vol 18, No 6, pp. 730–741.

Van Reenen, John (2021), "Going for Growth that's Sustainable and Equitable," https://cep.lse.ac.uk/pubs/download/cp604.pdf, Centre for Economic Performance, *CentrePiece Magazine* Summer .

Whittaker, Matthew (2019) "Follow the Money: Exploring the Link between UK Growth and Workers' Pay Packets,"Technical report, Resolution Foundation.

Williams, David M (2021) "Pay and Productivity in Canada: Growing Together, Only Slower than Ever," *International Productivity Monitor*, No 40, Spring, pp. 3-26

The Productivity-Median Compensation Gap in the United States: The Contribution of Increased Wage Inequality and the Role of Policy Choices

Lawrence Mishel and Josh Bivens[1]
Economic Policy Institute

Abstract

This article offers a narrative and supporting evidence on mechanisms that suppressed wage growth and generated a divergence of 43 percentage points (1.05 points per year) between net productivity and median hourly compensation growth between 1979 and 2017 in the United States. These dynamics reflect the strengthening of employers' power relative to white-collar and blue-collar workers. We offer empirical assessments of the impact of particular factors on wage growth and wage inequality. The three factors with the largest and best measurement impacts, i.e., excessive unemployment, eroded collective bargaining, and corporate-driven globalization — explain 55 per cent of the divergence. Other factors — a diminished overtime salary threshold, employee misclassification, employer-imposed non-compete agreements, and corporate fissuring-subcontracting and major-buyer dominance — explain another 20 per cent. Together, these policy-related factors can account for three-fourths of the 1979-2017 divergence between productivity and median hourly compensation growth.

Inequalities abound in the U.S. economy. A central driver of these inequalities in recent decades has been the widening gap between the hourly compensation of a typical (median) worker and labour productivity—the income generated per hour of

1 Lawrence Mishel is Distinguished Fellow at the Economic Policy Institue (EPI) and served as EPI President from 2002 to 2017. Josh Bivens is Director of Research at the Economic Policy Institute. This article is an abridged version of Mishel and Bivens (2021a). Emails: lmishel@epi.org, jbivens@epi.org. We gratefully acknowledge the assistance of people who provided information or commented on particular issues: Eileen Appelbaum, Danny Blanchflower, David Cooper, Daniel Costa, Nick Hanauer, Thea Lee, Thomas Lemieux, Mike Lipsitz, Celine McNicholas, Suresh Naidu, Heidi Shierholz, John Schmitt, Andrew Sharpe, Anna Stansbury, Evan Starr, Marshall Steinbaum, David Weil, Nathan Wilmers, and Ben Zipperer. We also thank three anonymous referees for their comments. Melat Kassa provided essential, excellent research assistance. We also are appreciative of the funding of the Economic Policy Institute's Unequal Bargaining Power initiative by the Nick and Leslie Hanauer Foundation, the William Flora Hewlett Foundation, and the Bernard and Anne Spitzer Charitable Trust.

work. This growing divergence has been driven by two other widening gaps, that between the compensation received by the vast majority of workers and those at the top of the wage distribution, and that between labour's and capital's share of income. This article presents an updated account of the United States productivity-pay divergence (Mishel and Gee 2012) and evidence that the divorce between the growth of median compensation and productivity has been generated primarily through intentional policy decisions designed to suppress typical workers' wage growth, namely the failure to improve and update existing policies, and the failure to thwart new corporate practices and structures aimed at wage suppression. Inequality will stop rising, and paychecks for typical workers will start increasing robustly in line with productivity, only when we enforce labour standards and embrace policies that reestablish individual and collective bargaining power for workers.

Between 1979 and 2017, the compensation of median workers trailed economy-wide (net) productivity growth by roughly 43 percentage points, or 1.05 percentage points per year. The effects have been felt broadly: During this time 90 per cent of U.S. workers experienced wage growth slower than the economy-wide average (0.99 per cent), while workers at the top (mostly highly credentialed professionals and corporate managers) and owners of capital reaped large rewards made possible only by this anemic wage growth for the bottom 90 per cent.

Sluggish median wage growth is not a political secret; it has been widely recognized across the political spectrum, even cited by both the Republican and Democratic Party platforms in 2016.[2] The root causes of the trend have frequently been misidentified, however. One prominent interpretation is that disappointing wage growth is an unfortunate result of apolitical market forces that one neither can nor would want to alter. Since labour markets are generally competitive and workers and employers have roughly balanced degrees of market power, this argument assumes, fundamental apolitical forces like technological change and automation, as well as globalization, have mechanically shifted demand away from non-college-educated and middle-wage workers. But the premier research cited in support of a competitive market-based explanation — predominantly focused on automation/technological change since the impact of globalization is frequently (though wrongly) considered to be minimal— has itself actually offered empirical metrics that demonstrate that automation/technological change fails to explain wage trends and wage inequality, especially in the period since 1995 (Mishel and Bivens

2 The Republican Party platform reads: "Our economy has become unnecessarily weak with stagnant wages. People living paycheck to paycheck are struggling, sacrificing, and suffering." The Democratic platform reads: "But too many Americans have been left out and left behind. They are working longer hours with less security. Wages have barely budged and the racial wealth gap remains wide, while the cost of everything from childcare to a college education has continued to rise."See https://www.presidency.ucsb.edu/documents/2016-republican-party-platform, and https://www.presidency.ucsb.edu/documents/2016-democratic-party-platform

2021b).

We need to look further for more convincing empirical explanations of why, during a period of rising productivity, hourly compensation for the bottom 90 per cent of all workers has risen so slowly in spite of overall income growth. Doing so requires explaining the key dynamics. The growing wedge between rising productivity and compensation growth for the typical worker financed the increased share of compensation going to top earners, especially those in the top 1 per cent and 0.1 per cent, along with a declining share of income going to labour. In addition, over the last four decades there has been a persistent disparity in the growth of earnings between those in the 90–99 per cent range and those in the middle. Further, wage disparities by gender, race, and ethnicity from the late 1970s, reflecting institutionalized gender and race discrimination, remain with us and have sometimes even worsened. Any accounting of where we are and what policies we need must address these issues.

This article offers a narrative and supporting evidence on the mechanisms that have suppressed wage growth since the late 1970s. We refer in this analysis to wage suppression rather than wage stagnation because it was an actively sought outcome — engineered by policymakers who invited and enabled capital owners and business managers to assault the leverage and bargaining power of typical workers, with the inevitable result that those at the top claim a larger share of income.

Six factors can collectively explain most of the growth of wage inequality and the erosion of labour's share that resulted in wage suppression over the last four decades (specifically 1979–2017):

- Austerity macroeconomics, including facilitating unemployment higher than it needed to be to keep inflation in check, and responding to recessions with insufficient force;
- Corporate-driven globalization, resulting from policy choices, largely at the behest of multinational corporations, that undercut wages and job security of non-college-educated workers while protecting profits and the pay of business managers and professionals;
- Purposely eroded collective bargaining, resulting from judicial decisions, and policy choices that invited ever more aggressive anti-union business practices;
- Weaker labour standards, including a declining minimum wage, eroded overtime protections, nonenforcement against instances of "wage theft," or discrimination based on gender, race, and/or ethnicity;
- New employer-imposed contract terms, such as agreements not to compete after leaving employment and to submit to forced private and individualized arbitration of grievances; and
- Shifts in corporate structures, resulting from fissuring (or domestic outsourcing), industry deregulation, privatization, buyer dominance affecting entire supply chains, and increases in the concentration of employers.

Concretely, our analysis attempts to account for the 43 percentage point (1.05 points per year) divergence between the growth of labour productivity (net of depreciation) and median hourly compensation (wage and benefit) growth between 1979 and 2017. This wedge excludes any impact of the differing measures of prices used to inflation-adjust productivity and compensation growth. Had median hourly compensation grown with net productivity it would have increased from $20.48 in 1979 to $33.10 in 2017 (2019$). In fact, median hourly compensation was $23.15 in 2017, a $9.95 shortfall from the net productivity benchmark.

We estimate that the first three factors—the impacts that are largest and best measured, i.e., excessive unemployment, eroded collective bargaining, and corporate-driven globalization—explain 55 per cent of the divergence between growth in productivity and median hourly compensation, and specific other factors included above a diminished overtime salary threshold, employee misclassification, employer-imposed noncompete agreements, and corporate fissuring-subcontracting and major-buyer dominance—explain another 20 per cent. Together, the factors [3] for which we have been able to assess their impact on the median wage can account for three-fourths of the divergence between productivity and median hourly compensation growth from 1979 to 2017.[4]

This article's analysis complements and points in the same direction as other recent research that has focused attention on worker power. For instance, Stansbury and Summers (2020) also argue that reduced worker power explains sluggish wage growth and a declining labour share of income. New empirical examinations of employer monopsony power have identified a growing (at least since the late 1990s) and pervasive employer ability to mark down wages from 20 per cent to 50 per cent and to exert more power over low-wage workers than others. This new monopsony literature provides a top-down empirical approach, estimating the aggregate potential employer power to suppress wages and then examining the contributing role of countervailing forces like unionization, high-pressure labour markets, and high values of minimum wages in explaining an aggregate net metric of employer power.In contrast, we provide a bottom-up empirical approach examining the impact of many specific factors and gauging their contribution to the overall divergence between productivity and median compensation growth.

Our research and other recent findings demonstrate that employer power is ubiquitous in labour markets, and that wages will be lower and wage growth suppressed absent institutions and policies that provide countervailing power.[5] In other words,

3 Other factors that we have not been able to empirically assess—increased wage theft and weak enforcement, anti-poaching agreements, increased discrimination, forced arbitration agreements, guestworker programs, and increased prevalence of employer-created "lawless zones" in the labour market where workers are deprived of effective labour protections because of their immigration status—have also contributed to wage suppression.

4 The growth of the wage inequality in the bottom half, the 50/10 wage gap, has been shown to result primarily from excessive unemployment and the deterioration of the minimum wage (Mishel and Bivens,2021a).

5 Joseph Stiglitz has long focused on power in markets, emphasizing both product market monopoly power and the weakening of employee power relative to employers. He recently provided an analysis similar to the framework of this article (Stiglitz,2021).

employer power is a constant of modern labour markets, but what has changed over the past generation or two is the erosion of institutions and policies — high-pressure labour markets, robust enforcement, unions, and meaningful minimum wages — that once provided that countervailing power.

The article proceeds as follows. The first step is to examine the wage trends that any explanation of wage suppression needs to explain and elaborating and analyzing the productivity-median hourly compensation divergence we seek to explain. The second section briefly assesses the conventional explanation of "skill-biased technological change"—namely, rapid automation workers while many lack the skills necessary for modern production systems. The third section identifies the six factors, from excessive unemployment and eroded collective bargaining to shifts in corporate structures, that we believe much better explain wage suppression. The fourth section reviews how this article fits into the overall literature on wage inequality and draws on the estimated impact of the various factors to establish how much they explain the overall divergence between productivity and median hourly compensation, The fifth and final section concludes.

Wage Trends and Patterns to be Explained

There are three disparities in growth of wages by workers' wage rankings that policymakers need to understand and economists need to explain: the one between the highest earners (the top 1 per cent and top 0.1 per cent) and other high-wage earners; the one between high-wage and middle-wage earners (the 95/50 or the 90/50 wage gaps); and the one between middle- and low-wage earners (the 50/10 wage gap). In addition, a theory about wage trends needs to explain the decline in the share of overall income accruing to labour, since this drop saps wage growth.

The rough summary of inflation-adjusted wage growth is as follows. Between 1979 and 2019 (the end of the last business cycle), inflation-adjusted annual wages at the very top have grown tremendously (Mishel and Kandra, 2020). Those in the top 1 per cent enjoyed 160 per cent growth (Chart 1), and those at the very top—the top 0.1 per cent—experienced growth of 345 per cent. Growth was much slower at the 95th percentile 63 per cent (using hourly wage data), slower still at the 50th (15 per cent), and a snail's pace at the 10th (3 per cent) though it is worth noting that growth rates at the middle and the bottom were not remarkably different since the late 1980s.[6]

Two key wage gaps have grown since the late 1980s: the one between the top and very top on the one hand and all other earners, including even those at the 95th percentile, on the other, and the gap between high earners and middle earners, illustrated by the ratio of wages at the 95th (or 90th) percentile and the median wage.

6 Economic Policy Institute's State of Working America Data Library. Wages by percentile and wage ratios https://www.epi.org/data/?subject=wage-percentil.

Chart 1: Cumulative Per Cent Change in Real Annual Wages, by Wage Group, 1979–2019

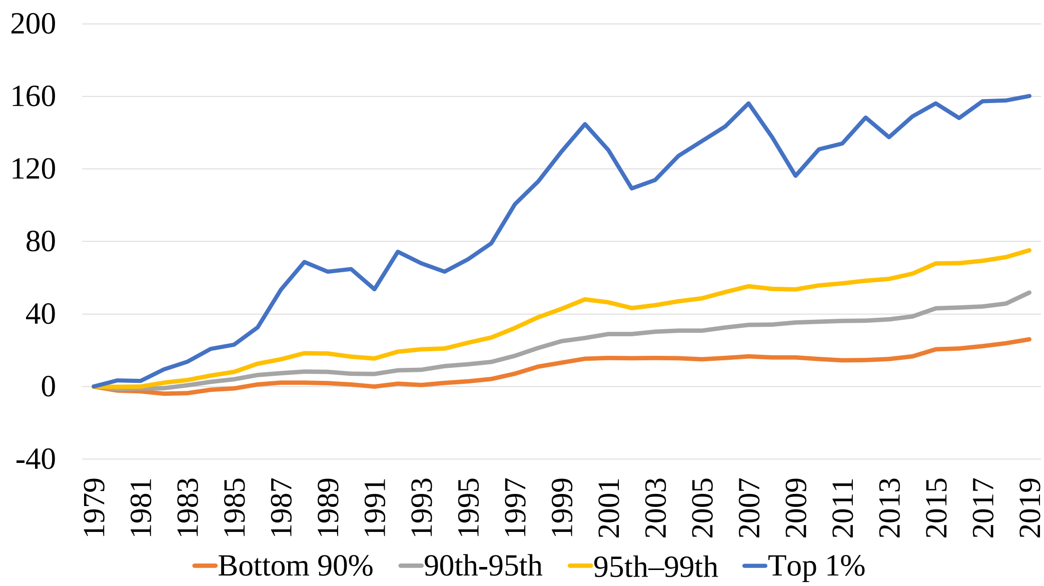

Source: EPI analysis of Kopczuk, Saez, and Song (2007, Table A3) and Social Security Administration wage statistics

The gap between the median earner and low-wage earners (50/10 wage gap) grew in the 1980s but has been stable since the 1987-88 (see Figures B and C in Mishel and Bivens 2021a).

Decomposing the gap between productivity and median hourly compensation

The last four decades have seen a systematic divergence between the growth of economy-wide productivity (the amount of income generated in an average hour of work) and the growth of hourly compensation (wages and benefits) for typical workers. We proxy the wages of "typical" workers as either wages for nonsupervisory workers (roughly 80 per cent of the private-sector workforce) or wages for the worker earning the median wage. Chart 2 shows the growth of economy-wide productivity (net of depreciation, measured with consumer prices) and the typical worker's hourly compensation since 1948,It uses the hourly compensation of private production-nonsupervisory workers because that is the only series available for the entire period since 1948. While productivity and a typical workers' compensation grew in tandem over the 1948-1979 period, they diverged thereafter, splitting entirely after 1979. In the latter period productivity growth decelerated significantly, but much more rapid deceleration (or even stagnation) occurred in growth of a typical worker's compensation. Net productivity grew 118.4 per cent from 1948 to 1979, accompanied by 107.5 per cent growth in a typical worker's compensation between 1979 and 2019 net productivity grew 59.7 per cent (1.18 per cent annually) further, but a typical worker's compensation (wages and benefits) grew only by 13.7 per cent (0.38 per cent annually).

This divergence was first pointed out

Chart 2: Gap Between labour Productivity and a Typical Worker's Compensation, 1948-2019

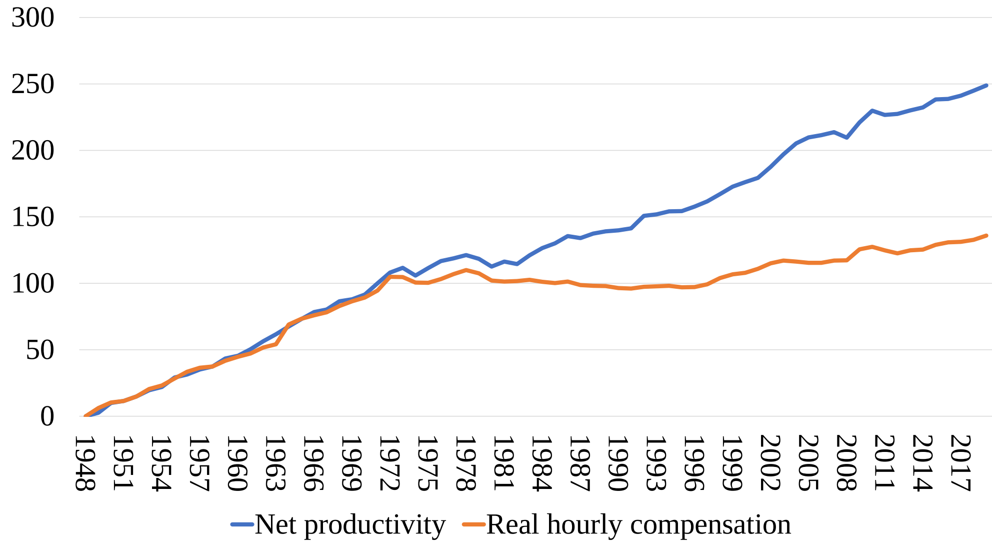

Source:

in the early 1990s (Mishel and Bernstein, 1994) to demonstrate that stagnant wages for the typical worker over the previous decade or so could not be explained solely by the slowdown of productivity growth. This section updates the Bivens and Mishel (2015) analyses of the wedges between typical workers' pay and productivity and the decomposition of the main factors generating it, drawing on previous work [7]

The top panel in Table 1 provides the basic trends required to decompose the divergence. The output data cover all sectors while the compensation and wage data are for all wage and salary workers (i.e., excludes self-employed). We focus on the gap between net productivity (productivity net of capital depreciation, which is a better metric than gross productivity for our purposes) and median hourly compensation (wages and benefits, in line 6) but provide data on median wages (line 5) and gross productivity (line 1) for completeness.[8] We decompose the divergence between net productivity (measured as the statistical agencies do, at output prices, shown in line 2) and median hourly compensation. This allows us to compute the contribution of the differing inflation rates in output (i.e. Net National Product) than in consumer goods and services as a factor. The definitions of the variables in Table 1 and the data sources are found in Appendix A.

Over the 1979-2019 period, net productivity grew 1.36 per cent annually while median hourly compensation grew just 0.38

7 See Mishel and Gee (2012) and the decomposition framework developed by the Centre for the Study of Living Standards: Sharpe, Arsenault, and Harrison 2008a; Sharpe, Arsenault, and Harrison 2008b; and Harrison 2009).

8 Net product encompasses the income that is distrusted to households.

Table 1: Contributions to Gap Between Median Hourly Compensation and Productivity Growth, 1979-2019 (Compound Annual Rate of Change)

A. Basic trends	**1973–1979**	**1979–1995**	**1995–2000**	**2000–2007**	**2007–2019**	**2000–2019**	**1979–2019**	**1979–2017**
1. Gross productivity	1.06	1.38	2.33	2.19	1.11	1.50	1.56	1.58
2. Net productivity (producer prices)	0.92	1.19	2.13	1.94	0.94	1.31	1.36	1.38
3. Net productivity (consumer prices)	0.53	0.96	1.65	1.84	0.89	1.24	1.18	1.18
4. Average hourly compensation	0.81	0.82	2.18	1.19	0.62	0.83	0.99	0.99
5. Median wage	-0.47	-0.04	1.41	0.47	0.37	0.41	0.35	0.30
6. Median compensation	0.13	0.07	1.09	0.65	0.33	0.45	0.38	0.33
7. Gross productivity–median compensation gap (1-6)	0.94	1.30	1.24	1.53	0.78	1.06	1.18	1.24
8. Net productivity–median compensation gap (2-6)	0.80	1.12	1.04	1.29	0.62	0.86	0.99	1.05
B. Explanatory factors for net productivity gap (annual growth)								
9. Inequality of compensation (4-6)	0.69	0.74	1.09	0.54	0.29	0.38	0.61	0.65
10. Loss in labor's share of income (3-4)	-0.28	0.14	-0.53	0.65	0.27	0.41	0.19	0.19
11. Divergence of consumer and output prices (3-2)	0.39	0.24	0.48	0.10	0.05	0.07	0.19	0.20
C. Relative contribution to net productivity gap (percent of gap)								
12. Inequality of compensation	86.2	66.2	105.0	41.9	47.6	44.5	62.2	62.6
13. Loss in labor's share of income	-35.2	12.7	-51.5	50.41	44.0	47.5	18.8	18.5
14. All inequality (12 + 13)	51.0	78.9	53.5	92.3	91.6	92.0	81.0	81.1
15. Divergence of consumer and output prices	49.0	21.1	46.5	7.7	8.4	8.0	19.0	18.9
Total	100.0	100.0	100.0	100.00	100.0	100.0	100.0	100.0

Source: Update of Bivens and Mishel (2015) Table 1. See Appendix A for data definitions and sources.

per cent annually, a sizeable divergence of 0.92 percentage points each year. The gap was primarily driven by factors associated with growing inequality—the decline of labour's share (since 2000) and growing inequality of compensation, as shown in rows 9 and 10. These inequality factors can explain 81 per cent of the growth of the productivity-median hourly compensation gap (row 14) with the remaining portion (19 per cent) due to difference in the growth of producer (used to measure productivity) and consumer (used to measure compensation) price growth.

In the most recent period from 2000 to 2019 the inequality factors can explain 92 per cent of the divergence and the difference in deflators was much less important (8 per cent).The net productivity-median hourly compensation divergence, exclusive of the price deflator differences (line 3-6), reflects the rising inequalities we seek to explain.

We focus on the 1979-2017 period because that best fits the various studies information available on the impact of various factors. Over 1979-2017 net productivity grew 56.2 per cent (1.18 per cent per year) while median hourly compensation grew 13.5 per cent,0.33 per cent per year a gap of roughly 43 percentage points or 1.05 percentage points per year.

We rely on this divergence of 43 percentage points to measure the extent of wage suppression. This divergence simply reflects the impact on the median hourly compensation of the growth of wage and benefit inequality and the loss of labour's share and thus measures the counterfactual of how much faster median hourly compensation could have grown had inequality not grown. A situation where productivity and median hourly compensation progress at the same pace is not offered as a description of how we expect the economy to work nor as a normative statement. Our analysis reviews the impact of particular factors on the growth of median wages (assumed to correspond to median compensation growth).

In the final section we examine the extent to which the cumulative impact of particular factors to gauge whether their cumulative effect can explain the divergence between net productivity and median hourly compensation growth, i.e. wage growth suppression

The decline in labour's share of income

One of the trends that alerted analysts to the erosion of worker bargaining power and the corresponding strengthening of employer bargaining power has been the erosion of labour's share of income in the 2000s. The distributional conflict between workers and employers (or capital and labour shares) is best examined in the corporate sector, where all income is divided between compensation going to workers and income accruing to owners of capital and avoids issues of having to decide whether "proprietor's income," or noncorporate businesses income is labour or capital (see Bivens 2019 for measurement details).

Labour's share fell from 82.4 per cent in 2000 to 77.9 per cent in 2007, the last year before the Great Recession. By 2016, when unemployment had reached levels comparable to what had prevailed in 2006 and 2007, labour's share remained roughly 2.5 percentage points below its 2007 level. The fall in labour's share from 82.4 per cent in 2000 to 75.5 per cent in 2016 is the equivalent of an 8.4 per cent across-the-board cut in compensation for every employee; equivalently, it would require an across-the-board compensation boost of 9.1 per cent to restore labour's share to its 2000 level. This shift toward greater capital income and returns is even more impressive given that real interest rates have fallen sharply in recent years, a development that should (all else equal) be accompanied by a lower return to capital (Farhi and Gourio,2018).

The Failure of Automation and Skill Gaps to Explain Wage Suppression or Wage Inequality

The predominant explanation offered by economists, pundits, policymakers, and the media to explain sluggish wage growth and growing inequality in the United States, at least until recently has been the skill-biased technological change hypothesis. It is assserted that a huge proportion of U.S. workers have "skills deficits," i.e., lack the skills necessary to deal with technological change. One version, focused on education wage gaps, argues that computerized automation has made more educated workers — generally referring to those with at least a four-year college degree - more valuable to employers and has correspondingly reduced the value of those without a college degree (Katz and Murphy, 1992; Goldin and Katz, 2007, 2008). This growing wage gap between college-educated and non-college-educated workers the college wage premium is used to explain rising wage inequality between high earners and the majority of earners who lack a four-year college credential (62 per cent of earners in 2019, down from 82 per cent in 1979.

A second version of the automation

story, frequently referred to as the "job polarization thesis," argues that technological change has increased the value of abstract reasoning, creativity, and expertise, judgment, resulting in the devaluation of skilled work done following well-understood rules and procedures (Autor, 2010).

These versions of skill-biased technological change portray the cause of wage suppression and growing wage inequality as due to a factor, automation, that is both inevitable (one can't stop technology's forward march) and desirable (after all, technological change is a key driver of rising living standards). Thus, the resulting economic adversity for some workers is the unfortunate byproduct of a dynamic that one would neither want to nor could change. Given this view, the only appropriate remedy is to adapt to automation, primarily by upgrading workers' skills and education and perhaps by providing a more adequate safety net for workers temporarily displaced.

Skill-biased technological change has always been a weak explanation for the wage trends since 1979, but is a prima facie implausible explanation for the trends since the mid-1990s or since 1999. None of the basic indicators of automation's impact and of skill deficits used to establish these narratives has been evident over the last 25 years. Consequently, there is no basis for considering automation-driven skill-biased technological change as a significant factor in wage suppression or the growth of wage inequality since the mid-1990s—and we assign it an impact of zero since 1995 in our analysis below.

Our critiques of both the skills narratives, focus particularly on their inability to explain wage trends since the mid-1990s.[9] These arguments are fully explored in Mishel and Bivens (2021b).

Omissions and Evidentiary Problems in the Skills-based Wage-gap Story

One problem with the automation narratives is that they fail to address the superlative wage growth of the top 1 per cent (and the top 0.1 per cent) and the corresponding upward shift of 6 percentage points of aggregate earnings to the top 1 per cent between 1979 and 2019 (Mishel and Kandra, 2020). The growth of wages for the top 1 per cent primarily reflects the growth of executive compensation and the expansion of the financial sector (and its high earners). Similarly, the narratives accord no attention to the erosion of labour's share of income. the data show that the link between automation and the decline of labour's share is at least as inconsistent with real-world data as is the link between automation and wage inequality (Stansbury and Summers, 2020).

9 Our discussion does not cover the 1979–1995 period specifically. However, we do not assign any impact of automation and skills gaps for this period. We remain skeptical that there was any impact, though, following the analysis in Mishel, Bernstein, and Schmitt (1997a) and Card and DiNardo (2002). Wage inequality surged in the early 1980s before there was much computer automation, for instance. Plus, much of the critique of automation's role in the period following 1995 also applies to the earlier period: there was a flat 50/10 wage gap since 1987; a spectacular growth for the top 1 per cent; and much of the growth of wage inequality was within education groups.

college wage premium

Autor, Goldin, and Katz (2020) note that "returns to a year of college rose by 6.5 log points, from 0.076 in 1980 to 0.126 in 2000 to 0.141 in 2017." Yet note the slowdown from the former period to the 2000–2017 period. In the former period, the log college wage premium rose 0.0325 percentage points each year, far faster than the 0.0088 percentage point increase each year between 2000 and 2017. This represents a 70 per cent reduction in growth. it is evident that the education wage gap has not driven wage inequality in the top half since 2000, or perhaps since mid-90s.

demand for college graduates

The substantial deceleration in the college wage premium, even as the supply of college graduates slowed, implies a dramatic slowing in the growth of relative demand for college graduates. As Autor, Goldin, and Katz (2020:5) note:

> "a puzzling slowdown in the trend demand growth for college equivalents starting in the early 1990s. Rapid and disruptive technological change from computerization, robots, and artificial intelligence is not to be found."

If automation's impact has been far less in the last 25 years than in earlier decades, it cannot explain the ongoing strong, even faster, growth of wage inequality in the top half, illustrated by the growth of the 95/50 wage gap.

wage gap in the bottom half

In the skills-gap story, the more education workers have, the more they are in demand and the higher their wages. Yet over the last three decades there have been no increases in the wage gaps between those with some college, those with a high school diploma, and those who left high school. Similarly, the wage gap between median (50th percentile) workers and low-wage (10th percentile) workers has been stable or declining since 1987 This is a long-standing critique of the education wage-gap hypothesis (Mishel, Bernstein, and Schmitt, 1997a; Card and DiNardo 2002; Acemoglu and Autor, 2012).

occupational employment polarization

Most strikingly, Autor (2010), Acemoglu and Autor (2012) and Autor (2014) document that job polarization has not been evident since 1999:

> "[G]rowth of high-skill, high-wage occupations (those associated with abstract work) decelerated markedly in the 2000s, with no relative growth in the top two deciles of the occupational skill distribution during 1999 through 2007, and only a modest recovery between 2007 and 2012. Stated plainly, the U-shaped growth of occupational employment came increasingly to resemble a downward ramp in the 2000s."(Autor, 2014:149-150)

Occupational employment patterns

The job polarization narrative relies on mapping occupational employment patterns to explain wage trends. But surprisingly, the polarization literature has never presented evidence that these occupational employment shifts directly affect wages. And indeed, they do not. Mishel, Shierholz, and Schmitt (2013) show that in the 1980s, 1990s, and 2000s, changes in occupational employment shares (whether employment in an occupation expands or contracts relative to other occupations) were not related to changes in relative wages by occupation. If occupational job polarization does not shape relative occupational wages, then it is certainly not much of an explanation for wage inequality.

A More Convincing Theory of Sluggish Wage Growth and Inequality: Policy-driven Wage Suppression

If forces unrelated to policy decisions, particularly automation, do not seem to be driving wage trends, what are the factors leading to wage suppression? Our answer is that there has been an intentional policy assault—including policy forbearance in the face of new anti-worker business practices—that diminished the institutional sources of leverage and bargaining power for typical workers in the labour market. The point was to suppress labour costs. This policy assault (acts of commission and omission, such as failing to update labour law or the value of the minimum wage) either directly undercut these institutional sources of power or accommodated employers' efforts to undercut them. Business forces were secure knowing that policymakers (legislators, executive branch officials, and judges) would not change legislation, enforcement priorities and effectiveness, or legal interpretations to countermand this assault on a typical workers' power in the labour market.

Why this policy and corporate assault began when it did, and why it was successful politically, are questions mostly outside the bounds of this article.[10] But a growing body of evidence shows that the specific policies launched in this attack can explain the overwhelming majority of wage suppression experienced in recent decades. This section identifies these policies and estimates their impact.

Austerity Macroeconomic Policy: Excessive Unemployment

The Federal Reserve Board's dual mandate is to pursue the maximum level of employment consistent with stable inflation. However, since 1979 the Fed's actions suggest that it took the inflation mandate more seriously, thereby tolerating (by failing to lower) or actually generating excessive unemployment for extended periods in the name of keeping inflation tame.

Hooper, Mishkin, and Sufi (2019: 25) note that "since the 1980s the Fed focused much more on avoiding labour market overheating in order to stabilize in-

10 See Hacker and Pierson (2011, 2020) for the political science explanations.

flation." Mishel and Bivens (2021a, Appendix B) examines the "The intentionality of macroeconomic policies", documenting the link of excessive unemployment to policy decisions.

Bivens and Zipperer (2018), analyzing the links between excess unemployment and wage growth, note that full employment (at least by the too conservative measure of matching actual unemployment to preexisting estimates of the "natural rate") was the norm after World War II but became the exception after 1979. Between 1949 and 1979, the cumulative difference between the actual unemployment rate and estimates of the unemployment rate consistent with stable inflation—the "natural rate" or the NAIRU, the nonaccelerating inflation rate of unemployment—was *negative* 15.3 percentage points, meaning that on average actual unemployment was 0.52 percentage points *below* the estimated NAIRU each year. In contrast, between 1979 and 2017 the cumulative difference was positive 35.7 percentage points or 0.94 points per year, meaning that actual unemployment was persistently above the estimated natural rate.

