ORIGINAL ARTICLE

Open Access

Uncertainty, redistribution, and the labor market since 2007

Casey B Mulligan

Correspondence:
c-mulligan@uchicago.edu
University of Chicago, Chicago, IL
60637, USA

Abstract

Uncertainty and its composition can affect the demand for social insurance, and thereby the labor market. This paper shows that small to medium-sized increases in uncertainty or risk aversion are enough to recommend an expansion of the safety net that would be broadly similar to the actual safety net expansions in the United States since 2007, which significantly depressed the labor market. Labor market effects of uncertainty through investment and insurance channels are also examined with employer and employee labor wedges measured from 2007 through 2013.

Keywords: Uncertainty; Redistribution; Taxes; Employment

Introduction

During the depths of the Great Depression, Franklin Delano Roosevelt famously said "The only thing we have to fear is fear itself". (White House 2012) Americans and even economists were fearful during the autumn of 2008 when large financial firms went bankrupt and financial markets were unable to operate normally. The federal government has been accused of creating significant uncertainty during the past several years – thereby impeding business planning and investment – by ceasing normal budgeting procedures, creating an extraordinary volume of new regulation for the health and financial sectors, and running up the national debt with no clear plan as to paying it back. All of this is alleged to depress the labor market.[1]

The purpose of this paper is to revive an old theory in which fear or uncertainty affects the efficiency of the labor market by increasing the amount of redistribution, and offer some metrics of the theory's quantitative importance. The redistribution approach stands in contrast to the many other explanations that are demand-driven in the sense that fear or uncertainty discourages business investment, and a lack of business investment is responsible for a lack of job opportunities.[2] The demand-driven explanations fail to generate the empirically observed "labor wedge" – a recession in which consumption and labor quantities both fall while real employer costs do not, which has been the pattern over the past couple of business cycles.[3]

Moreover, even if we assume that the entire investment drop was a consequence of fear or uncertainty, the vast majority of the decline in the quantity of labor remains unexplained unless fear or uncertainty affect the labor market through another mechanism. Real investment per capita in, say, 2011 was $2,000 (2005 prices) less than it was in 2007, before the recession began. $2,000 per capita is only about 4 percent of

Springer

GDP, and a 4 percent reduction in aggregate labor demand is expected to reduce the aggregate quantity of labor by about one percent.[4] But in fact the employment-population ratio was a full 7 percent less in 2011 than it was in 2007.

Using the principal-agent model of the "equity-efficiency tradeoff," this paper suggests that the demand for social insurance is an important, if not the primary, link between uncertainty and the quantity of labor because social insurance can produce a labor wedge on the supply side of the market, rather than primarily shifting the demand for labor as the investment-based explanations do.[5] Before the recession, uncertainty and fear were relatively low, and society had settled on a balance between the amount of redistribution or social insurance and the amount of labor market inefficiency in the forms of a distorted size and composition of the workforce. This pre-recession balance was a function of both the technological tradeoff between equity and efficiency, the degree of risk aversion, and political factors. With the greater amounts of uncertainty and fear that prevailed after 2007, the technological tradeoff was different and/or risk aversion was greater, and society changed the balance between equity and efficiency in the rational direction of less efficiency.

The size of this effect depends on the magnitude of the changes in social insurance program rules, which I have quantified elsewhere in terms of implicit marginal labor income tax rate changes (Mulligan 2012). Statutory marginal labor income tax rates facing the typical household head or spouse at the middle of the skill distribution increased about eight percentage points, and thereby reduced the after-tax return to working by almost fifteen percent, in less than two years. Marginal tax rate changes of this magnitude can easily reduce the quantity of labor by several percentage points, whereas the effects of uncertainty on labor through investment appear to be less.

