NBER WORKING PAPER SERIES

RISING LABOR PRODUCTIVITY DURING THE 2008-9 RECESSION

Casey Mulligan

Working Paper 17584
http://www.nber.org/papers/w17584

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2011

I appreciate several conversations with Lee Ohanian and Rob Shimer, and the comments of Gary Becker, Emir Kamenica, Kevin Murphy, Marcus Nunes, and Jesse Shapiro and participants at the Money and Banking Workshop at the University of Chicago. I also appreciate the financial support of the George J. Stigler Center for the Study of the Economy and the State. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Casey Mulligan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Rising Labor Productivity during the 2008-9 Recession
Casey Mulligan
NBER Working Paper No. 17584
November 2011
JEL No. E24,E32,J22

ABSTRACT

During the recession of 2008-9, labor hours fell sharply, while wages and output per hour rose. Some, but not all, of the productivity and wage increase can be attributed to changing quality of the workforce. The rest of the increase appears to be due to increases in production inputs other than labor hours. All of these findings, plus the drop in consumer expenditure, are consistent with the hypothesis that labor market "distortions" were increasing during the recession and have remained in place during the slow "recovery." Producers appear to be trying to continue production with less labor, rather than cutting labor hours as a means of cutting output.

Casey Mulligan
University of Chicago
Department of Economics
1126 East 59th Street
Chicago, IL 60637
and NBER
c-mulligan@uchicago.edu

A paycheck is an important reason why people devote time to, and exert effort for, working. Businesses willingly pay their employees because the workers help produce goods and services that can be sold to customers. Thus, three important indicators of the state of the labor market are the amount of time worked, the amount produced, and the ratio of production to time worked, known as labor productivity. A few indicators like these go a long way toward revealing causes of the 2008-9 recession, with some surprising conclusions.

Figure 1 displays monthly seasonally adjusted measures of time worked since January 2007.[1] The red and green series are civilian and nonfarm payroll employment, respectively, measured on the left axis in employees per thousand persons (civilian employment is shifted by 30 persons per thousand in order to be displayed on the same axis with nonfarm payroll employees). The blue series is an index of hours worked per person – a product of employees per person and weekly hours worked per employee – measured as the sum of private work hours per person (measured as the all-employees aggregate weekly hours index for all private industries, divided by population) and aggregate public work hours (estimated as public sector employment per person times private work hours per private sector employee).[2]

Although the two employment series are measured in different ways, their dynamics during the recession are pretty similar, as evidenced in the figure by the close agreement of the red and green series. To the extent that employment and hours per employee move together, the blue hours series has larger percentage changes than the other two series because it is the product of employment and hours per employee. All three series had their steepest declines in late 2008 and early 2009. Neither employment

[1] All of the source data used in this paper are seasonally adjusted, unless already measured annually.
[2] An Appendix explains how various sources agree on the amount of the aggregate hours drop.

series has had a prolonged increase between 2007 and the time I am writing, and the slight increase of the hours series since the end of 2009 still leaves hours per capita 9 percent below where it was when the recession officially began in December 2007. The purpose of my work is to explain, in the context of the wider economy, why time worked declined so much, why time worked has been so slow to recover to pre-recession levels, and to offer estimates of how the labor market might have been different with alternative public policies.[3]

A recession is, by definition, a period when time worked declines. A variety of explanations have been offered for previous recessions: adverse productivity shocks (Kydland and Prescott 1982), a surge in the demand for "liquidity" (Friedman and Schwartz (1963); Lucas, (2008)), a collapse in international trade (Crucini and Kahn 1996), and a stock market crash are among them. Do any of these explain the labor decline since 2007? This paper begins an answer to the question by examining the time series for labor productivity and real hourly wages, and finding that both series are higher now than they were in 2007.

In their studies of prior recessions, economists have closely examined the cyclicality of real wages – that is, whether real wages tend to fall during a recession and rise during an expansion (a pro-cyclical real wage), or instead rise during a recession and fall during expansions (a counter-cyclical real wage). John Maynard Keynes claimed in his *General Theory* that real wages were counter-cyclical: "… an increase in employment can only occur to the accompaniment of a decline in the rate of real wages." (Keynes 2008/1936, 15), although a counter-cyclical pattern is not readily observed in a number of recessions.[4] When it comes to the 2008-2009 recession, at least, I conclude that real wages and productivity are counter-cyclical.

---

[3] I focus on labor usage – the total hours in the nation that people are working – as opposed to the "labor force" that includes persons who are not working, but are looking for work. In most cases, as in Figure 1, my hours measures derive from hours tabulations by the Census Bureau or Bureau of Labor Statistics, which are tabulations of hours that employed persons are paid (and thereby include hours of paid vacation, sick days, etc.). In a few cases, noted below, it is important to distinguish paid work hours from hours actually worked, in which case I create my own hours-actually-worked series from the micro data from the Census Bureau's household survey. In a few cases, also noted below, employment is the only available labor usage measure.

[4] See (Abraham and Haltiwanger 1995) for a review of that literature.

High wages and low employment make some sense from an employer's perspective because wages are the main cost of employing a worker. But that raises the question of whether potential workers are all that eager to take the jobs that are available, or instead whether something prevents millions of people eager for work from bidding down wages. My next step adapts some of the labor market ingredients that are common to both "New Keynesian" and "real business cycle" models in order to decompose the labor reductions since 2007 into three types of potential "causes": labor distortions and labor preferences that raise productivity and reduce labor, productivity shocks that reduce labor and productivity, and wealth and intertemporal substitution effects that reduce labor and raise consumption. I conclude that labor fell more than output during the 2008-2009 recession while labor "supply" (defined more rigorously below) shifted to the left.

Analytically, this paper's decomposition is most like that of Katz and Murphy (1992), who look at changes over time in the relative amounts and productivity of skilled and unskilled labor in order to determine the relative importance of supply and demand shocks. In terms of substance, this paper is about the changes over time in the overall *levels* of labor and labor productivity, which raises the possibilities of tax distortions, wealth effects, and intertemporal substitution effects that would be less important for understanding one skill group's changes relative to another. In this regard, my analysis is more like that of Chari et al (2007), who also consider capital market fluctuations and total factor productivity. Gali et al (2007), Mulligan (2002) and Mulligan (2005) are three other papers using the supply-demand decomposition to quantify labor market distortions over time; Hall (1997) uses it to quantify labor preference shifts.

Section I displays the basic time series used to make the decomposition: aggregate labor, consumption, and productivity per hour. Four aggregate changes since 2007 help gauge the relative importance of various explanations for the recession and slow recovery. First, in contrast to the patterns of some previous recessions/depressions, output fell significantly less than did labor hours. That is, labor productivity (output per hour) increased. Section II shows that labor productivity increased by an amount that suggests that total factor productivity continued to follow its previous upward trend. This finding suggests that the recession and slow recovery cannot primarily be explained by

adverse productivity shocks such as interruptions to the production process created by unusually bad weather, strikes, labor immobility, or business sector attempts to economize on raw materials.

Second, as shown in Section III, aggregate labor compensation per hour is also significantly greater now than it was before the recession began. This finding suggests that wage-depressing impulses such as sectoral shifts in the direction of capital-intensive sectors, or perceived increases in employment costs (aside from wages themselves), cannot explain much of the recession. However, departures between the wage and productivity series since the middle of 2009 are consistent with some role for wage-depressing factors.

Third, consumption dropped significantly during the recession. As shown in more detail in Mulligan (2011), this drop contradicts some of the investment collapse models of the recession because at least some of the resources freed up from foregone investment opportunities would have been used for additional consumption.

Fourth, as shown in Section IV, growth accounting suggests that, on average, the use of production inputs other than labor hours actually increased during the recession, when labor hours fell. This factor substitution finding appears to contradict claims that the 2008-9 recession began because people were spending less, and that their low spending forced the businesses serving them to cut output. The spending-impulse theories may be a good description of the declines in manufacturing, residential construction, and perhaps a couple of other industries where all factors of production were used less, but are difficult to reconcile with the increased use of non-labor production factors that happened in the rest of the economy.

Sections V and VI explain how all of these findings are consistent with increases in labor market distortions, akin to marginal tax rate hikes, that occurred during the recession. The amount of the consumption drop helps quantify the amount of the labor market distortion that would be needed to explain what happened in the labor market since 2007. An Appendix compares these results to analogous calculations for previous recessions.

I. Quarterly Indicators of Aggregate Economic Quantities

Figure 2 displays four quarterly seasonally adjusted indices of real per capita consumption since the beginning of 2007.[5] Two of them are exclusive to the private sector: private nondurable consumption goods and private consumption services. The third (blue) series is public non-defense consumption. The fourth (black) series aggregates the three, chain-weighting by their contributions to total expenditure. My purpose here is to measure current living standards, so purchases of consumer durables are excluded, and non-defense public consumption is combined with the private series (because much of public non-defense consumption is publically provided health care, schooling, and housing similar to what is provided by the private sector).

The figure also displays a dashed black line to represent a sustainable trend: an estimate of how much total consumption could have increased if the economy would have continued to produce and grow as it had prior to the recession.[6] All of the series have declined four or five percent below the trend, and are lower in absolute terms than they were before the recession began. The most rapid deviations from trend occurred in 2008 for the private series, and after early 2009 for the public sector.

Overall, it is clear that per capita consumption dropped, and dropped much less in percentage terms than labor ultimately did (labor per person fell about ten percent). It appears that most of the 4-5 percent consumption drop below trend occurred in 2008, and since early 2009 consumption has resumed somewhat of an upward trend.[7] As shown below, the direction and amounts of the consumption changes are informative about the recession's causes.

---

[5] Real consumption is spending on consumption items and services, adjusted for inflation (that is, price changes) for those items and services, and therefore changes only if people change the number or quality of the items and services they purchase.

[6] Specifically, the "sustainable trend" increases at 0.6 percent per year, my estimate (discussed in Mulligan (2011)) of the rate of per capita total factor productivity growth. With some exceptions noted below, employee compensation per hour is thought to indicate marginal labor productivity to the extent that workers are paid their marginal product.

[7] The combined series closely follows private sector services because (a) private nondurables and services follow a similar pattern and (b) public non-defense consumption expenditure is small compared to private expenditures on non-durables and services (the latter is the single largest expenditure category).

Figure 3 displays three quarterly indexes of labor productivity: real GDP per hour, business sector output per hour, and real employee compensation per hour. The black dashed line is the prior trend for real GDP per hour of 0.6 percent per year. By the middle of 2009, all three measures were about 3 percent above what they were when the recession began, and therefore well above the prior trend. Almost four years have passed since the recession began, and labor productivity remains above trend by all three measures.[8]

By definition labor productivity is the ratio of output to hours worked, so it might seems almost automatic that the ratio would rise when its denominator falls, as it did in 2008 and 2009. However, as we see below, some of the severe recessions or depressions in the past have had output fall more than hours worked, and therefore labor productivity fell (see also Ohanian (2010)). Even in this recession, sectors most obviously depressed by lack of demand had their hours and output fall in roughly equal proportions. The time pattern of labor productivity indicates something about the causes of the recession, as we examine in more detail below.

II. Movements Along an Aggregate Marginal Productivity Schedule

Goods and services are produced with labor and other factors of production, so it is no surprise that output fell at about the time that labor did. But that still leaves open the question of whether the usage of other factors of production change in the same amount, or even in the same direction, as labor did. The aggregate Cobb-Douglas production function helps arrive at an answer. According to that function, and based on the observation that about 70 percent of national income accrues to labor, quarter *t*'s

[8] Real GDP per hour is a more comprehensive measure of productivity than business sector productivity because the former includes all sectors in its numerator and all workers in its denominator. As a comprehensive measure, GDP per hour is more readily integrated with models of aggregate labor supply in which total hours worked is a key variable. However, productivity by itself is often measured for the narrower business sector because output in the remaining sectors (government, non-profit institutions, and households) may be measured less reliably. Because the two measures experienced many of the same changes since 2007, henceforth I use the more comprehensive GDP per hour and note a few places where results would be slightly different if based on business sector output per hour.

aggregate output per hour y_t depends on the ratio of other inputs A_t to labor n_t, raised to the 0.3 power:[9]

$$y_t \equiv \left(\frac{A_t}{n_t} \right)^{0.3} \tag{0.1}$$

Holding constant the usage of other factors A_t, each unit reduction in log labor increases log labor productivity by 0.3. This downward-sloping relationship between labor and labor productivity is often called the "aggregate labor productivity schedule."

Changes in the usage or efficiency of other production factors, such as capital accumulation, technical change, or changes in capital utilization, would shift of the aggregate marginal productivity schedule. Inverting equation (0.1), and using data on output and labor hours, we can calculate the amount of the shift of that schedule, measured in the quantity dimension:

$$\Delta \ln A_t \equiv \Delta \ln n_t + \frac{\Delta \ln y_t}{0.3} \tag{0.2}$$

where ln denotes natural log and the difference operator Δ denotes changes from a benchmark quarter to quarter t. In other words, the "input residual" A_t calculated from equation (0.2) is the change in the usage of other factors that must have occurred in order for output to change as much as measured. [10]

Figure 4 illustrates how the actual changes in labor productivity from 2007 Q4 through 2011 Q2 can be decomposed using the marginal productivity schedule (0.1).

[9] More generally, when factor market are competitive, the magnitude of the elasticity of productivity with respect to labor input is equal to the share of output paid to non-labor production factors. I assume that share to be 0.3, and assume that the share is independent of the amount of labor and the amount of other inputs. As noted below, this paper's qualitative results are not sensitive to either of these assumptions.

[10] The concepts of labor productivity and the "input residual" are different from "total factor" or "multi-factor" productivity, which attempt to measure the change in output that cannot be explained by changes in labor and capital inputs. In practice, the multifactor productivity measures prepared by the Bureau of Labor Statistics do not account for changes in capacity utilization, which is sometimes the primary reason why the input residual changes over time.

Each data point in the Figure graphs the actual values of real GDP per hour and aggregate work hours (the same series as shown in Figure 1), measured on a logarithmic scale with the origin normalized to be 2007 Q4. The points are connected in chronological order.

Two of the points each have a straight line (with slope -0.3) drawn through them representing the marginal productivity schedule (0.1) applicable at those two dates. If (hypothetically), a single marginal productivity schedule applied at each date, then all of the data points would be on the same straight line with slope -0.3. In fact, each date is a different distance from any particular schedule, so the log input residual measures the horizontal distance from a schedule with slope -0.3 passing through the origin and the actual data.

By the end of 2009, labor quantity had declined eight quarters in a row, for total log labor change of -0.084. Productivity was essentially constant during 2008,[11] but eight quarters after the recession started, log labor productivity had risen 0.045. The marginal productivity schedule (0.1) attributes most, but less than all, of the 0.045 increase, namely 0.025, to the reduction in labor that occurred over that time. The other 0.020 of productivity change is attributed to an *increase* in the usage or efficiency of other factors. This finding is an important reason to doubt that a *reduction* in the usage or efficiency of other factors was a significant contributor to the recession.

Normally, the marginal productivity schedule tends to shift up over time; in the eight quarters prior to the recession it had shifted up 0.012 (0.0015 per quarter). In this regard, it is perhaps no surprise the marginal productivity schedule had also shifted up 0.020 between the end of 2007 and the end of 2009. For the same reason, it is perhaps unsurprising that the marginal productivity schedule continued to shift up after 2009.