This consistent excess unemployment was deeply damaging to wage growth. Using the lower bound of the Bivens and Zipperer estimates to assess the impact of excessive unemployment on median and 10th percentile wages in the 1979–2017 period, we find excessive unemployment had lowered the median hourly wage by 12.2 per cent. These estimated impacts of unemployment are far below those of Katz and Krueger (1999, Table 8), whose Phillips curve estimates using a 1973–1998 time series were double those of Bivens and Zipperer at the median and three times those at the 10th percentile.[11]

However, to err on the side of caution we make an adjustment to our estimates of the wage impact of higher unemployment to account for the "flattening" of the Phillips curve in recent years (a lessening of the relationship between unemployment and wage growth): We apply one impact for the 1979–2007 years and a lesser impact for the 2008–2017 years.[12] Taking this flattening of the Phillips curve into account we find that, if unemployment over 1979–2017 had averaged just the "natural rate" of 5.5 per cent rather than 6.3 per cent, median wages would have been 10.0 per cent higher in 2017. If the unemployment rate had been held even lower, say 5.0 per cent, median wages would have been 18.3 per cent higher by 2017.

Excessive unemployment had a larger impact on low-wage workers which included a disproportionately number of Black workers, lowering the 10th percentile wage by 11.6 per cent by 2017 from 1979 levels and raising the 50/10 wage gap by 2.7 percentage points. If our analysis uses

11 Note that the regression specification in Bivens and Zipperer (2018) (following Katz and Krueger (1999)) controls for a measure of inflation on the right-hand side which makes these estimates of the impact of excessive unemployment on real wages, i.e. a real wage Phillips curve.

12 Estimates of the wage impact of unemployment on the median and 10th percentile wage are from Bivens and Zipperer (2018, Chart 6). The impact of 1 percentage point higher unemployment lowers the median wage by 0.459 and 0.296 in the earlier and latter period and lowers the 10th percentile wage by 0.582 and 0.243 in the earlier and latter period.

5.0 per cent rather than 5.5 per cent as the full employment target, then the 10th percentile wage would have been 21.2 per cent higher in 2017 absent excessive unemployment.

Economic policy could have held unemployment to the average NAIRU rather than be one percentage point above NAIRU. The Volcker disinflation of the early 1980s was a mistake. Galbraith (1997) and DeLong and Summers (1988) have argued that the excessive unemployment was very costly and the benefits of lower inflation were overstated. Likewise, there could have been a much faster recovery from the 2007-2009 downturn. If overall public spending had simply matched the average rate of (per capita) growth in the 1980s, 1990s, and 2000s the pre-recession unemployment rate would have been restored in just 5-6 years instead of the actual eight years. (Bivens 2016).

Erosion of workers' rights to form unions and bargain collectively

The erosion of collective bargaining is the second largest factor that depressed wage growth in the middle and drove wage inequality over the last four decades.

The impact has been especially adverse for men because they were far more likely to be unionized in 1979 than women (31.5 per cent versus 18.8 per cent), so men had more to lose from the subsequent attack on unions and collective bargaining.

That collective bargaining leads to more equal wage outcomes was firmly established by research by Richard Freeman and James Medoff in the late 1970s and popularized in their important book, *What Do Unions Do?*, published in 1984 (Jake Rosenfeld's 2014 book, *What Unions No Longer Do*, provides an update of the issues).

More recent research has incorporated an assessment of the impact of unions on nonunion workers' wage — sometimes referred to as "spillover effects". The most recent research Fortin, Lemieux, and Lloyd (2021) provides an analysis that incorporates a spillover impact and provides additional insight because the results directly report on the impact of eroded collective bargaining on the wage gap between high-wage (90th percentile) and middle-wage (50th percentile) workers by gender. The authors have provided unpublished tabulations that provide the impact of deunionization on the median worker and median male worker.[13]

Some pundits and analysts, skeptical about the impact of weaker unions on wages or wage inequality, claim that the decline of unions reflects a decline in worker interest in unions or is due to globalization and automation, i.e., endogenous factors. Neither objection is well founded.

Kochan *et al.* (2018) examined the level of interest in joining a union among unorganized workers and found that the "demand for unions" has risen substantially since the late 1970s. Mishel, Rhinehart, and Windham (2020) assess the endogeneity of union decline and find that manufac-

13 Deunionization reduced the median hourly wage by 7.6 log points, or by 7.9 per cent (0.2 per cent annually).Between 1979 and 2017, the impact on men alone is larger, with deunionization lowering the male median wage by 10.9 log points, or 11.6 per cent (0.29 per cent annually).

turing employment decline can account for only a small part of it, perhaps 15–20 per cent. The authors point out that the share of workers covered by collective bargaining declined strongly across the private sector in sectors not heavily affected by globalization, including construction, transportation, communications, utilities, supermarkets, hotels, and mining. An Organisation for Economic Co-operation and Development (OECD, 2019b:15) analysis of the cross-country decline in collective bargaining across advanced nations found:

> "Contrary to a commonly held belief, the combined contributions of demographic changes and structural shifts, such as the shrinking of the manufacturing sector, are small and leave most of this declining trend [in collective bargaining] unexplained."

Managing globalization on capital's terms

Globalization has played a powerful role in disempowering workers and giving capital owners and managers a much-improved fallback position in their bargaining with workers. As such, the way that globalization has proceeded from U.S. workers' perspective has been profoundly shaped by intentional policy decisions that maximized its wage-suppressing effects.

Bivens (2017a:5) presents a summary of globalization's wage impacts, based on his own calculations and on the wider economics literature. He finds:

> "[T]he big damage is the permanent wage loss resulting from America's new pattern of specialization that requires less labour and more capital. Further, this wage loss is not just suffered by workers in tradeable goods sectors who are displaced by imports; it's suffered by all workers who resemble these workers in terms of credentials and labour market characteristics.. . . . The wage-suppressing effects of globalization hit all workers without college degrees, across the country."

Bivens (2013) found that the implied wage effects of trade expanded rapidly after 1995, as trade with lower-wage nations (particularly Mexico and China) picked up significantly. He also found that, by 2013, trade flows with low-wage nations were likely reducing wages for workers without a four-year college degree by roughly 5.6 per cent. For a non-college-degree worker making the median hourly wage and working full time, full year, the earnings reduction translated into just under $2,000 annually.

This estimate is nearly identical to what Autor, Dorn, and Hanson (2013) found in a regression-based investigation of the wage impacts of imports from low-wage countries. Their results indicate that each $1,000 in imports per worker from low-wage countries lower American wages by 0.7 per cent. Imports from all low-wage countries in 2016 stood at roughly $8,000 per worker, implying a wage reduction of roughly 5.6 per cent, or about $2,000 annually, for a full-time worker earning the median wage.

Weakened labour standards

Recent decades have seen the steady weakening of a number of key labour standards that once provided leverage and bargaining power for workers to improve job quality.[14] This part of the article discusses five specific areas of weakened labour standards.with particular emphasis on those that affected the median wage, A more complete discussion is found in Mishel and Bivens (2021a). The rapid erosion of the federal minimum wage's purchasing power is the most dramatic and most consequential. Other negative developments for workers protection are the erosion of overtime protection for salaried workers; weaker labour-standards enforcement and rising wage theft; the increased share of the workforce with no effective labour protections because of its immigration status, and more extensive misclassification of workers as independent contractors.

Erosion of the federal minimum wage

The failure to update the value of the minimum wage in line with wage or productivity growth is a premier illustration of policy choices, made on behalf of capital owners and corporate managers, that have had a huge impact on wage growth for low-wage workers and is the primary explanation for any growth in the wage gap between low- and middle-wage workers over the last four decades. Specifically, the failure to raise the federal minimum wage to an adequate level (defined for our purposes as $15 an hour by 2025) has lowered the wages of at least the bottom 22.2 per cent of earners and a full 31.0 per cent of earners if one includes those benefiting from state and local minimum wage increases since 2017.[15]

The growth of the minimum wage shapes the entire wage distribution of the bottom half, essentially setting the scale of the gap between the lowest-wage workers at the 10th percentile and the wages at the median.

The erosion of overtime protection among salaried workers

To be exempt from the minimum wage and overtime protections of the Fair Labor Standards Act under the "white collar" rule, a worker must be paid a salary (i.e., not be paid by the hour), must have bona fide "executive, administrative, or professional" duties (i.e., be an executive or a highly credentialed professional, or have supervisory duties), and earn above a specific salary threshold. Without a strong salary threshold, salaried workers who spend only a small share of their time actually doing exempt/"professional" work can be required to do hourly-worker-type duties (e.g., a store "manager" stocking shelves, unloading trucks, doing checkout at the cash register) for most of their work-time, including those beyond 40 hours in a week. Those hours beyond 40 are essen-

14 Shierholz (2021) reviews the evidence and importance of enforcement of labour standards.

15 This estimate is based on the impact of raising the minimum to $15 in 2025 and including the impact on those who received minimum wage increases at the state or local level since 2017. This estimate understates the share of earners affected since it ignores those in states that had a higher minimum than the federal threshold minimum wage in 2017 but did not increase it further since then.

tially unpaid.

The eroded share of the salaried workforce eligible for this overtime protection (i.e., receiving 150 per cent of regular hourly wages when working more than 40 hours a week) is another example of a labour standard that was substantially weakened in the last four decades. The share of the salaried workforce automatically eligible for overtime based on its pay was whittled down from roughly half (49.6 per cent) in 1975 to just 9.9 per cent in 2014 (Kimball and Mishel, 2016).

Analysis by the Department of Labor (2016) of the 2016 rule showed that raising the salary threshold increased hourly wages by for salaried workers directly affected by the rule.[16]

How much has the erosion since 1979 affected median wages? This depends on the impact on hourly wages of those affected and the share of middle-wage workers, say the middle fifth, affected by these overtime rules. The impact on the hourly wages of those affected by eroded protections would likely be about 1 per cent. But not all mid-level earners have been affected by changes in overtime protections. We assume that one-third of middle-wage earners lost 1 per cent of wages due to lost overtime protections, so the overall impact would be a 0.3 per cent reduction of hourly wages for the middle fifth.

Wage theft and weaker enforcement of labour standards

Many workers, particularly low-wage workers and the women and men of color who are disproportionately in this category, frequently fail to receive the wages they are owed. This is referred to as "wage theft" and reflects workers being paid below the minimum wage, not being paid for all hours worked, not being paid time-and-a-half though legally eligible for overtime, experiencing illegal deductions from pay, and having their tips stolen by employers or supervisors.

How extensive is wage theft? A 2008 study of 4,387 workers in low-wage industries in Chicago, Los Angeles, and New York; found that two-thirds of workers surveyed experienced at least one pay-related violation in any given week. The average violation amounted to 15 per cent of earnings (Bernhardt, Milkman, and Theodore 2009).

How much does wage theft affect wages in the middle and at the bottom? We do not have an estimate of aggregate wage theft across the wage spectrum or for middle-wage workers, so it is not possible to assess the impact of wage theft on the median wage. Among low-wage workers, Bernhardt, Milkman, and Theodore (2009) found that 68 per cent experience wage theft violations averaging 14.95 per cent of earnings. This translates into an average loss across all low-wage workers of 10.2 per cent. A speculative estimate is that if wage theft has doubled to the 10.2 per cent level implied by the Bernhardt, Milkman, and Theodore study, then it caused low-wage

16 Hourly wages increased 1.1 per cent for workers who occasionally worked overtime, 2.8 per cent for workers who regularly worked overtime and were newly covered by overtime protections. and 1.4 per cent for workers who regularly worked overtime and remained exempt.

workers' earnings to fall 5 per cent over the 1979–2017 period due to weaker wage standards enforcement, less access to legal recourse, and eroded unionization. For mid-level wages, theft of overtime pay, unpaid worktime, and the undercutting of prevailing wages likely also had an adverse impact.

Immigration policy that creates 'labour standard free zones'

Employers have increasingly hijacked immigration policy to create zones in the labour market where workers' ability to obtain enforceable basic labour standards is compromised by their immigration status. Note that the issue is not just the presence, or supply, of immigrants, but the legal situation that makes undocumented workers exploitable. In our economy, if you can be exploited, you will be. This exploitation of a sizable share of the workforce lowers the wages of migrants as well as those of the workers in their occupations and industries (Costa, 2019). Combining the estimates of unauthorized immigrants and guestworkers means that 6 per cent of the workforce is vulnerable to exploitation due to its legal status (Krogstad, Passel, and Cohn, 2019; Costa and Rosenbaum, 2017).

Research by Apgar (2015) comparing the wages of comparable migrant Mexican workers who were undocumented, had legal permanent resident status, or were temporary guestworkers in the H-2A and H-2B visa programs, found that unauthorized workers earned about 13 per cent less than legal permanent residents; temporary foreign workers (i.e., guestworkers) earned about 11 per cent less than legal permanent residents, and their wages did not significantly differ from unauthorized workers' wages.

The presence of exploitable migrant workers therefore undercuts labour standards in immigrant-intensive occupations and industries and thereby depresses wages and benefits of nonmigrants. Historical research highlights the intent of many employers in expanding the pool of workers in the United States who lack basic worker rights because of their immigration status. The conscious policy decisions to allow these circumstances clearly contribute to wage suppression.

Misclassifying employees as contractors

Employers in an array of industries have increasingly (and illegally) misclassified employees as independent contractors or are paying workers "off-the-books." This practice cheats workers of fringe benefits, social insurance protection (Social Security, unemployment insurance, workers' compensation), labour protections (regarding safety/health and race, age, and gender discrimination), and union rights. The point of this misclassification is to lower labour costs, and it undercuts labour standards and "undermines other, more responsible employers who face costs disadvantages arising from compliance with labour standards and responsibilities" (Weil, 2017:3).

It is difficult to quantify the extent of misclassification, since it is an illegal activity, and the extent to which it lowers wage and benefit costs. The fact that venture capitalists force this model on gig economy upstarts provides practical confirmation that the business strategy lowers labour costs and shifts risks to workers.

Uber, a prominent example of a firm whose business strategy is built on misclassifying rideshare drivers, acknowledged in its registration for an initial public offering (Uber 2019) that misclassification provides substantial cost savings.

Indications are that the practice has greatly increased. The last comprehensive federal estimate of independent contractor misclassification, a General Accounting Office (GAO) examination of tax year 1984, "found that 15 per cent of employers nationwide and across industrial sectors engaged in misclassification of a total of 3.4 million workers" (Carré, 2015:10).

Industry analyses provide information, too, on the numbers of workers affected.[17]

Precise estimates of the impact of rising misclassification are not possible with available data, but one can speculate about a range of possible impacts. To gauge the impact we assume that the 3.4 million misclassified workers found by GAO in 1984 (4.4 per cent of nonagricultural wage and salary employment) have risen to 9.0 million (a 7.4 per cent share)[18], and that misclassification lowers wages by either 15 per cent or 30 per cent. Further, we will assume that misclassification is either spread throughout the private nonagricultural wage and salary workforce or, more likely, targeted at the bottom two-thirds; in the former case the share of misclassified workers in total employment rises by 3.0 percentage points, while in the latter case it rises by 4.5 percentage points. The impact is likely to have been on both low-wage and middle-wage workers. These parameters provide a range of impacts: Misclassification lowered wages by between 0.5 per cent and 0.9 per cent if applied across the whole workforce and between 0.7 per cent and 1.4 per cent if affecting and applied to only the bottom two-thirds.[19] If one included all workers, including those in the public sector, then the estimated impacts would be proportionally less. We take a 1 per cent decline in the median wage as our ballpark estimate.

Failures to police or check new forms of employment 'contracts'

Employers are increasingly requiring em-

17 Ormiston, Belman, and Erlich (2020) estimate that in construction "between 1.30 and 2.16 million workers were misclassified or working in cash-only arrangements in an average month of 2017." The major rideshare companies, whose business model incorporates misclassification, have between 1 million and 2 million drivers. Other online demand firms also rely on misclassification. A newspaper investigation by Locke and Ordonez (2014) analyzing payroll records for government-backed construction housing projects across 28 states found that "companies using stimulus money routinely snubbed labor law and the Internal Revenue Service by treating workers as independent contractors in a clear violation of what's allowed." These companies "listed workers as contractors instead of employees in order to beat competitors and cut costs. . . .Scofflaws can save 20 per cent or more in labor costs by treating employees as independent contractors." Misclassification is common in trucking (Bensman 2009, cited in Carré 2015, Appel and Zabin (2019)) and in construction (Ormiston, Belman, and Erlich (2020).

18 The 7.4 per cent is the share of the nonagricultural private wage and salary workforce if 9 million workers are misclassified in 2017.

19 For instance, a rise of 3.1 percentage points of the entire private nonagricultural workforce yields a 0.5 per cent or 0.9 per cent wage reduction if misclassified workers are paid, respectively, 15 per cent or 30 per cent less than regular W-2 workers. Similarly, a rise of 4.6 percentage points among the bottom two-thirds of the private nonagricultural workforce yields a 0.7 per cent or 1.4 per cent wage reduction if misclassified workers are paid, respectively, 15 per cent or 30 per cent less than regular W-2 workers.

ployees to relinquish various rights when they accept employment or even after they are already employed. Noncompete and forced arbitration agreements are chief among these restrictions, and employers within various franchise chains also collude against employees through anti-poaching agreements. All of these agreements limit workers' options by limiting access to courts and the ability to readily find another job or even to know the basic terms of their employment arrangement. This works to suppress wages.

Noncompete agreements

Employers have increasingly required employees to sign noncompete agreements, which limit options for future employment and are now widespread. The practice suppresses worker mobility and suppresses wages, and it depresses firm entry and dynamism because employees are prohibited from starting their own firms (Starr, 2019b and 2020).

How widespread are noncompete agreements?,A 2017 national survey of private-sector American business establishments with 50 or more employees, found that it was "somewhere between 27.8 per cent and 46.5 per cent of private-sector workers are subject to noncompetes" (Colvin and Shierholz, 2019).

These data indicate that noncompetes have grown in their use since a survey of employees in 2014 showed just 18 per cent of the U.S. workforce covered by them, though 38 per cent were subject to one at some point in their careers (Starr, Prescott, and Bishara, 2020). The precise extent of the increased incidence of noncompetes is uncertain, however: The Colvin and Shierholz employer-based survey probably captures more noncompete use than the earlier employee-based survey, since many employees are unaware of having signed a noncompete agreement.

Noncompetes lower wages and mobility for both technical and low-wage workers, whether they reside in states where the contracts are enforceable or in those, such as California, where they are not. Moreover, "where non-competes are really common and highly enforced, the whole labour market suffers" (Starr 2019b), as wages, job mobility, and job satisfaction decline even among those not directly affected.

What is the impact of increased use of noncompete agreements on median wages, low wages, and various wage gaps? The best evidence regarding noncompetes and wage levels is the Lipsitz and Starr (2020) examination of the relationship between Oregon's 2008 ban on noncompetes for low-wage workers and the average hourly wages of hourly paid workers. The finding that the ban raised wages for hourly workers by 2.2 per cent to 3.1 per cent reflects the impact on those directly affected (about 14 per cent of hourly workers are subject to noncompetes) and the spillover effects on other comparable workers. It is important to note that these results are for hourly, not all, workers, and hourly workers comprised 67 per cent of Oregon earners.

Two indications in the Lipsitz and Starr research provide clues to the impact on the median worker. One is that the ban's impact was comparable across the 20th to 80th wage percentiles of hourly workers, suggesting that the impact on the median would be comparable to the 2–3 per cent average effect (if hourly workers comprise the bottom 67 per cent of earners, then the overall median is the 75th percentile of hourly workers). On the other hand, Lipsitz and Starr report that the impact of the noncompete ban was higher for two occupation groups with wage levels close to the overall median: In construction occupations and installation, maintenance, and repair occupations wages rose by 4.9 per cent and 4.3 per cent, respectively. The ban had basically no effect on a low-wage occupation, food service preparation. These differences across occupations reflect both the incidence and direct impact of noncompetes. In sum, these results suggest that the impact of noncompetes on the median is in the 4.3 per cent to 4.9 per cent range, there is little if any impact for the lowest-paid workers, and noncompetes actually narrow the wage gap (50/10) in the bottom half by depressing the median wage but not affecting the lowest-wage workers.

Assessing the impact of the increased use of noncompetes on median wage growth since 1979 requires quantifying that increased use. Unfortunately, there is no historical series on noncompete incidence. The agreements have been used for many years, especially among higher-wage professionals and executives, and use has increased as evidenced by the increased public and policymaker attention to the agreements, particularly for middle-wage or lower-wage (e.g., Jimmy John's sandwich shop workers) workers. If, say, the incidence among hourly workers has doubled since 1979 and the wage impact is roughly 4.5 per cent in recent years, then noncompetes have lowered the median wage by about 2.25 per cent.[20] It seems equally plausible to us, however, that "doubling" is an underestimate, since we know the incidence of forced arbitration agreements has enormously increased since the early 1990s (from 2 per cent in 1992 to more than 50 per cent in 2017), and firms insisting on forced arbitration also tend to insist on noncompetition agreements. So, we take the 2.25 per cent impact on median wage growth as a rough estimate.

Forced arbitration and class action waivers

The increasingly common employer requirement that workers sign arbitration agreements is another clear example of policy decisions, limiting workers' options to resist workplace exploitation. Forced arbitration is among a suite of agreements being forced on workers as a condition of employment (Colvin and Shierholz 2019). It suppresses claims, makes them less likely to succeed, and reduces awards. The consequence is the undermining of the enforcement of employment rights ranging from minimum wage and overtime pay to rights to equal pay and freedom from discrimination or harassment based on race, gender, or religion (Stone and Colvin 2015, Colvin 2018, and Deutsch *et al.* 2019).

The incidence of forced arbitration agreements took off after key Supreme Court decisions in 1991 and 2001 made

20 The impact might be lower if we base the assessment on the 2.2 per cent to 3.1 per cent overall wage effect identified by Lipsitz and Starr (2020) or if the rise in incidence is less than double.

clear "that an American employer may, with near total impunity, require an employee, as a condition of hiring and continued employment, to use private arbitration as the means of resolving public claims against the employer that involve a statutorily protected right" (Lipsky 2007:10). In 1992, just over 2 per cent of the workforce was covered by forced arbitration agreements, but that share rose to almost a quarter by the early 2000s. By 2017 the share was 56.2 per cent (Colvin 2018).

It is not possible to assess the wage impact of the spread of forced arbitration agreements. The practice is intended to and does undermine the enforcement of employment and civil rights workplace protections, further limiting employee options to resist employer exploitation.

Employer collusion and anti-poaching agreements

We do not know the extent of collusion among employers and how it has changed over time. It is, after all, illegal. There is research on explicit collusion in franchising, however, because this is a gray area in the law.[21]

Unfortunately, there is no systematic evidence of no-poaching agreements' impact on workers' pay and within-franchise job mobility. We do know that these agreements grew substantially over the 1996–2016 period, however, and disproportionately affect workers in low-wage industries and "potentially affect a large number of workers" (Krueger and Ashenfelter 2018).

Tolerating new business structures that disempower workers

In recent decades, employers have increasingly tried to build up concentrated power in product markets (as well as labour markets directly) and to leverage this increased product market power to augment their profitability and the pay of executives by lowering costs and suppressing wages. One mechanism has been to match market concentration with efforts to outsource key parts of their production or workforce to keep those costs from making a claim on the firm's income. In past years, policymakers might have used industry regulation such as in airlines and trucking and antitrust enforcement to keep these changes in check. But in recent decades, the pushback against these changes in business structure has been rare and muted.

Fissuring: contracting out/outsourcing, temping, and franchising

Perhaps the most pronounced way that employers have attempted to shape

21 Krueger and Ashenfelter (2018) report on anti-poaching agreements in franchises, agreements that limit a particular franchise of McDonald's, for instance, from hiring a worker from another McDonald's franchise.

22 David Weil's book, *The Fissured Workplace* (2014), as well as analyses by Appelbaum and Batt (2014), provide the details about what fissuring is and how it works to the advantage of employers.Weil points out that fissuring should be distinguished from contingent work or alternative work arrangements. "Fissured workplace arrangements can exist even though employment itself might be traditional (that is, ongoing and full time) when the worker is employed by a subcontractor, franchisee, or other business organization undertaking the work of a lead business."

labour market outcomes to their advantage through changes in business structure is the "fissuring" of workplaces.[22]

Fissuring is a corporate strategy that emerged from the focus on shareholder value. It raises profits in part by squeezing the costs of subcontractors, who in turn cut wages, and shifting risks onto other firms and workers. Fissuring probably has no impact on aggregate productivity, or on making the production of goods and services more efficient. Instead, its effects are overwhelmingly distributional, suppressing wages and profits among suppliers to the benefit of the contracting firms (Appelbaum and Batt, 2014).

But estimating the size of the fissured economy is a major challenge researchers are only now undertaking. However, it seems clear that somewhere between a fifth and a third of the economy is characterized by fissuring as a dominant force (Weil, 2019).

Growth in fissuring seems to have primarily occurred in business-to-business domestic outsourcing or subcontracting and not through use of independent contractors, staffing agencies, or franchising. Various studies confirm that workers in contractor firms earn less.

Fissuring, and particularly the outsourcing of particular tasks, is probably responsible for the fact that workers in the largest firms no longer receive higher pay than those in medium-sized firms. As Bloom *et al* .(2018) show, those in firms with more than 2,500 employees were not paid more in the 2007–2013 period than those in firms that had 1,000 to 2,500 employees, a sharp drop from what prevailed in the 1980s. This erosion of the quality of jobs in large firms affected a large swath of the workforce, as employment in firms exceeding 2,500 employees comprised 39.0 per cent of all jobs in 2014 compared to 37.0 per cent in 1999 and 35.3 per cent in 1979.[23]

A speculative estimate of the impact of fissuring is that a shift of 15 percentage points of employment into fissured workplaces earning 15.0 per cent less (Goldschmidt and Schmieder 2017) would yield an overall decline of wages of 2.25 per cent overall.

Product and labour market concentration, including dominant buyer

There has been increasing interest in two key changes in corporate structure in recent decades: product and labour market concentration (sometimes referred to as monopoly and monopsony).[24] It seems clear that there are many reasons for policymakers to be concerned about market concentration. Robust efforts (antitrust or regulation) to confront the malign effect of concentration should be part of the policy toolkit going forward.

Labour market concentration.

Though labour market concentration is definitely associated with lower wages, evi-

23 Analysis of the Census Bureau's Longitudinal Business Database.

24 We will not use the label "monopsony" here for labour market concentration, as modern labour economics has adopted the term "monopsony" to describe a wide range of influences—including but not limited to market concentration—that give employers power to set wages.

dence remains lacking that it has increased so as to greatly contribute to wage suppression.[25]

Product market concentration—monopoly power in product markets.

The growth of product market monopoly power over prices does not seem to have contributed to wage suppression, though firms may have leveraged their monopoly positions to suppress wages and profits in supplier chains. [26]

Rising monopoly power in product markets that has harmed U.S. households through excessive price growth seems unlikely to be a major channel through which concentration may be dragging on wage growth.

Product market concentration—dominant buyers squeezing suppliers..

In traditional conceptions of the harms done by product market concentration, firms' monopoly power is leveraged against consumers of their output, with prices being pushed above what would prevail in competitive markets. However, many real-world firms with substantial market share (Walmart and Amazon, for example) charge their own customers seemingly low prices while leveraging their market power instead against the firm's own suppliers, coercing them into providing supplies at low prices. This in turn squeezes both profits and wages for the supplier firms.

Path-breaking research by Wilmers (2018) has identified and quantified the impact of these "dominant buyers." Wilmers estimates that the share of nonfinance suppliers' revenue obtained from dominant buyers increased from 5 per cent in 1979 to 19 per cent in 2014 overall and from 6 per cent to 26 per cent in manufacturing and logistics. Wilmers argues that there was not only an increase in the role of dominant buyers but also an intensification of their wage impact. Wilmers estimates that the increase in dominant buyers lowered the growth of average annual earnings by 3.4 percentage points over the 1979 to 2014 period among publicly owned nonfinancial firms.

Deregulation of industries

Starting in the late 1970s, Congress deregulated various industries, including airlines, trucking, interstate busing, telecommunications, utilities, and railroads. Fortin and Lemieux (1997) showed that deregulation had a strong adverse impact on the wages of blue-collar workers for 9 per cent of the workforce. Card (1996) found a 10 per cent decline over 1980–1990 in the relative earnings of airline workers after deregulation. Deregulation also weakened the ability of employers to pay high wages and in many sectors, most notably trucking, led to a steep erosion of unionization (Viscelli, 2016). Unfortunately, we

25 Key papers assessing the effect of labour market concentration include Azar, Marinescu, and Steinbaum (2017); Benmelech, Bergman, and Kim (2018); Rinz (2018); and Naidu, Posner, and Weyl (2018). The evidence also shows that the average labor market is highly concentrated but the average worker is not employed in a concentrated labor market.

26 Key papers directly assessing the effect of product market concentration include Autor *et al.* (2017b); Barkai (2020); De Loecker, Eeckhout, and Unger (2020); Grullon, Larkin, and Michaely (2015); and an analysis by Goldman Sachs (Struyven 2018).

do not have an estimate of the impact of deregulation on wages.

Wage suppression and upward distribution to the top 1 per cent

The data are clear that wage growth for the vast majority of U.S. workers decelerated radically in the post-1979 era. This near-stagnation of median wages cannot be nearly fully explained by the slowdown in the economy's overall ability to pay higher wages (measured, for example, by growth in economy-wide productivity). Overall output and income growth did slow significantly post-1979, but growth for the bottom 90 per cent of wage earners slowed far more. This pattern left a large excess available for the top 10 per cent to grab, and most of it went to the top 1 per cent and, especially, the top 0.1 per cent. While this article does not undertake to directly explain the growth of wages at the very top those of the top 0.1 per cent and 1 per cent we would argue that this growth is the mirror image of wage suppression at the bottom. The forces that weighed on wage growth for the majority (excess unemployment, stagnation of the minimum wage, deunionization) largely do not slow wage growth for the top 1 per cent; instead, they allow more income (wages, and profits that are not going to typical workers' paychecks) to be claimed by the very top. In a sense, the wage suppression felt by the bottom 90 per cent was zero-sum (or even negative sum), as their loss financed a sharp redistribution of wages and incomes to the very top.'

Bivens and Mishel (2013) highlight the growth of Chief Executive Officer (and other executive) pay and the incomes of financial-sector professionals as the predominant source of the escalated wage and income growth at the top. This partly reflects the surge in the stock market which fuels executive and financial sector growth in pay. A vast body of research demonstrates that no plausible force coming from the interplay of competitive markets could account for the explosive income growth of these actors. These are not the "just deserts" of the high earners reflecting their heightened productivity.

Aggregate Impact of the Policy Choices Generating Wage Suppression

This section draws on the earlier assessments of the factors generating wage suppression to account for the divergence between the growth of net productivity and median hourly compensation over the 1979–2017 period

Relation to other literature

Our analysis builds on what Stansbury and Summers (2020) referred to as a "long history of progressive institutionalist work exemplified by Freeman and Medoff (1984), Levy and Temin (2007), and Bivens, Mishel, and Schmitt (2018)."

An important recent marker in this tradition was the keynote address by former chairman of the Council of Economic Advisers and Princeton economist Alan Krueger (Krueger, 2017) to the Federal Reserve Board Jackson Hole conference. Krueger said that certain economic models "give employers some discretion over wage setting" and, in a footnote, said, "Notice that I don't call these features 'imperfections.' They are the way the labour market works. The assumption of perfect competition is the deviation from the norm of 'imperfection' as far as the labour market is concerned."

The new monopsony literature reinforces our narrative in important ways and further highlights the need to identify the specific factors generating employer power over wages and ways this power has changed over time. A broad interpretation of employers' "monopsony" power does not hinge on labour market concentration (i.e., the proverbial one-company town), but instead diagnoses labour markets as being affected by employers' exercise of power that allows them to cut wages without fear of losing a large portion of their workforces—regardless whether the source of this power is market concentration or anything else. The focus needs to be on employer power relative to employees, so the erosion of countervailing power of employees is an essential dimension of understanding how employer power matters.

The emerging monopsony literature shows that employer power is ubiquitous in the modern U.S. labour market.[27]. The monopsony literature has identified a substantial amount of employer power such that employers are able to, as Bassier, Dube, and Naidu (2020) put it, "mark down" wages by anywhere from 20 per cent to 50 per cent. [28]

One way to interpret the evidence in the current paper is that employer power is the constant of modern labour markets, but what has changed over the past generation in the United States to generate anemic wage growth is the erosion of institutions and policies—high-pressure labour markets, unions, and binding minimum wages—that once provided countervailing power. Naidu and Sojourner (2020) have a similar interpretation:

This new monopsony literature provides a top-down analysis, which has primarily focused on estimating the aggregate scale of employer power. Some of the recent contributions have started to identify the underlying factors, examining the role of unionization, high-pressure labour markets, and high values of minimum wages, in explaining an aggregate metric of monopsony power. In contrast, our study is a bottom-up analysis examining the impact of myriad specific factors and gauging their contribution to the productivity–median compensation divergence over the past four decades.

Explaining the divergence between productivity and median hourly compensation growth

27 Webber (2015 and 2020); Dube, Giuliano, and Leonard (2019); Dube *et al.* (2020); Bassier, Dube, and Naidu (2020); Azar, Marinescu, and Steinbaum (2019); Langella and Manning (2020); Card *et al.* (2018); and the meta-analysis by Sokolova and Sorensen (2020)

28 There is some evidence on the time trend of employer monopsony power; two studies have shown that employer power increased since the late 1990s (Webber 2020; Langella and Manning 2020), though Bassier, Dube, and Naidu (2020, Table 6) show stability over the 2003–2012 period. One consistent finding of these studies is that employers are able to exert more power over low-wage than other workers, affirming that employer power generates wage inequalities.