A model of the equity-efficiency tradeoff

In the principal-agent model of the labor market, *ex ante* identical workers devote time and effort n to production. This includes time spent at work, time spent searching for new work, and effort devoted to enhancing and maintaining the productivity of those uses of time. The value produced by that effort is $y = n + v + \varepsilon$, where v and ε are the result of mean zero idiosyncratic random factors beyond the worker's control. For simplicity, the random factors act additively on the worker's effort, rather than multiplicatively as in Mirrlees (1971). My additive model is just a small adaptation of a special case of Holmstrom and Milgrom (1987) that has been applied to executive salaries by Rosen (1982), Garen (1994), and others.[6]

Some of the random factors, embodied in v, are widely observable. The others, embodied in ε, are not observed, except by the worker himself who can infer their combined value by subtracting his effort and v from the value produced. I normalize ε so that it is uncorrelated with v, and assume that the worker is risk-averse (more on this below). As a result, the worker optimally pools the result v of observed idiosyncratic random factors in a full-information insurance market: when v is observed he pays it to the insurance group, or receives $-v$ from the group in case its value is negative. In theory the insurance group is infinitely large and free from administrative costs so that total insurance premiums (from the group members with positive vs) exactly finance

total insurance awards (paid to group members with negative vs). In practice, the insurance group may be co-workers, family members, church members, etc.

Even with insurance against the observed random factors, the remaining value produced $n + \varepsilon$ is still random. The purpose of this paper is to examine possible ways of insuring the idiosyncratic ε risk, and deriving the effects of the amount and composition of uncertainty on the efficient amount of that insurance.

Suppose that a worker additionally enters into an imperfect information social insurance arrangement in which she pays a fixed fraction of her earnings $n + \varepsilon$ (net of payments with respect to full-information ν insurance) and receives an insurance benefit that is a linear function of net income. Her disposable income c is therefore a linear function of earnings, the intercept of which I denote as b and the slope I denote as μ:

$$c = (n + \varepsilon)\mu + b \tag{1}$$

The social insurance system has a budget constraint that relates benefits it pays to people having zero income to payments it receives from people having positive income:

$$b = \frac{1-\mu}{1+\phi}\bar{n} \tag{2}$$

where $\bar{n}$ is the average effort by members of the insurance system. $\phi \geq 0$ is an administrative or stigma cost reflecting the possibility that social insurance benefits might be worth less than they cost. The average idiosyncratic risk ε is zero but I explain below how the cross-section distribution might be related to aggregate "shocks".

Workers have a smooth utility function u defined over their value of consumption net of the costs of effort, with strictly positive first derivative and strictly negative second derivative in the relevant range. Their expected utility function is therefore:

$$\int u\left(\mu n + \mu \varepsilon + \frac{1-\mu}{1+\phi}\bar{n} - \gamma \frac{\eta}{\eta+1} n^{(\eta+1)/\eta}\right) dG(\varepsilon) \tag{3}$$

where G is the distribution of random factors ε (continuous, with all finite moments), and the constant $\eta > 0$ is the wage elasticity of labor supply. The constant $\gamma > 0$ is a preference parameter, so that the marginal rate of substitution between labor n and disposable income c has no wealth effects on labor. As a result, any worker facing this social insurance program will desire labor in the amount $n = (\mu/\gamma)^\eta$. The desired labor increases with μ because μ is the reward the worker receives for exerting effort; the remaining fraction $(1-\mu)$ of the results of effort goes to the social insurance system.

I assume that actual labor coincides with desired labor, which means that the social insurance system has no influence on the worker's labor decision aside from the parameters μ and b of the safety net benefit formula. As a result, average labor $\bar{n}$ is also equal to $(\mu/\gamma)^\eta$, which is generally different from the socially efficient labor $n^* = (1/\gamma)^\eta$ that equates the marginal rate of substitution to the social product of effort.

Because workers choose their effort before knowing the final outcome ε, the only source of randomness in (3) is the $\mu\varepsilon$ term. A smaller value for μ therefore means a lesser amount of risk faced by the worker, and more equal disposable income for workers who end up with different values of ε, but it also means less effort and thereby less aggregate income. This is known as the "equity-efficiency tradeoff" (Okun 1975), or less often as a "safety-efficiency tradeoff".