Figure 5 displays the quarterly measures of log labor input n_t and log input residual A_t, relative to their values for 2007 Q4. For each quarter, the log input residual is measured as the horizontal distance between the quarter's data point in Figure 4 and the marginal productivity schedule passing through the data point for the base quarter (recall equation (0.2)). The input residual can be interpreted as the change in non-labor inputs,

[11] As shown in Figure 3, business sector output per hour fell slightly during 2008.

or a change in input efficiency, needed to explain the change in output. Given that labor and the residual followed similar patterns during 2008, it seems that labor and non-labor inputs were falling in about the same proportions during that year. But in 2009 labor continued to fall while the residual quickly surged beyond its pre-recession values. The input residual continued to increase during 2010.

Thus, while labor's path during the beginning of the recession might be explained by a reduction in other inputs, the overall pattern since 2007 has been a significant *increase* in other inputs and a large reduction in labor. A key question is why production shifted so suddenly and so dramatically away from labor hours and toward the usage of other inputs.

III. On Average, Real Wages did not Fall

It is possible that aggregate hourly compensation changes do not accurately measure changes in the average person's reward to working, perhaps because of a change in the composition of the workforce that effectively gives more weight in the aggregate to higher paid workers during recessions than during expansions. For example, employment could drop the most for less skilled workers (as it has in past recessions; Bils (1985) and Solon, Barsky, and Parker (1994)), so that the average compensation (and output per hour) of those who remain employed rises even while none of those employees is compensated more than he was before the recession. Another possibility is that compensation is fixed for existing workers, but dropped during the recession for new hires, and it's the compensation of new hires that allocates labor in the marketplace.

Composition Bias is Relatively Small

The effect of changes in the composition of the workforce on aggregate hourly earnings and productivity is known as "composition bias," and the size of the bias depends on: (a) the percentage change in the size of the workforce, and (b) the degree to which persons leaving the workforce or joining the workforce have different hourly

wages than the rest of the employees. On the second point, job losses during the recession were certainly not random. Table 1 partitions the population by gender, schooling, age, and industry and displays percentage employment rate changes from 2007 to 2009, calculated from the CPS Merged Outgoing Rotation Groups public use microdata files. In order to help assess the size of the composition bias, the Table's right-hand column reports each groups average hourly earnings.[12] Some of the composition changes did serve to increase aggregate wage measures: young persons and persons with less schooling, who also tend to earn more per hour, had larger employment rate declines. On the other hand, the composition changes by gender tended to reduce aggregate wage measures because the gender with the higher average wage (male) is the one that experienced the larger employment rate decline.

The bottom of the table shows that the construction and manufacturing industries sharply reduced employment. Manufacturing is the larger of the two, and its average hourly wage was somewhat greater than average. Healthcare is an example of an industry that increased its employment since 2007, and its average hourly earnings were close to the national average. Overall, it appears that changing industry composition did little, if anything, to increase aggregate hourly earnings (see also (Bils 1985, 667)).

Based on the age, race, and schooling patterns shown in Table 1, and based on the likely effects of some of the public policy impulses discussed Mulligan (2011), I suspect that low wage persons in 2007 were disproportionately represented among those not employed in 2009, so that aggregate wage measures somewhat exaggerate the rate of wage increase and aggregate productivity measures somewhat exaggerate the rate of labor-quality-adjusted productivity growth. Nevertheless, workforce composition did not change enough to be the primary reason why aggregate real wage measures increased 3 percent from 2007 to 2009, rather than declining. In order for the composition bias to be as much as three percent, the roughly 6 million not working in 2009 were working in

[12] My statistics are weighted by the Current Population Survey "Final Weight," adjusted so that each month's survey is given equal weight in the annual averages. Unlike the labor compensation used for aggregate analysis, earnings measured from the Current Population Survey exclude the value of fringe benefits.

2007 would have had to have earned an average of about $5.50.[13] We know instead that the job losses included, for example, educated people and people over age 30, which are groups of people for whom wages almost always exceed the federal minimum. In fact, a wage of $5.50 would have violated the federal minimum wage for half of 2007, and have violated a number of state minimums for the entire year.[14] Thus, while changing composition alone likely increased aggregate wage and productivity measures between 2007 and 2009, the composition bias on the measured wage change is less than 3 percentage points, and likely much less.[15]

Wage Rates for Marginal Workers

For the purposes of understanding employer incentives to hire and potential employees' incentives to accept work, the wage rate for new hires and other workers with weak attachments to an employer may be more important than the average wage of all employees, and the latter may not evolve in the same way as the former. Nevertheless, my finding that, even after adjustments for composition, the average wage of all employees rose during the recession suggests that the average wage of marginal employees also rose (although perhaps in a different amount) because marginal workers are included in the overall average.

[13] 140 million workers earning an average of $20.20 per hour in 2007 plus 6 million workers earning an average of $5.50 per hour yields an overall average hourly earnings of $19.60 for the 146 million. In other words, the overall average understates the average for the 140 million by 3 percent.

[14] The United States Department of Labor (2008), (2007) calculations for 2006 and 2007 suggest that less than two million people earned at or below the federal minimum wage.

[15] As part of its work on multifactor productivity, the Bureau of Labor Statistics (2011) considers changes in age, education, and gender, and finds that those sources of composition bias changed 1.4 percent from 2007 to 2009. In their study of composition bias for business cycles between 1967 and 1987, Solon, Barsky, and Parker (1994) found that aggregate wage measures fell 0.6 percent for each percentage point increase in the national unemployment rate. The direction of this result by itself shows that something different happened during the 2008-9 recession, but if we take their composition-bias-corrected estimate of 1.0 percent (this is a 2007 weighted average of their separate male and female estimates) and thereby infer composition bias of 0.4 percent for each national unemployment rate percentage point, or a total 1.8 percent for the 4.7 point national unemployment rate change that occurred between 2007 and 2009. For the reasons I cited in the text (see also (Bils 1985, 684); Keane, Moffitt, and Runkle (1988)), I believe the Solon et al (1994) composition bias estimates may be exaggerated for my purposes, but still they imply that selection-bias-corrected real wages increased from 2007 to 2009.

To see this, suppose that a majority of the workforce, who number L, is paid pursuant to long term contract paying X, and will be employed under any contingency. The remainder M of the workforce is hired in a spot market at rate w. Average hourly earnings are:

$$\bar{w} = \frac{M}{L+M} w + \frac{L}{L+M} \frac{X}{L} \tag{0.3}$$

For small changes in w and M, the change in average hourly earnings is:

$$d\bar{w} = \frac{M}{L+M} dw - \left(\frac{X}{L} - w \right) d\left(\frac{M}{L+M} \right) \tag{0.4}$$

The first term on the right is the effect of the marginal workers' average wage change dw on the average change, holding the workforce composition constant. The second term is the composition bias. The change in the composition bias adjusted average wage therefore has the same sign as dw, and the amount of the change is proportional to the marginal workers' share of the workforce.[16]

IV. Was It Customer Demand? Factor Reduction and Factor Substitution by Industry

A reduction of labor usage during a recession with rising labor productivity is sometimes attributed to employer desires to reduce their production, which in turn derives from a lack of customer demand. If we assume that, as in a number of New Keynesian models (e.g., Woodford (2003)), non-labor inputs are fixed in the short run, then it follows from the diminishing returns to labor that reducing output involves reducing labor by an even greater proportion. The other input residual is constant by assumption.

[16] Young workers and recent graduates are examples of marginal workers, although some of them may also have experienced a relative reduction in the demand for their labor (e.g., the food service industry declined disportionately during the recession). Nevertheless, a number of measures of their real wage change is positive, rather than negative.

However, if at least some of the other inputs are adjustable in the short run, then reducing output would involve reducing those inputs too. In this case, the other input residual would decline, and perhaps enough that labor productivity is constant.

The assumption that inputs other than work hours are constant is convenient for many purposes, but a complete analysis of productivity statistics needs it relaxed because other inputs are subject to some of the same economic forces as work hours (Jorgenson 2009). In particular, an increase in other inputs together with a reduction in labor hours would suggest that employers may have perceived labor to be more expensive, and would be difficult to reconcile with the view that labor was cut merely as a means of reducing output.

Particular industries undoubtedly experienced a reduction in demand since 2007. We know, for example, that consumers cut their spending significantly, and one of the best ways to cut spending with minimum short-run impact on living standards is to reduce purchases of durable goods.[17] From the point of view of the manufacturers of those goods, customer demand fell. Residential construction is another industry that experienced reduced demand, as evidenced by the sharp decline in housing prices and increase in vacancy rates. To the degree that production inputs are adjustable, we expect manufacturing and residential construction to reduce their usage of all inputs, and not just work hours.

Table 2 displays measures of input and value-added (that is, output minus purchases of materials and services from outside the industry) changes from 2007-Q4 to 2009-Q4 for selected industries. Real value-added declines 0.17 log points in manufacturing, where labor hours fell 0.19 log points. Using the formula (0.2), I calculate the other input residual change to be -0.15 log points.[18] In other words, the value-added and labor hours data suggest that the manufacturing industry sharply reduced its usage of other inputs. The last row of the Table confirms this prediction with data

[17] As I explain in Mulligan (2011), even a recession that was the result of reduced supply would create an immediate reduction in the demand for durable goods.

[18] For the purposes of applying equation (0.2) to the manufacturing industries (as opposed to the entire economy), I used the labor share of 55% rather than 70% because manufacturing employee compensation was 55% of value-added in 2007 (United States Bureau of Economic Analysis 2011). Using a labor share of 70% would result in a -0.13 reduction in the other inputs residual for manufacturing.

from the Federal Reserve Board of Governors on capacity utilization in the manufacturing industry,[19] which changed -0.16 log points from 2007 Q4 to 2009 Q4.

The second column of Table 2 tells a similar story for the residential construction industry. Real value-added fell 0.42 log points while work hours and other inputs were reduced 0.44 and 0.36 log points, respectively. Thus, the manufacturing and residential construction industries reduced their usage of both labor hours and other inputs, and did so in about the same amount that they reduced value-added: just as one would expect if those industries had experienced a reduction in demand and were able to adjust the other inputs.

The third column of the table displays results for a familiar industry that experienced growing output and revenues throughout the recession: the mobile telecommunications industry (i.e., the sellers of wireless phone and wireless data services). On an annual basis, the industry's revenues and real value-added have been increasing every year for many years (United States Census Bureau 2011). According to the Federal Communications Commission (2011) and the industry association CTIA (2011), the industry's customers' mobile connections increased 12 percent from the end of 2007 to the end of 2009. Nevertheless, the industry cut its labor hours by 0.20 log points after increasing labor hours in 18 of the prior 20 years (United States Bureau of Labor Statistics 2010). As expected for an industry that cuts its labor hours without cutting output, the last row of the table suggests an increase in other inputs. The wireless telecommunications industry appears to be engaged in factor substitution: substituting other inputs for labor hours, rather than cutting all inputs in order to reduce output.

Admittedly, the wireless telecommunications industry is just a fraction of the overall economy, but so are the manufacturing and residential construction industries: the latter two industries combined for only 17 percent of the economy in 2007, and less now. The final column of the table therefore examines the entire business sector, apart from

[19] Capacity utilization is an example of a factor of production, other than labor hours, than can be adjusted in the short run.

manufacturing and residential construction.[20] Their real value-added fell 0.03 log points, while labor hours fell 0.08. The other input residual *increased* 0.09 log points. In this regard, most of the economy appears to be substituting other factors for labor hours, rather than reducing all factors.

Because efforts to reduce output should be associated with reductions in some of the other inputs, the fact that the other input residual increased on average for the entire economy (Figure 4) and for the non-manufacturing, non-construction parts of the business sector (Table 2), calls into to question the assertion that most industries cut their employment because of a lack of customer demand. A more obvious explanation for a substitution away from labor and toward other inputs is that businesses perceive labor to be more expensive than it was before the recession began. Below I use the consumption and productivity data to begin to quantify labor cost effects like this, and Mulligan (2011) separately quantifies labor cost effects using public policy measures.

V. Neither Wealth Effects nor Intertemporal Substitution Effects Explain the "Supply" Shift

Although my finding that the marginal productivity schedule (Figure 4) was stable or shifting up since 2007 rules out reductions in the usage or efficiency of complementary production factors as primary explanations for the sharp labor decline and prolonged recovery, that still does not tell us why the economy was on one part of that schedule in 2007, and a quite different part in the years thereafter. Economists often interpret movements along that schedule as changes in the supply of labor, or changes in the ability of the labor market to coordinate supply and demand (hereafter, "labor market distortions"), and interpret the schedule itself as the demand curve for labor.

An influx of immigrant workers into the economy is an example of a supply change. The influx would, in the short run, push down wages as people compete for jobs, and lower wages would induce employers to hire more. In the process of putting more

[20] Each quarter's log change in real value-added for (business – manufacturing – residential construction) is calculated so that it, together with the log changes for manufacturing and residential construction, average to the log change for the entire business sector, using each component industry's nominal value-added as weights.

people to work, output would increase but productivity would fall. That would be a movement along Figure 4's marginal productivity schedule in the direction of more labor and less productivity. A reduction in wealth is another supply change that would increase labor and reduce productivity.

A reduction in worker marginal tax rates – that is, an increase in the share of additional earnings a worker keeps after taxes – can also be interpreted as a supply change, and would also move the economy down the marginal productivity schedule. But a reduction in employer payroll taxes would, for similar reasons, also increase labor and reduce productivity, so taxes are often referred to as labor market distortions.

Consumption and Leisure have Moved in Opposite Directions

The wealth effect explanation for movements up the marginal productivity schedule says that people work less because they feel richer. The intertemporal substitution effect says that people work less in 2009 because they view 2009 as a relatively bad time to work and produce, either because the return to saving is low or because they expect future labor productivity to be even higher than it is now. Both the wealth and substitution effect theories imply that consumption is *high* during the recession (Barro and King 1984).

Figure 2 easily rejects the wealth and intertemporal substitution effect explanations for low aggregate labor because consumption expenditure has been low in this recession. Judging from the consumption drop, wealth and intertemporal substitution effects by themselves would be moving the economy down the marginal productivity schedule shown in Figure 4 – in the direction of more labor – so something else must be moving the economy up the schedule even more than the total change that combines the wealth and intertemporal substitution effects with other effects. In other words, if labor and productivity had remained constant while consumption had dropped, that itself would indicate an important change in the labor market because we expect adverse wealth effects to be associated with more labor and lower labor productivity.

A Labor Market Metric for Consumption Declines

Putting more structure on preferences for consumption and work permit me to quantify the size of the wealth and intertemporal substitution effects, and thereby the size of the leftward labor supply shift (or labor market distortion change) that would have occurred absent those effects. In particular, I assume that the month t marginal rate of substitution (*MRS*) between consumption and leisure is proportional to the ratio of real consumption per person, and proportional to work hours per adult:

$$MRS_t \sim \frac{c_t}{P_t}\left(\frac{n_t}{N_t}\right)^{1/\eta} \tag{0.5}$$

where c_t is aggregate real consumption of nondurables and services (including public non-defense consumption), P_t is population (adults and children), N_t is the adult population, and n_t is total labor time. η is a constant, assumed for the moment to be one, that can be interpreted as the Frisch elasticity of labor supply with respect to wages.

The *MRS* can be interpreted as the reservation wage of the marginal worker: the marginal worker is willing to work if and only if he is offered a wage that equals or exceeds his *MRS*. According to (0.5), this reservation wage increases with the marginal worker's living standard as measured by real consumption per capita: the lower is the living standard the lower the wage the marginal worker will accept. The dependence of the *MRS* on the amount of work hours n means that people are not willing to work still more hours unless offered a wage that is especially high. In other words, the equation version of (0.5) graphed in the [n,*MRS*] plane is a kind of labor supply function, with upward slope whose magnitude is determined by η, and with consumption shifting the function up (a "wealth effect").