This section examines the corporate and government policy levers that have suppressed wage growth, and concludes that they can account for the vast majority of wage suppression.

It can be difficult to assess causality and take interactions into account. But looking at the sum of the impact of the key factors supports the narrative that intentional policy decisions (either of commission or omission) have generated wage suppression. Analysts may differ on the assessment of particular factors, but our hope is that this compilation inspires further efforts, including ones for which we do not have sufficient empirical work to even make guesses.

How much needs to be explained? In our discussion of Table 1 we noted that between 1979 and 2017 net productivity (economy-wide productivity net of depreciation) grew 56 per cent while median hourly compensation (wages and benefits) grew 13 per cent, leaving a 43 per cent divergence.[29] By deflating both net productivity and the pay measures by the CPI-U-RS index, we have stripped out the influence of differing deflators (for productivity and compensation) from our calculation of the divergence, leaving only the changes in labour's share of income and changes in compensation inequality as drivers of the divergence (Bivens and Mishel 2015).

The impact of specific factors on the growth of the median wage is detailed in the first panel (Table 2) and draws on the discussion above. [30]

The share of the various factors in explaining the overall divergence of net productivity and real hourly compensation as presented in Table 2 is illustrated in Panel A of Chart 3 (examining growth, in percent, of factors) and Panel B of Chart 3 (examining growth, in dollars, of factors).

Austerity macroeconomic policy (excessive unemployment)

The impact of excessive unemployment caused by contractionary macroeconomic policy, promulgated to control inflation, (suppressing labour costs in the name of controlling inflation) reduced wages for the median worker by 10.0 per cent between 1979 and 2017. Adjusting for the "flattening" of the Phillips curve since 2008, as we do here, lessens the impact of higher unemployment on wage growth; without this adjustment the impact would have been 12.2 per cent.

Erosion of Collective Bargaining

The erosion of collective bargaining had an adverse impact by lowering the wages

29 If we used hourly compensation of production workers (82 per cent of payroll employment) as the pay measure rather than the median wage, the divergence would be even greater (45.1 percentage point divergence from a net productivity growth of 56 per cent and real hourly compensation of production workers growth of 11 per cent).

30 For purposes of the analysis we equate the impacts on median wages, as identified above, to be the same as the impact on median hourly compensation: this is not a consequential decision since the 13.0 per cent growth of median hourly compensation over the 1979–2017 period just slightly exceeded the 12.2 per cent growth of median hourly wages.

Chart 3: Factor's Contribution to Productivity-Median Compensation Divergence, 1979-2017

Panel A: Percentage Point Contribution

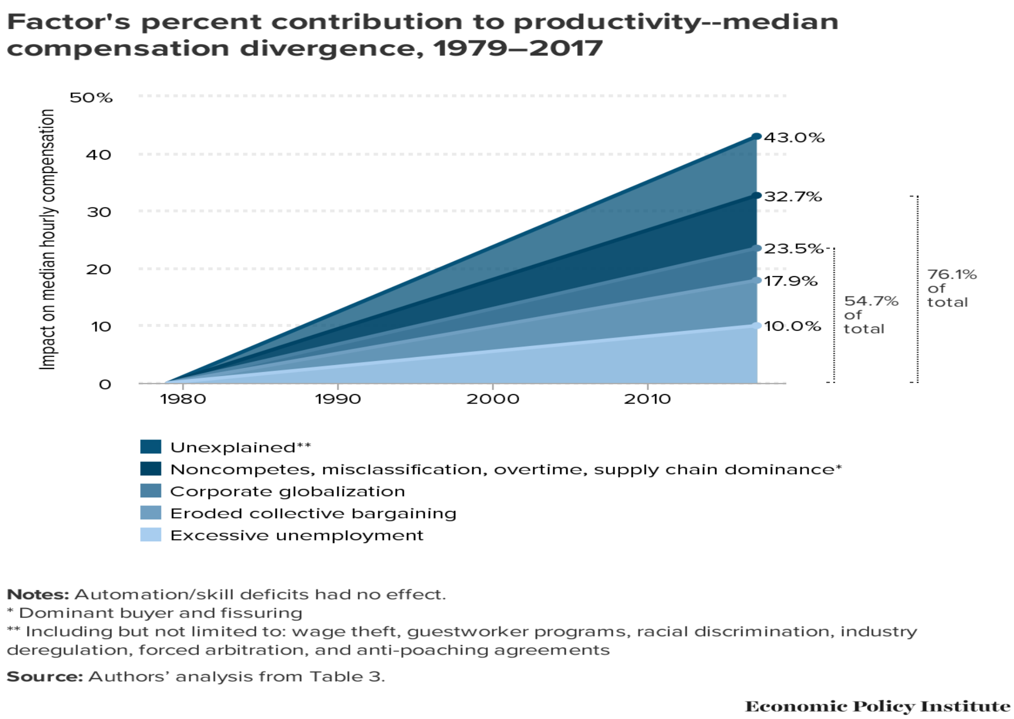

Panel B: Dollar Contribution

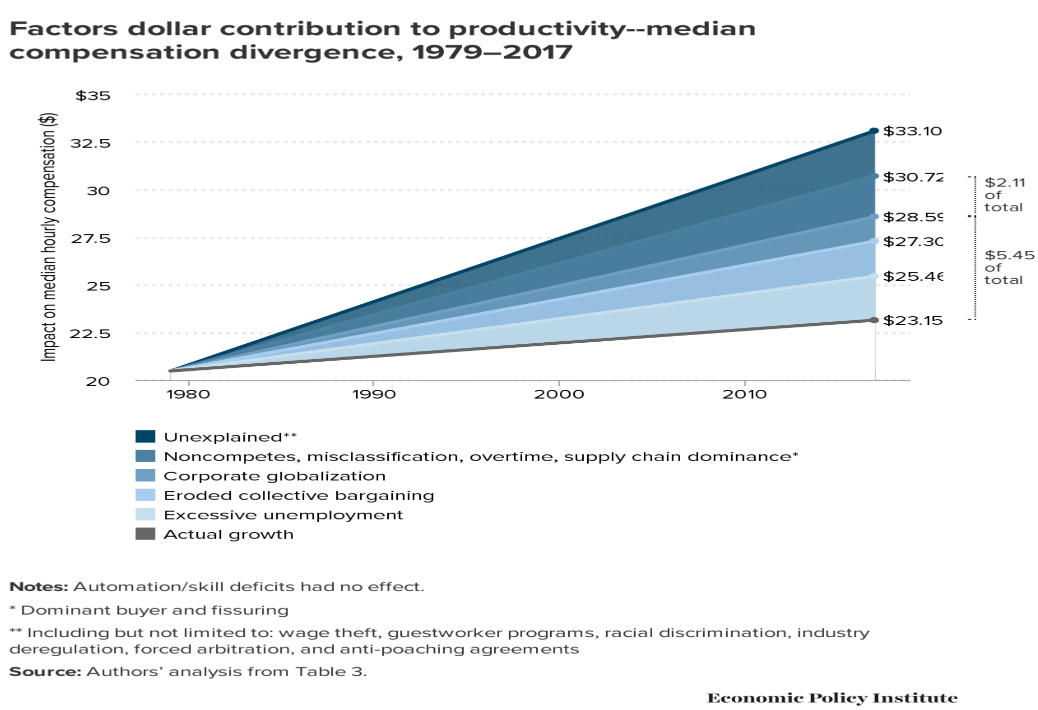

Note: Automation/skill deficits had no effect.
*Dominant buyer and fissuring
**Including but not limited to: wage theft, guest worker programs, racial discriminationm industry deregulation, forced arbitration and anti-poaching agreements
Source: Authors' analysis from Table 3

Table 2: Impact of Policy Areas on Median Wage Growth, 1979-2017 (percentage points)

1	**Excessive Unemployment***	10.0
2	**Erosion of Collective Bargaining**	7.9
3	**Globalization**	5.6
4	**Weaker Labour Standards**	
	Erosion of Minimum Wage	0
	Overtime Coverage for Salaried Workers	0.3
	Wage Theft	n.a
	Missclassification	1.0
	Increased Presence of Undocumented	n.a
	Guest Worker Programs	n.a
5	**Employer-Imposed Contract Terms**	
	Noncompetes	2.25
	Anti-Poacching	n.a
	Forced Arbitration	n.a
6	**Corporate Structure Changes**	
	Labour Market Concentration	0
	Dominant Buyer	3.4
	Fissuring	2.25
	Product Market Concentration	0
	Industry Deregulation	n.a
7	**Automation/skill-biased technological change**	
	1979-1995	n.a
	1995-2017	0

*Average relative to 5.5% unemployment

Sources: Authors' analysis based on the following studies: Excessive unemployment estimate based on Bivens and Zipperer (2018) and Katz and Krueger (1999); erosion of collective bargaining estimate based on Fortin, Lemieux, and Lloyd (2021), Stansbury and Summers (2020), Western and Rosenfeld (2011); globalization estimate based on Bivens (2013) and Autor, Dorn, and Hanson (2013); erosion of minimum wage estimate based on Fortin, Lemieux, and Lloyd (2021) and Autor, Manning, and Smith (2016); overtime coverage for salaried workers estimates based on analysis of Department of Labour (2016); noncompetes estimates based on Lipsitz and Starr (2020); dominant buyer estimates based on Wilmers (2018); fissuring based on Weil (2019); automation/skill-biased technological change from 1995–2017 estimates based on Autor, Goldin, and Katz (2020) and Autor (2017).
Note: n.a. means not available

of non-college-educated workers, particularly men, and has also lowered the wages and benefits of nonunion workers in sectors where collective bargaining had previously set wage patterns. We relied on an unpublished analysis of the Fortin, Lemieux, and Lloyd (2021) model to pinpoint at 7.9 per cent the impact of deunionization on the median wage of all workers (men and women combined) for the 1979–2017 period.

Globalization on capital's terms

Bivens (2013) found that, by 2013, trade flows with low-wage nations were likely reducing wages for workers without a four-year college degree by roughly 5.6 per cent. For a non-college-degreed worker making the median hourly wage and working full time, full year, this translates to about $2,000 annually. This estimate is nearly identical to what Autor, Dorn, and Hanson (2013) find in a regression-based investigation of the wage impacts of imports from low-wage countries.

Impact of the top three factors

As summarized in Table 3, together these three factors alone—excessive unemployment, eroded collective bargaining, and corporate-driven globalization—can account for a 23.5 per cent decline in the median wage growth from 1979 to 2017 and for 54.7 per cent of the divergence be-

Table 3: Impact of Policy Decisions on Median Wage Growth, 1979-2017 (percentage point)

	Impact on median wage (percentage points)	Share of divergence explained (percentage points)
Excessive unemployment; erosion of collective bargaining; globalization.	23.5	54.7
Other(Overtime, misclassification, fissuring, dominant buyer, noncompetes)	9.2	21.4
Automatom/skill deficits	0	0
Total explained	32.7	76.1
Unexplained*	10.3	23.9
Divergence btw productivity-median hourly compensation, 1979-2017	43	100

*Including but not limited to: wage theft, guest worker programs, racial discrimination, industry deregulation, forced arbitration and anti poaching agreements.
Source: Estimates presented in Table 2

tween net productivity and median hourly compensation. Chart 3 shows that excessive unemployment, eroded collective bargaining, and corporate-driven globalization lowered the growth of median hourly compensation by $5.45: absent these factors median hourly compensation would have risen to $28.59 rather than to $23.59.

Weaker labour standards

The failure to update the value of the federal minimum wage is a premier example of policy action shaping the wage structure and undermining the wages of the bottom third of earners (heavily women and minorities), or 46 million workers. The minimum wage's impact probably does not extend to the median, so we express that as zero in Table 3.

The erosion of other labour standards likely had an impact throughout the wage structure. Overtime protections for salaried workers declined precipitously and reduced median workers' wages by 0.3 per cent, while growing misclassification of workers as independent contractors lowered the median wage by 1.0 per cent. Other practices and policies, like lax protections against wage theft, the increased presence of undocumented workers and guestworkers, and more extensive racial discrimination have likely lowered wages, but we are not able to provide an empirical assessment.

Employer-imposed contract restrictions

Employers have increasingly required employees to relinquish various rights when they accept employment, or even after they are already employed, through agreements regarding noncompetition and forced arbitration. Employers within franchise chains have also colluded against employees through anti-poaching agreements, which limit workers' employment options. The effort to quantify the impact of these policies is still in the beginning stages. We estimate that noncompete agreements have reduced the median wage by 2.25 per cent, but we have not been able to derive estimates of the impact of forced arbitration (now covering more than half of nonunion employees) nor of anti-poaching agreements among franchisers.

Changes in corporate structure

Changes in corporate structure — from deregulation to fissuring to rising market concentration—likely pushed down wages by at least 5 per cent by 2017. A speculative gauge of the impact of a shift of 15 percentage points of employment into fissured workplaces where wages are 15 per cent less would imply an overall decline of wages of 2.25 per cent and probably an even larger decline at the median. Wilmers' (2018) estimated that the increase in dominant buyers lowered annual earnings by 3.4 per cent over the 1979–2014 period among workers in publicly owned nonfinancial firms.

There is likely to be some double counting when aggregating the fissuring and dominant buyer factors, but it is also likely that the unassessed components of corporate structures exerted at least as much downward wage pressure to offset it.

Automation/skill-biased technological change

As detailed earlier (and in Mishel and Bivens (2021b)), automation and skill-biased technological change are prima facie implausible explanations of the wage suppression or wage inequality experienced at least since 1995. Given the deceleration of the salient indicators of automation and automation's impact on key labour market metrics (relative demand for college education, occupational polarization), we assign no impact in Table 2 to automation in driving the productivity–pay divergence for the 1995–2017 period. We are also skeptical that there was any impact in the earlier 1979-95 period, following the analysis in Mishel, Bernstein, and Schmitt (1997a) and Card and DiNardo (2002).

Conclusion

In all, the policy-driven factors delineated in Tables 2 explain a vast share of the divergence between productivity and median hourly compensation.[31]

The best-measured impacts, those for excessive unemployment, eroded collective bargaining, and corporate-shaped globalization, can account for 23.5 percentage points (or 55 per cent of the total) of the 43 percentage points productivity–median compensation divergence. The harder-to-measure impacts of other factors (lowering of the overtime threshold for salaried workers, misclassification, noncompete agreements, and changes in corporate structures like fissuring) can collectively account for another 9.2 per cent of the erosion of the median wage and explain another 21.4 per cent of the divergence. These sum up to explaining about three-fourths (76.1 per cent) of the divergence (Table 2 and Chart 1). This is an understated conclusion since there are many additional policy factors that we have not been able to empirically assess: wage theft, guestworker programs, racial discrimination, industry deregulation, exploitable immigrants, forced arbitration, and anti-poaching agreements.

31 See Mishel and Bivens (2021a) for a discussion on the suppressed wage growth at the 10th percentile and the increase in the 50/10 percentile wage gap

It is possible that summing these estimates overstates the aggregate impact if there is some endogeneity or interactions (e.g. higher unemployment causes greater loss of unions). On the other hand, there are reasons to believe that the impact of these factors is larger than the sum of their individual effects. One way of understanding what has happened is to gauge all the ways that an individual workers' options to obtain better employment conditions or to affect their current employment have been increasingly foreclosed—limiting both exit and voice. When workers want to improve their conditions of work, they have increasingly limited options to organize a union, rely on adequate and enforceable government standards (e.g., the minimum wage, safety and health, overtime, anti-discrimination, correct classification), or make employers accountable through litigation. Exit is more limited because of anti-poaching agreements, noncompetes, and generally higher unemployment, and the downward pressure on their wages is intensified by globalization, fissuring, and dominant buyer power. Increasingly, resistance is futile.

Our analysis, admittedly, does not rely only on pure causal estimates of each factor or guarantee that there are no interactions between the different determinants of wage growth. We offer our estimates as an informed summation of what existing research finds, and, crucially, we argue that any interaction effects are at least as likely to amplify the effects of policy levers as they are to dampen them.

So, while we conclude that these policy levers can account for the vast majority of the rise of wage inequality and the productivity-median hourly compensation divergence we are not wedded to any specific number. We imagine that any quibbles with our estimates would still leave policy factors explaining a clear majority. We look forward to others offering alternative estimates of these factors and associated aggregations or to offering an alternative narrative with corresponding empirical evidence. But the simple existence of possible interactive effects does not mean that our results will clearly shrink – they may well rise–or that policy levers did not drive wage suppression.

The lessons here are simple. Wage growth has been greatly directed by policy decisions and is a political variable. It responds — robustly — to big policy changes. But for decades these policy decisions have gone in the wrong direction. Policymakers can deliver prosperity to the vast majority of U.S. workers based on faster wage growth. Whether workers obtain a fair share of the economy's gains in the future will depend not so much on abstract forces beyond their control but on demanding that their political representatives restore bargaining power to workers, individually and collectively.

References

Acemoglu, D. and D. Autor (2012) "What Does Human Capital Do? A Review of Goldin and Katz's The Race between Education and Technology," *Journal of Economic Literature,* Vol. 50, No. 2, pp. 426–463.

Apgar, L.A. (2015) "Authorized Status, Limited Returns: The Labor Market Outcomes of Temporary Mexican Workers," Economic Policy Institute.

Appel, S. and C. Zabin (2019) "Truck Driver Misclassification: Climate, Labor and Environmental Justice Impacts," *UC Berkeley Labor Center.*

Appelbaum, E. and R. Batt (2014) "Private Equity at Work: When Wall Street Manages Main Street," *Russell Sage Foundation.*

Autor, D (2014) "Polanyi's Paradox and the Shape of Employment Growth." Working Paper 17040. National Bureau of Economic Research.

Autor, D (2010) "The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings," Center for American Progress and Brookings Institution, The Hamilton Project.

Autor, D., C. Goldin, and L. F. Katz (2020) "Extending the Race Between Education and Technology." *AEA Papers and Proceedings,* Vol. 110, pp. 347-51.

Autor, D., Dorn, D. and G.H. Hanson (2013) "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, Vol. 103, No. 6, pp. 2121–2168.

Azar, J., I. E. Marinescu, and M. Steinbaum (2017) "Labor Market Concentration," *National Bureau of Economic Research,* Working Paper No. 24147. https://doi.org/10.3386/w24147.

Azar, J. I. Marinescu, and M. Steinbaum (2019) "Measuring Labor Market Power Two Ways," textitAEA Papers and Proceedings Vol. 109.

Bassier, I., A. Dube, and S. Naidu (2020) "Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies," *Working Paper 27755*, National Bureau of Economic Research.

Benmelech, E., N. Bergman and H. Kim (2018) "Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?" *National Bureau of Economic Research,* Working Paper 24307.

Bernhardt, A., Milkman, R., Theodore, N., Heckathorn, D., Auer, M., DeFilippis, J., González, A.L., Narro, V., Perelshteyn, J., Polson, D. and M. Spiller (2009) "Broken Laws, Unprotected Workers: Violations of Employment and Labor Laws in America's Cities," Center for Urban Economic Development, National Employment Law Project, and UCLA Institute for Research on Labor and Employment.

Bivens, J (2016) "Why Is Recovery Taking So Long—and Who's to Blame?" Economic Policy Institute.

Bivens, J (2013) "Using Standard Models to Benchmark the Costs of Globalization for American Workers Without a College Degree," Economic Policy Institute.

Bivens, J (2017a) "Adding Insult to Injury: How Bad Policy Decisions Have Amplified Globalization's Costs for American Workers," Economic Policy Institute.

Bivens, J. and B. Zipperer (2018) "The Importance of Locking in Full Employment for the Long Haul," Economic Policy Institute.

Bivens, J. and L. Mishel (2015) "Understanding the Historic Divergence Between Productivity and a Typical Worker's Pay: Why It Matters and Why It's Real," Economic Policy Institute.

Bivens, J. and L. Mishel (2013) "The Pay of Corporate Executives and Financial Professionals as Evidence of Rents in Top 1 Percent Incomes," *Journal of Economic Perspectives*, Vol 27, No. 3, pp. 57–78.

Bivens, J., L. Mishel and J. Schmitt (2018) "It's Not Just Monopoly and Monopsony: How Market Power Has Affected American Wages," Economic Policy Institute.

Bloom *et al.* (2018) "Is the Large Firm Wage Premium Dead or Merely Resting?" *AEA Papers and Proceedings,* Vol. 108, No. 1, pp.317–22.

Card, D.*et al.* (2018) "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics,* Vol. 36, No. S1, pp: 13-70.

Card, D, and J. DiNardo (2002) "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics*, Vol. 20, No. 4, pp. 733–783.

Card, D (1996) "The Effect of Unions on the Structure of Wages: A Longitudinal Analysis," *Econometrica*, Vol. 64, No. 4, pp. 957–979.

Carré, F (2015) "(In)dependent Contractor Misclassification," Economic Policy Institute.

Colvin, A.J.S. (2018) "The Growing Use of Mandatory Arbitration: Access to the Courts Is Now Barred for More Than 60 Million American Workers," Economic Policy Institute.

Colvin, A. J.S. and H. Shierholz (2019) "Noncompete Agreements: Ubiquitous, Harmful to Wages and to Competition, and Part of a Growing Trend of Employers Requiring Workers to Sign Away Their Rights," Economic Policy Institute.

Costa, D and J. Rosenbaum (2017) "Temporary Foreign Workers by the Numbers: New Estimates by Visa Classification," Economic Policy Institute.

Costa, D (2019) "Employers Increase Their Profits and Put Downward Pressure on Wages and Labor Standards by Exploiting Migrant Workers," Economic Policy Institute.

De Loecker, J., J. Eeckhout and G. Unger (2020) "The Rise of Market Power and the Macroeconomic Implications," *Quarterly Journal of Economics,* Vol. 135, No. 2, pp: 561-644.

DeLong, J.B. and L. Summers (1988) "How Does Macroeconomic Policy Affect Output?" *Brookings Papers on Economic Activity*, Vol. 2, pp. 423-494.

Deutsch, R., Hamaji, K., Nicolas, E., McNicholas, C., Poydock, M. and H. Shierholz (2019) "Unchecked Corporate Power: Forced Arbitration, the Enforcement Crisis, and How Workers Are Fighting Back," Center for Popular Democracy and Economic Policy Institute.

Dube, A. *et al* (2020) “Monopsony in Online Labor Markets” *American Economic Review: Insights 2*, No. 1, pp: 33-46.

Dube, A. L. Giuliano and J. Leonard (2019) “Fairness and Frictions: Impact of Unequal Raises on Quit Behavior,” *American Economic Review,* Vol. 109, No. 2, pp:620-63.

Farhi, E. and F. Gourio (2018) “Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia,”*Brookings Papers on Economic Activity,* Fall, pp. 147-250.

Fortin, N. and T. Lemieux (1997) “Institutional Changes and Rising Wage Inequality: Is There a Linkage?” *Journal of Economic Perspectives*, Vol. 11, No. 2, pp. 75–96.

Fortin, N., Lemieux, T. and N. Lloyd (2021) “Labor Market Institutions and the Distribution of Wages: The Role of Spillover Effects,” Working Paper 28375, National Bureau of Economic Research.

Freeman, R.B. and J.L. Medoff (1984) *What Do Unions Do?*, Basic Books.

Galbraith, J (1997) "Time to Ditch the NAIRU," *Journal of Economic Perspectives*, Vol. 11, No. 1, pp. 93-108.

Goldin, C. and L. Katz (2007) “The Race Between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005,” Working Paper 12984. National Bureau of Economic Research.

Goldin, C, and L.F. Katz (2008) *The Race Between Education and Technology*, Harvard University Press.

Grullon, G., Y. Larkin and R. Michaely (2015) “Are U.S. Industries Becoming Concentrated?” *Swiss Finance Institute Research Paper 19-41*, forthcoming in Review of Finance, posted online May 31, 2015.

Hacker, J.S. and P. Pierson (2011) *Winner-Take-All Politics: How Washington Made the Rich Richer—and Turned Its Back on the Middle Class.* Simon and Schuster.

Hacker, J.S. and P. Pierson (2020) *Let Them Eat Tweets: How the Right Rules in an Age of Extreme Inequality*, Liveright Publishing.

Harrison, P (2009) "Median Wages and Productivity Growth in Canada and the United States," CSLS Research Note 2009-2, Centre for the Study of Living Standards, http://www.csls.ca/events/cea2009/sharpe-etal.pdf.

Hooper, P, Mishkin, F.S. and A. Sufi (2019) “Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or Is It Just Hibernating?” Working Paper 25792, National Bureau of Economic Research.

Katz, L.F. and K.M. Murphy (1992) “Changes in Relative Wages, 1963–87: Supply and Demand Factors,” Quarterly *Journal of Economics* Vol. 107 (February): pp 35–78.

Katz, L.F. and A. Krueger (1999) “The High-Pressure Labor Market of the 1990s,” *Brookings Papers on Economic Activity Vol 1.*

Kimball, W. and L. Mishel (2016) "Nearly One-Third of Salaried Workers Will Be Guaranteed Overtime Under the New Rule, but That Falls Short of the Half Covered in the Mid-1970s," Economic Snapshot. Economic Policy Institute, June.

Kochan, T.A., Kimball, W., Yang, D. and E.L. Kelly (2018) “Worker Voice in America: Is There a Gap Between What Workers Expect and What They Get?" *ILR Review*, Vol. 72, No. 1, pp. 3–38.

Krogstad, J.M., Passel, J. and D. Cohn (2019) “5 Facts About Illegal Immigration in the U.S.,” Pew Research Center.

Krueger, A.B (2017) “The Rigged Labor Market,” *Milken Institute Review*, April 28.

Krueger, A.B. and O. Ashenfelter (2018) “Theory and Evidence on Employer Collusion in the Franchise Sector,” Working Paper 24831. National Bureau of Economic Research.

Langella, M. and A. Manning (2020) “The Measure of Monopsony,” Video, 56:28. Marshall Lecture, Annual Congress of the European Economic Association.

Levy, F. and P. Temin (2007) “Inequality and Institutions in 20th Century America," Working Paper 13106, National Bureau of Economic Research.

Lipsitz, M. and E. Starr (2020) “Low-Wage Workers and the Enforceability of Non-Compete Agreements.” *Management Science*, pp. 1-19.

Lipsky, D.B (2007) “Conflict Resolution and the Transformation of the Social Contract,” Labor and Employment Relations Association Presidential Address, January.

Locke, M. and F. Ordonez (2014) “Taxpayers and Workers Gouged by Labor-Law Dodge,” *McClatchy Washington Bureau,* September 4.

Mishel, L. and J. Bernstein (1994) “Is the Technology Black Box Empty? An Empirical Examination of the Impact of Technology on Wage Inequality and the Employment Structure, ” *Technical Paper,* Economic Policy Institute.

Mishel, L. and J. Bivens (2021a) “Identifying the Policy Levers Generating Wage Suppression and Wage Inequality,” Economic Policy Institute Report,May 13.https://www.epi.org/unequalpower/publications/wage-suppression-inequality.

Mishel, L. and J. Bivens (2021b) “The Failure of Automation and Skill Gaps to Explain Wage Suppression or Wage Inequality,” Economic Policy Institute Report, May 20.

Mishel, L. and K-F. Gee (2012) “Why Aren’t Workers Benefiting From Labour Productivity Growth in the United States?” *International Productivity Monitor*, Number 23, Spring.

Mishel, L. and J. Kandra (2020) “Wages for the Top 1 per cent Skyrocketed 160 per cent Since 1979 While the Share of Wages for the Bottom 90 per cent Shrunk: Time to Remake Wage Pattern with Economic Policies That Generate Robust Wage-Growth for Vast Majority,” Working Economics (Economic Policy Institute blog), December.

Mishel, L., Shierholz, H. and L. Schmitt (2013) “Don’t Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality,” Economic Policy Institute.

Mishel, L., Bernstein, J. and J. Schmitt (1997a) *"The State of Working America 1996–1997"* An Economic Policy Institute book. M.E. Sharpe.

Mishel, L., Rhinehart, L. and L. Windham (2020) “Explaining the Erosion of Private-Sector Unions: How Corporate Practices and Legal Changes Have Undercut the Ability of Workers to Organize and Bargain,” Economic Policy Institute.

Naidu, S., E.A.Posner and E. G. Weyl (2018) “Antitrust Remedies for Labor Market Power,” *Coase-Sandor Working Paper Series in Law and Economics,* No. 850.

Ormiston, R., D. Belman, and M. Erlich (2020) "An Empirical Methodology to Estimate the Incidence and Costs of Payroll Fraud in the Construction Industry," *National Alliance for Fair Contracting.*

Rinz, K. (2018) “Labor Market Concentration, Earnings Inequality, and Earnings Mobility,” *U.S. Center for Administrative Records Research and Applications, Census Bureau,* Working Paper 2018-10.

Sharpe, A., Arsenault, J. and P. Harrison (2008a) “Why Have Real Wages Lagged Labour Productivity Growth in Canada?” *International Productivity Monitor*, No. 17, Fall, pp. 17–26.

Sharpe, A., Arsenault, J. and P. Harrison. (2008b) "The Relationship between Productivity and Real Wage Growth in Canada and OECD Countries: 1961-2006," CSLS Research Report 2008-8, http://csls.ca/reports/csls2008-8.pdf.

Shierholz, H. (2021) “The Importance of Enforcement in Designing Effective Labor Market Tools.” in *Combating Inequality: Rethinking Government’s Role*, edited by Olivier Blanchard and Dani Rodrik. MIT Press.

okolova, A. and T. Sorensen (2020) “Monopsony in Labor Markets: A Meta-Analysis,” *Washington Center for Equitable Growth.*

Stansbury, A.M. and L.H. Summers (2020) “Declining Worker Power and American Economic Performance.” *Brooking Papers on Economic Activity.*

Starr, E., Prescott, J.J. and N.D. Bishara (2020) “Noncompete Agreements in the U.S. Labor Force.” in Journal of Law and Economics, online version May 7, 2020.

Starr, E (2019b) "The Use, Abuse, and Enforceability of Non-Compete and No-Poach Agreements: A Brief Review of the Theory, Evidence, and Recent Reform Efforts," Economic Innovation Group.

Stiglitz, J.E (2021) "Fostering More-Competitive Labor Markets," in *Inequality and the Labor Market: The Case for Greater Competition*, edited by Sharon Block and Benjamin Harris. Brookings Institution Press.

Stone, K. V.W., and A. J.S. Colvin (2015) "The Arbitration Epidemic: Mandatory Arbitration Deprives Workers and Consumers of Their Rights", Economic Policy Institute

Struyven, D. (2018) “Superstar Firms, Super Profits, and Shrinking Wage Shares,” *US Economics Analyst,* February 4.

Uber Technologies (2019) “Form S-1 Registration Statement, United States Securities and Exchange Commission,” April 11.

Viscelli, S (2016) *The Big Rig: Trucking and the Decline of the American Dream,* University of California Press.

Webber, D.A. (2020) “Labor Market Competition and Employment Adjustment over the Business Cycle,” *Journal of Human Resources,* Vol. 56, No. 2.

Webber, D. A. (2015) “Firm Market Power and the Earnings Distribution” *Labour Economics,* Vol. 35, pp. 123-134.

Weil, D (2017) “Lots of Employees Get Misclassified as Contractors. Here’s Why It Matters,” *Harvard Business Review*, July 5.

Weil, D (2019) " The Present and Future of Work in the Fissured Workplace Context," *Russell Sage Foundation Journal of the Social Sciences*, Vol. 5, No. 5.

Wilmers, N (2018) "Wage Stagnation and Buyer Power: How Buyer-Supplier Relations Affect U.S. Workers' Wages, 1978 to 2014," *American Sociological Review*, pp. 1–28.

Appendix A: Data Definitions and Sources

Average hourly compensation: compensation equals wages and salaries and benefits. Benefits are the sum of Health Benefits and 'Non-health benefits' ['Total compensation' (Nominal, NIPA 6.2A-D line 1) - 'Health benefits' Wages and salaries (NIPA 6.3A-D line 1) deflated by CPI-U-RS. Compensation is divided by 'Total hours' (NIPA 6.9B-D) to obtain real hourly compensation.

Average hourly compensation (production/non-supervisory workers): Average hourly earnings (AHE) of production/non-supervisory workers is from: BLS, series ID = CEU0500000008 and for years 1947–1963, series ID = EEU00500006 for historical AHE. Production/non-supervisory worker compensation obtained by multiplying real average hourly earnings by the compensation-wage ratio.

Compensation-to-wage ratio: the ratio of real compensation (wages and salaries plus health and non-health benefits) to real wages and salaries.

Consumer prices: measured as changes in the CPI-U-RS index, https://www.bls.gov/cpi/research-series/r-cpi-u-rs-allitems.xlsx

Gross productivity: gross output provided in unpublished BLS Total Economic Productivity (TEP) "Labor productivity", which matches GDP in NIPA data. Hours worked from TEP data. Gross productivity growth is growth of gross output per hour.