Absent labor market distortions and other determinants of the marginal rate of substitution, the marginal rate of substitution would equal marginal labor productivity. Given the Cobb-Douglas assumption (0.1), marginal productivity is proportional to labor

productivity, so that both average and marginal productivity changes have the same log changes over time. As explained by Mulligan (2005), changes in the gap between (0.5) and average productivity y_t are therefore measures of changes in the combined effect of changes in labor market distortions and other (omitted) determinants of the marginal rate of substitution. Denoting that gap as (1-τ_t), its changes can be calculated as:

$$\Delta \ln(1-\tau_t) \equiv \left[\Delta \ln(c_t / P_t) + \frac{1}{\eta} \Delta \ln(n_t / N_t) \right] - \Delta \ln y_t \qquad (0.6)$$

In words, each log point that consumption declines is a log point that distortions must increase in order to explain a given path for labor and productivity.

With the only the data presented in this paper, one cannot determine whether the gap (1-τ_t) captures preferences or distortions. Henceforth, for the purposes of brevity, I refer to -ln (1-τ_t) as the amount of labor market "distortion."

An appendix to this paper examines sensitivity of the results to alternative assumptions about functional form, elasticity magnitude, and consumption concepts.[21] However, only weak assumptions are needed to conclude that equation (0.6) is correct at least in terms of the qualitative effects of consumption, labor hours, and productivity on the labor market distortion.[22] Given that consumption and labor clearly fell during the recession, and labor productivity clearly rose, we must conclude that the distortion increased; the only question is the amount of that increase.

[21] We also expect the *MRS* to depend on additional demographics, such as the age of the population, but demographics change very little at a business cycle frequency as compared to the amounts of change in consumption, labor, and productivity. For this reason, my analysis does not consider any additional demographics.

[22] The marginal rate of substitution function (0.5), which corresponds to the utility function $u(c/P, n/N) = \ln(c/P) - \gamma \frac{\eta}{\eta+1}\left(\frac{n}{N}\right)^{\eta/(\eta+1)}$, can be interpreted as a first order approximation to the marginal rate of substitution function associated with any utility function $u(c/P,n/N)$ with those two arguments, and having the properties that (a) consumption and leisure are normal goods, (b) the indifference sets are convex, and (c) consumption and the *MRS* can have the same trend without any trend for labor per adult.

VI. Labor Market Distortions since 2007

Figure 6 graphs quarterly changes in the labor market distortion, together with its supply or "reservation wage" component (the square bracket term in equation (0.6)) and its productivity component y_t. Distortions increased throughout 2008 and 2009. During 2008, much of the increase can be described as falling consumption and labor in the face of fairly constant productivity. In other words, by the end of 2008 living standards decline – consumption per person had fallen two percent – and this by itself reduced the reservation wage schedule by two percent. With a lower reservation wage, and constant productivity, we expect people to work more, not less. In fact, the elderly did work more after 2007, and likely for this reason, but nonetheless Figure 1 shows that the average person was working less by the end of 2008. With people working less, we further suspect that reservation wages were low: an additional three percent lower as of the end of 2008. Thus, the reservation wage or MRS value for 2008 Q4 shown in Figure 6 is -0.051 (see the blue curve). In order to rationalize these outcomes in the face of productivity that was essentially constant, the distortion term must have changed about 5 percent in the direction of greater distortion; +0.050 is graphed in Figure 6 as the 2008 Q4 value for the distortion (see the red curve).

During 2009, productivity increased and the reservation wage fell further, which implies that labor distortions increased further. The total distortion change from 2007 Q4 to 2009 Q4 was 0.159. To put this in perspective, to explain the labor market events with, say, an across-the-board labor income tax hike, the amount of the hike would be about 16 percentage points.[23] Since 2009, productivity, the reservation wage, and the distortion have been fairly constant.

As shown in the Appendix, the 1981-82 recession's residuals were quite different: the input residual fell while the labor supply residual was constant. In addition to having

[23] More precisely, from a base tax rate of zero, the hike would have to be 15 percentage points because 0.15 = $e^{-0.16}$: recall that the distortion in Figure 6 is measured as $-\ln(1-\tau)$.

much larger reductions in labor and output, the 1930's "Great Depression" was also different: both the input and labor supply residuals fell.

Wage measures are not part of equation (0.6) but, to the degree that wages can be reliably measured, could be used to decompose the overall distortion – the difference between log productivity and log *MRS* – into a difference between log productivity and the log wage and a difference between the log wage and the log *MRS*. The former difference is sometimes associated with employer-side distortions, such as employment costs they incur apart from wages, and the latter difference associated with employee-side distortions such as income taxes owed by employees on their wage income. However, given that the gap between wage changes and productivity changes is small by comparison with the overall distortion (compare Figures 3 and 6), I do not present such a quantitative decomposition, except to note that most of the distortion since 2007 appears to be employee-side distortions rather than employer-side distortions.

The calculation of the distortion for any particular date can be illustrated in a diagram like Figure 4 by adding a *MRS* or reservation wage schedule (0.5) to the marginal productivity schedules already shown. Figure 7 therefore takes Figure 4 and expands the scales of the axes to make room for a reservation wage schedule, and deletes all of the observed values of labor and productivity except for 2007 Q4 and 2009 Q4. Due to their familiar geometric shapes and conceptual relations with demand and supply, marginal productivity schedules are labeled as "demand" and *MRS* schedules are labeled as "supply." In a labor market without distortions, the quantity of labor is determined by the intersection of supply and demand.

The 2009 Q4 demand schedule (solid greed) is the schedule satisfying the formula (0.2), and taking on the other inputs value *A* that fits the actual labor and productivity for 2009 Q4. The 2007 Q4 supply schedule is the schedule satisfying the formula (0.5) for a fixed amount of consumption per person. The 2009 Q4 supply schedule is the same as the 2007 Q4 supply schedule, except that it is shifted vertically in the amount of the log per capita consumption change over those eight quarters, -0.030. In other words, the reservation wage schedule was about 3 percent lower in 2009 Q4 than eight quarters earlier. As illustrated in Figure 7, the actual 2009 Q4 values for labor and productivity do

not lie on the 2009 Q4 supply schedule, and the distortion is measured as the vertical distance between the supply schedule and the actual values, which turns out to be 0.159.

Figure 7 can be used to simulate hypothetical values for labor and labor productivity if labor distortions and the labor supply function had remained unchanged since the beginning of the recession yet consumption, population, and the input residual had followed their actual values. This hypothetical outcome is labeled HA in the Figure, and is at the intersection of the 2009 Q4 supply and demand schedules. If instead of falling 0.084 log points, labor had *risen* 0.038 log points (as at the hypothetical outcome HA), the outcome would have been exactly on the 2009 Q4 supply schedule and the distortion term would have been zero. In other words, the actual labor supply distortion not only prevented an increase in labor that would have been consistent with the consumption drop, but actually reduced labor.[24] In this sense, the labor supply distortion is responsible for more than 100% of the employment decline since 2007. As future work seeks to explain the origins of growing labor distortions since 2007, Figure 6's red series is a guide as to how much their growth must have been.

VII. Conclusion: Productivity Patterns Begin to Reveal the Recession's Causes

Employment, hours, and consumption per person declined significantly in 2008 and 2009, while real wages and labor productivity rose. Since 2009, none of these variables have returned to their pre-recession values. This paper decomposes the hours and productivity changes into three types of "causes":

[24] If 2009 Q4 real consumption per consumption per capita had been the same as in 2007 Q4, the productivity residual followed its actual values, and the labor supply distortion had not changed over time, then the supply schedule would be unchanged and log labor would have increased 0.015 log points, as shown in the Figure as the hypothetical outcome HC. Mulligan's (2011) general equilibrium model combines the marginal productivity (0.1) and marginal rate of substitution equations (0.5) together with equations describing the evolution of consumption and capital in order to simultaneously determine the effects of distortions on all four variables.

- changes in other production inputs that change labor and output in the same direction,
- (unmeasured) labor distortions and labor preferences that raise productivity, reduce labor, and reduce consumption, and
- wealth and intertemporal substitution effects that reduce labor and raise consumption.

The macroeconomic concept of the "marginal productivity schedule" relating wages or output per hour worked to the number of hours worked – a concept shared by real business cycle models, New Keynesian models, and even Keynes' (2008/1936) *General Theory* – helps to isolate the first group of causes. Output declined enough less than work hours that it appears that other production inputs (aside from work hours) tended to *increase* during the recession. When viewed through the lens of *any* model in which aggregate output is a function of labor hours and other inputs with an elasticity of output with respect to hours of about 0.7,[25] the recession and slow recovery cannot primarily be explained by, or even associated with, adverse productivity shocks such as interruptions to the production process created by unusually bad weather, strikes, labor immobility, or business sector attempts to economize on raw materials.

Perhaps the most commonly cited theories of the 2008-9 recession are that the housing collapse, stock market crash, and/or the banking crisis caused people to spend less, so that the businesses serving those spenders experienced less demand for their products. Rather than cutting prices to induce customers to continue buying the quantities that they did before the recession began, those businesses decided to cut output. These spending-impulse theories may be a good description of the declines in manufacturing, residential construction, and perhaps a couple of other industries, and can explain why labor productivity increased. But such theories cannot explain the apparently increased use of other production inputs that occurred on average for the whole economy.

[25] Between 2007-Q4 and 2009-Q4, log real GDP declined 0.039 and log labor hours declined 0.082, which means that other inputs increased as long as the elasticity of real GDP with respect to hours is greater than 0.48 (=39/82), even if that elasticity varies with the amount of hours. In order to conclude that other inputs increased in the non-manufacturing non-residential-construction parts of the business sector, the elasticity of real value-added with respect to hours must only exceed 0.36 (see the final column of Table 2).

Factor substitution motivated by a perceived increased cost of labor can explain why output and work hours fell while the usage of other inputs increased. That increase can come on the employer-side, as with an anticipated and often discussed employer tax credit for new hires, which amounts to an implicit tax hike on the payroll employers have before the tax credit goes into effect.[26] Another example is that health care reform or some other forthcoming employer regulation will create employer liabilities based on the number of employees they had in the past. Yet another possibility is that, thanks to the banking crisis, employers find payroll management more costly. In all these employer-side examples, labor productivity growth is consistent with a drop in labor demand, as evidenced by lower labor compensation per hour, to which marginal workers respond by not working. For this reason, the finding in Figure 3 that both labor productivity and labor compensation per hour are greater now than they were before the recession began suggests that perceived employer costs are not the primary reason for the sharp drop in labor usage between 2007 and 2009. On the other hand, the gap between productivity and hourly compensation changes since 2009 is consistent with the hypothesis that employment is recovering slowly in part because employers perceive employment costs that exceed measured labor compensation.

Employee-side distortions, shifts in labor supply, or a failure of nominal wages to fall enough to clear the labor market, are all consistent with rising productivity and real wages during the recession. The last part of this paper used the theory of labor supply to quantify the combined amount of the distortions or supply shifts. The labor supply theory, also shared by real business cycle and New Keynesian models, says that workers' reservation wages increase with living standards and with the amount of time worked. With consumption and work time so much lower than they were when the recession began, the theory says that people would be more willing to work now than they were then,[27] unless something else were significantly reducing their reward to work, or reducing their willingness to work. Pinpointing such "distortions" is beyond the scope of

[26] The implicit tax accrues on payrolls before the credit goes into effect, but is not collected until afterwards, which is why it is excluded from the usual measures of labor compensation at the time that the implicit tax accrues.

[27] On the third group of causes, it is well known (e.g., Barro and King (1984); Hall (1997)) that previous business cycles do not appear to be wealth or intertemporal substitution effects because both labor and consumption decline. The 2008-9 recession is hardly unique in this regard.

this paper, but it does find that the distortions were, at their peak, as large as a 16 percentage point increase in the labor income tax rate.

VIII. Appendix: Productivity, Labor, and Residuals in Prior Downturns

This appendix examines the input and labor supply residuals for four previous postwar recessions, and for the Great Depression of the 1930s.[28] The changes during the 2008-9 recession were much less in magnitude than during the Great Depression, but often larger in magnitude than the changes during recessions since then. Of particular interest is that 2008-9 recession is quite different from (a) both the 1981-82 recession and the Great Depression in terms of the direction of the input residual, and (b) the 1981-82 recession in terms of the direction of the labor supply residual.

Figure 8 is the 1980s version of Figure 3 for the years 1981-83, except that real GDP per capita and real employee compensation is divided by a slightly different measure of hours (see below) than used in Figure 3. Note that the NBER dates the business cycle peak in 1981-Q3 and the trough in 1982-Q4. Both real GDP per hour and business sector output per hour decline significantly, although the two series disagree about the exact amount of the decline. Both series remained below pre-recession levels until the recovery was underway.

Real employee compensation per hour actually rises during the 1981-82 recession, and does not fall until the recovery. Both productivity and real wages are likely subject to a counter-cyclical composition bias (see section III in the body of this

[28] See also Ohanian (2010), who reports averages for all postwar recessions, with similar conclusions.

paper), but composition bias probably cannot explain why average hourly earnings increased rather than decreased.[29]

The input residual A_t is measured in the quantity dimension – that is, it is a horizontal distance in Figure 7 – while the labor supply residual is measured in the price dimension (a vertical distance in the Figure). In order to examine the two residuals in the same units, I transform the other input residual into a "productivity residual" by multiplying it by non-labor's share 0.3. Figure 9 is a scatter diagram of the labor supply residuals and productivity residuals for each of five recessions that began in 1974, 1981, 1990, 2001, and 2008. The chart's origin indicates the values for the quarter designated as the NBER business cycle peak, and the axes measure each residual as a gap from its value at the NBER business cycle peak.

The 1980s (green) and 2008-9 (red) recessions show two very different patterns. The 1980s recession was characterized by a reduction in the productivity residual, and little change in the labor supply residual.[30] The recovery from that recession involved a large increase in the productivity residual, but again little change in the labor supply residual by the seventh quarter after the NBER peak (the former was 1983-Q2; the business cycle trough was 1982-Q4).[31] The 2008-9 recession began with a reduction in both residuals, but it was the labor supply residual that dropped the most. Later in the recession, the productivity residual increased beyond its pre-recession values, while the labor supply residual continued to fall. A full recovery has not happened yet, but Figure 6 shows that the partial recovery so far that has occurred may have involved reductions in

[29] Employment declined three percent from peak to trough. As an upper bound on the composition bias, suppose that the bias was the result of workers exiting the workforce between peak and trough in the amount of three percent of peak employment, and whose average wage was one-half of the overall average (according to the Bureau of Labor Statistics (2011), average hourly earnings in the private sector were $7.44, half of which would be close to the federal minimum wage of $3.35). In this case, the business cycle peak wage would underestimate the average wage of those who would be working in the trough by about 1.5 percent.

[30] Despite the 1980s reduction in time worked, the labor supply residual did not fall because: (a) productivity fell, and (b) consumption per person hardly fell.

[31] The labor supply residual for 1981-82 would decline about 0.02 more if real hourly wages were used in the calculation, rather than real GDP per hour, but still much less than the labor supply residual drops shown for the other four recessions.

the productivity residual,[32] as opposed to the increases associated with many of the previous recoveries.

The recession beginning in 1990 may be the one most similar to the 2008-9 recession, because the productivity residual increased significantly from peak to trough, while the labor supply residual fell. Growth in the productivity residual stopped during the 1990s recovery, while the labor supply residual returned near to its pre-recession value.

The Great Depression of the 1930s was unique in its magnitude, and therefore not shown in Figures 8 and 9. Table 3 offers a comparison of the early 1930s to the economy since 2007. In this recession, the productivity residual has increased. The productivity residual fell more than 5 percent 1929-33 (Cole and Ohanian 1999), which is many times more than it did in the 1980s recessions. Both the 2008-9 recession and the Great Depression had labor distortions of more than 0.15 appear within a year or two, but after four years the Great Depression labor distortion (Mulligan 2005) was much larger than it is now. Although it is not clear whether the Great Depression was just an amplified version of the 1970s recession – with both the labor supply residual and the productivity residual falling – it is qualitatively different from the 2008-9 recession when the productivity and labor supply residuals moved in opposite directions.

IX. Appendix: Sensitivity Analysis

Many of the conclusions in this paper derived from the empirical finding that aggregate work hours dropped significantly more than output. Fortunately, aggregate work hours have been measured in various ways, depending on the sampling method, the concept of work time, and the means of eliciting work time. Regarding sampling

[32] In order to maximize comparability across recessions, hours worked for the 2008-9 recession are measured differently in Figure 9 than in the rest of the paper, because the basis for my preferred measure of hours (all employee hours in all private industries) is not available prior to 2003. For the purposes of Figure 9 (and Figure 8's real GDP and labor compensation per hour series), aggregate hours are the sum of private work hours per person (measured as the all-employees aggregate weekly hours index for the business sector) and aggregate public work hours (estimated as public sector employment times private work hours per private sector employee).

method, the U.S. government conducts a monthly survey of employers (the "establishment survey") and another monthly survey of households (the "household survey"). The establishment survey asks employers about their employees' paid work time (that is, including paid sick days and paid vacation days) during the reference week. The Bureau of Labor Statistics projects sample measures to the entire economy for that month, and those aggregate hours can be divided by total persons aged 16 and over (regardless of employment status) to arrive at an estimate of weekly hours worked per adult.

The household survey asks respondents about their work status (and the status of other adults in the household) during the reference week and, if at work during that week, hours worked. These hours worked, counting zeros for persons not at worked can be averaged across respondents to arrive at another measure of hours worked per adult. The two differ somewhat in terms of the population covered (the establishment survey does not include agriculture or self-employed workers).[33]

The Bureau of Labor Statistics combines the establishment and household surveys to form industry- and sector-specific measures of hours worked, which (as noted in the main text) I use for some of my productivity estimates.

But sometimes surveys can be misleading about hours worked, because there is a tendency for people to report round numbers like "40 hours" or "35 hours" even when actual hours worked are not a round number (over 40 percent of employed persons in the household survey reported that they worked 40 hours in the reference week, as opposed to a mere 0.4 percent who reported 39 hours of work). It is logically possible that a number of employed people were working more hours in recent years, but continue to report the round number of 40.

[33] Although the establishment survey measures jobs, and the household survey measures workers (some workers have more than one job), my establishment and household measures are not different in this regard because I summed hours in the establishment survey and divided by persons, not jobs. With the exception of the civilian employment series shown in Figure 1, I take averages from the household surveys, rather than totals, and therefore make no adjustment for the survey's January update of its population controls (United States Bureau of Labor Statistics 2011).

Since 2003 the Census Bureau has supplemented its household survey with a time diary study – the "American Time Use Survey" (ATUS) – dedicated to measuring time use. Participants in that survey are ask to account for all of their waking hours in a specific day with their various activities that day, including eating, watching TV, working, traveling, caring for children, etc. The diary study therefore has no automatic bias toward finding that masses of people work exactly eight hours every day for exactly 5 days per week. The Bureau of Labor Statistics averages daily hours worked across respondents, days of the week, and survey months to arrive at an annual average daily hours worked per person aged 15 and over. Figure 10 below displays changes in hours worked per person for each of the three measures. They closely agree: the indices for 2009 range from 91.1 to 91.7, while the 2010 indices range from 90.5 to 90.7.

The establishment survey and ATUS measures shown in Figure 10 are not available for previous recessions, and those previous recessions were too short to have a large enough sample from the household survey to accurately estimate hours changes change from peak to trough. The BLS does form an index of hours of production and non-supervisory for all private industries from the establishment survey, and has an index of hours of all business sector employees that is part of their productivity calculations (and is formed from both the establishment and household surveys). These two series disagree as to the amount that work hours dropped from peak to trough, by almost two percentage points. The household and establishment surveys for that time period disagree almost as much as to the peak to trough percentage employment change.

Table 1 2007-2009 Employment Rate Changes by Demographic Group and Industry

	Employment per person, % chg	Earnings per hour among those at work, 2007
by gender		
Men	-7%	21.17
Women	-4%	17.53
Both	-6%	19.60
by schooling		
Less than HS Grad	-13%	11.97
High School Grad	-8%	15.62
Beyond HS	-4%	23.23
by age		
16-24	-12%	11.77
25-39	-6%	18.85
40-64	-4%	21.83
65+	3%	19.46
by industry		
Construction	-19%	18.76
Manufacturing	-13%	20.21
Education	0%	21.22
Healthcare	3%	19.54
Other	-5%	19.35

Note: Each industry employment rate is measured as the ratio of industry employees to total adults

Source: Author calculations from CPS-MORG microdata.

Table 2. Factor Reduction and Factor Substitution by Industry

	Log Changes from 2007-Q4 to 2009-Q4			
	Factor Reduction Industries		Factor Substitution Industries	
	Manufacturing	Residential Construction	Wireless Telecommunications	Business Sector, excluding manufacturing & residential construction
real value added	-0.17	-0.42	0.09	-0.03
labor hours	-0.19	-0.44	-0.20	-0.08
labor's share, 2007	0.55	0.67	0.32	0.70
other inputs residual	-0.15	-0.36	0.22	0.09
Addendum: capacity utilization	-0.16	NA	NA	NA

Residential construction labor share, and hours per employee, assumed equal to the corresponding values for the entire construction industry.
Wireless telecommunications value-added changes are 2007-2009, rather than 2007-Q4 to 2009-Q4

Sources: BEA. Components of Value Added by Industry as a Percentage of Value Added
Census Bureau. 2009 Services Annual Survey Data: Information Services
BLS. Labor Productivity and Costs. Industry Employment and Hours
Board of Governors of the Federal Reserve System

Table 3. Residuals for the 2008-9 Recession and the Great Depression

	residual changes (log points, in the price dimension)	
	productivity	**labor supply**
2007 Q4 - 2009 Q4	+0.018	-0.16
2007 Q4 - 2011 Q2	+0.035	-0.13
1929-30	-0.053	-0.17
1929-33	-0.150	-0.46

Sources:
Cole and Ohanian (1999)
Mulligan (2005)