Health benefits: Nominal health benefits (NIPA 6.11A-C line 30, NIPA 6.11D line 32) are inflation adjusted by a constructed health inflator: CPI-U-RS times PCE Health (NIPA 2.5.4 line 37) / PCE total (NIPA 2.5.4 line 1).The PCE health deflator is used because it more fully captures health costs than the CPI health deflator which only captures out-of-pocket costs.

Median hourly wage: Economic Policy Institute analysis of Current Population Survey (CPS) Outgoing Rotation Group (ORG) microdata, as presented in the State of Working America Data Library.

Median hourly compensation: real median hourly wage multiplied by the compensation-wage ratio (see: Average hourly compensation (production/non-supervisory workers)

Net productivity: Gross output converted to net output by multiplying the ratio of net-to-gross domestic product ratio [NDP (NIPA 1.7.6 Line 4), GDP (NIPA 1.7.6 line 1)]. Net productivity presented deflated by output prices, i.e. implicit price deflator of Net Domestic product (NDP) and also by consumer prices (CPI-U-RS).

National Income and Product Accounts (NIPA): Department of Commerce, Bureau of Economic Analysis: https://apps.bea.gov/iTable/iTable.cfm?reqid=19step=2reqid=19step=2isuri=11921=survey

Wedges between median hourly compensation and productivity

- **Total net productivity-median hourly compensation divergence**: this is what the decomposition explains. It is the difference between growth of real net productivity (at output prices) and real median hourly compensation.
- **Inequality of compensation**: difference between growth of real average hourly compensation and real median hourly compensation.
- **Loss in labour's share of income**: difference between growth of net productivity growth (at consumer prices) and real average hourly compensation.
- **Divergence of consumer and output prices**: difference between net productivity growth (at output prices) and net productivity growth (at consumer prices).

The Evolution of the Productivity-Median Wage Gap in Canada, 1976-2019

Andrew Sharpe and James Ashwell
Centre for the Study of Living Standards [1]

Abstract

The median wage is a key metric to assess developments in the standard of living of the population. Productivity gains are passed on to workers as real wage gains. But in recent decades the proportion of labour productivity gains that are being passed on to the typical or median worker has fallen in many advanced countries, a process known as decoupling. The article uses an accounting framework developed by the Centre for the Study of Living Standards to quantify the importance of the factors affecting the relationship between productivity and real median wages. It presents results for the 1976-2019 period in Canada. A key finding is that the annual gap between labour productivity growth and real hourly median wage growth fell from 1.36 percentage points per year in 1976-2000 to 0.46 points in 2000-2019. This was due to slower growth in wage inequality, the end of the decline of the labour share and an improvement in workers terms of trade. Productivity growth was relatively stable between periods. In the 1976-2000 period, the bargaining power of workers fell dramatically due to high unemployment, falling unionization rates and a rising import share. After 2000, these trends reversed or stabilized, improving the bargaining power of workers.

How does one assess developments in the standard of living of the population? Since labour income or wages is by far the most important source of income, trends in wages for the typical or median worker appears an obvious metric. Indeed, *The Economist* (April 10-16, 2021) concludes "It is right to judge economic progress by the purchasing power of median wages, not profits or share prices." In the long run, wages are determined by productivity growth. Productivity gains are passed on to workers as real wage gains. But in recent years the proportion of labour pro-

1 Andrew Sharpe is Executive Director of the Centre for the Study of Living Standards (CSLS). James Ashwell, now a law student at McGill University, was a summer student at the CSLS when this research was undertaken. This article draws on Ashwell (2021). The authors would like to thank IPM Co-Editor Bart van Ark, Bruno Rainville and other officials at the Employment and Social Development Canada (ESDC), and four anonymous referees for useful comments. The CSLS thanks the ESDC for financial assistance for this project. Email: andrew.sharpe@csls.ca.

ductivity gains that are being passed on to the typical or median worker has fallen in many advanced countries, a process known as decoupling.[2]

The Centre for the Study of Living Standards has developed an accounting framework or methodology to quantify the importance of the factors affecting the relationship between productivity and real median wage growth.

The main objective of this article is to update these estimates on decoupling in Canada and the factors behind it to 2019. The article also seeks to provide a narrative to explain the reasons for the gap between productivity and median wage growth over the 1976-2019 period and in particular why this gap fell from 1.36 percentage points per year in the pre-2000 period to 0.46 points in the post-2000 period.

The bargaining power of labour largely determines the ability of labour to share in the overall productivity gains of the economy (Summers and Stansbury, 2020). This bargaining power is affected by labour market conditions, as proxied by the unemployment rate, by the strength of collective bargaining institutions, as proxied by the unionization rates and by international trade developments related to globalization, as shown by the merchandize imports share of GDP. This article shows that after 2000 the average unemployment rate was lower, the unionization rate, after falling sharply in the 1976-2000 period, fell at much slower rate, and the imports share fell from its 2000 peak. These trends reversed the downward pressures on bargaining power of labour that existed in the last quarter of the 20th century that resulted in the emergence of a large gap between labour productivity and median wages growth,

This article has five sections. The first section briefly reviews the literature of the decoupling of productivity and median wage growth in Canada. The second section presents the accounting framework which underpins the analysis of the decoupling issue. The results are presented in section three with special attention to the large fall in the gap after 2000. Section four examines the reasons for this development, The fifth and final section concludes.[3]

Review of Empirical Estimates for Canada[4]

The productivity-wage gap has become an important object of study in economics, and estimates of its magnitude have been conducted in a number of countries. For Canada, Sharpe *et al.* (2008a) quantify this gap and develop the decomposition accounting approach used in this article. Median earnings barely grew from 1980 to 2005, increasing 0.01 per cent per year, av-

2 In this article the term productivity refers to labour productivity. The term median wages refer to real median wages unless otherwise specified.

3 See Appendixes 1 and 2 in Ashwell (2021) to compare our results with Uguccioni (2016) and Williams (2021), who also study Canada. Appendix 3 compares the trends identified in this article with those of the United States identified by Mishel and Bivens (2021). The Data Appendix available at http://www.csls.ca/ipm/41/IPM_41_Data_Appendix.xlsx contains all the data used in this report along with some supplemental series pertaining to wages and productivity.

4 For a review of estimate of decoupling for the United States, the U.K. and other countries, see (Ashwell, 2021)

erage earnings grew 0.36 per cent per year, while productivity grew 1.27 per cent per year. Increased mean-median inequality explained 28 per cent of the gap and the decline of the labour share explained 20 per cent. Increases in supplementary income explained a further 20 per cent and the loss in labour's terms of trade explained 33 per cent.

Uguccioni, Sharpe and Murray (2016) use the same methodology to update the numbers for productivity and wages to 2014. They find that productivity grows from 1976-2014 by 1.12 per cent per year while median earnings grew at 0.09 per cent per year, for a productivity-wage gap of 1.03. Fifty per cent of the gap can be explained by increased mean-median earnings inequality, 30 per cent by the decreasing labour share, and 20 per cent by labour's terms of trade.

Williams (2021) investigates the same essential question as the previous authors but modifies the approach slightly to account for recent debates about measurement and variable selection. Williams argues that depreciation and taxation costs must be accounted for when considering productivity and wage trends, and he develops measures of net productivity and net labour share to integrate these considerations to his analysis. Williams finds that before accounting for taxes and depreciation, labour productivity rises from 1961-2019 by 1.65 per cent per year, whereas it rises 1.47 per cent per year after including them.

Instead of median wages, Williams prefers average compensation measures to evaluate transmission of productivity gains to workers. He argues that this allows a clearer picture of how labour overall is faring, but does not provide insight as to the distribution of those gains among workers. Deflated using the CPI, average wages grew 1.59 per cent per year over the 1961-2019 period, and using Williams' preferred implicit consumption deflator they grew 1.73. The labour share of GDP before accounting for depreciation and taxes therefore decreases over the same period by 0.19 per cent per year, but after accounting for those factors and calculating labour's share of NDP it decreases just 0.01 per cent per year. Labour productivity growth advanced 1.67 per cent per year, while net productivity growth was 1.47 per cent. These growth rates imply that there was no decoupling between productivity and average wages (not median wages) over the last 60 years in Canada.

Accounting Approach to the Decomposition of the Productivity-Median Wage Gap

Basic models of the labour market predict that workers overall will be paid an amount roughly equal to the marginal economic value that they provide to employers. As workers generate more value, their compensation should therefore rise accordingly. "Productivity" measures the amount of value that workers provide their employers, in terms of dollars per hour worked. Workers are generally paid on a per hour basis, so growth in productivity should be equal to growth in hourly pay. We measure the growth of productivity in per cent change per year for a given period, and measure the growth in median hourly pay over the same period. Subtracting the rate of growth in wages from the rate of growth in productivity, we obtain the gap between the two variables in percentage points. To understand this gap, we break it down into four components, also given in percentage points, which add up to the overall productivity-wage gap.[5]

The first component is the labour share. When productivity goes up, part of the benefits of that increase go to workers, and the other part of it goes to capital. Roughly speaking, capital is anything used in production other than the labour of workers. The proportion of economic benefits going to labour has historically been steadily around 50-60 per cent in Canada, with the other 40-50 per cent going to capital. These proportions are referred to as the labour share and the capital share, and while they have historically remained at similar levels, these shares do change over time. If a higher proportion of the benefits of productivity growth goes to capital, then the labour share becomes smaller, and workers obtain less than they normally would from the increased productivity, generating a productivity-pay growth gap.

The second component of the gap is called "labour's terms of trade", and it relates to price changes. Because of inflation, "real wages" and productivity must be calculated using constant dollar amounts which account for the differences in price changes over time. However, the prices of goods and services consumed by workers may not rise at the same speed as the prices of goods and services produced by workers. The adjustment of nominal wages must therefore be done either on the basis of consumer prices or producer prices.

The factor by which we multiply the nominal wage series in order to adjust for price changes is called a price index, which can be derived from the price levels in the economy as a whole, or from the price levels of goods and services consumed by workers. The first type of price index is called a GDP deflator, and the second is a measure of the prices of consumption goods, the most widely used example of which is the Consumer Price Index (CPI). To understand how workers' living standards are changing, we deflate their wages using the CPI since that deflator reflects the costs of living as experienced by workers. To deflate the output of the economy as a whole, we use the GDP deflator, because that deflator includes the prices of everything that is produced, rather than just the prices of goods and services which are consumed domestically. If the prices of consumer goods changes at a different rate than prices overall, there will be a gap between the growth of real wages from the point of view of workers and from the point of view of their employers. This difference can contribute to the overall gap between productivity and median real wages. The term we use for this component of the gap is "labour's terms of trade".

The third contributor to the productivity-wage gap is called the SLI/Self-employment component. This component is essentially the difference between the rates of growth of average compensation and average wages. The compensation measure is more exhaustive, as it includes supplementary labour income

5 For a formal presentation of the acoounting framework, see Ashwell (2021) and Sharpe *et al.* (2008).

(SLI) and an estimate of labour income for the self-employed, in addition to wages. SLI refers to compensation that employees receive from their employers beyond their regular wages, salaries and commissions, such as contributions to pension plans and to employment insurance. The labour component of self-employment income is estimated ("imputed") because there is no way of directly measuring how much of the income of the self-employed can be characterized as labour income as opposed to capital income, since these workers tend to invest both their time and their capital into their endeavors.

The final component is wage inequality, as proxied by the difference in growth rates of average and median wages. The average hourly wage is obtained by adding up wage income of all workers and dividing it by the number of total hours worked in a year. The median hourly wage is the wage received by the worker in the very middle of the wage distribution. Put another way, the median wage is the wage of the worker for whom the number of people who earn more is equal to the number of workers who earn less. If workers in the top of the distribution enjoy faster wage growth than everyone else, then the average wage will rise faster than the median wage.[6] The inequality component measures the difference between the rates of growth between median and average wages. With this fourth component in place, we can fully explain the gap between productivity and median hourly wages.[7]

Empirical Results[8]

The Labour Productivity and Median Wage levels

Chart 1 shows trends in the absolute levels of labour productivity, defined as real output per hour expressed in 2012 dollar, and real hourly median wages, also expressed in 2012 dollars, in Canada from 1976 to 2019. In 1976, the median wage in Canada was $16.40 per hour. Labour productivity was $37.60 per hour. In other words, the median worker received 43.5 pent cent of the amount of output produced in an hour of work.

By 2019, the median wage had grown to $17.40 and the level of labour productivity to $60.20. The median wage was now only 28.8 per cent of the average level of labour productivity. This development reflected the relative growth rates of the median wage and labour productivity over the period. Indeed, the median wage only increased 6 per cent from 1976 to 2019 while productivity was up by 60 per cent.

The median wage/productivity ratio will rise when the growth rate of the median wage exceeds that of the productivity growth. It will fall when the median wage advances less rapidly than productivity, as was the case in the 1976-2019 period. Equitable sharing of productivity gains for

6 It is worth noting that if the poorest workers benefit faster than everyone else, then the same phenomenon is observed: faster growth in the average than in the median. For this reason it is important to dig deeper into how the benefits of productivity are distributed among workers as done in Ashwell (2021).

7 For an algebraic presentation of the framework, see Ashwell (2021) and Sharpe *et al.* (2008).

8 See the online Appendix for data sources and variable definitions.

Chart 1: Median Real Wage and Labour Productivity in Canada, 1976-2019

Panel A: Absolute Level in 2012 dollars

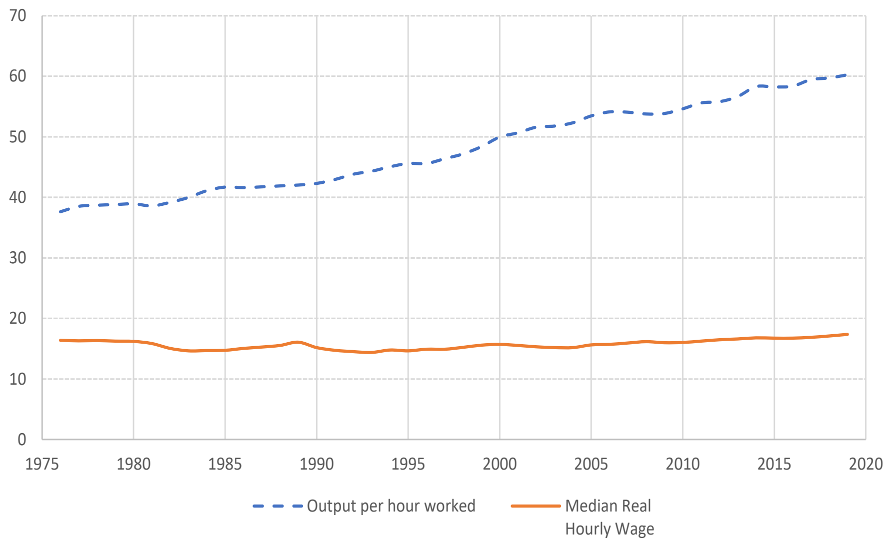

Panel B: Relative level (median real wage to output per hour)

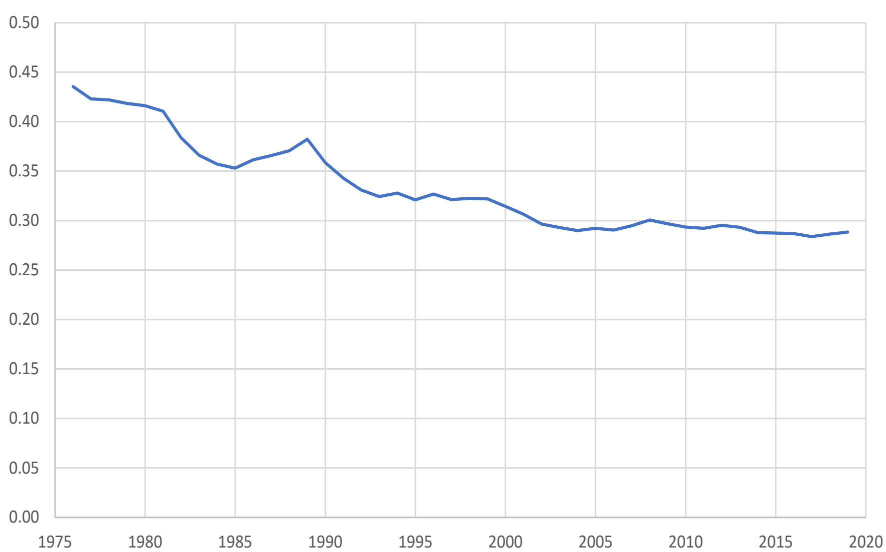

Source: Statistics Canada, See Data Appendix T3 and T10 for more details.

workers is defined as the same growth rates of median wages and productivity. Equitable sharing of productivity growth does not mean that workers receive the total value of the output they produce.

It can be noted that the absolute level of the median wage can rise even when the median wage/productivity ratio is falling,

The focus of this article is on the gap in growth rates of the median wage and productivity. It is this growth rate differential that determines the path of the median wage/productivity ratio.

The Productivity Median-Wage Gap Growth, 1976-2019

Panel A of Chart 2 shows the gap between growth in productivity and median wages in Canada for the overall 1976-2019 period, for four cyclically neutral peak-

Chart 2: Productivity Gap and Median-Wage Growth, 1976-2019 and Sub-periods

Panel A: Productivity-median hourly wage gap (percentage outputs per year)

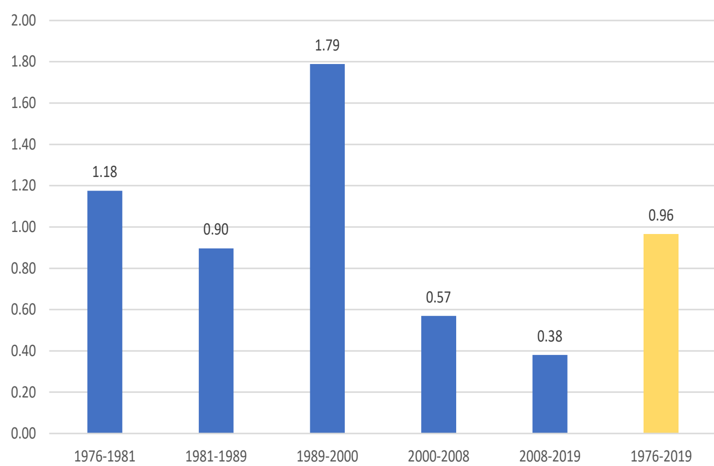

Panel B: Productivity and median wages, (average annual rate of change)

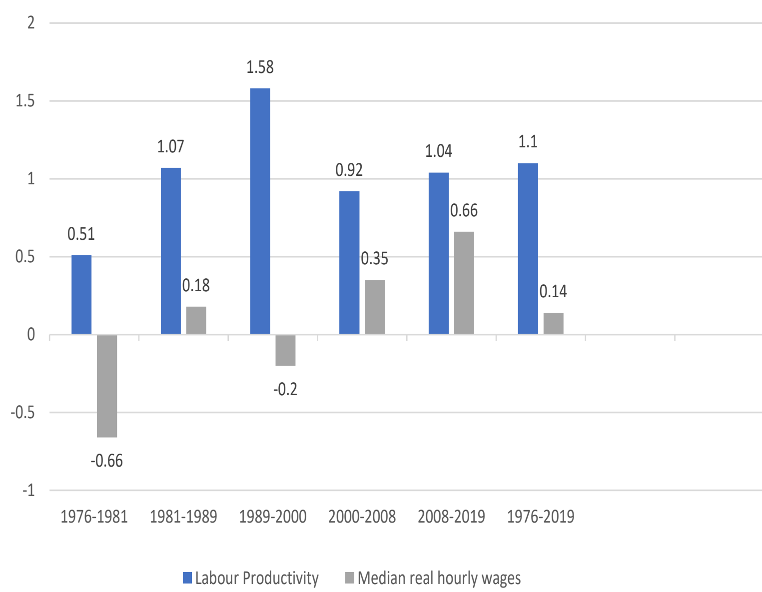

Source: Center for the Study of Living Standards (CSLS).

to-peak business cycles (1981-1989, 1989-2000, 2000-2008, 2008-2019) and for the incomplete business cycle at the start of the period of data availability (1976-1981). Panel B shows the growth rate for productivity and median wages.[9]

Chart 3 shows the time series between 1976 and 2019 for labour productivity and real median wages and three other series that can be used used to identify the factors

9 See Appendix Table 1 for the growth rates for the components of the gap for all periods. Appendix Table 2 for the absolute contributions of the components to the gap in all periods, and Appendix 3 for the relative contributions. See Appendix Chart 1 for the labour shares, Appendix Chart 2 for the labour terms of trade, Appendix Chart 3 for the difference between total compensation and wage, and Appendix Chart 4 for the inequality component as expressed by the ratio of average to median wages.

Chart 3: Productivity, Median Average Wages (CPI deflated), and Compensation (GDP deflated), 1976=100

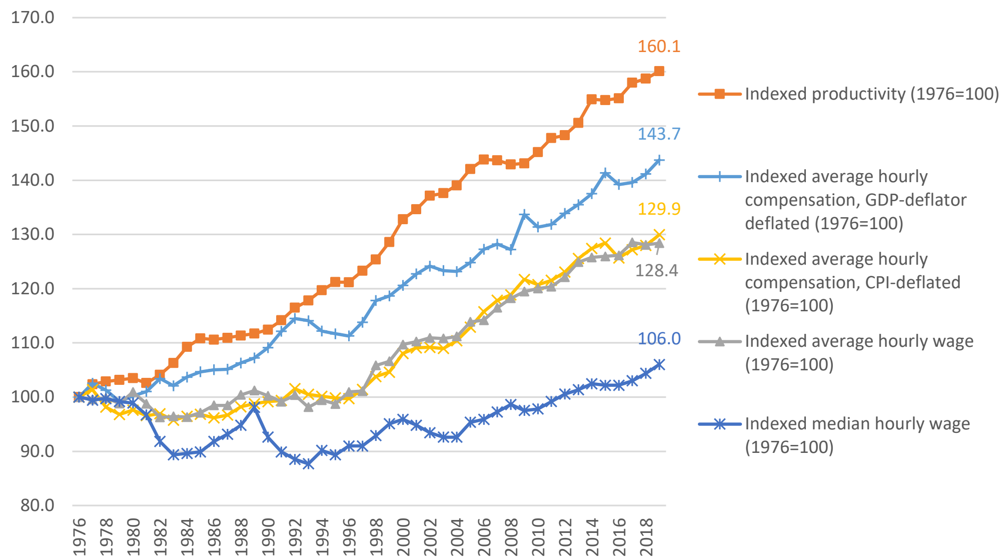

Source: Statistics Canada, see Data Appendix for details.

explaining the productivity-median wages gap. These series are labour compensation deflated by the GDP deflator, labour compensation deflated by the CPI, and average wages. The overall gap for the complete 1976-2019 period was 0.96 percentage points.

Based on the accounting framework presented earlier in the article, growing inequality between average and median worker wages accounted for 48 per cent of this gap, while a decrease in labour's terms of trade and in the labour share of GDP each account for roughly 25 per cent. Differences between the growth rate of supplementary labour income and the labour component of self-employed income and wage income account for the remaining part of the gap.

Change in the Productivity-Median Wage Gap Between 1976-2000 and 2000-2019

The size of gap between labour productivity growth and median wage growth in Canada in the first two decades of the 21st century was one third that of the last quarter of the 20th century: 0.46 percentage versus 1.36 points. The median worker has still not been fully benefiting from labour productivity growth, but he or she is doing much better. In 1976-2000, median wage growth was negative, resulting in the median worker receiving no benefit from labour productivity growth of 1.19 per cent per year. In contrast, in 2000-2019 median pay grew 0.53 per cent per year, slightly more than half the rate of productivity growth (0.99 per cent). In the 2013-2019 period the situation improved further, with median wage growth rising to three quarters of productivity growth.

This section provides an in-depth examination of the relative improvement in median pay relative to productivity first from an accounting perspective and then

Table 1: Basic trends (per cent annual growth)

Factors	1976-2000	2000-2019	Difference between periods	Percentage contributions
A. Labour productivity	1.19	0.99	-0.2	22.2
B. Labour share of nominal GDP	-0.4	-0.06	0.34	-37.8
C. Average real hourly compensation (GDP deflator)	0.78	0.93	0.15	-16.7
D. Average real hourly compensation (CPI deflator)	0.32	0.98	0.66	-73.3
E. Average real hourly wages	0.39	0.83	0.44	-48.9
F. Median real hourly wages	-0.17	0.53	0.7	-77.8
G. Productivity– median wage gap	1.36	0.46	-0.9	100
A. Real output per hour worked, constant 2012 dollars. Source: Statistics Canada, see Data Appendix T1 for details				
B. Total nominal labour compensation divided by total nominal GDP. Growth rate here shows change in that fraction. Total labour compensation includes imputed labour income for self-employed. Source: Statistics Canada, see Data Appendix T5 for details				
C. Total labour compensation (including imputed labour income of self-employed and SLI) divided by total hours worked, deflated with implicit GDP deflator. Source: Statistics Canada, see Data Appendix T6 for details				
D. Total labour compensation (including imputed labour income of self-employed and SLI) divided by total hours worked, deflated with CPI. Source: Statistics Canada, see Data Appendix T6 for details				
E. Average annual income from wages, salaries and commissions (excl. self-employed), divided by average hours worked (PA, incl. self-employed), deflated with CPI. Source: Statistics Canada, see Data Appendix T10 for details				
F. Median annual income from wages, salaries and commissions (excl. self-employed), divided by median hours worked, deflated with CPI. Source: Statistics Canada, see Data Appendix T10 for details				
G. Labour productivity (A) minus median hourly wages (E)				

Sources: The Centre for the Study of Living Standards (CSLS).

in terms of the fundamental factors driving the relationship. The growth rates of the components of the productivity-median wage relationship in 1976-2000, 2000-2019 and between the periods are first discussed, and the contribution of the four components of the decomposition examined. Developments in Canada between the two periods are compared with those in the United States and the United Kingdom.

1976-2000

From 1976 to 2000, labour productivity, defined as output per hour worked in the total economy advanced at a 1.19 per cent average annual rate (Table 1). In contrast, real hourly median wages actually fell 0.17 per cent per year. This resulted in a 1.36 percentage point annual gap between the growth rates of productivity and median pay. Three factors contribute roughly equally to this gap. First, the labour share of nominal income fell from 69.5 per cent of gross value added in 1976 to 63.1 per cent in 2000, a 0.40 per cent average annual rate of decline. This development accounts for 0.40 points or 30 per cent of the gap (Table 2 and 3).

Second, wage inequality rose as real hourly average wages grew 0.39 per cent per year, compared to -0.17 per cent for median wages, a difference of 0.56 points or 41 per cent of the total 1.36 point gap.

Third, the CPI rose at a much faster rate than the GDP deflator in 1976-2000, 4.78 per cent per year versus 4.30 per cent, a difference of 0.46 per cent per year. This meant that the consumer wage rose more slowly than the producer wage. Average hourly compensation deflated by the CPI rose 0.32 per cent per year compared to 0.78 per cent for average hourly compensation deflated by the GDP deflator. The difference of 0.46 points accounted for 34

10 The fourth factor, non-wage labour market income, made only a very small contribution to the gap. This

Table 2: Explanatory Factors for Productivity-Median Wage Gap (percentage points)

Factors	1976-2000 (A)	2000-2019 (B)	Difference between periods (C)=(B)-(A)	Percentage Contributions (D)=(C)/-0.90*100
A. Inequality	0.56	0.3	-0.26	28.9
B. Labour's share of income	0.4	0.06	-0.34	37.8
C. Labour's terms of trade	0.46	-0.05	-0.51	56.7
D. SLI/Self-employment	-0.06	0.14	0.2	-22.2
G. Sum of factors	1.36	0.46	-0.9	100
A. Average real hourly wages (E) minus Median real hourly wages (F)				
B. Total nominal labour compensation divided by total nominal GDP (C)				
C. Average real hourly compensation deflated with GDP deflator (C) minus Average real hourly compensation deflated using CPI (D)				
D. Average real hourly compensation (D) minus average real hourly wage (E)				
G. Sum of all factors				

Sources: The Centre for the Study of Living Standards (CSLS).

Table 3: Explanatory Factors for Productivity-Median Wage Gap (percentages)

Factors	1976-2000	2000-2019	Difference between Periods
Inequality	41.3	65.9	24.6
Labour's share of income	29.5	13.2	-16.3
Labour's terms of trade	33.9	-10.4	-44.3
SLI/Self-employment	-4.7	31.4	36.1

Sources: The Centre for the Study of Living Standards (CSLS).

per cent of the gap between productivity and median pay growth.[10]

2000-2019

After 2000, the productivity-median wage growth gap fell by two thirds (0.90 points) from 1.36 points to 0.46 points. All three factors that made large positive contribution to the gap in 1976-2000 made smaller contributions, or even negative contributions in 2000-2019 (Table 2). After falling significantly in the last quarter of the 20th century, the labour share stabilized in the first two decades of the 21st century of income more or less stabilized (63.1 per cent in 2000 versus 62.4 per cent in 2019). This factor now only contributed 0.06 percentage points or 13 per cent to the much smaller productivity-median wage gap of 0.46 percentage points.

Wage inequality continued to grow after 2000, but the pace was around one half that of the pre-2000 period. Median wages lagged average wages by 0.30 points from 2000 to 2019 (0.53 per cent versus 0.83 per cent), down from 0.56 points in 1976-2000. This represented about two thirds of the 0.46 point productivity-median wage gap.

In contrast to the pre-2000 period when the CPI inflation exceeded that of the GDP deflator, after 2000 CPI growth was 0.05 percentage points less that GDP deflator growth (1.88 per cent versus 1.93 per cent). This means that average hourly compensation deflated by the CPI rose 0.05 percentage points more per year compared to av-

factor encompasses supplementary labour income and the labour component of self-employed income and is included in total labour compensation. Its rate of growth is reflected in the difference between compensation growth and wage growth. As wage growth exceeded compensation growth (0.39 per cent per year versus 0.32 per cent) in 1976-2000, non-wage income grew at a slower rate year than wage income. This factor actually reduced the productivity-median wage gap by 0.07 percentage points or 5 per cent.

erage hourly compensation deflated by the GDP deflator (0.98 per cent versus 0.93 per cent). Instead of contributing significantly to the gap as it did in 1976-2000, labour's terms of trade, defined as the ratio of the trends in CPI inflation to overall economy inflation as expressed by the GDP deflator, improved after 2000 and reduced the productivity-median wage gap.

The fourth factor, non-wage labour market income, which had reduced the gap slightly in 1976-2000, now made a moderate absolute contribution to the gap in 2000-2019, and an important relative contribution. As noted, its rate of growth is reflected in the difference between compensation and wage growth. Compensation growth exceeded wage growth in 2000-2019 (0.98 per cent per year versus 0.83 per cent) since non-wage income such as employer contributions to social programs such as CPP outpaced the growth of wage income. This factor actually boosted the productivity-median wage gap by 0.14 percentage points or 31 per cent,

Change between 1976-2000 and 2000-2019

The explanation for the fall in the gap between productivity and median wage growth in Canada between 1976-2000 and 2000-2019 can be approached from two perspectives, first developments in the two variables themselves and second developments in the variables affecting median wages, as discussed above.

Two developments directly explain the fall in the gap by 0.90 points from 1.36 points to 0.46 points after 2000, slower productivity growth and much faster median wage growth. After advancing at 1.19 per cent per year in 1976-2000, labour productivity growth fell off to 0.99 per cent in 2000-2019 a fall on 0.20 points, or 22 per cent of the fall in the gap. The more dramatic development was the turnaround in median wages. After falling 0.17 per cent per year in 1976-2000, median wages advanced at a 0.57 per cent average annual rate in 2000-2019, an improvement of 0.70 percentage points, or 78 per cent of the fall in the gap,

As was discussed for the 1976-2000 and 2000-2019 periods, four factors mediate the difference between the growth rates of productivity and median wages, the labour share, labour's terms of trade, wage inequality and non-wage income. To understand changes between periods one looks at the changes in the absolute contributions of these four factors to the change in the gap (Table 2). The largest contribution to the fall in the gap was made by labour's term of trade, which has fallen significantly in the pre-2000 period and then slightly improved after 2000. This factor experienced a 0.51 point turnaround between periods and thus accounted for 57 per cent of the change in the gap, The relative stabilization of the labour's share after 2000, after falling before 2000 contributed 0.26 points to the fall in the gap or 38 per cent. Even though wage inequality was still rising after 2000, it was advancing at a slower pace. This change in growth rates between periods meant that the contribution of this factor fell 0.26 points, accounting for 29 per cent of the 0.90 point fall in the gap, the final factor is the growth of non-wage income, after making a negative contribution to the gap in 1976-2000 and a positive contribution in 2000-2019, the difference between these two contributions was

Table 4: The Progressive Positive Improvement of the Economic Situation of the Median Worker in Canada (average annual per cent or percentage point change)

Year & change	Median Wage (%)	Prod-Median Wage Gap	Median wage/ Prod Growth (%)
1976-2000	-0.17	1.36	-14
2000-2019	0.53	0.46	54
change	0.7	-0.9	68
2000-2008	0.35	0.57	38
2008-2019	0.66	0.38	63
Change	0.31	-0.19	25
2008-2013	0.55	0.5	52
2013-2019	0.75	0.28	73
Change	0.2	-0.22	21

Sources: CSLS estimates based Appendix Table 1.