Figure 1. Employment and Hours per Person, Jan-2007 to Aug-2011

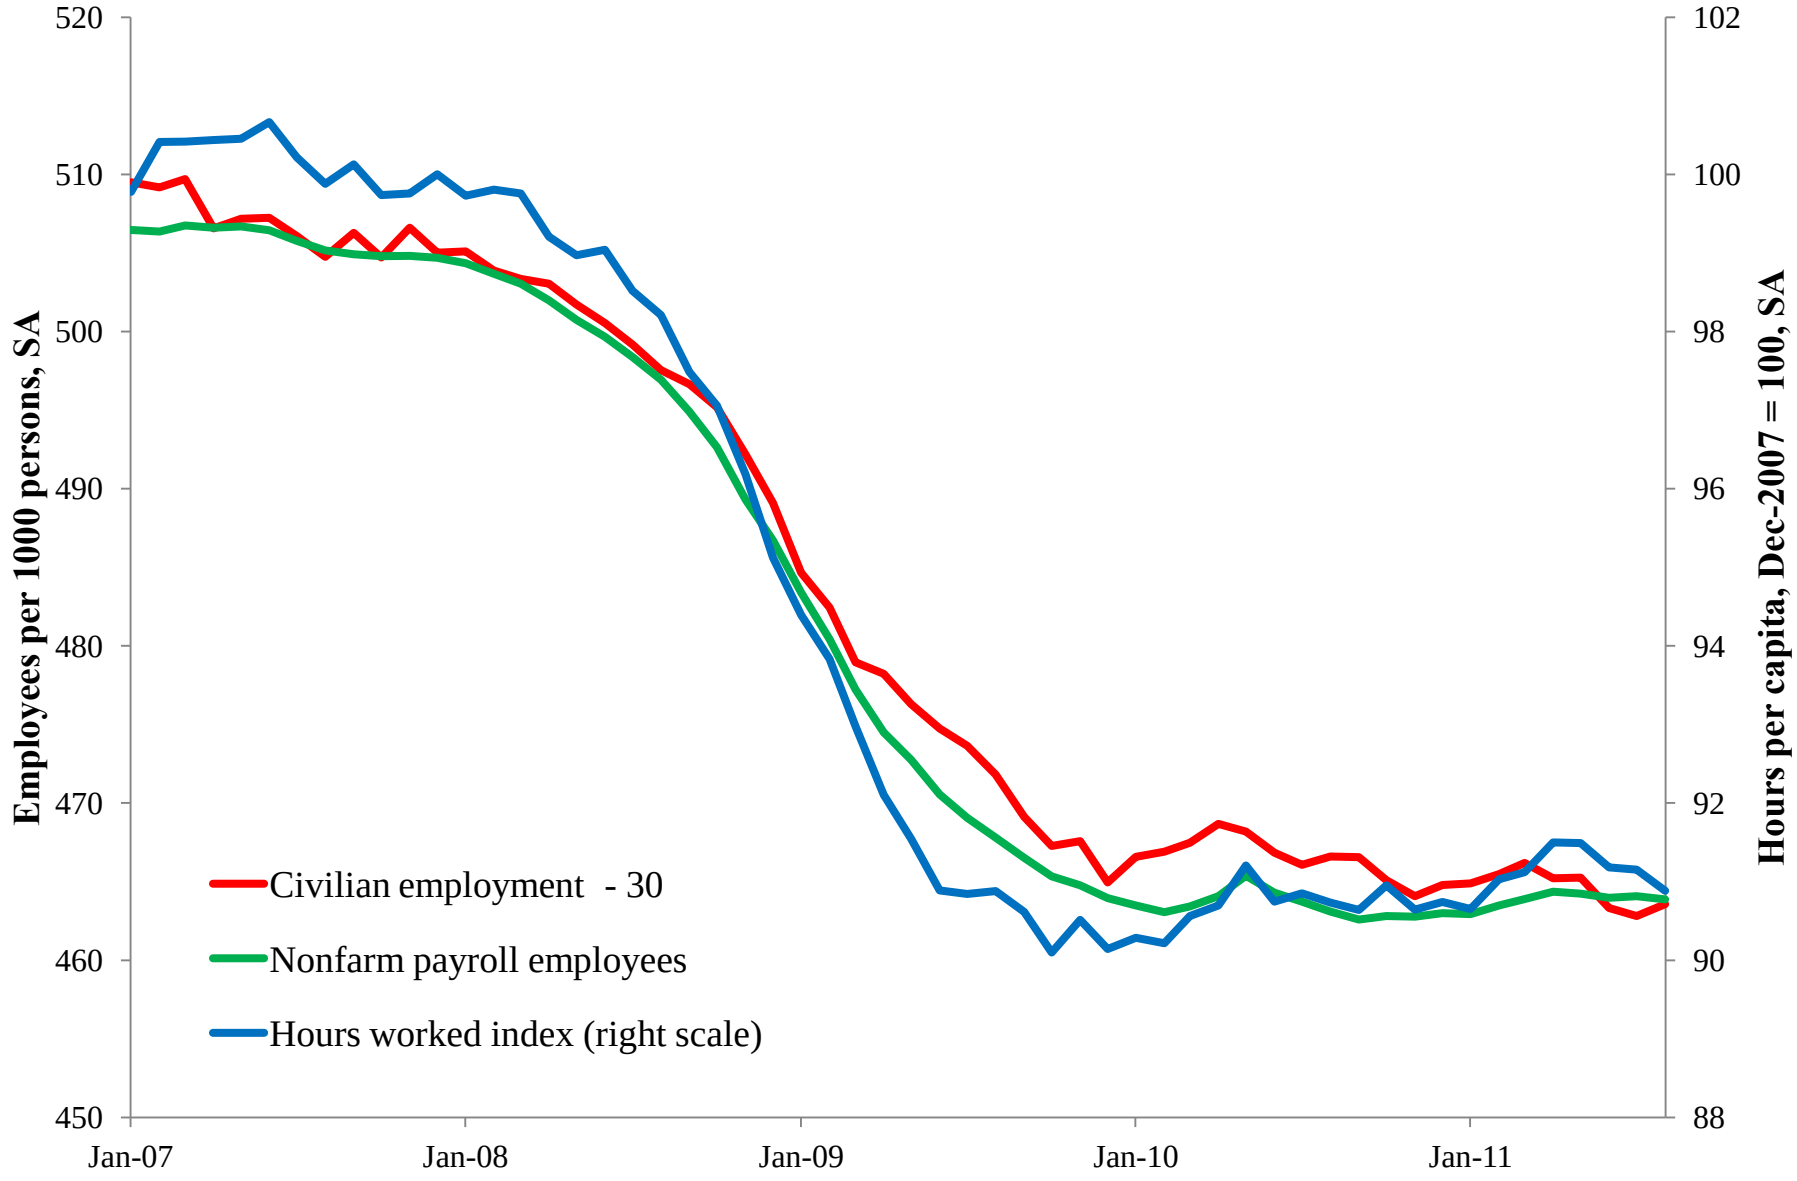

Figure 2. Real Consumption per Person, 2007-Q1 to 2011-Q2

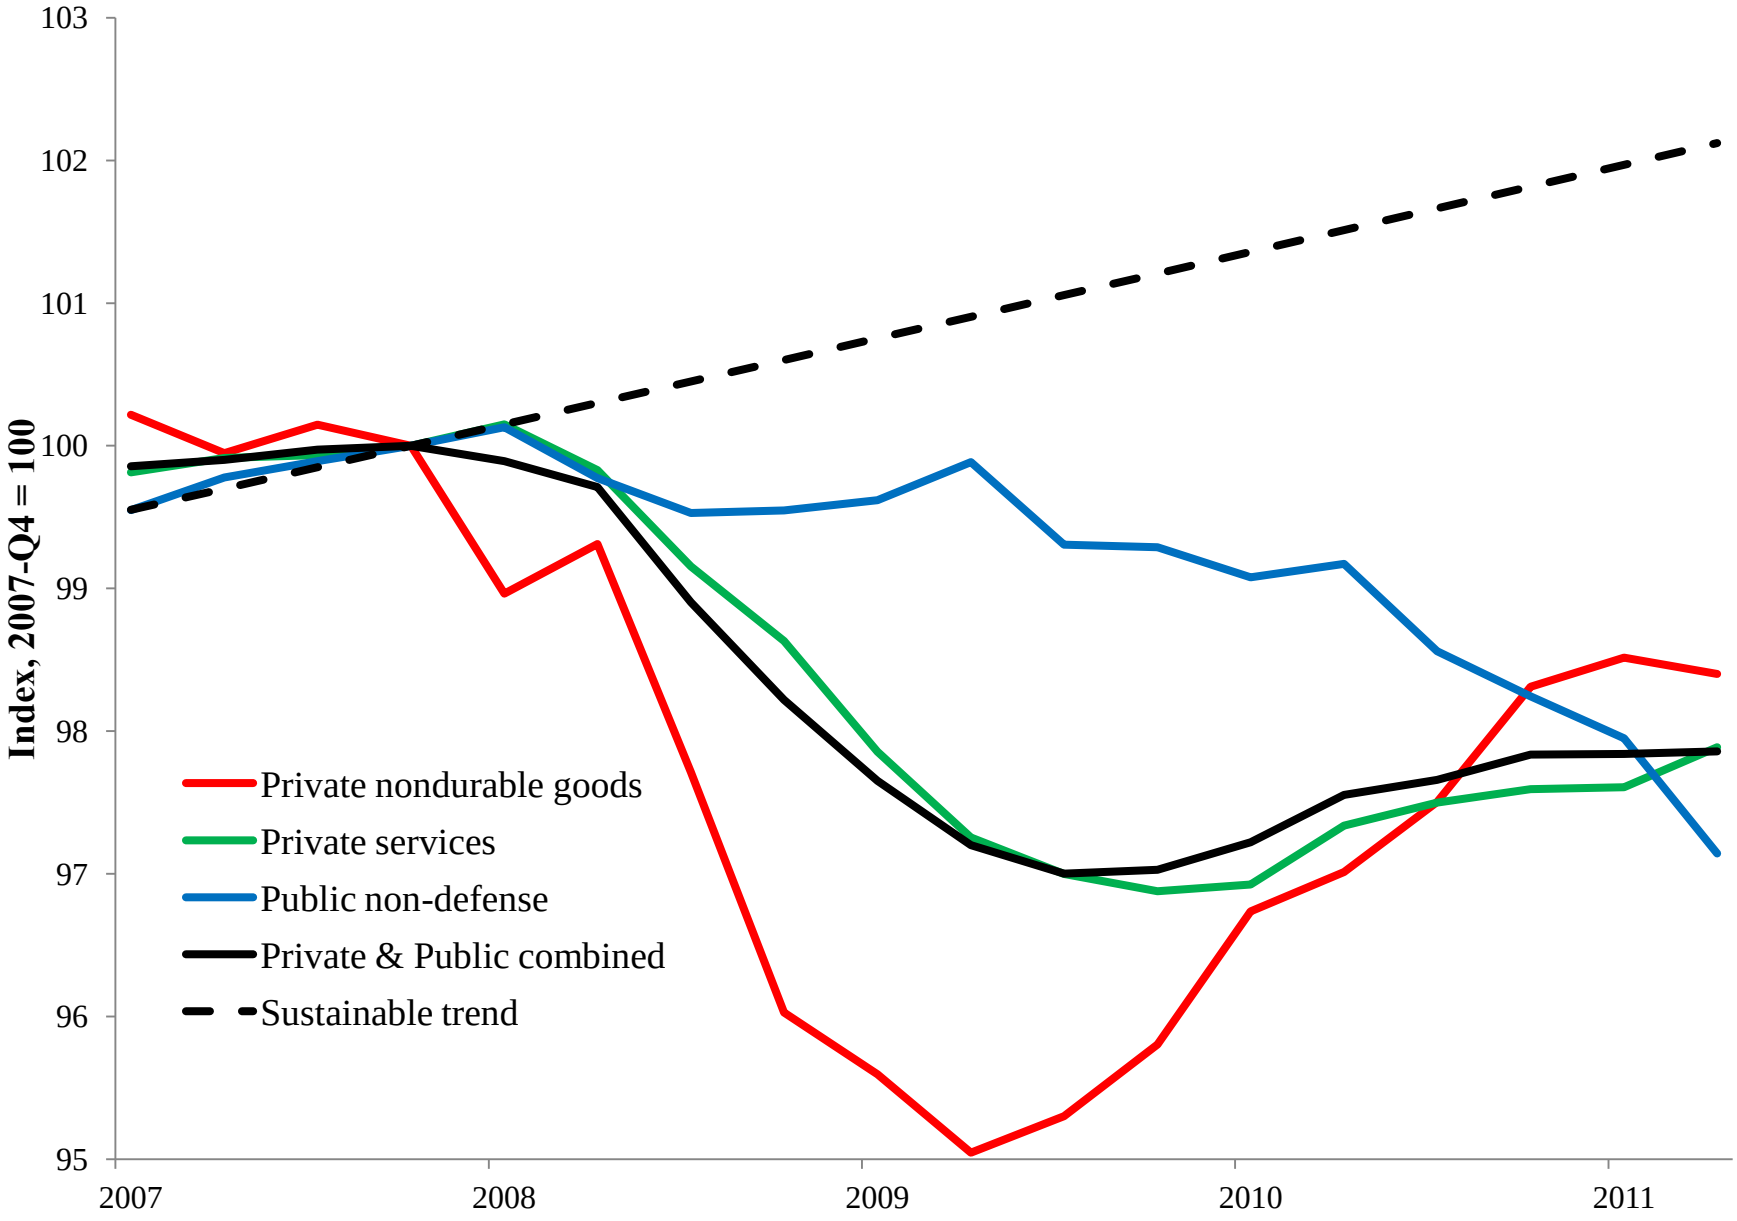

Figure 3. Labor Productivity, 2007-Q1 to 2011-Q2

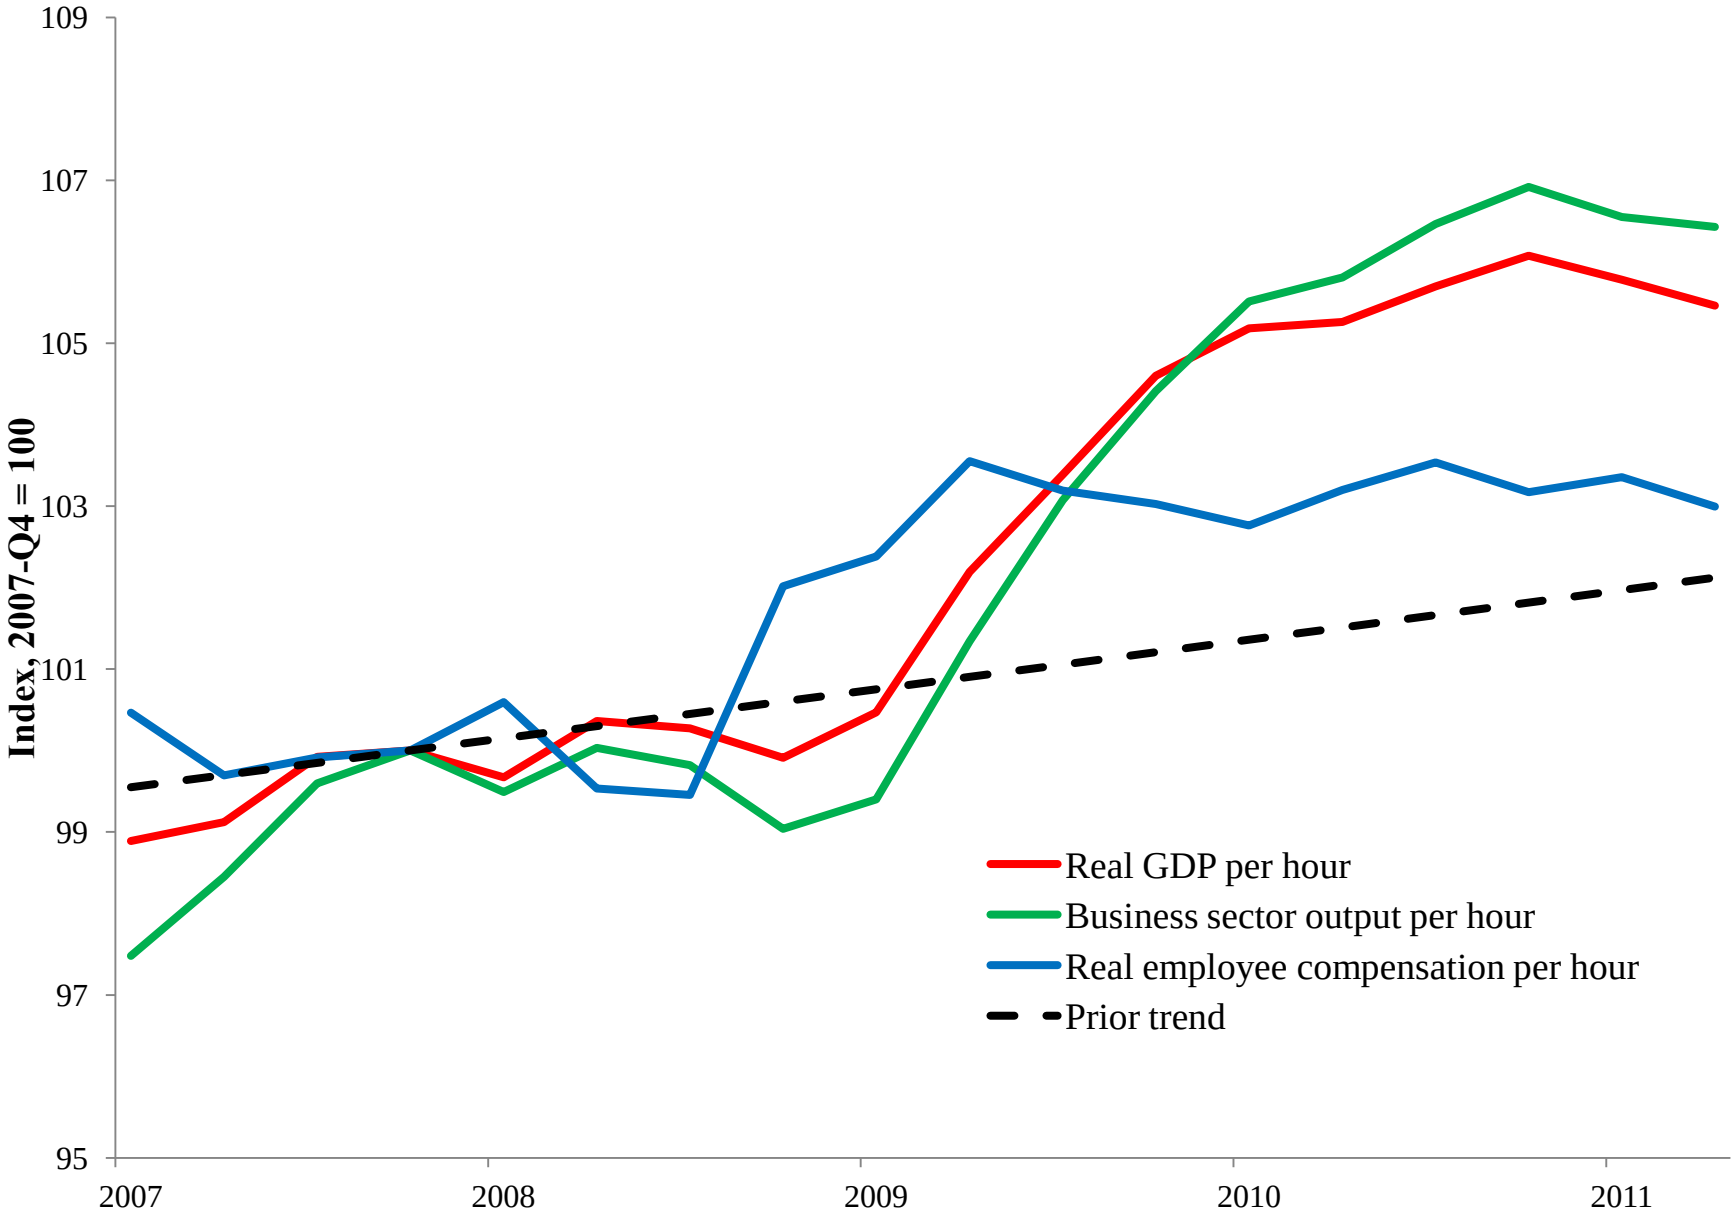

Figure 4. Marginal Productivity Schedules, 2007 Q4 to 2011 Q2

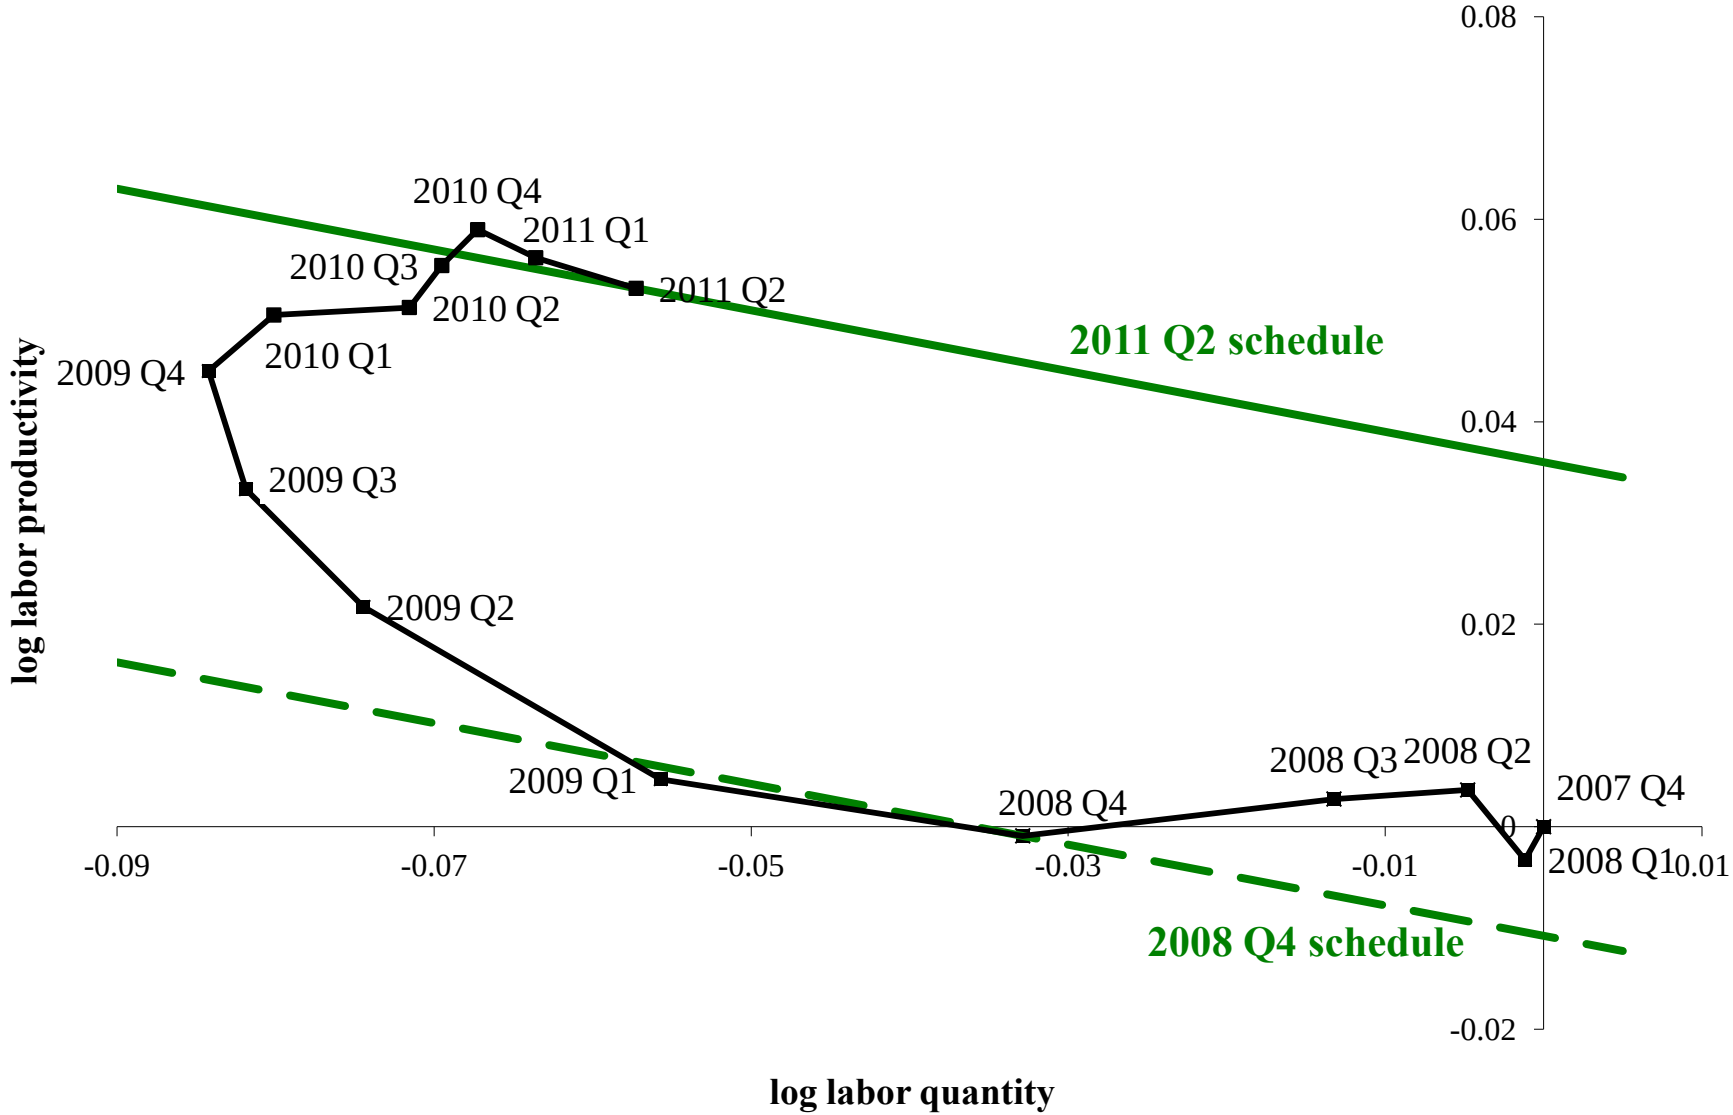

Figure 5. Labor and Other Inputs, 2007-Q1 to 2011-Q2

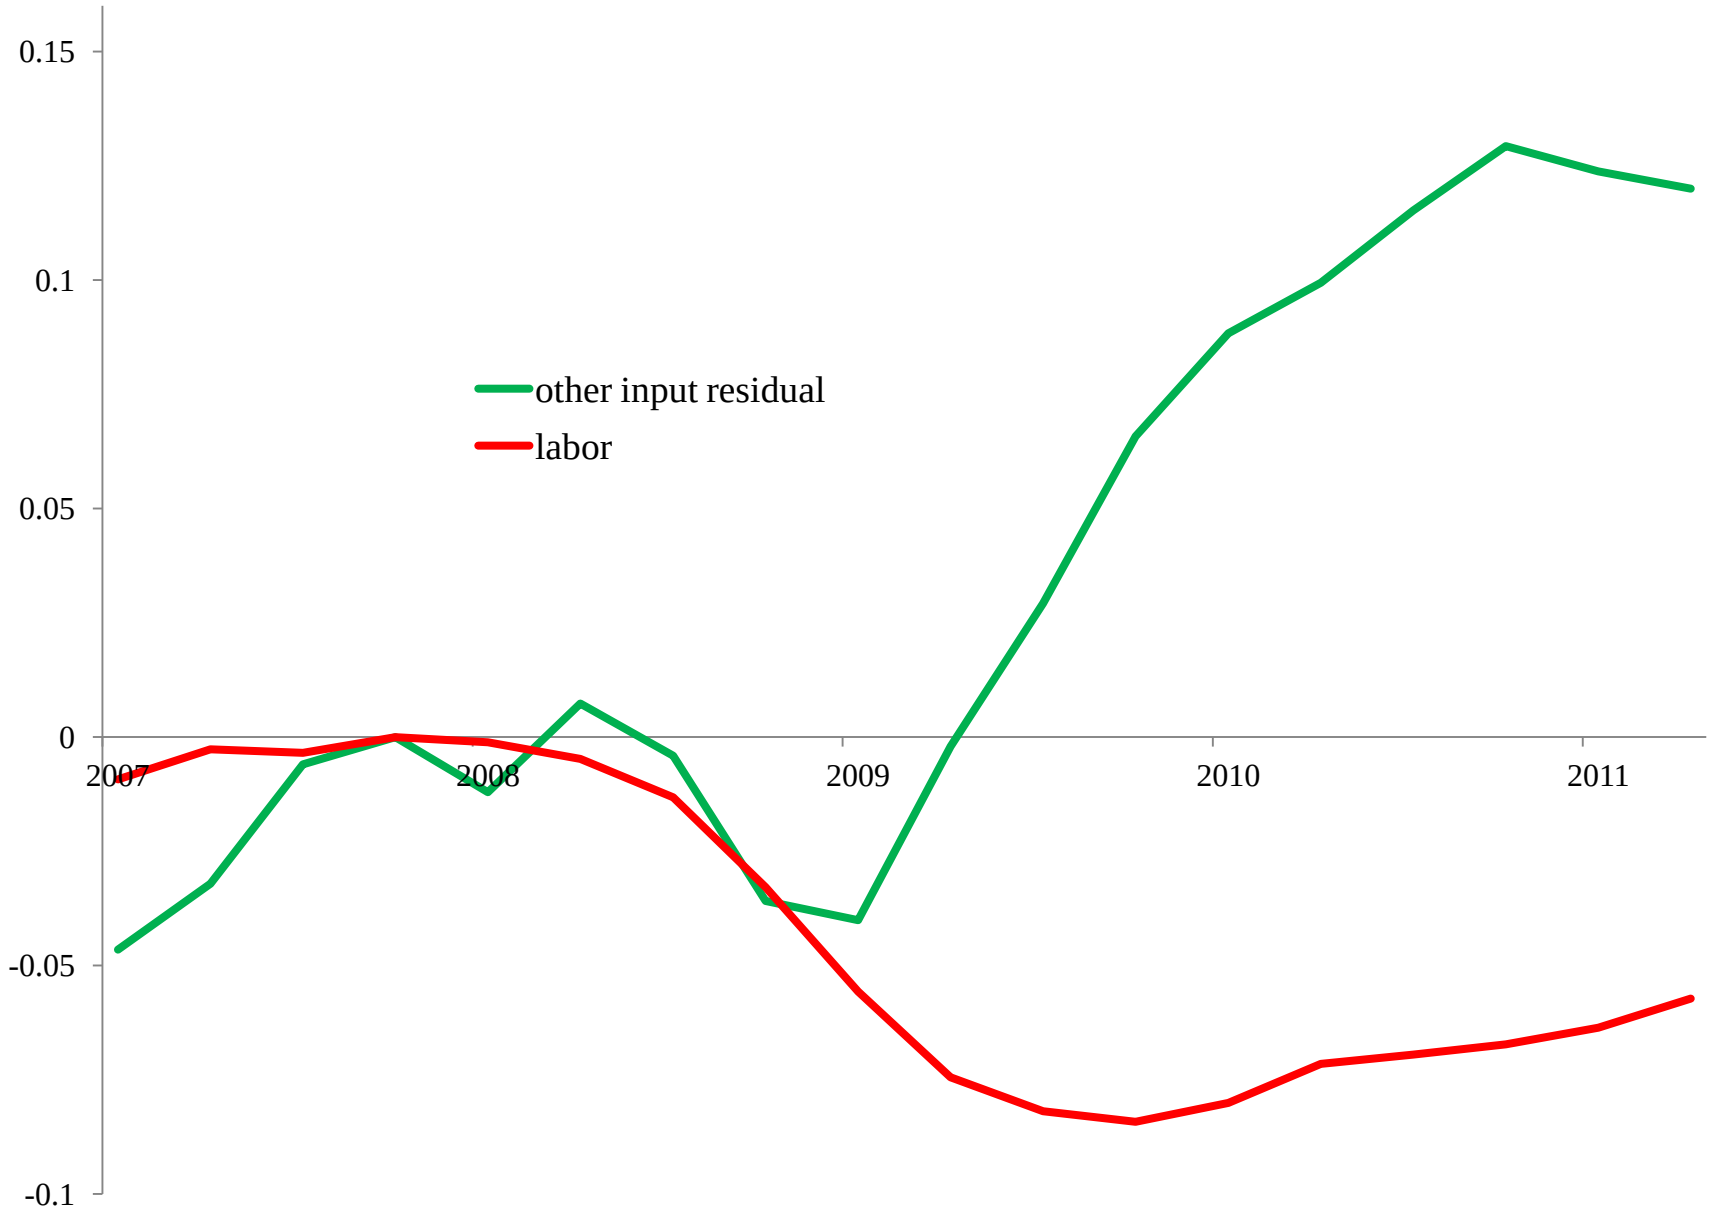

Figure 6. Supply, Demand, and Distortions since 2007-Q4

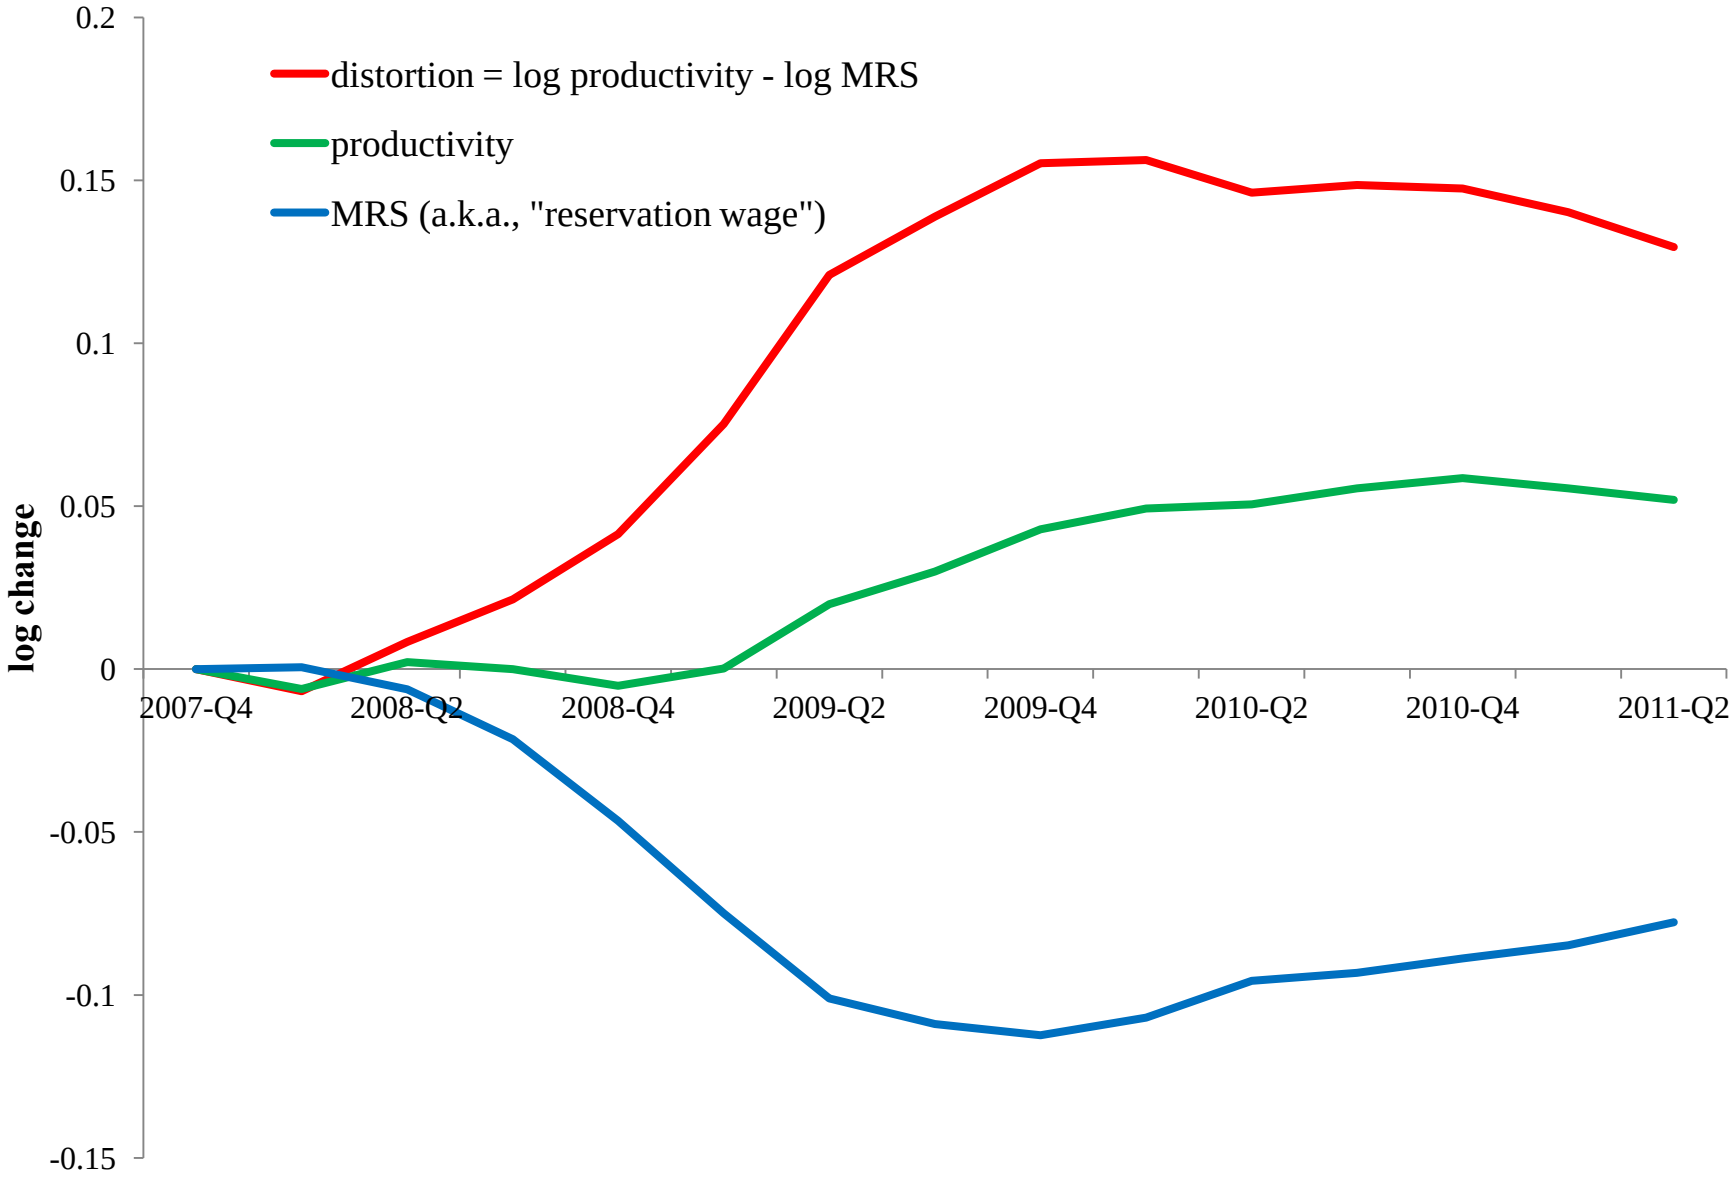

Figure 7. The 2009 Q4 Distortion, and Two Hypotheticals

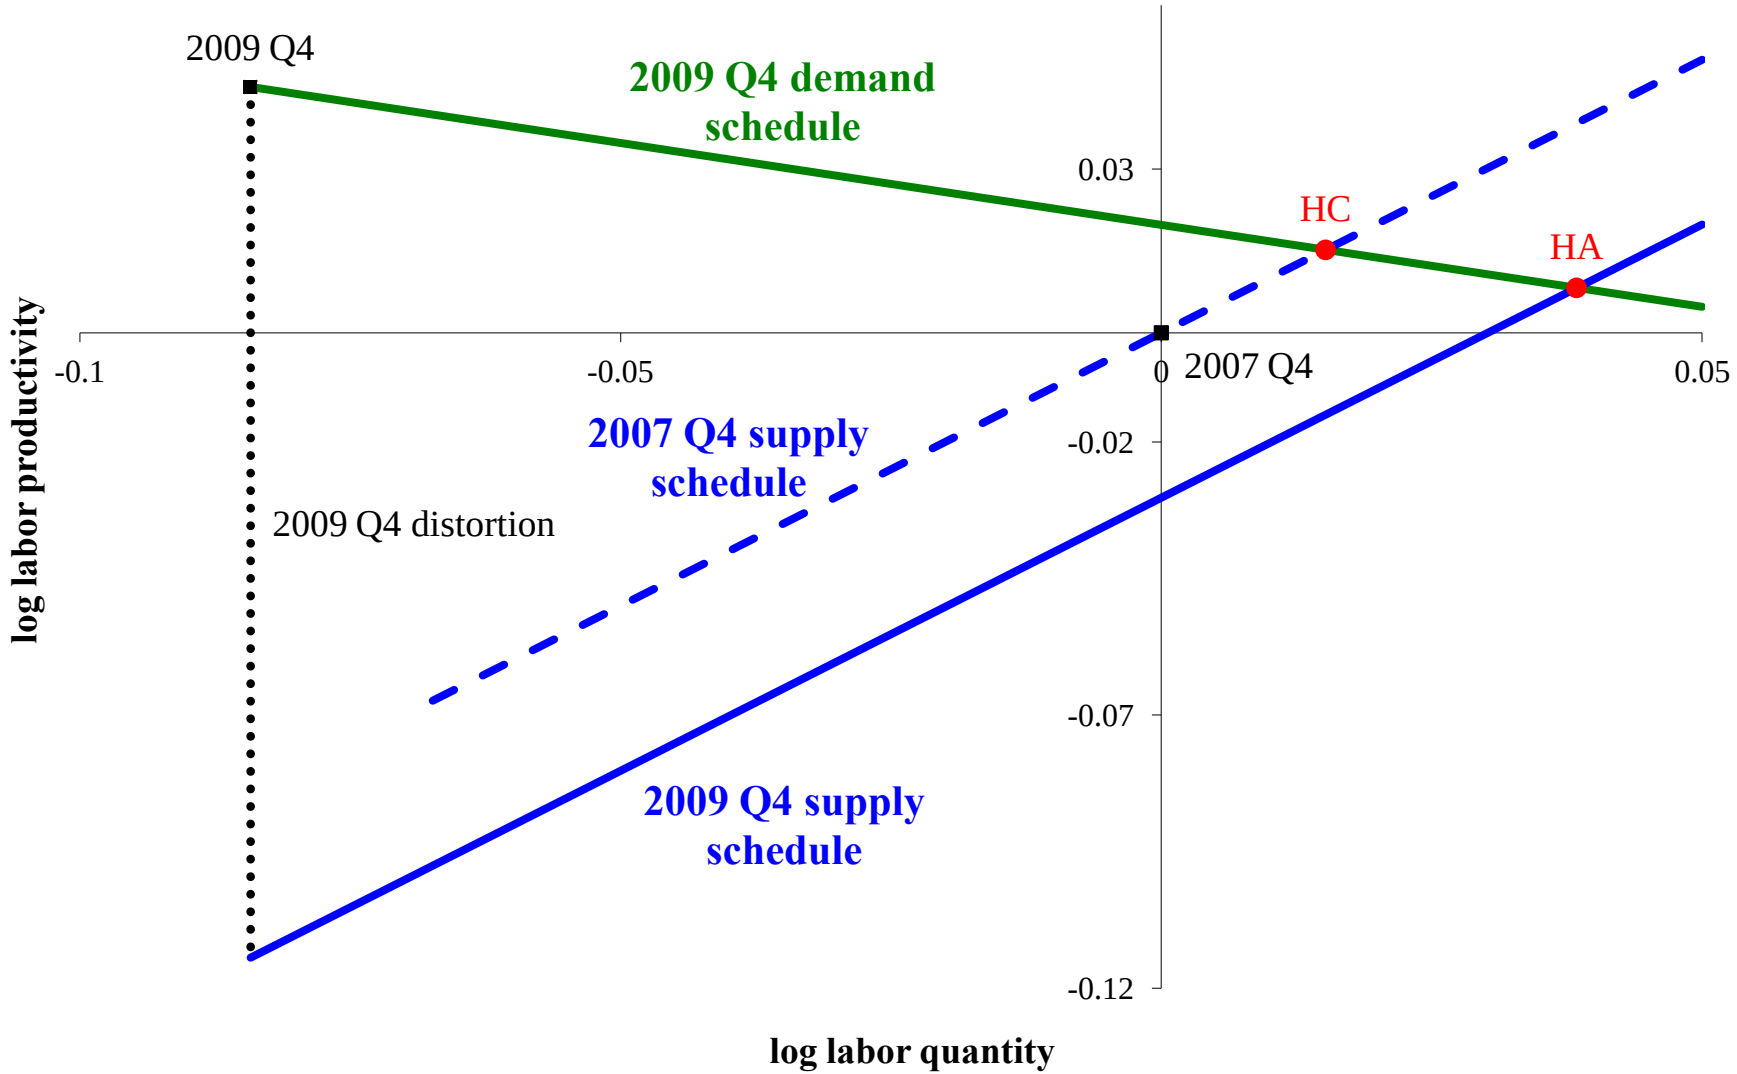

Figure 8. Labor Productivity, 1981-Q1 to 1983-Q4

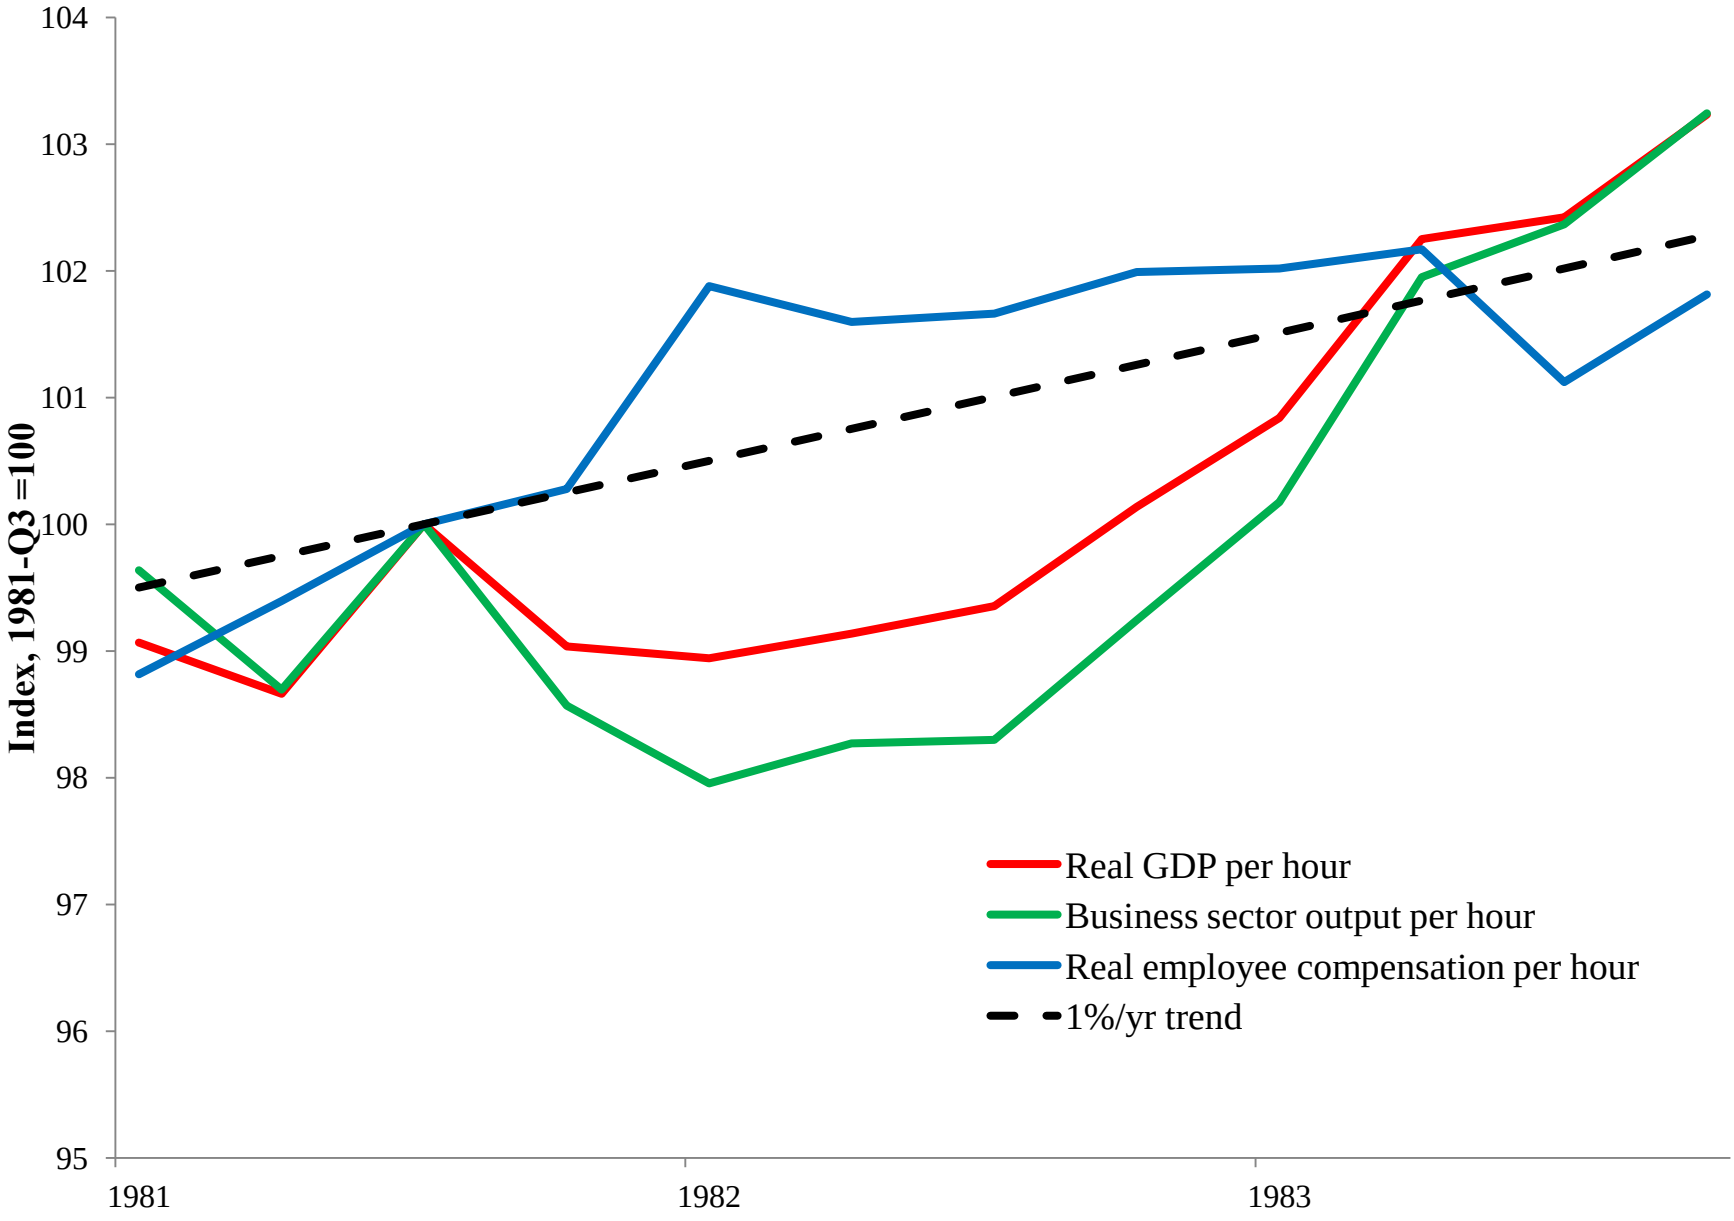

Figure 9. Labor Supply and Productivity Residuals in 5 Recessions

0-12 quarters from NBER peak

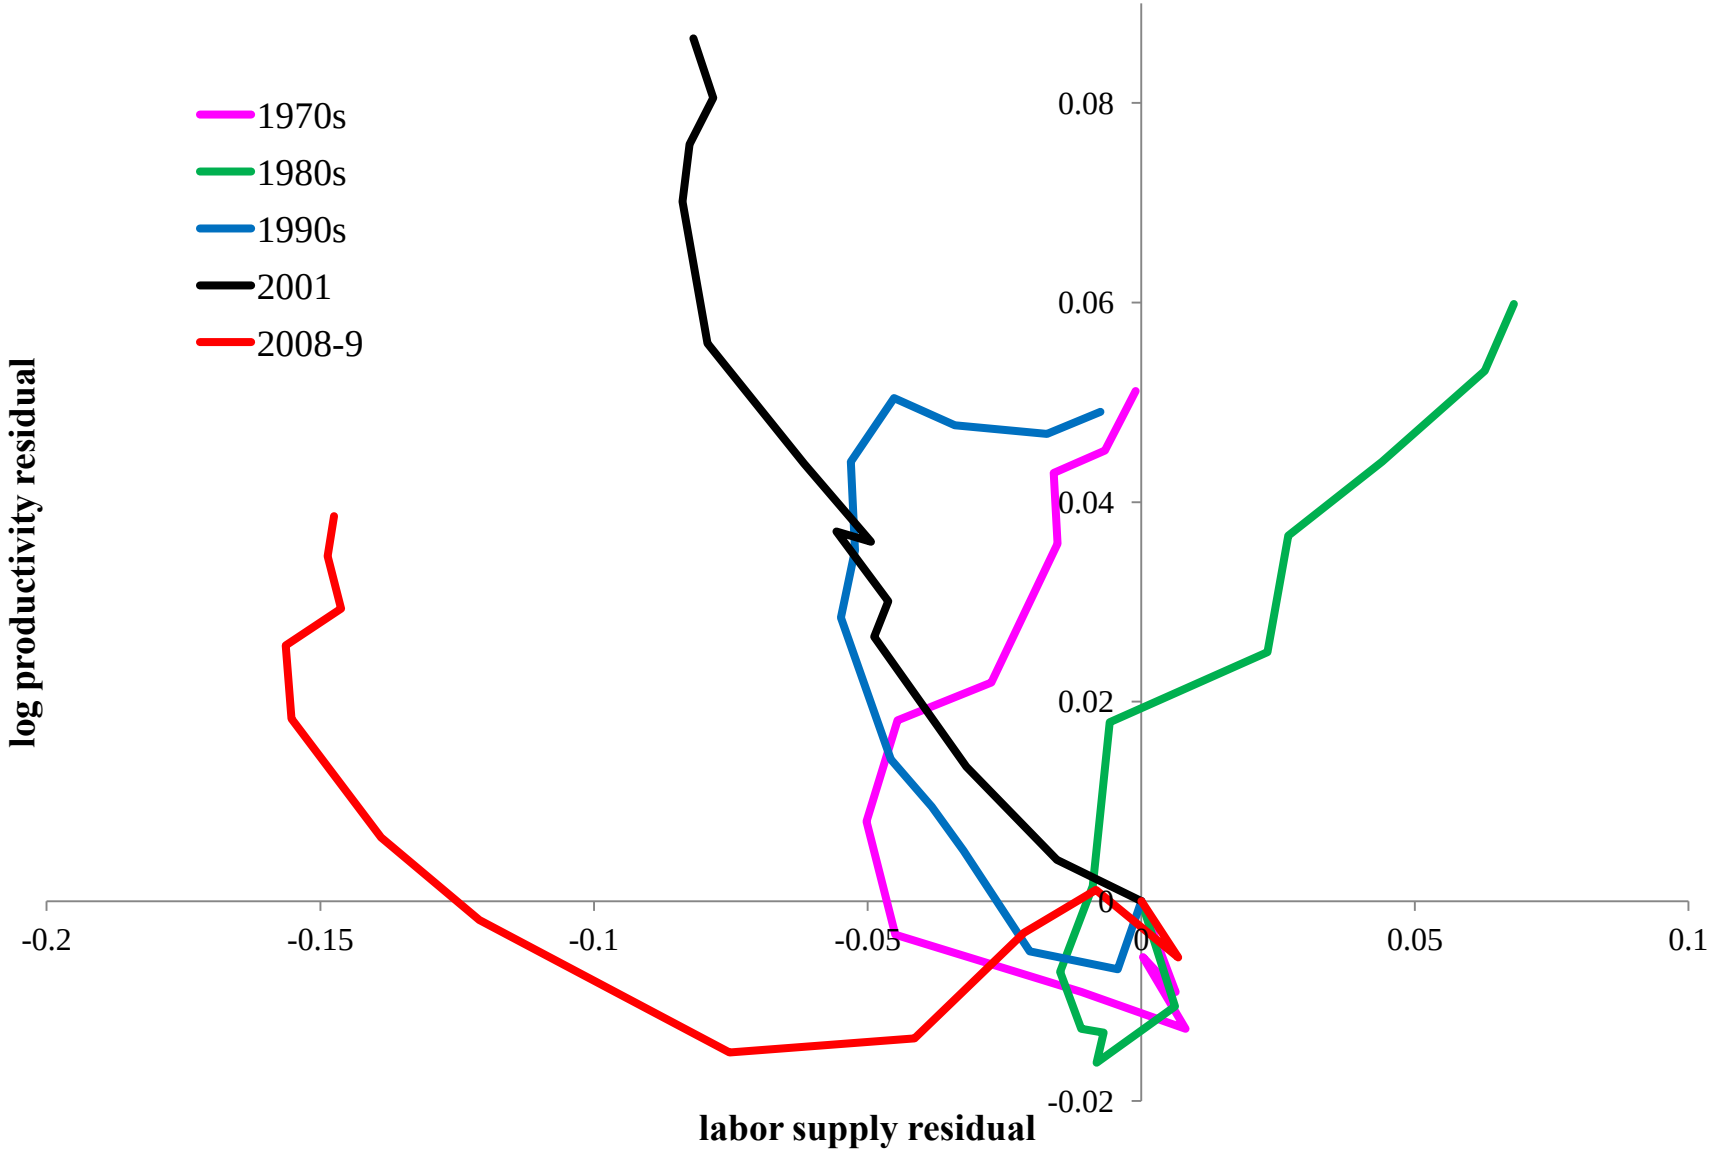

Figure 10. Work Hours per Person from Three Sources

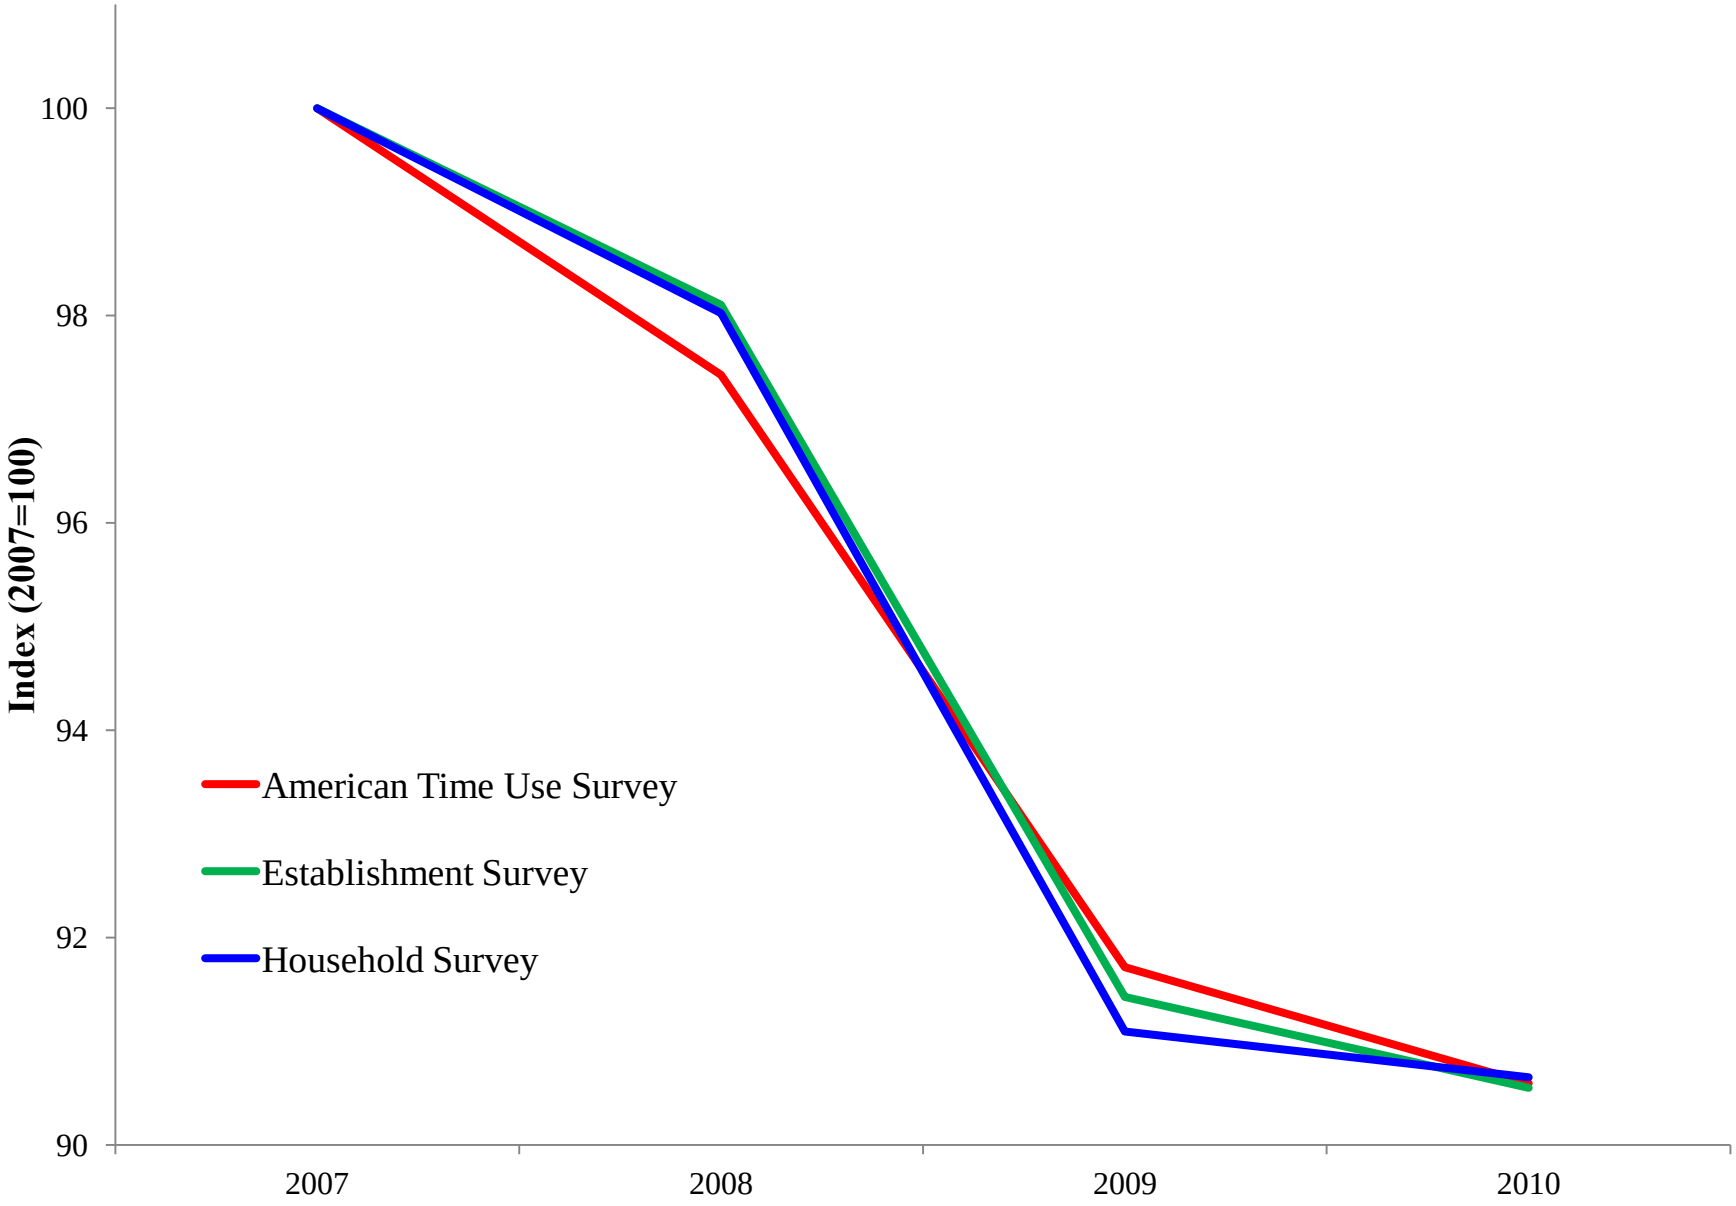

X. References

Abraham, Katharine G., and John C. Haltiwanger. "Real Wages and the Business Cycle." *Journal of Economic Literature* 33, no. 3 (September 1995): 1215-1264.

Barro, Robert J, and Robert G. King. "Time Separable Preferences and Intertemporal Substitution Models of Business Cycles." *Quarterly Journal of Economics* 99, no. 4 (November 1984): 817-39.

Bils, Mark J. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* 93, no. 4 (August 1985): 666-89.

Chari, V. V, Patrick J. Kehoe, and Ellen R. McGrattan. "Business Cycle Accounting." *Econometrica* 75, no. 3 (April 2007): 781-836.

Cole, Harold L., and Lee E. Ohanian. "New Deal Policies and the Persistence of the Great Depression: A General Equilibrium Analysis." *Journal of Political Economy* 112, no. 4 (August 2004): 779-816.

Cole, Harold L., and Lee E. Ohanian. "The Great Depression in the United States from a Neoclassical Perspective." *Federal Reserve Bank of Minneapolis Quarterly Review* 23, no. 1 (Winter 1999): 2-24.

Crucini, Mario J., and James Kahn. "Tariffs and Aggregate Economic Activity: Lessons from the Great Depression." *Journal of Monetary Economics* 38, no. 3 (December 1996): 427-67.

CTIA. *Wireless Quick Facts.* 2011. http://www.ctia.org/advocacy/research/index.cfm/aid/10323 (accessed October 24, 2011).

Federal Communications Commission. "15th Mobile Wireless Competition Report." *Commercial Mobile Radio Services (CMRS) Competition Reports.* June 27, 2011. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-11-103A1.pdf (accessed October 24, 2011).

Friedman, Milton, and Anna J. Schwartz. *A Monetary History of the United States, 1867–1960.* Princeton, NJ: Princeton University Press, 1963.

Gali, Jordi, Mark Gertler, and J. David Lopez-Salido. "Markups, Gaps, and the Welfare Costs of Business Fluctuations." *Review of Economics and Statistics* 89 (February 2007): 44-59.

Hall, Robert E. "Macroeconomic Fluctuations and the Allocation of Time." *Journal of Labor Economics* 15, no. 1, Part 2 (January 1997): S223-50.