0.20 points. Unlike the other factors that were working in the same direction to reduce the gap 0.90 points between periods, this factor worked to increase the gap.

Developments Since 2000 in Canada

Just as the gap between productivity and median wage growth fell, and sharing of productivity gains for the median worker improved, between the 1976-2000 and 2000-2019 periods in Canada, the same trends are observed after 2000 in the 2000-2008 and 2008-2019 subperiods (Appendix Table 1). The gap between productivity and median wage gap growth fell from 0.57 points in 2000-2008 to 0.38 points in 2008-2019 due largely to the virtual end the upward trend in wage inequality with the pick-up in median wage growth (0.66 per cent versus 0.35 per cent). The share of productivity gains going to the median worker rose from 38 per cent to 63 per cent.

A disaggregation of the 2008-2019 period into 2008-2013 and 2013-2019 sub-periods shows a continued improvement of the economic situation of the median worker in the more recent sub-period. The gap between productivity and median wage growth fell from 0.50 points in 2008-2013 to 0.28 points in 2013-2019 due largely to the virtual end of the upward trend in wage inequality with the pick-up in median wage growth (0.75 per cent versus 0.55 per cent). The share of productivity gains going to the median worker rose from 52 per cent to 73 per cent between period.

In summary, a comparison of the economic situation of the median worker in Canada shows a progressive improvement over time (Table 4). Between the last quarter of the 20th century and the most recent 2013-2019 sub-period, the rate of growth of median pay has risen from -0.17 per cent per year to 0.75 per cent, the productivity-median wage gap, although still positive has shrunk from 1.46 points to 0.28 points, and the proportion of productivity growth that the median worker received has shifted from nothing (-14 per cent) to 73 per cent.

Comparison with Developments in the United States

The fall in the gap between productivity and median wage growth identified in Canada after 2000 has also been observed

Table 5: A Comparison of the Evolution of the Productivity Median Wage Gaps in Canada, US and UK

Countries	Pre-2000			Post-2000			Change between periods		
	Productivity (A)	Median wage (B)	Gap (C)	Productivity (D)	Median Wage (E)	Gap (F)	Productivity (G)=(D)-(A)	Median Wage (H)=(E)-(B)	Gap (I)=(F)-(C)
US-Mishel (1973-2000)	1.49	0.13	1.36	1.5	0.41	1.09	0.01	0.28	-0.27
GSS (1976-2000)	1.2	0	1.2	1.5	0.7	0.8	0.3	0.7	-0.4
UK (1981-1996 and 2007-2019)	2.38	1.51	0.87	0.2	0.17	0.03	-2.18	-1.34	-0.84
Canada	1.19	-0.17	1.36	0.99	0.53	0.46	-0.2	0.7	-0.9

Note: GSS (Greenspon, Stansbary and Summers) .

in the United States, although to a smaller degree than in Canada, and also in the United Kingdom (Table 5).

Mishel and Bivens (2021)) report that the gap in the United States from 1973 to 2000 was 1.36 points, based on annual gross productivity growth of 1.49 per cent and media wage growth of 0.13 per cent. These growth rates are similar to those in Canada. In the 2000-2019 period, US productivity growth was virtually unchanged at 1.50 per cent, and median wage growth picked up to 0.41 per cent, reducing the growth gap to 1.09 points, This means that the fall in the gap between periods was 0.27 points in the United States about one third of the 0.90 fall in Canada. This smaller fall reflected the much weaker pick-up in median wage growth in the United States than in Canada and the failure of productivity growth to fall off after 2000, as it did in Canada.

Greenspon, Stansbury and Summers (2021) finds results comparable to those of Mishel and Bivens. For the 1976-2000 period, they report labour productivity growth of 1.20 per cent per year with no change in median wages over the period, resulting in a growth gap on 1.20 points per year. In the 2000-2019 period, productivity growth accelerated to 1.50 per cent per year, but median wages picked up even more to 0.70 per cent, with a reduction in the gap to 0.80 points with a fall in the gap of 0.40 points between periods,

As in Canada. the median worker in the United States has failed to fully benefit from labour productivity growth since the 1970s, but the gap has fallen after 2000. According to Mishel and Bivens, the median worker was receiving only 9 per cent of productivity growth in the 1973-2000, a share that rose to 27 per cent after 2000. With stronger median wage growth after 2000, Greenspan, Stanbury and Summers find a much greater sharing of productivity gains for the median worker, from none in 1976-2000 to 50 per cent in 2000-2019. It is interesting to note that around 50 per cent of productivity growth went to the median worker in Canada after 2000. While this represents an obvious improvement over the pre-2000 situation in both countries, it falls well short of the full sharing of the benefits of productivity growth.

Mishel and Bivens (2021:Table 1) use the same framework to decompose the factors explaining the divergence between productivity growth and median wages for the United States that this article uses for Canada. They find that in the 1973-2000 period it was the large increase in wage inequality that accounted for the lion's share of the gap (o.6 points of x per cent), with no contribution from a falling labour share. After 2000 the contributions of the two factors were similar. The fall in the gap between periods can be explained by the slower pace of increase of wage inequality.in the post-2000 period as well as the virtual end of faster consumer prices growth relative to producer prices.[11]

Comparisons with Developments in the UK

In the most recent study on decoupling in the UK (Teichgräber and Van Reenen, 2021), the pre-2000 and post-2000 productivity-median wage gap comparisons requires an analysis of the 1981-1996 and 2007-2019 periods.[12] In the first period labour productivity advanced at a 2.38 per cent per year, compared to 1.51 per cent for median wage growth, a gap on 0.87 points. In the second period, productivity growth collapsed to 0.21 per cent year, but median wage growth also fell to 0.17 per cent. A gap of only 0.04 points. In other words, the gap between productivity and median wage growth fell 0.84 points between periods. But the median worker was enjoying much fast wage growth in the first period when wage growth was strong even though he or she was only receiving 63 per cent of productivity gains, compared to the 81 per cent in the second period when wage growth was virtually non-existent. This is explained by the simple fact that 63 per cent of 2.48 is much greater than 83 per cent of 0.21. It is thus both the magnitude of the productivity growth rate and the sharing of this growth rate that determines the rate of growth of median wages. In the UK the problem for workers in the 21st century (or at least since 2007) has been the collapse of labour productivity growth (Oulton, 2019). In contract, in Canada and the United States, the problem has been the continued unequal sharing of the benefits of productivity growth.

Explaining the Progressive Relative Improvement of Median Wages in Canada[13]

This article has shown that the median worker in Canada fared very poorly during the last quarter of the 20th century, with the median wage falling despite pos-

11 Like Canada, the productivity-median wage gap has fallen in the United States since 2000. Between the 2000-2007 and 2007-2017 periods, the annual gap fell from 1.73 points to 0.74 points (Mishel and Bivens, 2021: Table 1). But the reasons for this development differ fundamentally between countries. As was noted earlier in Canada the fall in the gap between these periods was driven by the pick-up in median wages. In contrast, in the United States almost all the decline was caused by the massive fall-off in labour productivity growth from 2.19 per cent in 2000-2007 to 1.11 per cent. In 2007-2019. Median wage growth in the United States was actually slower in 2007-2019 than in 2000-2007 (0.33 per cent versus 0.65 per cent.

12 Growth rates are also available for the 1997-2007 period, which overlaps both centuries.

13 For a review of the literature on the factors affecting wages, see Ashwell (2021).

14 See the articles in Banting, St Hilaire, Sharpe (2001) for discussion of economic conditions in Canada bin the 1980s and 1990s.

Chart 4: Unemployment Rate in Canada, 1976-2019

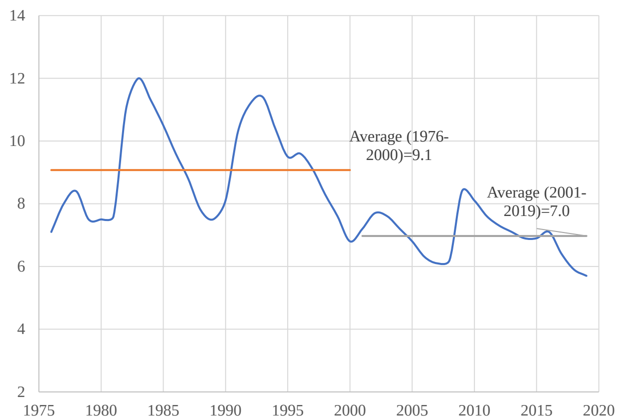

Source: Statistics Canada.

itive productivity growth.[14] The situation has progressively improved during the first two decades of the 21st century, although by the latter part of the period the median worker was still not receiving the full benefits of productivity growth. What explains this relative improvement?

At a general level, it is bargaining power that determines the proportion of productivity growth or gains that goes to workers (Summers and Stansbury, 2020; Dufour and Russell, 2015). Wage growth depends on both the rate of productivity growth and bargaining power so the latter is not the only story to explain wage growth. Weak productivity growth leads to weak wage growth even though bargaining power may remain unchanged, as has been seen in the UK in the 2007-2019 period. But in Canada productivity growth actually fell off after 2000 (0.99 per cent per year in 2000-2019 versus 1.19 per cent in 1976-2000) so stronger median wage growth is not due to an improved productivity performance. This means that the fall in the productivity-median wage gap in Canada is due to higher growth of the median wage reflecting improved bargaining power of the median worker.

Worker bargaining power is affected by demand and supply conditions in the labour market. Three key determinants of these conditions are the tightness of the labour market, as captured, for example, by unemployment rate, the collective bargaining power of workers as represented by the unionization rate, and globalization as shown by the import share.[15] Mishel and Bivens (2021) show that these three factors account for most of the gap between productivity and median wage growth in the United States over the 1979-2019 period. This article will show that developments on these three factors can account for the 0.90 point fall in the gap in the productivity-median wage growth in Canada between 1976-2000 and 2000-2019.

15 The nature of technological change may also be a factor affecting bargaining power, particularly for low-level skill groups who are affected by skill-biased technological change.

Chart 5: Unionization Rates in Canada, 1976-2019

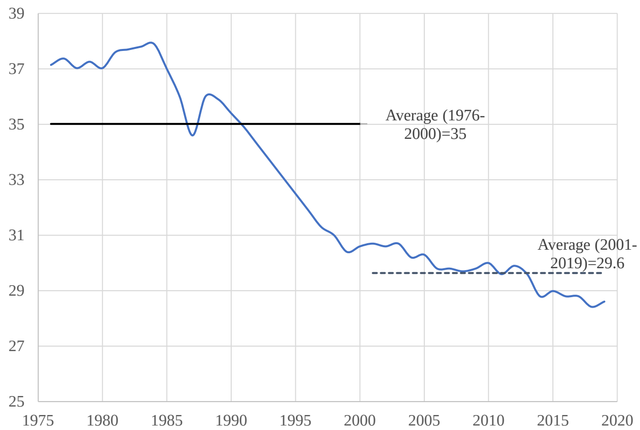

Source: Statistics Canada.

Labour market tightness

There are many measures of the degree of tightness in the labour market, but the most widely used is the official unemployment.[16] Chart 4 shows the evolution of the unemployment rate in Canada from 1976 to 2019. The downward trend after 2000 is readily apparent. The unemployment rate averaged 9.1 per cent from 1976 to 2000, then fell 2.1 points to average 7.0 per cent in 2001-2019. The recessions in the early 1980s and early 1990s resulted in large spikes in the unemployment rate, which took many years to unwind. The recession in 2008-2009 saw a much smaller run-up in the unemployment rate. In the second half of the 2010s the unemployment rate was fell below 7 per cent, reaching a low of 5.5 per cent in 2019. A key explanation for this improvement in macroeconomic conditions has been the transition from a regime of high real interest rates in the 1980s and 1990s to much lower rates in the 2000s and 2010s.

The high unemployment rate of the last quarter of the 20th century thus contributed to the inability of workers to maintain their purchasing power, with the median wage falling in real terms, despite productivity growth of over 1 per cent per year. The lower unemployment rate after 2000 allowed the workers to obtain higher wages, although still below the pace of productivity advance. Only in the second half of the 2010s when the unemployment rate fell below 7 per cent did median wage growth begin to approach productivity growth (0.75 per cent per year versus 1.03 per cent in 2013-2019).

Collective bargaining power

The rate of unionization is a well known measure of collective bargaining power. Chart 5 shows the evolution of the unionization rate in Canada from 1976 to 2019. The downward trend is readily apparent, especially before 2000. The unionization rate averaged 35.0 per cent per cent from

16 Additional measures of slack in the labour market are job vacancies and labour underutilization measures that include discouraged worker and involuntary part-time workers. These measures are strongly correlated with the unemployment rate.

Chart 6: Merchandise Imports as Share of GDP in Canada, 1976-2019

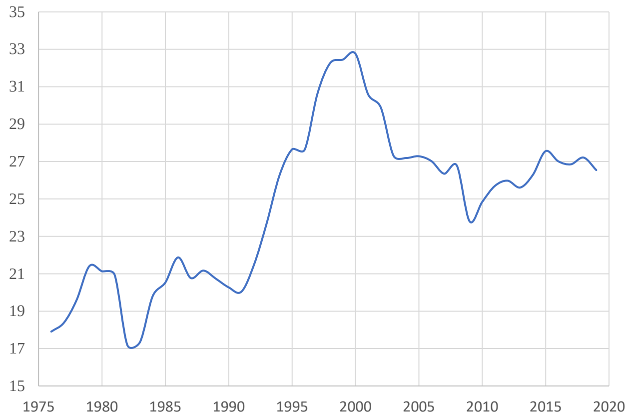

Source: Statistics Canada

1976 to 2000, then fell 5 points to average 29.6 per cent in 2001-2019. There was a strong union movement in Canada in the 1950s, 1960s and 1970s. For a variety of reasons union fortunes started to fall in the mid-1980s, with the unionization rate plummeting from 38 per cent in 1984 to 30.5 per cent in 2000.[17] Since 2000, there has been a continued downward trend although at a much slower pace with the unionization rate, down a further 3 points to 28 per cent by 2019.

The sharply falling unionization rate in pre-2000 period is consistent with the inability of the median worker to maintain his or her purchasing power.[18] The much smaller declines in the unionization rate after 2000 means the effect of deunionization, such as the adoption of two-tier wage schemes on pay increases, were less allowing, more opportunity for median real wage growth.

Globalization

Globalization can have myriad impacts on the bargaining power of labour through various channels.[19] A well-used indicator of

17 For an analysis of this decline, see (Morissette, Schellenberg, Johnson, 2005).

18 Dufour and Russell (2015) find union membership to be positively correlated with productivity-wage transmission in Canada at the 10 per cent significance level, and Card *et al.* (2004) show evidence that the decline in unionization in the United States and UK has contributed to higher income inequality in those countries. Mishel and Bivens (2021) argue that policy-driven erosions of bargaining power are the primary causes of workers' .

19 For example, Autor *et al.* (2020) and Schwellnus *et al.* (2017) have documented, how globalization has allowed successful firms to become dominant at a scale which was previously impossible and become so-called "superstar firms". This dominance can translate into monopsony power, meaning that large firms responsible for employing large shares of employees in a particular market have disproportionate power in that market. The scale of these firms can also translate into monopoly power in the product market, meaning they can raise consumer prices and obtain higher profits for capital in the form of monopoly rents, and thus reduce the labour share in that industry (Autor, Katz, Kearney, 2006). Autor *et al.* (2006) and Goldin and Katz (2007) propose the "skills-biased technological change" theory whereby technological changes and automation have led to higher demand for highly-skilled labour and lower demand for workers in the middle and bottom of the earnings distribution, thus increasing inequality. The increasing globalization of production in general and manufacturing in particular has also been proposed as a contributor to rising within-country inequality (Katz Murphy, 1992). Helpman (2016) summarizes the relevant literature on this hypothesis and finds that globalization and trade have had a "significant, yet modest" impact on wage inequality.

globalization is the import share, showing inability to capture the full gains of productivity. Chart 6 shows the evolution of the share of merchandise imports in GDP in Canada from 1976 to 2019. In the pre-2000 period this share was on a strong upward trend, rising from 17 per cent in 1976 to 32 per cent in 2000, with the lion's share of the increase taking place after 1990. After the 2000 peak the import share fell and by 2019 was at 27 per cent, still well above the pre-1990 level.

The jump in the import share from 1990 to 2000 was due to the implementation of the Canada-US Free Trade Agreement (FTA) starting in 1990 and the North American Free Trade Agreement (NAFTA) starting in 1994. These agreements boosted Canada's trade with the United States and Mexico, giving employers the opportunity to relocate production facilities in the country that minimizes labour costs while maintaining market access to all three countries. A number of manufacturers relocated production from Canada to lower-wage locations in the U.S. South and Mexico. The threat of relocation in the manufacturing sector also reduced the bargaining power of workers in wage negotiations and had negative spillover effects on wages in other sectors. The fall in the growth rate of median wages from 0.18 per cent per year in the 1981-1989 period to -0.20 per cent in 1989-2000 is consistent with this reduction in worker bargaining power. By the 2000s after firms had time to adjust to the new trade regimes, some of the downward pressure on wages from lower trade barriers may have dissipated. This is consistent with the pick-up in median wage growth to 0.35 per cent per year in 2000-2008, a turnaround of 0.55 points from 1989-2000.

In addition to the implementation of the FTA and NAFTA in the 1990s, the accession of China to the WTO lead to increased imports from China. The rise in the China's share of total Canadian imports was limited in the 1990s, from 1.8 per cent of total imports in 1990 to 3.4 per cent in 2000 (Murray, 2017:Table x) so this development accounted for little of the overall rise in the import share of GDP in the 1990s. The globalization associated with increase imports from China was concentrated in the first decade of the 21st century when imports from China rose from 3.5 per cent of total Canadian imports in 2000 to 13 per cent in 2010. Since then, the share has exhibited limited further progress at 14 per cent in 2015. This leveling off the negative effect of China trade on bargaining power of Canadian workers with the stabilization of the China import share in the 2010s is consistent with the pick-up of median wage growth to 0.66 per cent in 2008-2019 from 0.35 per cent in 2000-2008.

Conclusion

Reprising the quotation from *The Economist* at the beginning of this article that it is the median wage, not profit or share price, that is to be used to judge economic progress of an economy or society, one must conclude that progress in Canada over the 1976-2019 has been meager. Both profits and share prices have done well, but the median wage has advanced at only 0.14 per cent per year despite labour productivity growth of 1.10 per cent per year. About one half of this gap of 0.96 points between productivity and median wage growth is due to growing wage inequality, as manifested by faster growth of average versus median wages, with one quarter due to the fall in labour share of income and a second quarter arising from the fall in labour's terms of trade, as reflected in the faster increase in the CPI compared to the GDP deflator. The median worker's limited economic progress reflects their weak bargaining power to obtain wage increases from employers, compared to workers in the top half of the wage distribution who did considerably better. This weak bargaining power of the median worker in turn reflects a number of factors, especially high unemployment, falling unionization rates, and globalization leading to increased competition from imports.

Despite this dismal overall assessment of the economic progress of Canadians over the last near half century, a more nuanced picture emerges when the period is broken into sub-periods. In particular, the first two decades of the 21st century have experienced much more economic progress, with the median wage advancing 0.46 per cent per year compared to a fall of 0.17 per cent per year in the last quarter of the 21st century. With productivity growth relatively stable between periods at around 1 per cent, the pick-up in median wages reduced the gap between productivity and median wage growth by two thirds from 1.36 points to 0.46 points between 1976-2000 and 2000-2019. Since 2000, the median worker has received about one half of the gains from productivity growth, still very far from a full and equitable sharing, but a dramatic turnaround from the pre-2000 period when the median worker received no benefit from productivity growth. In the 1976-2000 period, the bargaining power of workers fell dramatically due to high unemployment, falling unionization rates and a rising import share. After 2000, these trends reversed or stabilized, improving the bargaining power of workers.

The situation of the median worker looks even better the closer one approaches the present. Median wage growth was higher over the 2008-2019 business cycle than over the 2000-2008 cycle (0.66 per cent per years versus 0.35 per cent). It was also higher during the second part of the most recent business cycle than the first half (0.75 per cent in 2013-2019 versus 0.55 per cent in 2008-2013. Indeed, in this most recent period the worker median was receiving nearly three quarters of the benefits of productivity growth. It is no coincidence than during this most recent period labour demand was strong, with the unemployment rate falling below 7 per cent and reaching a low of 5.5 per cent in 2019, the first time since the 1960s that a rate this low had been achieved.

A fully employed economy characterized by strong demand for the skills of workers in the bottom half of the wage distribution is the key to ensure that the median worker receives an equitable share of the real income benefits generated productivity growth. The current situation of robust labour demand and widespread labour shortages, as evidenced by the record numbers of job openings, has boosted median wage growth If this situation continues in the years to come, the economic progress in the Canadian economy and society will significantly outpace the dismal overall performance recorded over the 1976-2019 period, especially the 1976-2000 period.

References

Ashwell, J (2021). "A strengthening Position at the Bargaining Table? Understanding the Productivity Median-Wage Gap in Canada, 1976-2019," CSLS Research Report 2021-09, December.

Autor, D., D. Dorn, L. F. Katz and C. Patterson and J. Van Reenen (2020). The Fall of the Labor Share and the Rise of Superstar Firms. The Quarterly Journal of Economics, 135(2), May, 645-709. Retrieved from https://doi.org/10.1093/qje/qjaa004.

Autor, D. H., L. F. Katz and M. F. Kearney (2006). The Polarization of the US Labour Market. American Economic Review, 96(2), 198-194. https://doi:10.1257/000282806777212620.

Banting, K., A. Sharpe, F. St-Hilaire (2001). The Longest Decade: Canada in the 1990s, Thr Review of Economic Performance and Social Progress, Vol. 1, June (Ottawa Center for the Study of Living Standard and Montreal: Institute for Research on Public Policy).

Goldin, C. and L. Katz (2007). The Race Between Education and Technology: The Evolution of US Educational Wage Differentials,1890 to 2005. National Bureau for Economic Research. Retrieved from https://www.nber.org/papers/w12984.

Greenspon, J., A. Stansbury, L. H. Summers (2021). "Productivity and Pay in the United States and Canada," *International Productivity Monitor,* in Symposium on the Decoupling of Productivity and Pay in the United States, the United Kingdom and Canada, Fall, No. 41.

Helpman, E (2016). Globalization and Wage Inequality. National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w22944.

Katz, L. F. and K. M. Murphy (1992). Changes in Relative Wages, 1963–1987: Supply and Demand Factors. Quarterly Journal of Economics, 107, 35-78.

Mishel, L. and J. Bivens (2021). "The Productivity-Median Compensation Gap in the United States: The Contribution of Increased Wage Inequality and the Role of Policy Choices," *International Productivity Monitor,* in Symposium on the Decoupling of Productivity and Pay in the United States, the United Kingdom and Canada, Fall, No. 41.

Mishel, L. and J. Bivens (2015). Understanding the Historic Divergence Between Productivity and a Typical Worker's Pay: Why It Matters and Why It's Real. Economic Policy Institute. Retrieved from https://files.epi.org/2015/understanding-productivity-pay-divergence-final.pdf.

Mishel, L. and J. Bivens (2021). Identifying the Policy Levers Generating Wage Suppression and Wage Inequality. Washington, D.C.: Economic Policy Institute.

Morissette, R., G. Schellenberg and A. Johnson (2005). Diverging Trends in Unionization. Statistics Canada Perspectives, April, Catalogue no. 75-001-XIE.

Oulton, N (2019). "The UK and Western Productivity Puzzle: Does Arthur Lewis Hold the Key?," *International Productivity Monitor,* Centre for the Study of Living Standards, No. 36, Spring, pp.110-141.

Schwellnus, C., A. Kappeler and P. A. Pionnier, (2017). The Decoupling of Median Wages from Productivity in OECD Countries. International Productivity Monitor, 32(Spring), 44-60.

Sharpe, A., J. F. Arsenault and P. Harrison (2008a). Why Have Real Wages Lagged Labour Productivity Growth in Canada? International Productivity Monitor, 17(Fall), 16-27.

Sharpe, A (2002). Productivity Concepts, Trends And Prospects: An Overview. In K. Banting, F. St.-Hilaire and A. Sharpe (Eds.), The Review of Economic Performance and Social Progress (Vol. 2, pp. 31-56). Centre for the Study of Living Standards and the Institute for Research on Public Policy. Retrieved from http://www.csls.ca/repsp/2/andrewsharpe.pdf.

Stansbury, A. and L. Summers (2018). Productivity and Pay: Is the Link Broken? Washington, DC: Peterson Institute for International Economics. Retrieved from https://www.piie.com/system/files/documents/wp18-5.pdf.

Stansbury, A. and L. H. Summers (2020). The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy. Cambridge, MA: National Bureau of Economic Research. doi:10.3386/w27193.

Teichgräber, A. and J. V. Reenen (2021). "Have Productivity and Pay Decoupled in the UK?", *International Productivity Monitor,* in Symposium on the Decoupling of Productivity and Pay in the United States, the United Kingdom and Canada, Fall, No. 41.

Uguccioni, J. and A. Sharpe (2016). Decomposing the Productivity-Wage Nexus in Selected OECD Countries, 1986-2013. Centre for the Study of Living Standards. Retrieved from http://www.csls.ca/reports/csls2016-16.pdf.

Uguccioni, J., A. Sharpe and A. Murray (2016). Labour Productivity and the Distribution of Real Earnings in Canada, 1976-2014. Centre for the Study of Living Standard. Retrieved from http://www.csls.ca/reports/csls2016-15.pdf.

Williams, D. M (2021). Pay and Productivity in Canada: Growing Together Only Slower Than Ever. International Productivity Monitor.

Is Egypt Really More Productive than the United States? The Data behind the Penn World Table

Robert Inklaar and Pieter Woltjer
Groningen Growth and Development Centre
University of Groningen[1]

Abstract

A new feature in recent versions of the Penn World Table (PWT) is data on comparative levels of total factor productivity (TFP) across countries. TFP is defined as the efficiency with which inputs are transformed into outputs, and differences across countries can be due to factors such as better technology or better resource allocation. Yet, surprisingly, in PWT version 10.0, a number of low-income countries have a TFP level well above that of the United States. In this article we discuss the case of Egypt in 2017. PWT then reports a productivity level that is 23 per cent higher than that of the United States despite having an income level of only one fifth of the US level. We trace this anomalous outcome to the underlying data on comparative inputs. A fully satisfactory answer to the question in the title is elusive at this point, but the analysis highlights the data challenges that affect TFP level estimates, alongside more familiar modeling and measurement challenges.

One of the benefits of the development of the System National Accounts, and subsequent global measurement effort, is comprehensive and consistent cross-country data on consumption, investment and production. These data, in turn, can be used to systematically account for the sources of economic growth (Solow, 1957). Growth accounting leads to estimates of total factor productivity (TFP) growth, defined as the growth in output that cannot be accounted for by growth of factor inputs, capital and labour. Similarly, the development of measures of comparative levels of prices and output for a large set of countries (Kravis, Heston and Summers, 1978), opens the door for development accounting, which aims to assess how much of the

1 Robert Inklaar is Professor of Economics at the University of University of Groningen. At the time of writing, Pieter Woltjer was a postdoctoral researcher at the University of Groningen. The authors thank three anonymous reviewers and the editors, Bart van Ark and Andrew Sharpe, for their thoughtful comments and suggesstions. Any remaqining errors are our own. Email: r.c.inklaar@rug.nl; p.j.woltjer@gmail.com.

differences in income levels across countries can be accounted for by differences in input levels. The residual variation is variation in TFP levels across countries.

A result from the development accounting literature is that approximately half of the variation in GDP per worker can be accounted for by variation in factor inputs and the other half due to variation in TFP levels.[2] An implication of this result is that TFP levels show less variation than GDP per worker levels. For example, data for 2017 from the Penn World Table (PWT) version 10.0 show that the GDP per worker level of a country in the 95th percentile of the country distribution is 25 times higher than that of a country in the 5th percentile. For TFP levels, the multiple is only 3.6.[3]

This substantial cross-country variation makes understanding why TFP levels differ an important research question. Two broad (proximate) explanations can be relevant, namely that individual firms are less productive in one country than another, for example because of differences in technology adoption (Comin and Hobijn, 2010), or the allocation of resources between firms may be less efficient (Jones, 2011). Regardless of which of these explanations is most important, it implies that a high-productivity economy is more efficient in meaningful ways.

To the extent that high-productivity economies are also high-income economies, this implication seems to fit many people's priors. But in PWT, 10.0 some countries with income levels that are comparatively low exhibit TFP levels that exceed the TFP level of the United States. The most extreme example is Egypt, which in 2017 had a TFP level in PWT that was 23 per cent higher than in the US level, despite a GDP per capita level that is only one-fifth of the US level. Yet concluding that Egypt's firms are more technologically advanced or its economic system more successful in allocating resources to productive firms and industries than the United States may strike observers as implausible. Put simply, if Egypt's economy were truly so efficient, why are Egyptians not richer?[4]

Before drawing that conclusion, it is important to realize that measured TFP levels are the outcome of choices (and constraints) regarding the model, measurement and data. Modelling choices are about the underlying economic model and its assumptions, measurement choices are about how concepts of output and input are defined and measured, while data choices are about the approximations and assumptions that are necessary to operationalize the output and input concepts that the earlier choices prescribe. In this article, we emphasize choices and constraints regarding data; below we provide some discussion and examples of model and measurement choices.

In this article, we give an overview of the model, measurement and data choices that

2 See, for example, the surveys by Caselli (2005) and Hsieh and Klenow (2010).

3 See Feenstra, Inklaar and Timmer (2015) for the more recent version of the Penn World Table and section two from a summary

4 Whenever we talk of an economy being more or less productive or efficient, this should be read as a statement about total factor productivity (TFP), unless otherwise noted.

underlie TFP level data in the Penn World Table, version 10.0 (Feenstra, Inklaar and Timmer, 2015). We then zoom in on the case of Egypt as a marked outlier in terms of measured TFP level and we focus on data for 2017. We compare Egypt to a group of countries in the Middle East and North Africa region, because cultural, climatic or geographic factors may have a similar impact on outputs and inputs in these countries. We also compare Egypt to a group of countries at a similar income level, as income may affect outputs and inputs similarly. We find that, compared to these regional and developmental peer groups, Egypt has a low employment rate, low investment rate and a high price level of capital.

We compute counterfactual TFP levels, replacing observed Egyptian values by average values from the two peer groups and these counterfactual TFP levels are notably lower than observed ones. However, it is the combination of the three factors (employment, investment and capital prices) that leads to Egypt's status as a TFP outlier. This is clearly a conundrum for data users. Throwing out Egypt or any other country raises the question what the criterion should be and no simple criterion for the underlying data presents itself.

Earlier versions of PWT have reported letter grades to provide a sense of data quality but following the same grading logic would give Egypt the second-highest data-quality grade. Likewise, of the set countries with a statistical capacity number, Egypt is in the top 10 per cent. At the same time, accepting these numbers is hard as well, as discussed above. From the perspective of developers of the Penn World Table, there is no clear solution to this conundrum, other than to outline why it is such a conundrum.

The choice to focus on Egypt in 2017 should not be taken to imply that this is the only problematic case. For 59 of the 66 years for which PWT reports TFP level estimates for Egypt, the TFP level exceeds the US level. And beyond Egypt, there are 45 other countries for which PWT reports TFP levels that are higher than in the United States for one or more years. Many of those countries, such as Belgium, the Netherlands and Taiwan, have high income and labour productivity levels, so high TFP levels are no surprise. But this list also includes countries, such as Gabon and Jordan with notably lower income and labour productivity levels. We could also use the relationship between income level and TFP level to identify outliers and, again, there is a broader range of countries and years that deviates substantially from the cross-county pattern and could be classified as outliers. Despite this longer set of problematic/remarkable cases, focusing on the case of Egypt can be useful to illustrate how researchers may use the PWT data to gauge the plausibility of figures they are interested in. Those with a specific interest in data for one country may especially benefit from such a diagnostic approach before deciding how to proceed.

As discussed earlier, our emphasis in this article is on data choices and constraints, leaving aside modelling and measurement choices. The typical model underlying TFP level estimates is the Solow model and its assumptions on constant returns to scale, perfect competition, Hicks neutral technical change and no complementarities between inputs. Many of these modeling choices have been criticized, primarily in the context of growth accounting (Hulten, 2010). A general assessment of how changing these assumptions would affect estimated TFP levels is hard to give, but the impact could be substantial. Alternatively, different modelling choices may primarily affect the interpretation. As shown in Basu *et al.* (2020), traditional 'Solow residual' measures of TFP may still be relevant for consumer welfare, even if Solow-model production-side assumptions are not satisfied. This is because the Solow residual still reflects the trade-off between the output available for consumption versus the effort, in terms of labour hours supplied and deferred consumption that is used for investment, that is needed to produce that output.