Jorgenson, Dale W. "Introduction." In *The Economics of Productivity*, by Dale W. Jorgenson, 9-28. Northampton: Edward Elgar, 2009.

Katz, Lawrence F, and Kevin M. Murphy. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107, no. 1 (February 1992): 35-78.

Keane, Michael, Robert Moffitt, and David Runkle. "Real Wages over the Business Cycle Estimating the Impact of Heterogeneity with Micro Data." *Journal of Political Economy* 96, no. 6 (December 1988): 1232-66.

Keynes, John Maynard. *The General Theory of Employment, Interest and Money*. New Dehli: Atlantic Publishers and Distributors, 2008/1936.

Kydland, Finn, and Edward C. Prescott. "Time to Build and Aggregate Fluctuations." *Econometrica* 50, no. 6 (November 1982): 1345-70.

Lucas, Robert E. Jr. "Bernanke is the Best Stimulus Right Now." *Wall Street Journal*, December 2008.

Mulligan, Casey B. "A Century of Labor-Leisure Distortions." *NBER working paper*, no. 8774 (February 2002).

Mulligan, Casey B. "Means-tested Subsidies and Economic Performance since 2007." *NBER working paper*, no. 17445 (September 2011).

Mulligan, Casey B. "Public Policies as Specification Errors." *Review of Economic Dynamics* 8, no. 4 (October 2005): 902-926.

Ohanian, Lee E. "The Economic Crisis from a Neoclassical Perspective." *Journal of Economic Perspectives*, Fall 2010: 45-66.

Solon, Gary, Robert Barsky, and Jonathan A. Parker. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *Quarterly Journal of Economics* 109, no. 1 (February 1994): 1-25.

United States Bureau of Economic Analysis. "Components of Value Added by Industry as a Percentage of Value Added." April 26, 2011. http://www.bea.gov/industry/gpotables/gpo_action.cfm?anon=1014136&table_id=27013&format_type=0 (accessed October 24, 2011).

United States Bureau of Labor Statistics. "BLS series id MPU490017." *Major Sector Multifactor Productivity Index.* September 23, 2011. http://www.bls.gov/mfp/ (accessed October 26, 2011).

—. "Characteristics of Minimum Wage Workers: 2006." February 28, 2007. http://www.bls.gov/cps/minwage2006.pdf (accessed October 15, 2011).

—. "Characteristics of Minimum Wage Workers: 2007." March 24, 2008. http://www.bls.gov/cps/minwage2007.pdf (accessed October 15, 2011).

—. "Employment from the BLS household and payroll surveys: summary of recent trends." *Economic News Release: Employment Situation.* October 7, 2011. http://www.bls.gov/web/empsit/ces_cps_trends.pdf (accessed October 25, 2011).

—. "Historical Hours and Earnings, Table B-2." October 2011. ftp://ftp.bls.gov/pub/suppl/empsit.ceseeb2.txt (accessed October 22, 2011).

—. "Industry Employment and Hours." *Labor Productivity and Costs.* October 29, 2010. http://www.bls.gov/lpc/ipr_aphin.xls (accessed October 24, 2011).

United States Census Bureau. "Information Sector Services." *2009 Service Annual Survey Data.* July 12, 2011. http://www2.census.gov/services/sas/data/51/2009_NAICS51.xls (accessed October 24, 2011).

Woodford, Michael. *Interest and prices: foundations of a theory of monetary policy.* Princeton, NJ: Princeton University Press, 2003.