Measurement choices also have an important impact on estimated TFP levels. For example, Lagakos *et al.* (2018) show that workers in high-income countries accumulate more human capital on the job than those in lower-income countries, an effect that is not accounted for in traditional measures of human capital (as used in PWT). As a result, factor input variation in PWT would be understated while TFP variation is overstated. Similarly, the productive use of subsoil assets, such as oil or iron ore, is typically not included as part of factor inputs, even though rents from their extraction are an important contributor to GDP in resource-rich countries in (predominantly) the Middle East and Africa. As Freeman, Inklaar and Diewert (2021) show, this omission of subsoil assets as a factor input leads to an overstatement of TFP levels in those countries.

It is important to note that TFP is — by construction — a residual and any measurement or data problem in output or inputs will be reflected in that number. Given the conceptual and practical challenges in measuring input of human and produced capital, this could prompt users to rely on measures that are less sensitive to such problems, such as comparative labour productivity. Such a choice risks throwing away the baby with the bath water as the TFP estimates for Egypt in 2017 (and a set of other countries and years) are outliers to a broader pattern of factors input use and productivity that fits more closely with economic intuition and theories. But, again, caution may be in order when zooming in on specific countries.

The article contains six sections. We first give a brief introduction to the Penn World Table in its current form in Section 1. In Section 2, we introduce a general development accounting framework that we use to measure TFP levels and that can be used to identify outliers. Section 3 covers the measurement of output and inputs in PWT. Section 4 presents the results of a development accounting analysis for 2017, followed in Section 5 by an in-depth analysis of the case of Egypt. Section 6 provides some concluding remarks.

The Penn World Table

The Penn World Table has a long history, originating in the pioneering work by Irving Kravis, Robert Summers and Alan Heston at the University of Pennsylvania to develop measures of comparative price and income levels that started in the

1960s.[5] The core feature of PWT has always been to combine National Accounts data on GDP, divided into consumption, investment and net exports, with data on comparative price levels for those same expenditure categories from the International Comparison Program (ICP) produced by the World Bank (2020). The result is a measure of "real GDP" that allows for comparisons of comparative income levels across countries, rather than only over time as in country National Accounts. And while ICP comparisons have been done at substantial intervals (5–6 years or more apart) and for an initially small group of countries, PWT has always provided annual data for global comparisons.

The article that introduced PWT version 5, Summers and Heston (1991), remains one of the most highly cited research papers in economics, in part due to its ubiquitous use in the literature on cross-country growth regressions but also as a standard dataset for measures of comparative income levels for most countries in the world since 1950.

With the release of PWT version 8.0 in 2013, the development of the database moved to the University of Groningen and the University of California, Davis. Feenstra, Inklaar and Timmer (2015) launched the "Next Generation of the Penn World Table", which introduced a series of measurement innovations:

1. Rather than relying on a single benchmark/reference year for comparative price levels, use each price benchmark. For example, PWT 7.0 and 7.1 were based on relative price data from ICP 2005 for the year 2005. For other years, relative prices are estimated based on inflation of each country relative to the reference country, the United States. In comparison, since PWT 8.0, relative prices for 1970 are based directly on data from ICP 1970 (for participating countries). This approach means that new releases of ICP do not lead to potentially major shifts in comparative income rankings going back in time, an approach that has since been adopted for more recent years by the World Bank (2020).

2. In earlier versions of PWT (and currently still in ICP), no explicit information was available about the relative prices of exports and imports. But in a world with many differentiated products and incomplete passthrough of exchange rate movements into prices, this is a substantial omission. As demonstrated in Feenstra and Romalis (2014), relative prices of imports and exports do vary substantially. Feenstra, Inklaar and Timmer (2015) provide a conceptual framework demonstrating the importance of accounting for these prices differences to draw sensible conclusions about the productive capacity of different economies.

3. PWT has traditionally emphasized measures of GDP, with a split by major expenditure category. However, for many questions, it is important to not only account for relative output, but also for relative inputs and productivity. Measures of comparative inputs and productivity were

5 A comprehensive history of their work and the development of the International Comparison Program (ICP) can be found here: https://www.rug.nl/ggdc/productivity/pwt/related-research-papers/heston$_{i}cp_{m}emoir_{2}017.pdf$.

introduced in PWT 8.0 and refined subsequently; we discuss this in more detail below.[6]

In 2021, version 10.0 of PWT was released, covering data for 183 economies and the period 1950–2019. The main data table, as well as a range of supporting datasets and documentation is available at www.ggdc.net/pwt.

Development Accounting

The tool we rely on to identify outliers in TFP levels is development accounting, which is typically used to assess the degree to which variation in observed per-capita factor inputs — capital and labour — can account for variation in output per capita. As we show in this section, one other outcome of such an analysis is to highlight the average relationship between factor inputs or productivity and output, i.e. countries with higher levels of output per capita tend to have higher levels of inputs per capita and higher productivity. Using this result, we can identify countries that fall outside this average range for more detailed scrutiny. The remainder of this section introduces the development accounting conceptual framework, but can be skipped without loss of continuity.

As detailed in Caselli (2005), the typical starting point in development accounting is an aggregate production function for country *m:*

$$Y_m = A_m f(K_m, L_m) = A_m f(K_m^{\alpha} L_m^{1-\alpha}) \tag{1}$$

A country's GDP, Y, is produced using production function f with input of capital K with input of capital L and labour and productivity level A. In equation (1) we assume a constant-returns to scale Cobb-Douglas production function with a constant output elasticity of capital α for expositional simplicity; in the implementation we rely on a translog production function.

Let a lower-case variable denote a quantity divided by country population, P_m, and let us express per capita quantities relative to the United States, to be indicated by a $\sim$.

This means that relative GDP per capita can then be expressed as

$$\tilde{y_m} \equiv \frac{Y_m/P_m}{Y_{US}/P_{US}}$$

Based on equation (1) and this notational convention, we can decompose a country's GDP per capita level relative to the United States into the contribution from differences in factor inputs and differences in productivity levels:

As discussed in Hsieh and Klenow (2010), this accounting for differences in GDP per capita levels answers the hypothetical question: by how much would GDP per capita increase if one of the factor inputs or productivity were to increase, holding constant the other two elements. This can be a sensible hypothetical when comparing growth over a short period of time as it is plausible to assume that the economy has not yet moved from one steady state to another. Yet when comparing across countries, it seems more plausi-

6 Some earlier versions of PWT did include measures of comparative capital stocks. In much of the literature on development accounting, researchers estimate their own measures of capital and productivity (Caselli 2005).

ble that the comparison is between countries in a (Solow model) steady state, i.e., where the investment response to the level of technology has worked itself out. Hsieh and Klenow (2010) argue that a more sensible hypothetical in a cross-country context would be based on:

$$\tilde{y_m} = \tilde{A}_m^{\frac{1}{1-\alpha}} \left(\frac{\tilde{k_m}}{\tilde{y_m}}\right)^{\frac{\alpha}{1-\alpha}} \tilde{l_m} \quad (2)$$

$$log\tilde{y_m} = (\frac{1}{1-\alpha})log\tilde{A_m} + (\frac{\alpha}{1-\alpha})log\frac{\tilde{k_m}}{\tilde{y_m}} + \log\tilde{l_m} \quad (3)$$

This equation rearranges the production function in intensive form, with the expression in logs in the second row. The benefit of this expression is that it accounts for the endogenous response of investment, and thus capital stocks, to differences in human capital and productivity. This follows the logic of the Solow growth model, in which the capital/output ratio of a country is constant in the steady state.[7] This is particularly relevant in a cross-country context, where differences in steady states are likely a larger factor in accounting for income differences than different positions relative to the steady state. Put differently, this decomposition does justice to the idea that an important reason for low capital levels in low-income countries is that productivity and human capital levels are lower.

Output and input levels in equation (3) are expressed in per-capita terms. As we also discuss in Section 4, labour input is an estimate of total hours worked, adjusted for the impact of schooling levels, $l = \frac{L}{P} = \frac{(N*H^a)*h}{P}$,where N is the number of workers. H^a is the average number of hours worked [8] and h is an index of the average years of schooling with an assumed rate or return to schooling. Capital input is based on capital stocks by asset, weighted using rental price weights.[9]

Equations (1)–(3) assume a fixed output elasticity of capital α. In PWT we follow Jorgenson and Nishmizu (1978), Schreyer (2007), Feenstra *et al.* (2015) and Inklaar and Diewert (2016) and assume a translog production function. From PWT 9.1 on, we also rely on a multilateral index, meaning the factor output and inputs are compared relative to a hypothetical average country based on all countries in our sample $c = 1, ..., C$. Given the translog production function we assume, the multilateral Törnqvist input index can be expressed as:

$$logQ_m = \alpha\left[logK_m - \overline{logK}\right] + (1-\alpha_m)\left[logL_m - \overline{LogL}\right] \quad (4)$$

with $\alpha_m \equiv$

$$\frac{1}{2}(\frac{r_m K_m}{R_m K_m + W_m L_m} +$$

7 In the Solow model, the parameters define a steady-state level for the capital/worker level and that implies a corresponding steady-state GDP/worker level. If productivity increases, then the marginal product of capital is higher at the initial capital/worker level, leading to new investment and a rise in the capital/worker and GDP/worker level. The capital/output ratio is the same in the initial and the new steady state.

8 Data on average hours worked are not available for all countries. When not available, we assume that average hours worked in the country equal those in the United States, so that relative TFP estimates (US=1.00) are not affected. See Section 6 for some further discussion of this assumption.

9 See Inklaar *et al.*. (2019) and the discussion in Section 3.

$$\frac{1}{c}\sum_{c=1}^{C}\frac{r_cK_c}{r_cK_c+W_cL_c})$$

the two-country average share of capital income in GDP, and $\overline{logK}$ the cross-country average of capital input levels, $\overline{logK} \equiv \frac{1}{c}\sum_c logK_c$. Equation (4) gives the input index relative to a hypothetical average country, but that index can be expressed relative to any reference country, such as the United States. This implementation of α implies assuming constant returns to scale, so that total income equals total cost, and perfect competition in factor markets so that inputs are used up to the point where marginal product equals marginal costs.

To identify extreme values of relative TFP and the factor output and inputs, we will assess the role of each term in equation (3) in accounting for income differences by estimating the following regressions:

$$\frac{1}{1-\alpha_m}log(\tilde{A_m}) = \beta^A log(\tilde{y_m}) + \varepsilon_m^K \quad (5)$$

$$\frac{\alpha_m}{1-\alpha_m}log(\frac{\tilde{k_m}}{\tilde{y_m}}) = \beta^A log(\tilde{y_m}) + \varepsilon_m^K \quad (6)$$

$$log(\tilde{l_m}) = \beta^A log(\tilde{y_m}) + \varepsilon_m^A \quad (7)$$

These equations use the expression in the second row of equation (3), so after taking logs. That expression states that the log of relative GDP per capita, $\tilde{y_m}$ is equal to the contribution from productivity differences, $\frac{1}{1-\alpha_m}log(\tilde{A_m})$, the contribution from differences in the capital/output ratio, $\frac{\alpha_m}{1-\alpha_m}log(\left(\frac{\tilde{k_m}}{\tilde{y_m}}\right)$ and the contribution from differences in labour input, $log(\tilde{l_m})$. To assess how much each of these three factors contributes to the overall variation in GDP per capita, we run regressions 5–7.

Since the sum of the dependent variables equals the independent variable, the coefficients β^A, β^K and β^L add up to one and can inform us of the relative importance of each term in accounting for cross-country income differences. This approach for assessing the contributions to income differences was first used in Inklaar *et al.* (2019). Compared to the variance decomposition of (Caselli, 2005), this approach has as a benefit that covariances between inputs need not be separately accounted for. We will use these equations here not to assess the β_s but to identify outliers, i.e. countries that are far outside the typical cross-country relationship between inputs, productivity and income levels. But, first, implementing equations 5-7 requires data on relative output and input levels.

Measurement of Output and Inputs

Current price GDP

We estimate real GDP by dividing GDP at current prices, in national currency, by purchasing power parities (PPPs) to correct for differences in prices across countries. "Real" in this context should thus be read as "in units comparable across countries". Nominal GDP data is readily available from the National Accounts as published by the United Nations. The primary contribution of PWT is in the estimation of PPPs at the level of consumption, investment, the trade balance and GDP for a long period of time.

As discussed in the previous section, the more recent versions of PWT - including

PWT 10.0, which we use in this article – use all available PPP benchmark data for estimating the PPP time series. So, if a country participated in, for example, the ICP comparison for 1980, then the PPPs for 1980 are based on data from that comparison. As the relative prices for final consumption and trade are now based on linked benchmark data, the annual changes in the price levels and real GDP are no longer (automatically) consistent with growth rates as reported in the national accounts, since (in general) consecutive PPP benchmarks are not consistent with national inflation.[10]

To facilitate research into economic growth of a single country over time, PWT also includes a GDP volume series with the growth rates over time identical to those in the National Accounts. In this article, we use the data for 2017, the year in which the most recent ICP benchmark comparison was held and the reference year for PWT 10.0. Using data for a benchmark year helps focus on the role of input data for productivity. The methods for estimating output and capital prices for non-benchmark years do not play a role.

Human capital

In PWT version 8.0, a human capital index based on the average years of schooling from Barro and Lee (BL, 2013) was introduced. The years of schooling were weighted using assumed rates of return to education, based on Mincer equation estimates by Psacharapoulos (1994).[11] This followed the approach of Caselli (2005) and assumes an average return on the first four years of 13.4 per cent, a rate of 10.1 per cent on years 5 to 8 and 6.8 per cent on every year of schooling beyond 8 years.

In PWT 9.0 the source for the years of schooling was revised to address criticism by De La Fuente and Domenech (2006) and Cohen and Soto (2007), who argue that the Barro and Lee data used source data inconsistently. The Barro and Lee (BL) data was supplemented with years of schooling data compiled by Cohen and Leker (CL, 2014).

We opt for either BL or CL depending on whether data for a country are only available from one of these sources, or whichever is closer to the level or trend over time in De La Fuente and Doménech (2006) and years of schooling data from UNESCO.[12]

The assumption of fixed rates of return to education across both time and space may underestimate actual differences in educational attainment between countries.[13] The human capital index could instead be based on country-specific weighting factors

10 This is partly due to index number reasons, as PPPs rely on expenditure shares for multiple countries while national inflation data uses only home-country expenditure shares. Yet most of the inconsistency cannot be readily traced to a clear source, see Inklaar, Marapin, Woltjer and Timmer (2021). As a result, though, these real GDP estimates are less suitable to measure changes over time in a single country.

11 In a Mincer equation, differences in wages are explained by differences in individual characteristics, such as education. The return to education is then given by the coefficient of education on wages.

12 The source of years of schooling data is listed for each country and year in the labour detail file available for download in the additional data and programs sections on http://www.ggdc.net/pwt.

13 We follow the standard implementation of Caselli (2005), though see Lagakos, Moll, Porzio, Qian and Schoellman (2018) for a broader view of human capital in a development accounting context.

using observed wage and employment data, as implemented in, for example, the Total Economy Database produced by the Conference Board (2021). Human capital could also vary between other dimensions, such as gender, work experience and occupation. Unfortunately, data on wages by employment category are not available for many of the lower-income countries in the PWT dataset or these do not span the full time series of the dataset. There is also no consensus about how much the quality of education differs.[14] The current PWT approach ensures that the widest range of countries can be incorporated in our growth and development accounting exercises and improves transparency.

Capital stocks and services

In PWT 9.1 we addressed two important shortcomings in the measurement of capital input. First, we estimated initial capital stocks based on better data and an improved procedure that does more justice to country-specific experiences. Second, we implemented a capital services methodology in accordance with standard productivity measurement theory. By doing so, we account for more of the cross-country variation in income levels. Inklaar *et al.* (2019) provide a full description of the estimation procedure. Below we provide a short summary and a discussion of the potential issues and extensions.

The quantity of capital input K_i for each of the nine assets i distinguished in PWT is typically not directly observable.[15] Instead, it is based on estimated net capital stocks N_i which are in turn based on the total accrued investment l_i depreciated over time using the Perpetual Inventory Method (PIM):

$$N_{i,t} = (1-\delta)N_{i,t}{-_1} + l_{i,t} \quad (8)$$

We next estimate the rental prices for each asset and take account of the differences in investment patterns, particularly evident between poor and rich economies. Following the framework of Jorgenson and Nishimizu (1978) — and more recently discussed in the OECD (2009) capital manual — the asset rental price at time t can be approximated as:

$$r_{i,t} = P_{i,t}^{N}\, i_t + P_{i,t}^{N}\, \delta_i - P_{i,t}{-_1^N} \frac{1}{5}\left(\sum_{T=t-4}^{t} \hat{P_{i,T}}\right) \quad (9)$$

where i_t is the required rate of return on capital, P_i^N is the purchase price of asset i, δ_i is the geometric depreciation rate and $\hat{p}$ is the percentage change in prices. To address volatile asset prices, we use a five-year moving average to estimate the change in asset prices. Assuming that the flow of capital inputs from a particular asset is proportional to the stock of that asset, N_iK_i, we can express the income flow from asset i (the capital compensation for asset i) as r_iN_i and estimate relative capital input for equation (4) as:

14 One line of evidence for this is based on comparing wages of immigrants to the United States depending on whether they were educated in a low-income or a high-income country (Schoellman, 2012). See Hsieh and Klenow (2010) on the difficulty of (fully) accounting for quality differences in education based on within-country estimates of the return to education.

15 The assets distinguished are residential structures, non-residential structures, transport equipment, information technology equipment, communication equipment, other machinery, software, cultivated assets and other intellectual property products.

$$LogK_m = \sum_i \frac{1}{2}(v_{i,m} + v_{i;})(logN_{i,m} - \overline{logN_i}) \quad (10)$$

where $V_{i,m} \equiv \frac{r_{i,m}N_{i,m}}{\sum_i r_{i,m}N_{i,m}}$ is the share of asset i in total capital compensation in country m, $v_{i;} = \frac{1}{c}\sum_c v_{i,c}$ is the cross-country average compensation share and $\overline{logN_i} = \frac{1}{c}\sum_c logN_{i,c}$ the cross-country average capital stock. As in equation (4), equation (8) is a multilateral Törnqvist index. For each asset, the capital stock of a country is compared to a geometric average of all countries and the differences in each capital asset, $\overline{logN_{i,m}} - \overline{logN_i}$ are weighted using the share of that asset in capital compensation.

In the standard Jorgensonian approach to rental prices, the required rate of return on capital is chosen to exhaust the income left after subtracting labour income from GDP. This gives an internal rate of return on capital and an important advantage is that this return sets 'pure profits' to zero and is thus consistent with the maintained assumption of perfect competition. An important drawback, in a global context, is that in some countries the rents from extracting natural resources like oil and gas is a sizeable fraction of GDP. For those countries, computing the internal rate of return based on the income that does not flow to labour would substantially overestimate the required rate of return on assets.[16] So instead, we determine the income flowing to capital as nominal GDP minus labour income minus natural resource rents.[17]

Outliers from development accounting

This concise overview of the main variables that are used for development accounting highlights that numerous assumptions and choices on measurement and are necessary to compile the output and inputs data to assess relative TFP across country. We now take a more in-depth look at data from PWT 10.0 for the year 2017. As discussed above, we choose to focus on data for the most recent PPP benchmark comparison.

Using the resulting data from PWT 10.0, we then estimate equations 5–7 for the year 2017, including data for the 114 countries with the required information for all variables. Chart 1 shows three scatter plots, for productivity, produced capital and human capital, against GDP per capita. The regression line is also plotted. The outlying levels of relative TFP in Panel A could reflect either extremely high or low productivity but could also result from measurement errors in either the inputs or output. Countries whose price level of GDP is identified as an outlier, for instance due to hyperinflation, are already excluded from Chart 1, so potential measurement errors are limited to capital output, labour input, the share of labour in GDP and nominal GDP.

Panel B demonstrates that most of the outliers identified in Panel A also show extraordinarily low or high relative levels of

16 Ideally, natural resources should be recognized as production factors in their own right. That is beyond the scope of this article but see Freeman, Inklaar and Diewert (2021).

17 Natural resource rents are from the World Development Indicators.

Chart 1: Relationship between GDP Per Capita, TFP, Produced Capital, and Human Capital in 114 Countries, 2017

Panel A: Total Factor Productivity

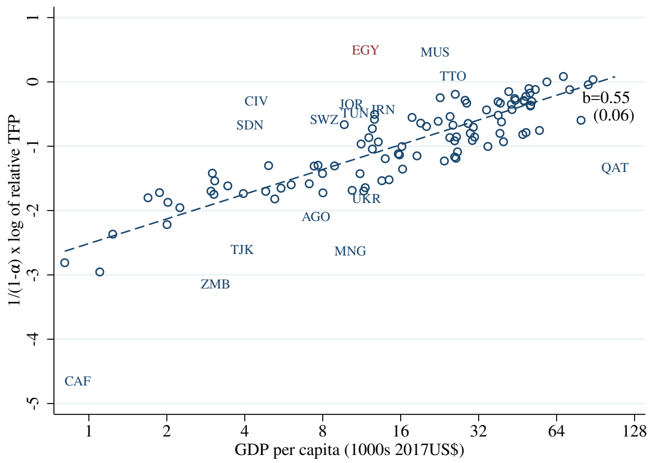

Panel B: Produced Capital

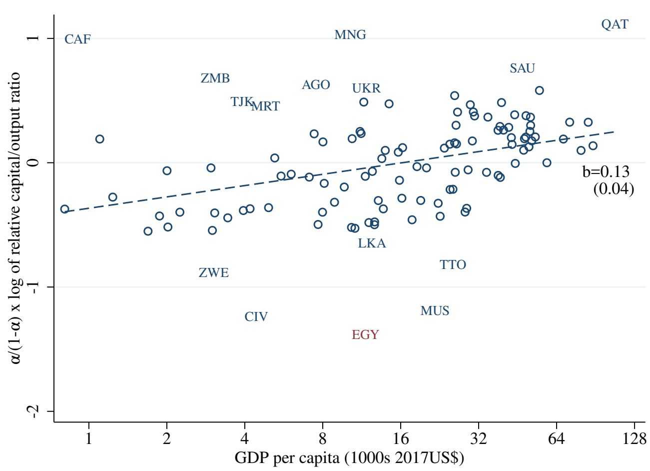

Panel C: Human Capital

Source: Penn World Table, PWT 10.0, Feenstra *et al.* (2015).

Note: Shown are the 114 countries in PWT 10.0 for which TFP estimates can be made; omitted are countries that did not participate in ICP 2017 as well as those already designated as outliers in PWT. The line shows the OLS line of best fit. Also shown is the slope coefficient and associated robust standard error.

capital to output. For Cote d'Ivoire (CIV), Egypt (EGY) and Mauritius (MUS), the very low level of capital intensity could explain at least part of the high relative productivity since relative human capital does not appear to be excessively low for these countries as shown in Panel C. For the Central African Republic (CAF), the low share of labour in GDP may be responsible for both the low observed level of relative TFP and high levels of capital to output in Panels A and B respectively.

The countries that are identified here as outliers, by deviating substantially from the main cross-country pattern, depend on the year of the analysis. The analysis could have been done for all benchmark years or all years in PWT. The qualitative results are very similar, in the sense that data for many countries fit the broader pattern of inputs, productivity and output. The list of countries that are outliers in one or more years would be longer.

Likewise, as discussed in the introduction, Egypt, Mauritius and Trinidad and Tobago (TTO) are not the only countries with a lower income level, yet a TFP level that exceeds the US level. This is to say that we could have picked from a longer list of countries for the case study that is to follow. Yet, the approach for the case study would be very similar, so it is a diagnostic tool that can be applied more broadly. More generally, we would recommend that users who are interested in a particular country should follow similar steps to see to whether their country fits the broader cross-country patterns or, if not, whether there are specific variables for which country observations are remarkable.

Case Study: Egypt

As a first step, Table 1 computes relative output, inputs and TFP compared to different groups of countries. This follows equation (4) and varies the set of countries C in the comparison. The first row shows the data directly from PWT with all 114 countries as reference group but, as in PWT, expressed with US=1.00. The second row is based on a multilateral comparison with 12 countries in the Middle East and Africa (MENA) region and the third row uses a group of 19 countries that is within 20 per cent of Egypt's level of GDP per capita. By expressing output, input and TFP relative to each reference country (group), we can highlight where data for Egypt are atypical.

The comparison versus the United States is like the standard presentation in PWT, though the numbers in Table 1 are a bilateral comparison, rather than a multilateral comparison with all countries. The MENA comparison group is chosen as regional factors, such as climate, geography and culture, are more similar within this group than with the overall world and it may be that these factors influence output and inputs in similar ways. The United Nations Arab Human Development Report project is one example of the usefulness of such a regional perspective. There is also important diversity in this region, for example, as some countries in the region rely heavily on oil and gas production while others do not. This is one reason why we also consider a second reference group, based on income level. The income reference group is chosen because, from Chart 1, we know that countries at similar income levels have more similar levels of inputs.

Table 1: Development Accounting for Egypt with Varying Reference Countries, 2017

Reference countries	y	l	Av. hours	HC	k	A
All, US=1.00	0.20	0.55	1.00	0.70	0.07	1.23
MENA=1	0.50	0.70	0.99	1.00	0.18	1.68
Similar income=1	0.99	0.71	0.96	0.97	0.38	2.07

Source: Penn World Table, 10.0, Feenstra *et al.* 2015)

Note: The table show levels for Egypt relative to reference countries. y: GDP per capita, l: employment/population ratio, Av. hours: average hours worked per worker, HC: human capital (years of schooling with assumed rates of return), k: capital/population ratio., A: total factor productivity. Total factor productivity is computed using equation (4) with varying set of countries. Row 1 uses all 114 countries in PWT, and expresses output, inputs and productivity relative to the United States. Row 2 uses Egypt plus 12 countries in the Middle East and North Africa (MENA) region with MENA=1 (BHR, IRN, IRQ, ISR, JOR, KWT. MAR, MLT, QAT, SAU, TUN). Row 3 uses Egypt plus 19 countries that are within 20 per cent of Egypt's GDP per capita level with group income=1 (ARM, BRA, BRB, CHN, COL, ECU, FJI, IDN, IRN, IRQ, JOR, LKA, MNG, NAM, PER, PER, TUN, UKR, ZAF

The first row expands on what we learned from Table 1. Egypt's GDP per capita level, y, is 20 per cent of the US level and its TFP level A is 23 per cent higher. Egypt's employment-to-population ratio is considerably lower, at 55 per cent of the US level; its human capital level stands at 70 per cent and its level of capital per head of the population is only 7 per cent of the US level.

Looking at rows 2 and 3 makes clear that the level of human capital in Egypt is comparable to that in the MENA region and the Similar income group, implying similar average years of schooling in the population. The most substantial differences appear for employment per capita l and capital per capita k. Egypt's employment to population ratio is only 70 per cent of the average of the MENA and Similar income groups and its capital to population ratio is only 18 per cent of the MENA average and 38 per cent for the income group. For these reasons, Egypt's TFP level relative to the MENA group is 68 per cent higher, despite an income level of only half of the MENA group. Egypt's average income level is, by construction of the group, very close to the income group, but its TFP level is 107 per cent higher.

Of note is that PWT does not have information on average hours worked in Egypt, so for that reason this is not a contributing factor to TFP differences. As discussed above, PWT does account for differences in average hours worked where available, but data is typically more abundant for higher-income countries. If not available, TFP calculations are done assuming the same number of average hours worked. This assumption is also made for many countries in the MENA and Similar income group, leading to small differences across the rows. From the work of Bick, Fuchs-Schündeln and Lagakos (2018), we know that lower-income countries tend to work more hours per adult. They do have estimates for average hours worked in Egypt and those numbers imply that the average Egyptian workweek at 48 hours is much longer than the average US workweek of 39 hours. If we would use these numbers, Egypt's TFP level relative to the United States would be 1.12 rather than 1.23, though the adjustment relative to the other groups would be smaller since the income differences are smaller.

Table 2 looks more closely at the

Table 2: Employment, Working-age Population and Total Population in the MENA Region (in per cent).

Country	N/P	N/P^A	P^A/P
Egypt	27	40	66
Iran	30	39	76
Iraq	23	37	61
Israel	44	61	72
Jordan	22	33	65
Kuwait	57	72	79
Morocco	30	41	73
Qatar	76	87	86
Saudi Arabia	39	52	75
Tunisia	30	40	76

Source: : PWT for employment and population, WDI for working-age population.

Note: N is employment, P is population and P^A is the working-age population, i.e., the population aged 15–64.

low employment-to-population ratio $l = N/P$ by dividing this ratio into the employment-to-working-age-population ratio, N/P^A and the share of working-age population, P^A/P for a set of countries in the MENA region. The rate N/P^A is low at 40 per cent, though five other countries in the table are close to or below this participation rate. Where Egypt stands out most is in its relatively low share of the working-age population, which is due to it having a large share of young people. Yet, Jordan and Iraq are similarly young and have similar employment-to-population ratio, which shows that Egypt's numbers are not beyond belief.

Table 3 examines capital input in more detail for the same group of countries. The aim is to understand the low level of capital input in Egypt. The final column corresponds to the capital input variable used for the 'produced capital' panel of Chart 1. Egypt clearly has the lowest level (0.34) of this group of countries, with Jordan (0.60) and Tunisia (0.58) closest. Starting from the first column, we can see that Egypt has a low investment rate, at 15 per cent of GDP. Only Iraq's investment rate, at 16 per cent, is close. This low investment rate is a longer-run feature of the Egyptian economy, as its nominal capital-output ratio is very low, at 1.74. Here again, Iraq's ratio is similar, at 1.72. Such a low investment rate could be a sign of under-recording of investment; also in a global comparison, there are few countries with investment that are so low. Such an investigation is beyond our scope, but in principle, cross-checks on data on imports or firm-level surveys could be useful.

The extent to which this low nominal investment rates translates into low levels of capital input depends on the relative prices for capital versus output. Comparing the GDP price and (capital) stock price columns shows that Egypt is almost the only country in the region for which the relative capital stock price is higher than the relative price of output; Tunisia is the only other country that break this pattern. When moving from capital stock prices to capital services prices, all countries show an increase in relative prices (US=1.00). The main factor is that all countries have a higher internal rate of return on produced capital than the United

Table 3: Investment and Capital in the MENA Region in 2017

Country	Investment (% of GDP)	K/Y (nominal,stock) (US = 1.00)	GDP Price (US = 1.00)	Capital Stock Price (US = 1.00)	Capital Services Price (US = 100)	IRR	K/Y (real services) (US = 1.00)
Egypt	15	1.74	0.17	0.20	0.78	0.28	0.34
Bahrain	28	3.29	0.48	0.30	0.74	0.17	1.14
Iran	20	3.77	0.41	0.26	0.86	0.09	0.75
Iraq	16	1.72	0.44	0.31	0.66	0.14	1.17
Israel	21	2.66	1.10	0.82	1.38	0.13	0.89
Jordan	20	2.47	0.40	0.31	0.85	0.17	0.60
Kuwait	27	2.20	0.53	0.30	0.69	0.09	1.42
Morocco	29	3.25	0.41	0.25	0.41	0.10	1.26
Qatar	45	2.63	0.56	0.32	0.64	0.16	1.77
Saudi Arabia	24	2.53	0.43	0.26	0.47	0.12	1.63
Tunisia	19	3.08	0.32	0.36	0.67	0.11	0.58

Source: Penn World Table 10.0 (Feenstra et al. 2015).

Note: Investment is gross fixed capital at current prices; K/Y (nominal, stock) is the current-cost net capital stock in local currency units over GDP; GDP Price is the purchasing power parity of GDP over the nominal exchange rate (XR) (US=1.00), Stock Price is the PPP for investment goods, weighted using the share of each asset in the current-cost net capital stock, over XR; Services Price is relative rental price (equation (7)) for each asset, weighted by the share of each asset in capital costs, over XR; IRR is the internal rate of return, the return that equates capital cost to GDP minus labor costs minus natural resource rents; K/Y (real, services) is capital services input (PWT variable ck, US=1.00) over real GDP (PWT variable CGDPo) relative to US real GDP.

States (0.07), which means capital costs are higher.[18] The increase is largest in Egypt, the country with the highest internal rate of return of this group.[19] So, in summary, Egypt has a low level of capital input, in part because the country devotes a relatively small share of its resources to investment purposes and a result of the high capital prices is that those resources buy relatively few capital goods.

These figures suggest two possibilities. First, it could be that all these statistics are a true reflection of Egypt's economy. This seems hard to accept, since it implies an improbably high TFP level for Egypt's economy. The second possibility is that Egyptian statistics are substantially mismeasured. This may be an attractive conclusion if the alternative is to accept that Egypt's economy is more productive than the US economy.

But an objective basis for such a conclusion is hard to find (other than that these data imply an improbable outcome). While Egypt has a low employment rate, a low investment rate and high capital prices compared to the two groups of comparison countries, they are not so far away from plausible measurements that they can be easily dismissed. For example, other coun-

18 See Inklaar *et al.* (2019) for an analysis showing that the internal rate of return tends to be higher in countries with lower income levels.

19 This high internal rate of return is needed to reconcile the high observed share of capital income in GDP (64 per cent according to PWT) with the relative low level of capital. Note that the bank lending rate, as shown in the World Bank's World Development Indicators (WDI), also shows Egypt with the highest lending rate (18 per cent), with lower rates in United States (4 per cent) and other countries in the region (though data coverage is incomplete).

tries with low investment rates are Brazil (15 per cent), Uruguay (16 per cent), Portugal (17 per cent) and Poland (18 per cent). It is also hard to find objective measures of the quality of the statistical system. The World Bank provides a 'statistical capacity' indicator, which is based on the frequency with which important data collection (e.g. an agricultural census) or revisions (to, e.g. consumption baskets for inflation) takes place. On this measure, Egypt in 2017 receives a score of 83 (out of 100), much higher than the average score of 53 for the MENA region and in the top 10 per cent of developing countries.

Several previous versions of PWT also included a letter grade (A–D) as an indicator of data quality.[20] This grade was based on three factors: 1) did the country participate in one or more official PPP benchmarks' 2) what was the inconsistency between consumer inflation and the change in consumption PPPs between benchmarks; and 3) how high is the country's income level. Income level is included because resources available for the statistical system are assumed to increase with income level. Factors 1 and 2 are especially geared primarily at PPP measurement.[21]

For this article, we replicated the letter grading using only factors 2 and 3, because all but three countries in PWT have participated in at least two official PPP benchmarks. The degree of inconsistency is measured between the two most recent PPP benchmarks, for 2011 and 2017, which is likely most relevant for the analysis of 2017 data. Following the earlier methodology, the inconsistency results are divided into five bins, with low inconsistency reflected in placement in a higher bin. Income levels are divided in six bins. The overall indicator is computed by giving the inconsistency bin score twice the weight of the income bin score as this final indicator is grouped into four bins. Following this procedure gives Egypt a grade of B, the second highest, while many MENA countries score much lower. For example, Iraq and Jordan have grade D and Tunisia grade C.

The correlation between these grades and the statistical capacity indicator is positive, but at a value of 0.36 not very high. This could mean that there are various dimensions to data quality or that these indicators do not capture data quality very well. But as Egypt scores high on both indicators, there is no (ex-ante) reason to doubt Egypt's statistics more than those of many other countries around the world.

Counterfactual TFP levels for Egypt

From Table 1–3 we have learned that the main reasons for the high Egyptian TFP levels is the low employment rate, the low investment rate and the high price of capital. To see the impact of these factors on TFP levels, we present in Table 4 three counterfactual Egyptian TFP levels for each country group. In each counterfactual, one factor is set equal to the unweighted geometric average of the country group. The first row of Table

20 See, for example, the documentation to PWT 6.1 (https://www.rug.nl/ggdc/docs/appendix$_{p}wt_{6}1.pdf$).

21 See also Inklaar *et al.* (2021) on inconsistency between inflation and PPP changes.

Table 4: Counterfactual TFP levels for Egypt (US=1.00)

	MENA	Similar Income
Baseline	1.23	1.23
N/P	1.06	1.06
K/Y(nominal)	0.99	0.93
Stock PPP	0.91	0.94

Source: Penn World Table 10.0 (Feenstra *et al.* 2015).

Note: The baseline TFP level (US=1.00) is the same as in Table 2. Subsequent rows recompute TFP, setting one of the three factors equal to the (unweighted geometric) average of the MENA or Income reference groups; see Table 2 for the country lists. Row N/P changes Egypt's employment-to-population ratio; K/Y (nominal) changes Egypt's ratio of the current-cost net capital stock in local currency units over GDP; Stock PPP changes the PPP for investment goods, weighted using the share of each asset in the current-cost net capital stock, over XR.

4 shows that if Egypt would have had the same employment-to-population ratio as the average MENA country or the average country at similar income level, Egypt's TFP level would have been only 6 per cent higher than that of the United States rather than 23 per cent higher. TFP would even be 1 to 7 per cent lower if Egypt had the same nominal capital-to-output ratio as the two country groups and 6 to 9 per cent lower if the capital stock price were the same.[22]

These are large adjustments and in most of these counterfactuals, Egypt no longer has a TFP level that is higher than that of the United States. In terms of ranking these factors, the most impactful seems to be the PPP for capital goods, followed by the nominal capital/output ratio and the employment/population ratio. But note that Egypt would still be an outlier if only one of these variables were changed. Recall from Chart 1 that countries with lower income levels tend to have lower TFP levels. From that relationship, the predicted TFP level for Egypt would be only 61 per cent of the US level.

Conclusions

As discussed in the introduction, understanding why some countries have higher TFP levels than others is an important goal of empirical development economics. From the literature we know that choices regarding the conceptual model for development accounting and regarding the measurement of outputs and inputs are important in telling us how large TFP differences across countries are. What we have illustrated in this article is how the data choices and constraints are likewise very important.

We have highlighted one country in one year, Egypt in 2017, with a very high TFP level compared to what we would expect given Egypt's income level and the average cross-country relationship between TFP and income level. That relationship predicts a relative TFP level of 61 per cent of the US level in 2017, while the model, measurement and data of PWT 10.0 show a level of 123 per cent. We have used regional and income-level comparison groups

22 The capital services price of Egypt differs by less from the other countries than the capital stock price, so this adjustment overstates the impact changing capital prices.

of countries to illustrate that this exceptional TFP level is due to a low employment rate, a low investment rate and a high price of capital. Yet none of these variables is inherently implausible in these reference groups. Estimating counterfactual TFP levels for Egypt using values from either reference group for the employment-population ratio, the nominal capital stock or the price level of capital would lead to lower estimated TFP levels, but none of those individual counterfactual estimates gets close to the 61 per cent that would be the predicted value solely based on Egypt's income level and measured TFP levels for other countries.

One reading of these results follows the argument above: "See, we knew that Egypt could not be more productive and efficient than the United States and the country turns out to have these crazy output and input figures, so best to ignore this result or deem this measurement approach to be invalid."

But with equal justification, this conclusion can be questioned: "So you are saying that Egypt must be employing more people than all sources say? And their investment levels are understated? And the price of capital is also mismeasured? Why would you distrust all these figures?"

Egypt does not have a particularly weak statistical system, at least as judged by the World Bank's Statistical Capacity indicator or from replicating the data quality letter grades that was provided in some earlier versions of PWT. Neither of these data quality indicators speaks directly to the reliability of Egyptian National Accounts and price measurement — and we are not aware of any indicator that does. We do not see an interpretation that fits both the broader cross-country pattern (countries with lower income levels tend to be less productive) and the fact that none of these data points is inherently implausible in comparative perspective. That leaves us as PWT data developers with little choice but to present the numbers as they are and leave it to individual users to decide how to interpret the numbers in a way is suitable for their purpose.

While the title of this article suggested that we render a definite verdict on Egypt's TFP level, the broader goal of this article has been to show how a data user might proceed when faced with some figures in PWT (or other databases) that strike them as implausible. The development accounting framework is a useful guide to distinguishing outliers from regular patterns in the data. And especially if a user is interested in analyzing a particular country, we offered a diagnostic approach that may be useful for a more in-depth analysis. TFP, being a residual, will always be sensitive to measurement problems in output or inputs, so user beware.

References

Barro, R. J. and L. Jong-Wha (2013) "A New Data Set of Educational Attainment in the World, 1950-2010," *Journal of Development Economics,* Vol. 104, pp. 184–98.

Basu, S., L. Pascali, F. Schiantarelli, and L. Serven (2020) "Productivity and the Welfare of Nations," mimeo.

Bick, A., N. Fuchs-Schündeln and D. Lagakos (2018) "How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications," *American Economic Review,* Vol. 108, No. 1, pp. 170-99.

Caselli, F. (2005) “Accounting for Cross-Country Income Differences” in Aghion, P. and S. N. Durlauf (eds.) *Handbook of Economic Growth,* Elsevier, Vol. 1A, pp. 679–741 (also as NBER Working Paper 10828).

Cohen, D. and M. Soto (2007) “Growth and Human Capital: Good Data, Good Results,” *Journal of Economic Growth,* Vol. 12, No. 1, pp. 51–76.

Cohen, D. and L. Leker (2014) “Health and Education: Another Look with the Proper Data," mimeo, Paris School of Economics.

Comin, D. and B. Hobijn (2010) “An Exploration of Technology Diffusion,” *American Economic Review,* Vol. 100, No. 5, pp. 2031-59.

De la Fuente, A. and R. Doménech (2006) “Human Capital in Growth Regression: How Much Difference Does Quality Data Make?” *Journal of the European Economic Association,* Vol. 4, No. 1, pp. 1–36.

Feenstra, R. C., R. Inklaar and M. P. Timmer (2015) “The Next Generation of the Penn World Table,” *American Economic Review,* Vol. 105, No. 10, pp. 3150–382.

Feenstra, R. C. and J. Romalis (2014) “International Prices and Endogenous Quality,” *Quarterly Journal of Economics,* Vol. 129, No. 2, pp. 477–527.

Freeman, D., R. Inklaar and W. E. Diewert (2021) “Natural Resource and Missing Inputs in International Productivity Comparisons,” *Review of Income and Wealth,* Vol. 67, No. 1, pp. 1–17.

Hsieh, Chang-Tai and P. J. Klenow (2010) “Development Accounting,” *American Economic Journal: Macroeconomics,* Vol. 2, No. 1, pp. 207–23.

Hulten, C. R. (2010) “Growth Accounting,” in Bronwyn H. Hall and Nathan Rosenberg (eds.), *Handbook of the Economics of Innovation,* Elsevier, Vol. 2, pp. 987–1031.

Inklaar, R. and W. E. Diewert (2016) “Measuring Industry Productivity and Cross- Country Convergence,” *Journal of Econometrics,* Vol. 191, No. 2, pp. 426–33.

Inklaar, R., P. Woltjer and D. Gallardo-Albarran (2019) “The Composition of Capital and Cross-Country Productivity Comparisons,” *International Productivity Monitor*, No. 36, Spring, pp. 34-52.

Inklaar, R., R. Marapin, P. Woltjer and M. Timmer (2021) “Inconsistencies in Comparing Relative Prices over Time: Patterns and Facts,” *GGDC Research Memorandum,* No. 189.

Jones, C. I. (2011) "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics,* Vol. 3, No. 2, pp. 1-28.

Jorgenson, D. W. and M. Nishimizu (1978) “U.S. and Japanese Economic Growth, 1952-1974: An International Comparison,” *Economic Journal,* Vol. 88, No. 352, pp. 707–726.

Kravis, I. B., A. Heston and R. Summers (1978) “Real GDP Per Capita for More Than One Hundred Countries,” *Economic Journal,* Vol. 88, pp. 215–42.

Lagakos, D., B. Moll, T. Porzio, N. Qian and T. Schoellman (2018) “Life Cycle Wage Growth across Countries,” *Journal of Political Economy,* Vol. 126, No. 2, pp. 797–849.

OECD (2009) "Measuring Capital: OECD Manual 2009," Second Edition.

Psacharopoulos, G. (1994) “Returns to Investment in Education: A Global Update,” *World Development,* Vol. 22, No. 9, pp. 1325–43.

Schoellman, T. (2012) “Education Quality and Development Accounting,” *Review of Economic Studies,* Vol. 79, No. 1, pp. 388–417.

Schreyer, P. (2007) “International Comparisons of Levels of Capital Input and Multi-Factor Productivity,” *German Economic Review,* Vol. 8, No. 2, pp. 237–54.

Solow, R. M. (1957) “Technical Change and the Aggregate Production Function,” *Review of Economics and Statistics,* Vol. 39, No. 3, pp. 312–20.

Summers, R. and A. Heston (1991) “The Penn World Table (mark 5): An Expanded Set of International Comparisons, 1950–1988,” *Quarterly Journal of Economics,* May, Vol. 106, No. 2, pp. 327–68.

The Conference Board (2021) “Total Economy Database,” April.

World Bank (2020) "Purchasing Power Parities and the Size of World Economies: Results from the 2017 International Comparson Program," World Bank, Washington, D.C. https://openknowledge.worldbank.org/handle/10986/33623.

Chaos Before Order: Productivity Patterns in U.S. Manufacturing

Cindy Cunningham, Sabrina Wulff Pabilonia, and Jay Stewart
U.S. Bureau of Labor Statistics
Lucia Foster and Cheryl Grim
Center for Economic Studies, U.S. Census Bureau
John Haltiwanger
University of Maryland
Zoltan Wolf
New Light Technologies[1]

Abstract

Within-industry productivity dispersion is pervasive and exhibits substantial variation across countries, industries, and time. We build on prior research that explores the hypothesis that periods of innovation are initially associated with a surge in business start-ups, followed by increased experimentation that leads to rising dispersion potentially with declining aggregate productivity growth, and then a shakeout process that results in higher productivity growth and declining productivity dispersion. Using novel detailed industry-level data on total factor productivity and labour productivity dispersion from the Dispersion Statistics on Productivity dataset along with novel measures of entry rates from the Business Dynamics Statistics and productivity growth data from the Bureau of Labor Statistics for U.S. manufacturing industries, we find support for this hypothesis, especially for the high-tech industries. An increase in entry rates in a two-year period t is associated with an increase in dispersion and decrease in aggregate productivity growth in two-year

1 Cindy Cunningham and Sabrina Pabilonia are Research Economists at the U.S. Bureau of Labor Statistics, Jay Stewart is a Senior Research Economist at the U.S. Bureau of Labor Statistics, Lucia Foster is the Chief Economist and Chief of the Center for Economic Studies at the U.S. Census Bureau, Cheryl Grim is the Assistant Center Chief for Interdisciplinary Research of the Center for Economic Studies at the U.S. Census Bureau, John Haltiwanger is the Dudley and Louisa Dillard Professor of Economics at the University of Maryland and a Schedule A part-time employee at the U.S. Census Bureau, and Zoltan Wolf is a Senior Data Scientist at New Light Technologies. The authors would like to thank two anonymous referees, Lucy Eldridge, Andrew Sharpe, T. Kirk White, and participants at the Sixth World KLEMS Conference in March 2021 for comments. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau or the U.S. Bureau of Labor Statistics. The Census Bureau has reviewed this data product for unauthorized disclosure of confidential information and has approved the disclosure avoidance practices applied to this release (Approval ID: CBDRB-FY21-261). Corresponding author: Pabilonia.Sabrina@bls.gov.

period $t+1$ and a decrease in dispersion and increase in aggregate productivity growth in two-year period $t+2$.

Within-industry productivity dispersion is large and exhibits substantial variation across countries, industries, and time (Bartelsman and Doms, 2000; Syverson, 2011). Many factors have been shown to be related to this dispersion, including frictions and distortions that vary across these same dimensions (e.g., Decker, Haltiwanger, Jarmin, and Miranda, 2020). These frictions and distortions, such as barriers to entry, costs of adjusting factors of production, establishment-specific markups, and regulations preventing the equalization of marginal products, may inhibit productivity-enhancing reallocation. This would suggest that increasing within-industry dispersion is associated with slower productivity growth.

An alternative hypothesis is that periods of rising within-industry dispersion may reflect innovation and experimentation. This hypothesis is based on seminal research by Gort and Klepper (1982) and Jovanovic (1982). These papers hypothesize that periods of innovation are initially associated with a surge in firm entry, followed by increased experimentation that yields rising dispersion potentially with declining aggregate productivity growth and then a shakeout process, where successful businesses grow and unsuccessful ones exit, which eventually results in higher productivity growth and declining productivity dispersion.

To explore this latter hypothesis, Foster, Grim, Haltiwanger, and Wolf (2021) looked at the dynamic relationship between entry rates (an indirect measure of innovation), within-industry labour productivity (LP) dispersion, and LP growth using firm-level data for the entire U.S. private sector, where LP is defined as output per job. They find that a surge in firm entry in a four-digit NAICS industry during a three-year period is followed by an increase in within-industry dispersion and a temporary slowdown in industry-level LP growth in the next period. In the subsequent period, there is a fall in dispersion and a rise in LP growth. These relationships are stronger in high-tech industries, where the pace of innovation is presumably faster.

In this article, we build on Foster *et al.* (2021) by exploiting novel, detailed industry-level data on within-industry total factor productivity (TFP) and LP dispersion from the Dispersion Statistics on Productivity (DiSP) data, along with new measures of establishment and firm entry rates from the Business Dynamics Statistics (BDS) data for U.S. manufacturing industries. We combine these data with the official U.S. TFP and LP growth measures

2 The DiSP (developed jointly by BLS and the Census Bureau) is public-use data available at https://www.bls.gov/lpc/productivity-dispersion.htm and https://www.census.gov/disp. Restricted-use microdata is available for qualified researchers on approved projects in the Federal Statistical Research Data Centers (FSRDCs) (http://www.census.gov/fsrdc). The BDS is available at https://www.census.gov/programs-surveys/bds.html. Industry productivity growth data are available at https://www.bls.gov/lpc/tables_by_sector_and_industry.htm. The public-use data and STATA code to

from the Bureau of Labor Statistics (BLS) to examine the relationships between entry, dispersion, and productivity.[2]. To abstract from business cycle dynamics and to focus on the hypothesis, we examine low-frequency variation (average annual growth rates over two-year periods) and include industry and period effects. Relative to Foster *et al.* (2021), a primary contribution of this article is the use of dispersion and growth measures of TFP, which are better metrics for examining the innovation hypothesis.

We find support for the hypothesis that innovation is an important driver of within-industry TFP dispersion and aggregate TFP growth, especially for high-tech industries, using entry rates as a proxy for innovation. A surge in entry in a high-tech industry over a two-year period results in an increase in within-industry TFP dispersion in the next two-year period, followed by an increase in TFP growth in the two subsequent two-year periods. We also find evidence that the increase in dispersion in the first two-year period following a surge in entry is accompanied by negative TFP growth. Relatedly, we find evidence of the reverse, declining TFP dispersion and faster TFP growth in the second two-year period. In addition, we find the relationships between entry and TFP dispersion are stronger when we focus on high-tech industries. For non-tech industries, we find a small decrease in TFP growth, but with an additional lag and no subsequent increase in the following period. We find broadly similar results for LP measures of dispersion and growth.

The article proceeds as follows. In the first main section, we describe the data and present descriptive statistics. The main results are in section two. Concluding remarks are in section three.

Data and Descriptive Statistics

This article uses detailed industry-level data on productivity growth, establishment and firm entry rates, and establishment level productivity dispersion from three public-use data sources: BLS Industry Productivity Statistics, Business Dynamics Statistics (BDS), and Dispersion Statistics on Productivity (DiSP). In addition, we construct additional dispersion measures from the restricted-use data underlying DiSP.[3] Throughout the article, we use industry-level measures for all 86 four-digit NAICS industries in the manufacturing sector. To mitigate business cycle influences, we construct our measures for non-overlapping two-year periods to examine the longer-term relationships between entry, productivity dispersion growth, and productivity growth.

BLS produces the official U.S. measures of LP and TFP growth for four-digit NAICS manufacturing industries (Bureau of Labor Statistics, 2020).The industry LP measures are defined as the ratio of the growth in real sectoral output—the total value of goods and services sold outside

replicate the analyses based on the public-use data are available at https://doi.org/10.5281/zenodo.5770628

3 The experimental data product DiSP was first released in September 2019. Industry-level BDS data were first released in September 2020.

the four-digit NAICS industry — to the growth in hours worked by all persons in the industry.[4] For most industries, real output is derived by deflating sales revenue using industry-level BLS implicit price indexes. Output is also adjusted to remove resales and to account for changes in finished goods and work-in-process inventories. Data for industry output measures are primarily from economic censuses and annual surveys of the U.S. Census Bureau. Data on hours worked come from BLS surveys.[5] The industry TFP measures are defined as the ratio of the growth in real sectoral output to the growth in the cost-weighted combined inputs utilized in producing that output. Inputs include capital, labour hours, energy, materials, and purchased business services.

Although the BLS productivity data for detailed industries in the manufacturing sector are available annually beginning in 1987, we restrict our main analyses to growth in productivity and dispersion over the 1997–2017 period, because the DiSP data start in 1997.[6] The BLS productivity growth rates exhibit considerable year-over-year variation for many manufacturing industries (see Online Appendix Table A1 for four-digit NAICS industry productivity means and coefficients of variation).[7]

http://www.csls.ca/ipm/41/IPM_41_Data_Appendix.xlsx

For this reason, we use the BLS industry productivity indexes to construct non-overlapping average annual growth rates for two-year subperiods from 1997 to 2017 (1997–1999, 1999–2001,. . . , 2015–2017).[8]

DiSP is a newly developed public-use dataset from the Bureau of Labor Statistics and the Census Bureau (2020). This dataset, which is constructed primarily from establishment level data, includes several measures of within-industry dispersion in LP and TFP — the interquartile range (IQR), interdecile (90–10) range, and standard deviation for all 86 four-digit NAICS industries in the manufacturing sector from 1997 to 2016. LP is the log of real output per hour, where output is based on the value of shipments adjusted for resales and changes in inventories and the deflator is the BLS implicit price deflator for that industry.[9] TFP is the log of real output per unit of all factor input costs, where the factors are capital, labour hours, energy, and materials. These measures are available with and without activity weighting, where the activity weights for LP are an establishment's hours share (the share of

4 For very detailed industries, sectoral output is very close to gross output. For more-aggregated industries, sectoral output is closer to value added. For more details on the importance of removing intrasectoral transactions for aggregate industry productivity measurement, see Kovarik and Varghese (2019).

5 For more information on the construction of hours measures, see https://www.bls.gov/lpc/iprhours.htm.

6 The dispersion series will be expanded backward to 1976 as well as forward in future releases.

7 The online appendix tables are posted at http://www.csls.ca/ipm/41/IPM_41_Productivity_Dispersion.pdf.

8 We use standard growth rate measures calculating the ratio of indexes in the current (2-year) period to the prior (2-year) period and then annualizing. For example, $LP_{1997-1999} = (\text{index}_{1999}/\text{index}_{1997})^{0.5} - 1) * 100$.

9 To make the dispersion measures comparable across industries and over time, we normalize each establishment's productivity level each year by subtracting the mean productivity of that establishment's four-digit industry.

a plant's hours of the total hours in its industry) and for TFP are an establishment's share of combined inputs.[10] In addition, we use 90–50, 50–10, 75–50, and 50–25 measures of dispersion from the restricted-use data underlying the DiSP product to consider skewness in the within-industry distribution of productivity.

For our main analysis, we calculate average annual growth rates for LP and MFP dispersion in each of the two-year subperiods in our sample using activity-weighted IQR dispersion measures. (In the last period, we use a one-year growth rate, because the series ends in 2016.) The within-industry IQR dispersion measure describes how much more productive an establishment at the 75th percentile of the productivity distribution is than one at the 25th percentile. Activity-weighted measures should more closely correspond to the BLS aggregate productivity measures. BLS published productivity growth rates can be thought of as changes in the first moment of the underlying distribution of productivity among establishments, where the weights are appropriately defined, while changes in dispersion from DiSP measure changes in the second moments of that distribution.[11]

On average, throughout this period and using the unweighted measures, Cunningham *et al.* (2021) find that establishments at the 75th percentile are 2.4 times more productive than establishments at the 25th percentile when looking at LP and 1.7 times as productive when looking at TFP.[12] However, they also find significant variability in the IQR dispersion measure across industries and a slight increase in dispersion over time. We use the IQR measures for our main analyses because they are less sensitive to outliers; however, we also include a robustness check using the interdecile dispersion measure.

Our entry rates come from the BDS, which the Census Bureau (2020) significantly redesigned and expanded with the release of the 2018 data in September 2020. This novel public-use dataset compiled from the Longitudinal Business Database includes the distribution of firms and establishments by age (based on when they first report positive employment) within detailed industries, allowing us to identify the number of establishment births or firm startups.[13] We construct entry rates (both establishment-based and firm-based) for each four-digit NAICS industry as the simple average of annual entry rates for each two-year subperiod, where the entry rate is the number of establishments aged zero (births) divided by the average count of active establishments in year t and year

10 See Cunningham *et al.* (2021) for a detailed description of these new dispersion measures.

11 Recall, activity weights are applied at the establishment level. They give a higher weight to establishments with more activity when calculating productivity dispersion for an industry.

12 As described in Cunningham *et al.* (2021), unweighted measures use inverse propensity score weights at the establishment level to correct for sample selection issues for the Annual Survey of Manufactures. Activity weighting is the product of the inverse propensity weight and an activity weight.

13 In instances where the number of births in an age bin is not disclosed because there were only 1–2 firm births, we set the number of births equal to 1. Results are essentially the same if we were to set births at 2 firms in the undisclosed age bins.

t-1.[14] Our hypothesis is that increases in entry rates lead to growth in dispersion but with a lag. We construct entry rates for three lagged two-year subperiods. For example, the first-period lagged entry rates corresponding to the average annual growth rates for the 1997–1999 subperiod are the average of entry rates in 1996 and 1997. Thus, our entry rate data cover the 1992–2015 period.

Table 1 shows summary statistics for our data. The average value of the two-year average annual BLS industry LP growth rates was 1.6 per cent for the 1997–2017 period. Over the same period, TFP grew on average 0.4 per cent per year. Dispersion growth is the growth rate in the IQR for LP and TFP dispersion. The LP dispersion growth rate was 0.6 per cent on average, while the TFP dispersion growth rate was 1.5 per cent on average; however, there was considerable variation in aggregate productivity and productivity dispersion growth across industries and time (see the minimum and maximum values). Entry rates were 6.1 per cent on average (establishment and firm). The negative means of the changes in entry rates indicate that, on average, entry rates were falling in the manufacturing sector.

In our analysis, we differentiate between high-tech and non-tech industries, because the former have been an engine of productivity growth, especially over the earlier years in our sample period (Brill, Chansky, and Kim, 2018). We classify 16 of the 86 industries in our sample as high-tech based on the share of jobs held by STEM workers (including engineers, IT workers, scientists, and managers of these workers). The industry is considered high-tech if the share of these workers in the industry exceeds 2.5 times the national average, as determined by Wolf and Terrell (2016).[15] For our main regressions, we use establishment entry rates, which are consistent with our establishment-based dispersion measures. However, both establishment and firm entry rates are relevant in this context because the Gort and Klepper (1982) experimentation stage arguably involves both establishment and firm-level entry. Importantly, establishment-entry rates include the contribution of both firm-level entry and new establishments of existing firms.

We begin our analysis by illustrating graphically the relationships between (1) establishment entry rates and TFP dispersion growth and (2) establishment entry rates and TFP growth for the two high-tech industries that were the top contributors to the marked TFP slowdown that occurred around 2005: semiconductor and other electronic component manufacturing and computers and peripheral equipment manufacturing (Brill, Chansky, and Kim,

14 See https://www.census.gov/programs-surveys/bds/documentation/faq.html for more details on the construction of the entry rates.

15 The high-tech industries include: petroleum and coal products; basic chemical; resin, synthetic rubber, and artificial and synthetic fibers and filaments; pharmaceutical and medicine; industrial machinery; commercial and service industry machinery; engine, turbine, and power transmission equipment; other general purpose machinery; computer and peripheral equipment; communications equipment; audio and video equipment; semiconductor and other electronic components; navigational, measuring, electromedical, and control instruments; manufacturing and reproducing magnetic and optical media; electrical equipment manufacturing; aerospace products and parts.

Table 1: Summary Statistics, All Four-digit NAICS Industries in the Manufacturing Sector

Variable	Years	N	Mean	Std. dev.	Min	Max
Productivity growth						
BLS labour productivity (LP)	1997–2017	860	1.6	6.0	-24.4	38.4
BLS total factor productivity (TFP)	1997–2017	860	0.4	4.0	-11.1	28
Dispersion growth						
LP dispersion	1997–2016	860	0.6	8.6	-33.9	79.4
TFP dispersion	1997–2016	860	1.5	13.1	-63.8	118.5
Entry rate						
Establishment entry rate	1992–2015	1,032	6.1	2.5	1.5	21.1
Firm entry rate	1992–2015	1,032	6.1	2.7	1.2	23.2
Entry rate (per cent change)						
Establishment entry rate	1992–2015	946	-0.5	24.7	-63.6	371.7
Firm entry rate	1992–2015	946	-0.6	27.2	-62.4	486.5

Note: Productivity and dispersion growth are calculated as non-overlapping two-year-average annual growth rates, except in the last period dispersion is a one-year growth rate because this series ends in 2016, e.g., $LP_{1997-1999} = (index_{1999}/index_{1997})^{0.5} - 1) * 100$. Entry rates are two-year-average rates, i.e., $entry_{1999-1998} = (entry_{1999} + entry_{1998})/2$. LP (TFP) dispersion is the interquartile range of within-industry log real output per hour (log real output per unit of combined inputs), activity weighted. Min and max statistics are for industry by period (two-year) variation.

Source: Authors' tabulations based on BLS Industry Productivity Statistics, Dispersion Statistics on Productivity, and Business Dynamics Statistics.

Chart 1: Semiconductor and Other Electronic Component Manufacturing, 1991–2017

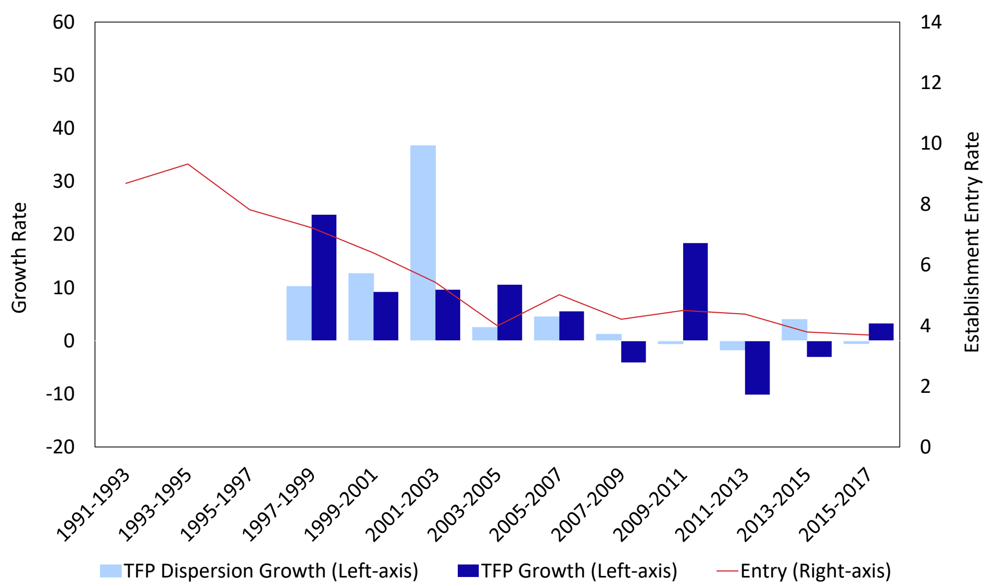

Note: Productivity and dispersion growth are calculated as non-overlapping two-year-average annual growth rates. Entry rates are two-year-average rates. TFP dispersion is the interquartile range of within-industry log real output per unit of combined inputs, activity weighted.

Source: Authors' tabulations based on BLS Industry Productivity Statistics, Dispersion Statistics on Productivity, and Business Dynamics Statistics.

Chart 2: Computer and Peripheral Equipment Manufacturing, 1991–2017

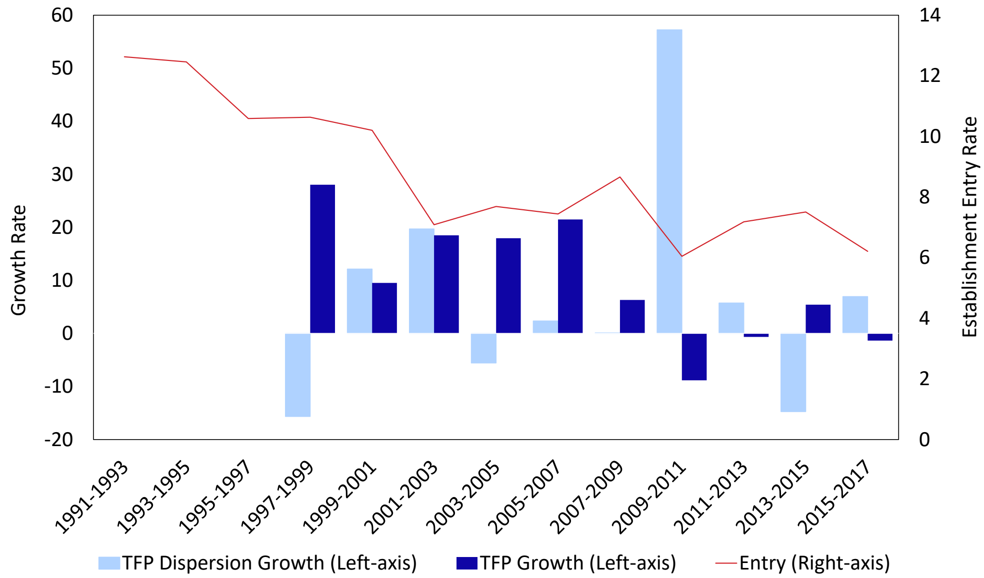

Note: Productivity and dispersion growth are calculated as non-overlapping two-year-average annual growth rates. Entry rates are two-year-average rates. TFP dispersion is the interquartile range of within-industry log real output per unit of combined inputs, activity weighted.

Source: Authors' tabulations based on BLS Industry Productivity Statistics, Dispersion Statistics on Productivity, and Business Dynamics Statistics.

Chart 3: Grain and Oilseed Manufacturing, 1991–2017

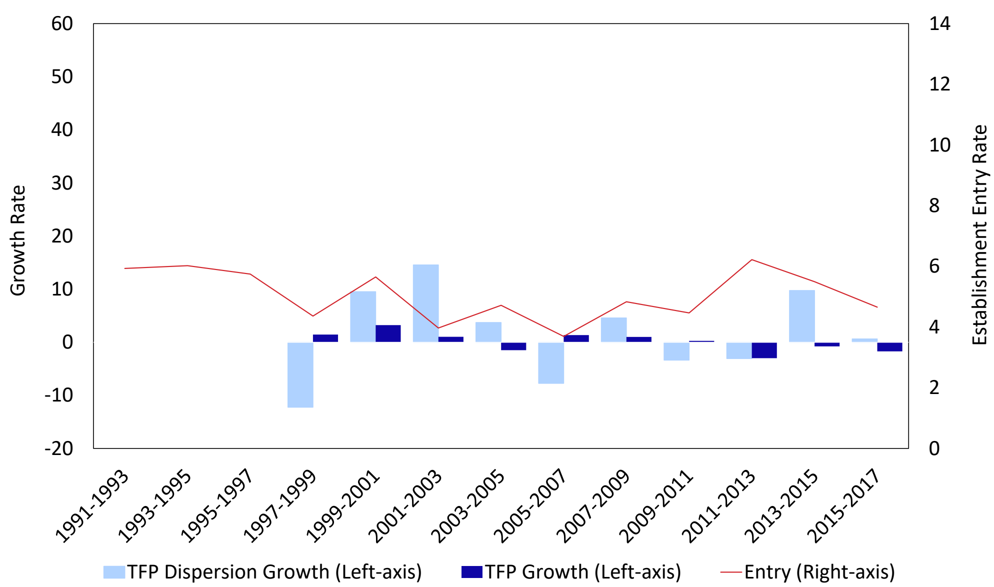

Note: Productivity and dispersion growth are calculated as non-overlapping two-year-average annual growth rates. Entry rates are two-year-average rates. TFP dispersion is the interquartile range of within-industry log real output per unit of combined inputs, activity weighted.

Source: Authors' tabulations based on BLS Industry Productivity Statistics, Dispersion Statistics on Productivity, and Business Dynamics Statistics.

2018). We then consider a non-tech industry, grain and oilseed manufacturing, where we do not necessarily expect to see innovations that lead to entry.

In Chart 1, we see high entry rates in semiconductor and other electronic component manufacturing in the early 1990s followed by high growth in dispersion between 1997 and 2003, especially in 2001–2003, when dispersion grew by 37 per cent. Around 2003, entry rates became relatively stable at around 4 to 5 per cent, with little change in dispersion from one period to the next after that. We see TFP grew from 1997 to 2007 and was especially high in 1997–1999, several periods after a surge in entry. Growth was modest but still positive in 2003–2005 and 2005–2007, following a large spike in dispersion in 2001–2003. In two out of the four periods following the Great Recession, TFP growth was negative.

Chart 2 shows the relationships for computer and peripheral equipment manufacturing. Again, we see that entry rates are initially very high through 2001, exceeding 10 per cent. Thereafter, entry rates are consistently below 8 per cent, except during the Great Recession when the entry rate rose to about 8.7 per cent. Dispersion rises and falls with a large increase during the Great Recession, but there is no obvious pattern that it follows changes in entry; however, TFP growth is very high until the Great Recession, following several periods of relatively high entry rates by a lag.

Chart 3 illustrates the relationships for grain and oilseed manufacturing. Here, we see much lower entry rates that hover between 4 and 6 per cent. Movements in dispersion do not appear to be tied to movements in entry, and there is little growth in productivity.

Empirical Model and Results

We explore the relationships between entry, productivity dispersion, and aggregate productivity growth by estimating panel models of the following form:

$$\begin{aligned} Y_{i,t} = {} & \alpha + \lambda_t + \lambda_i + \sum_{k=1}^{3}[\beta_k Entry_{i,t-k} \\ & + \delta_k Entry_{i,t-k} * Tech_i] + \epsilon_{i,t} \end{aligned} \tag{1}$$

where $Y_{i,t}$ is either average annual within-industry productivity dispersion growth or aggregate industry productivity growth where productivity is measured as LP or TFP. The subscript i denotes the industry, while the subscript t denotes time in two-year subperiods. Entry is either the establishment or firm entry rate, which enters the equation with one-, two- and three-period lags, thus covering a total of six years. Tech is a binary variable equal to one if the industry is high tech and zero otherwise. The parameters of interest, β_k and δ_k, represent the associations between entry and growth, allowing for differences by industry type (high tech or not). The parameter α is a constant term. The model also includes period effects (λ_t) and industry effects (λ_i). The parameter ϵ is a random error term. We estimate the models by ordinary least squares and cluster the standard errors at the industry level.

We estimate both the productivity and dispersion models in growth rate specifications. Differences in levels of productivity are difficult to interpret. For productivity dispersion, levels are more readily interpretable. However, there are industry-

Table 2: Productivity Growth, IQR Dispersion Growth, and Establishment Entry Rates (1997–2017)

	Labour Productivity		Total Factor Productivity	
	Dispersion (1)	Productivity (2)	Dispersion (3)	Productivity (4)
Lag 1 Entry	1.00***	0.45**	0.00	0.01
	(0.33)	(0.21)	(0.44)	(0.07)
Lag 2 Entry	-0.36	-0.20	-0.35	-0.20**
	(0.27)	(0.24)	(0.42)	(0.09)
Lag 3 Entry	-0.31	-0.15	0.33	-0.05
	(0.37)	(0.17)	(0.40)	(0.09)
Lag 1 Entry x Tech	-1.60	-1.59***	2.90*	-0.67**
	(1.24)	(0.51)	(1.49)	(0.30)
Lag 2 Entry x Tech	0.91	1.30*	-4.24**	0.78*
	(1.27)	(0.70)	(1.63)	(0.46)
Lag 3 Entry x Tech	1.33**	1.39**	0.85	0.83
	(0.57)	(0.69)	(1.94)	(0.57)
Joint Hypothesis Tests:				
Lag 1 Entry + Lag 1 Entry x Tech	-0.60	-1.14**	2.91**	-0.66**
	(1.21)	(0.49)	(1.41)	(0.31)
Lag 2 Entry + Lag 2 Entry x Tech	0.55	1.10	-4.59***	0.58
	(1.25)	(0.66)	(1.60)	(0.45)
Lag 3 Entry + Lag 3 Entry x Tech	1.02**	1.24*	1.18	0.79
	(0.51)	(0.69)	(1.90)	(0.58)
Observations	860	860	860	860
R-squared	0.08	0.28	0.09	0.34

Note: Robust standard errors in parentheses are clustered at the industry level. Controls also include a constant, period effects, and industry effects. $*p < 0.10$, $**p < 0.05$, $***\ p < 0.01$.
Source: Authors' tabulations based on BLS Industry Productivity Statistics, Dispersion Statistics on Productivity, and Business Dynamics Statistics.

specific differences in trends in productivity dispersion. The growth rate specifications control for these differences, which are outside the scope of our analysis, in a parsimonious manner.

Our main results are presented in Table 2. The first three rows in Table 2 present the associations for non-tech industries. The second three rows are the differential associations for high-tech industries. The last three rows, which are calculated by summing the associations for non-tech industries and the differential associations for high-tech industries, are the associations for high-tech industries.

We begin with the discussion of the results using TFP dispersion (measured as the IQR) and growth, as these reflect our more important and novel results. These results are in columns 3 and 4 of Table 2. For high-tech industries, a one-percentage-point increase in the establishment entry rate is associated with a 2.9-percentage-point increase in TFP dispersion growth in the next period (column 3). In contrast, a one-percentage-point increase in the establishment entry rate is associated with a 0.7-percentage-point decrease in TFP growth in the next period (column 4).[16] In the second period after entry, dispersion growth falls dramatically (a 4.6-percentage-point decrease) while TFP growth rises (a 0.6-

16 As a robustness check, we also examine the relationship between entry and the 90–10 dispersion statistics. The patterns are similar for TFP, although statistical significance is not as strong (Online Appendix Table A2). We also looked at the relationships using dispersion statistics that were not activity weighted (Online Appendix Table A3). Results are not as strong without activity weighting.

percentage-point increase). The latter estimate is not statistically significant at conventional levels but the difference between high-tech and non-tech industries is about 0.8 of a percentage point and is statistically significant in the second period after entry.

For non-tech industries, we find little relationship between entry, dispersion, and growth (entry is associated with a small drop in TFP growth two periods later, with no subsequent growth). As a sensitivity analysis, we used the longer aggregate productivity series back to 1987, but we still did not find productivity growth for non-tech industries in the third period following an increase in entry (see Online Appendix Table A4).

Turning to LP results, column 1 shows the relationship between LP dispersion and entry, controlling for differences by industry type. For non-tech industries, we find a one-percentage-point increase in the establishment entry rate is associated with a one-percentage-point increase in the growth rate of LP dispersion in the following period. For high-tech industries, we find entry is associated with an increase in dispersion only three periods later. Column 2 shows the relationship between aggregate LP growth and entry. We find that a surge in entry is associated with a small increase in LP growth among non-tech industries in the next period. The results do not show significant changes in LP growth for higher-order lags of entry. However, in high-tech industries, a one-percentage-point increase in entry leads to a 1.1-percentage-point decrease in LP growth one period later and to over 1.2-percentage-points higher LP growth two subsequent periods later. The differences between high-tech and non-tech are large and statistically significant. The results for LP are broadly consistent with those for TFP but less systematic.[17]

Table 3 presents results using firm entry rates instead of establishment rates, which are largely similar to those in Table 2. The coefficient estimates are consistent with the innovation hypothesis, though not always statistically significant at conventional levels. As in Table 2, results in Table 3 are more systematic using TFP dispersion and growth measures for high-tech industries.

Lastly, we consider whether there are stronger relationships between entry and dispersion growth for different parts of the productivity distribution. For example, we may expect to find larger effects of entry among establishments above the median if more productive establishments are able to benefit more from innovations or if innovation induces entry of many establishments with relatively similar productivity levels. In Table 4, we present estimates of the relationship between entry rates and the dispersion growth for the 75–50 and 50–25 ranges of the productivity distribution. We focus on the TFP results for this exercise.[18] For high-tech industries, entry

---

17 The weaker results for LP are not inconsistent with the findings by Foster *et al.* (2021) who focused on LP dispersion, growth, and firm entry. Foster *et al.* (2021) used four-digit NAICS data for the entire private sector, while the current article is restricted to the manufacturing sector. The primary value added of the current paper is the use of TFP dispersion and growth measures at the detailed industry level within manufacturing.

18 Results for LP are presented in Online Appendix Table A5. Results using the 90–50 and 50–10 ranges for both TFP and LP are presented in Online Appendix Table A6.

Table 3: Productivity Growth, IQR Dispersion Growth, and Firm Entry Rates, 1997–2017

	Labour Productivity		Total Factor Productivity	
	Dispersion (1)	Productivity (2)	Dispersion (3)	Productivity (4)
Lag 1 Entry	0.90***	0.31*	-0.36	0.00
	(0.31)	(0.16)	(0.39)	(0.06)
Lag 2 Entry	-0.43*	-0.18	-0.35	-0.11
	(0.25)	(0.20)	(0.41)	(0.08)
Lag 3 Entry	-0.24	-0.09	0.26	-0.02
	(0.30)	(0.16)	(0.31)	(0.07)
Lag 1 Entry x Tech	-1.86	-0.78	1.63	-0.13
	(1.29)	(0.52)	(1.85)	(0.27)
Lag 2 Entry x Tech	1.41	0.82	-4.60**	0.34
	(1.40)	(0.75)	(1.83)	(0.37)
Lag 3 Entry x Tech	1.04	1.32**	2.22	0.92*
	(0.65)	(0.66)	(1.59)	(0.47)
Joint Hypothesis Tests:				
Lag 1 Entry + Lag 1 Entry x Tech	-0.96	-0.47	1.27	-0.13
	(1.27)	(0.50)	(1.80)	(0.27)
Lag 2 Entry + Lag 2 Entry x Tech	0.98	0.65	-4.95***	0.23
	(1.38)	(0.72)	(1.82)	(0.36)
Lag 3 Entry + Lag 3 Entry x Tech	0.80	1.23*	2.48	0.90*
	(0.63)	(0.66)	(1.54)	(0.48)
Observations	860	860	860	860
R-squared	0.08	0.27	0.09	0.34

Note: Robust standard errors in parentheses are clustered at the industry level. Controls also include a constant, period effects, and industry effects. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.
Source: Authors' tabulations based on BLS Industry Productivity Statistics, Dispersion Statistics on Productivity, and Business Dynamics Statistics.

initially leads to an increase in dispersion among both below- and above-median establishments, but the relationship is significant only for the lower part of the IQR (50–25). However, dispersion falls significantly both below and above the median in the second period but more dramatically among more productive establishments. In the third period, dispersion in the upper part of the support increases significantly. For non-tech industries, we find asymmetric effects, with entry leading to lower dispersion in the 75–50 range, but higher dispersion in the 50–25 range three periods later. Again, results are similar when we consider the relationships between firm entry rates and dispersion growth. We interpret these results as providing suggestive evidence that entry yields not only changes in overall dispersion but also changes in the shape of the dispersion.

In closing this section, it is instructive to observe that underlying the dynamic relationships we have uncovered are highly persistent processes. Productivity (LP and TFP), dispersion (LP and TFP), and entry levels all exhibit substantial persistence within industries.[19] Our findings highlight that these persistent processes relate to each other in complex and interesting ways.

19 The average AR1 coefficient for LP (TFP) productivity levels is 0.61 (0.54) for high-tech industries and 0.57 (0.45) for non-tech industries. The average AR1 coefficient for LP (TFP) dispersion levels is 0.42 (0.23) for high-tech industries and 0.30 (0.36) for non-tech industries. The average AR1 coefficient for entry rates for establishments is 0.61 for high-tech industries and 0.56 for non-tech industries. Table A7 in the Online Appendix presents estimates from an AR1 model for establishment entry for each manufacturing industry.

Table 4: 75–50 and 50–25 TFP Dispersion Growth and Entry Rates, 1997–2017

	Establishment Entry		Firm Entry	
	75–50 **(1)**	**50–25** **(2)**	**75–50** **(3)**	**50–25** **(4)**
Lag 1 Entry	-0.39	0.50	-0.48	0.03
	(0.53)	(0.75)	(0.40)	(0.61)
Lag 2 Entry	0.11	-1.27	0.35	-1.34
	(0.56)	(0.79)	(0.42)	(0.91)
Lag 3 Entry	-0.94*	0.97*	-1.13**	1.23***
	(0.49)	(0.56)	(0.48)	(0.44)
Lag 1 Entry + Lag 1 Entry x Tech	3.34	3.04	1.45	1.85
	(2.22)	(1.86)	(1.56)	(2.00)
Lag 2 Entry + Lag 2 Entry x Tech	-7.20***	-1.54	-9.75***	-0.30
	(2.72)	(1.50)	(3.31)	(2.30)
Lag 3 Entry + Lag 3 Entry x Tech	4.51***	-1.02	7.67***	-0.83
	(1.69)	(1.66)	(1.90)	(1.38)
Joint hypothesis tests:				
Lag 1 Entry + Lag 1 Entry x Tech	2.95	3.54**	0.97	1.88
	(2.20)	(1.76)	(1.54)	(1.95)
Lag 2 Entry + Lag 2 Entry x Tech	-7.09***	-2.80**	-9.41***	-1.64
	(2.69)	(1.36)	(3.30)	(2.14)
Lag 3 Entry + Lag 3 Entry x Tech	3.57**	-0.05	6.54***	0.41
	(1.65)	(1.61)	(1.79)	(1.35)
Observations	859	859	859	859
R-squared	0.10	0.08	0.12	0.08

Note: One observation is missing for the TFP regressions because the productivity levels at the different points in the distribution were the same in one period, and thus the percent change was undefined. Robust standard errors in parentheses are clustered at the industry level. Controls include a constant, period effects, and industry effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Source: Authors' tabulations based on BLS Industry Productivity Statistics, Dispersion Statistics on Productivity, and Business Dynamics Statistics.

We regard our findings as suggestive rather than definitive. Gort and Klepper (1982) examine lags over many years using business registry data that tracked entering, exiting, and continuing firms after 46 specific product innovations (e.g. electric shavers or windshield wipers). They found long and varying lags in the responses to innovations, but they did not relate these dynamics to either productivity dispersion or growth, which we explore in this article. We have imposed a relatively simple lag structure to investigate the timing of the relationships between entry, productivity dispersion growth (a proxy for experimentation), and productivity growth. Exploring the long and variable lags from the suggestive evidence in Gort and Klepper (1982) from business registry data will require longer time series and likely a more disaggregated analysis.[20]

Conclusion

This article uses novel detailed industry-level data on TFP and LP dispersion in establishment level productivity levels from the DiSP along with new measures of establishment and firm entry rates from the BDS to examine the relationships between productivity growth, productivity dispersion growth, and entry for U.S. manufacturing industries. We test the hypothesis that pe-

20 Given these issues, generating cumulative effects from Tables 2 and 3 would be incomplete. We also note that because we used standard growth rates, the cumulative effect is not the simple sum of the lagged effects.

riods of innovative activity in an industry are initially associated with a surge in entry of new firms or establishments that is followed by an increase in experimentation that leads to rising within-industry dispersion with potentially declining productivity growth. Under this hypothesis, there is then a shakeout process, where the successful businesses grow and thrive while the unsuccessful ones exit, causing productivity dispersion to decline and productivity growth to rise.

We find the strongest support for this hypothesis using the high-tech industries and measures of TFP dispersion and TFP growth. An increase in entry rates is initially associated with an increase in TFP dispersion and a decline in TFP productivity growth for high-tech industries. This is followed in subsequent periods by a decline in TFP dispersion and an increase in TFP growth for high-tech industries (especially relative to TFP growth for non-tech industries).

Overall, these results lend support to the hypothesis that rising within-industry dispersion at least partly reflects innovation and experimentation. Future work using the restricted-use micro-productivity data could explore the reasons we observe a stronger relationship between entry and productivity dispersion for the upper half of the productivity distribution. Future research using the restricted-use micro-productivity data could also explore whether high entry increases dispersion because the new establishments are more disperse than the existing ones or they change the productivity levels of the incumbent firms. A more disaggregated analysis, such as at the 6-digit NAICS level or detailed product class, would also permit greater flexibility in exploring the variable lags in the entry, experimentation, and productivity growth dynamics suggested by Gort and Klepper (1982). Finally, it would be interesting to explore how measures of innovation such as patenting relate to dispersion and productivity growth.[21]

Given the recent trend of low entry rates prior to the pandemic, we may expect to see slower productivity growth in the years to come. However, the surge in new business applications in the second half of 2020 and the first three quarters of 2021 suggests the possibility of a new round of productivity growth (Dinlersoz, Dunne, Haltiwanger, and Penciakova, 2021; Haltiwanger, 2021).

References

Bartelsman, E.J. and M. Doms (2000) "Understanding Productivity: Lessons from Longitudinal Microdata," *Journal of Economic Literature,* Vol. 38, No. 3, pp. 569–594.

Brill, M., Chansky, B., and J. Kim (2018) "Multifactor Productivity Slowdown in U.S. Manufacturing," *Monthly Labor Review,* July, pp. 1–15.

Bureau of Labor Statistics, U.S. Department of Labor (2020) "Industry Productivity Statistics," https://www.bls.gov/lpc/tables_by_sector_and_industry.htm.

Bureau of Labor Statistics, U.S. Department of Labor and the Census Bureau, U.S. Department of Commerce (2020) "Dispersion Statistics on Productivity," https://www.bls.gov/lpc/productivity-dispersion.htm and https://www.census.gov/disp.

Census Bureau, U.S. Department of Commerce (2020) "Business Dynamics Statistics," https://www.census.gov/data/datasets/time-series/econ/bds/bds-datasets.html.

21 Goldschlag and Perlman (2017) describe new measures of innovative activity planned for the BDS.

Cunningham, C., Foster, L., Grim, C., Haltiwanger, J., Pabilonia, S. W., Stewart, J. and Z. Wolf (2021) "Dispersion in Dispersion: Measuring Establishment Level Differences in Productivity," *IZA Discussion Paper,* No. 14459.

Decker, R., Haltiwanger, J., Jarmin, R. S. and J. Miranda (2020) "Changing Business Dynamism and Productivity: Shocks vs. Responsiveness," *American Economic Review,* Vol. 100, No. 12, pp. 3952–3990.

Dinlersoz, E., Dunne, T., Haltiwanger, J. and V. Penciakova (2021) "Business Formation: A Tale of Two Recessions," *American Economic Review Papers and Proceedings,* Vol. 111, pp. 253–257.

Foster, L., Grim, C., Haltiwanger, J., and Z. Wolf (2021) "Innovation, Productivity Dispersion, and Productivity Growth." in *Measuring and Accounting for Innovation in the 21st Century*, edited by Carol Corrado, Jonathan Haskel, Javier Miranda, and Daniel Sichel, pp. 103–36, Chicago: University of Chicago Press.

Goldschlag, N. and E.R. Perlman (2017) "Business Dynamics Statistics of Innovative Firms," Center for Economic Studies Discussion Paper No. 17-72.

Haltiwanger, J. (2021) "Entrepreneurship during the COVID-19 Pandemic: Evidence from the Business Formation Statistics," NBER Working Paper, No. 28912.

Gort, M. and S. Klepper (1982) "Time Paths in the Diffusion of Product Innovations," *Economic Journal,* Vol. 92, No. 367, pp. 630–653.

Jovanovic, B (1982) "Selection and the Evolution of Industry," *Econometrica,* Vol. 50, May, pp. 649–670.

Kovarik, T. and J. Varghese (2019) "Intrasectoral Transactions: The Most Important Productivity Statistic You've Never Heard of," *Monthly Labor Review,* September.

Syverson, C (2011) "What Determines Productivity?" *Journal of Economic Literature,* Vol. 49, No. 2, pp. 326–365.

Wolf, M. and D. Terrell (2016) "The High-Tech Industry, What Is It and Why It Matters to Our Economic Future," *Beyond the Numbers: Employment and Unemployment,* Vol. 5, No. 8, U.S. Bureau of Labor Statistics, May 2016.

Response to Review Article by Bert Balk on *Measurement of Productivity and Efficiency: Theory and Practice*

Robin C. Sickles
Rice University
Valentin Zelenyuk
University of Queensland[1]

We would like to thank the Editors of *the International Productivity Monitor* (IPM), and Andrew Sharpe in particular, for commissioning a review article of our book, *Measurement of Productivity and Efficiency: Theory and Practice* (Sickles and Zelenyuk, 2019), which appeared in the Spring 2021 issue of the IPM. It is a great honour for us and we are very appreciative for this opportunity to reach out to the wide audience of the IPM. We also would like to wholeheartedly thank Professor Bert Balk for undertaking this challenging task.

The reading of Professor's Balk review article (Balk, 2021) was for us quite reminiscent of the challenges we had while writing the book itself. Indeed, in the process we recognized early on that to write a good chapter for a book takes about half a year or so, and writing two such chapters roughly doubles the time (under constant returns to scale), not to mention the extra time to interconnect them. As combinatorial math tells us, the complexity of the interconnections among the chapters increases dramatically with the number of chapters. And, to be frank, our initial goal of combining the 17 chapters spanning several major fields (with many sub-fields) in the area of productivity and efficiency analysis was not achieved at the level of perfection we had hoped. In such a dynamic field one is always trying to catch up to a fast moving target of advances in theory and statistical methods.

Our focus on completeness and coverage of such a broad topic as productivity and efficiency was the reason the book grew to over 800 pages, which was about double what the publisher agreed to initially. Editorial oversight required us to cut or condense some of the topics and we apologize

1 Robin Sickles is Reginald Henry Hargrove Chair of Economics Emeritus, Department of Economics, Rice University, Valentin Zelenyuk is ARC Future Fellow and Professor, School of Economics and Centre for Efficiency and Productivity Analysis, University of Queensland, Emails: rsickles@rice.edu and v.zelenyuk@uq.edu.au.

to the readers, including Professor Balk, if something they wanted to see is not there.

This also led to a reduction in citations in the author index counts for several leaders in the field of productivity and efficiency, such as Knox Lovell and Shawna Grosskopf, among others. Due to the publisher's policy, for citations where only one author is mentioned and the other authors are hidden under "*et al.*", it is only the first author who is included in the author index for such citations. Alphabetizing co-authored works in economics, the common tradition, does have its costs. In fact, we tried to convince the publisher that it would be fairer to mention all co-authors in the author index but they did not modify their editorial policy.

Professor Balk (2021:139) opines that:

> "For the first part, especially Chapters 1-7, I am not so certain that there has been much progress. A global comparison with Balk (1998) reveals that theoretical differences are almost negligible."

With all respect to the book of Balk (1998), which we find as a very useful source for ourselves, we also find a substantial amount of material we cover in Chapter 1 through 7 to be very different (and the difference is far from negligible) from material in Balk (1998). Because these seven chapters contain a bit more than 40 per cent of our book, we feel obliged to clarify and thank the IPM Editors for giving us the right and the opportunity to respond to Professor Balk's review.

It is important to clarify that the topics covered in Chapters 1-7 (as well as others) are not unique either to our book or to Balk (1998), simply because they are the foundational material that is expected to be covered in any graduate level textbook for this field. Indeed, without such material it would be hard or even infeasible to follow much of the rest of the book. We also note that the similarity of this material in both books is likely due to both in part following the classic book of Färe and Primont (1995), who in turn followed Shephard (1953, 1970), as we all did, acknowledging it at many places. Even given this, *we have a lot of new material in those seven chapters*. Moreover, the classic material is cast in different perspectives and with different emphases than in the other books. To make our explanations brief, we will focus on just a few examples.

In regard to Chapter 1, its distinctive feature is in the much deeper coverage of the axioms of production theory and how these reflect in properties of distance functions, while the distinctive feature of Chapter 2 is to cover duality for production theory in much greater detail than Balk (1998). For Chapter 3, one of the important distinctions from Balk (1998) is that it covers many types of efficiency measures that were not even mentioned in Balk (1998), e.g., hyperbolic - type measures, additive and multiplicative Russell-type measures, slack-based efficiency measures, etc. We also discussed their advantages, caveats and practical issues. Chapter 3 also presents a comprehensive discussion of axioms for the efficiency measures — material based on the stream of literature in *the Journal of Economic Theory*, started by Färe and Lovell (1978) and followed up by Zieschang (1984), Bol (1986), Russell (1990), Dmitruk and Koshevoy (1991), among others, which we think is far from negligible. Another important distinction of Chapter 3 is the coverage of the most general profit efficiency framework developed most recently, by Färe *et al.* (2019), which unifies the previous developments by including other efficiency measures as special cases of this general measurement framework.

Chapter 4 and Chapter 7 — on index numbers are perhaps the closest to material covered in Balk (1998), which is natural because Professor Balk is undoubtedly one of the best experts on those topics in the world. Yet, even here, besides presenting the material in somewhat different perspectives, we have included new and important material that was discovered after 1998. We invite the readers to explore this new material.

The material we covered in Chapter 5 — on Aggregation is completely different from Balk (1998) and, to our knowledge, provides novel and the most comprehensive coverage of that topic. Here it is worth noting that Professor Balk is correct that we have omitted the topic of reallocation of resources in the context of aggregation. We do mention it in passing, in Sections 5.6 and 5.7. Chapter 5 is already quite long and extending it to cover 'reallocation' would require yet another chapter. Such a chapter in fact appeared a few months after our book, as Mayer and Zelenyuk(2019).

Finally, the material covered in Chapter 6 — on Functional Forms—also has a very different level of detail and coverage than what has appeared in earlier books, whether Balk (1998) or Färe and Primont (1995) or any other of which we are aware.

In summary, based on our accounting, on the order of 50 per cent of the material in our Chapters 1-7 is new to Balk (1998). One could, of course, still claim that this new material we synthesize there (which is based on the leading journals that publish in this field, such as *Journal of Economic Theory, Journal of Econometrics, Operations Research*, etc.) is "almost negligible". In our opinion, it is very important, although we leave it to the reader to evaluate how well we have discussed these previous works and their theoretical and empirical contributions.

We would like to use the rest of this reponse to point out some new and interesting publications that came out recently, as well as some that are in the pipeline. Two of these are the recent handbooks edited by Ray, Chambers and Kumbhakar (2020) and ten Raa and Greene (2019). Both have a great collection of many chapters from many giants of the field. To some extent these books overlap in coverage with each other and other works (including our book), which is natural. Yet often they cast those same topics (among others) at different angles or different degree of detail and so, we think, both deserve special attention on their own. We also hope these books will be reviewed in *the International Productivity Monitor* soon.

Two books we are eagerly looking forward to see are by Kumbhakar and Parmeter (under contract with Cambridge University Press) that focuses on Stochastic Frontier Analysis and by Simar and Wilson that focuses on statistical aspects of DEA. We also hope these two books will be reviewed by the IPM, along with other reviews, and this will help the readers navigate in the expanding and enriching literature in the field of productivity and efficiency analysis.

Again, we thank Professor Bert Balk for dedicating his efforts for reviewing our book and to the editors of IPM to publish it and for the opportunity to respond.

Last, yet not least, we believe that for any book, there is still always room for improvements—at least because the knowledge discoveries go on and thus material can become dated. Hence, we enthusiastically encourage productivity scholars to continue developing new and refining old theoretical and empirical approaches to the measurement of productivity and efficiency. Indeed, as the famous wisdom says (usually attributed to business consulting guru Peter Drucker):

> "If you can't measure it, you can't improve it."

Given the challenges facing the world economy, improvements in living conditions in the world by way of shared productivity growth can only be realized when its measurement has been adequately assessed. Ultimately, that is the purpose of our book.

References

Balk, B. M. (1998) *Industrial Price, Quantity, and Productivity Indices: The Micro-Economic Theory and an Application* (Boston, MA: Kluwer Academic Publishers.)

Balk, B. M. (2021) "Productivity Measurement with Data Envelopment Analysis and Stochastic Frontier Analysis: A Review Article on Measurement of Productivity and Efficiency: Theory and Practice," *International Productivity Monitor*, Vol. 40, Spring, pp.134–39.

Bol, G. (1986) "On Technical Efficiency Measures: A Remark," *Journal of Economic Theory*, Vol. 38, No. 2, pp. 380–86.

Dmitruk, A. V. and G. A. Koshevoy (1991) "On the Existence of a Technical Efficiency Criterion," *Journal of Economic Theory,* Vol. 55, No. 1, pp. 121–44.

Färe, R., X. He, S. K. Li and V. Zelenyuk (2019) "A Unifying Framework for Farrell Profit Efficiency Measurement," *Operations Research*, Vol. 67, No. 1, pp. 183-97.

Färe, R. and C. K. Lovell (1978) "Measuring the Technical Efficiency of Production," *Journal of Economic Theory*, Vol. 67, No. 19, pp. 150–62.

Färe, R. and D. Primont (1995) *Multi-Output Production and Duality: Theory and Applications"* (New York, NY: Kluwer Academic Publishers.)

Mayer, A. and V. Zelenyuk (2019) "Aggregation of Individual Efficiency Measures and Productivity Indices," in ten Raa, T. and W. H Greene, editors, *The Palgrave Handbook of Economic Performance Analysis*, Springer International Publishing, pp. 527–57.

Ray, S. C., R. Chambers and S. Kumbhakar (2020) *Handbook of Production Economics* (Singapore: Springer Singapore.)

Russell, R. R. (1990) "Continuity of Measures of Technical Efficiency," *Journal of Economic Theory*, Vol. 51, No. 2, pp. 255–67.

Sickles, R. C. and V. Zelenyuk (2019) *Measurement of Productivity and Efficiency: Theory and Practice* (Cambridge: Cambridge University Press.)

Shephard, R. W. (1953) *Cost and Production Functions* (Princeton, NJ: Princeton University Press.)

Shephard, R. W. (1970) *Theory of Cost and Production Functions, Princeton Studies in Mathematical Economics* (Princeton, NJ: Princeton University Press.)

Ten Raa, T. and W. H. Greene, (2019) *The Palgrave Handbook of Economic Performance Analysis,* Springer International Publishing.

Zieschang, K. D. (1984) "An Extended Farrell Technical Efficiency Measure," *Journal of Economic Theory*, Vol. 33, No. 2, pp. 387–96.