The equity-efficiency tradeoff can be analyzed quantitatively by considering a measure of equity or safety S that is negatively related to the standard deviation of disposable income $S = (1 + s_c)^{-1} = (1 + \mu s_\varepsilon)^{-1}$, where s_ε is the standard deviation of the random factor ε. By eliminating μ from the safety equation and the average labor equation $\bar{n} = \mu^\eta n^*$, we have a single equation for the tradeoff:

$$S = \frac{1}{1 + s_\varepsilon \left(\frac{\bar{n}}{n^*}\right)^{1/\eta}} \tag{4}$$

More safety S achieved with more social insurance (1-μ) means less labor efficiency $\bar{n}/n^*$, and vice versa.[7] Figure 1 displays the tradeoff in a simple diagram, with safety measured on the vertical axis and efficiency on the horizontal axis. Points to the northwest on the frontier shown in Figure 1 correspond to more social insurance—that is, low values for the self-reliance rate μ—and points to the southeast correspond to less social insurance. The solid curve is the equity-efficiency frontier, because combinations of equity and labor efficiency beyond it are not possible since equalizing outcomes reduces incentives to supply labor.

Possible changes in the equity-efficiency tradeoff, and the optimal degree of social insurance

Both safety and labor efficiency are desirable, and overall efficiency strikes a balance between the two. The optimal amount of social insurance can be described as the value of μ—a point on the equity-efficiency tradeoff shown in Figure 1—that maximizes worker expected utility.[8]

$$\int u\left(\mu\varepsilon + \frac{1 + \phi\mu}{1 + \phi}(\mu/\gamma)^\eta - \gamma\frac{\eta}{\eta + 1}(\mu/\gamma)^{\eta+1}\right)dG(\varepsilon) \tag{5}$$

It is straightforward to prove that the optimal amount of social insurance depends on the amount of risk the worker faces, as embodied in the distribution function G.

Figure 1 The Equity-Efficiency Frontier.
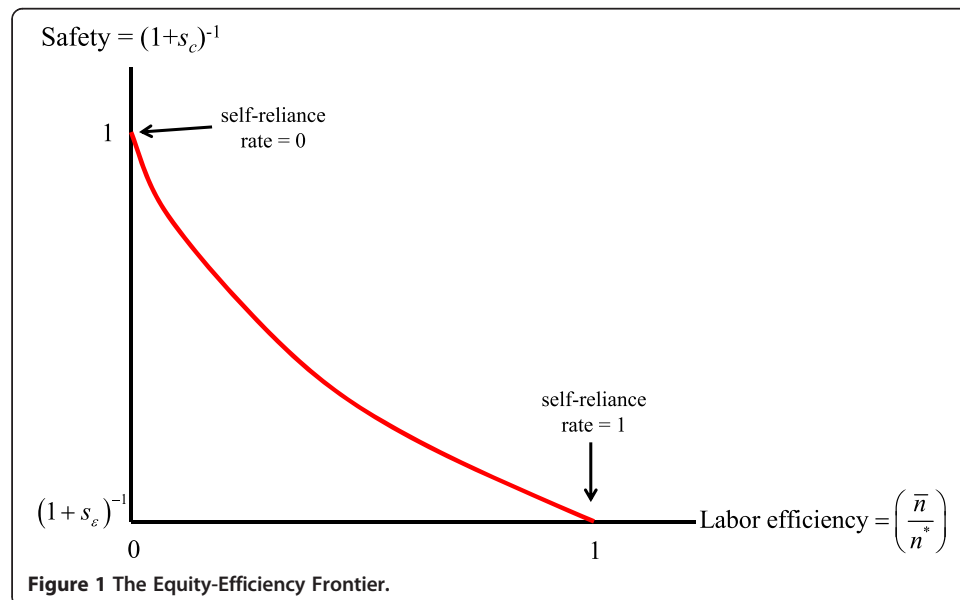
Proposition 1 The optimal amount of social insurance (1-μ) increases with the standard deviation s_ε of the random factor ε, holding constant higher-order moments of the distribution function G.

Proof Because u is smooth, I prove the proposition by using its Taylor expansion (in the neighborhood of ε = 0) in the worker expected utility expression:

$$u\left(\frac{1+\phi\mu}{1+\phi}(\mu/\gamma)^\eta - \gamma\frac{\eta}{\eta+1}(\mu/\gamma)^{\eta+1}\right) + \int \sum_{k=2}^\infty \frac{u^{(k)}}{k!}(\mu\varepsilon)^k dG(\varepsilon)$$
$$= u\left(\frac{1+\phi\mu}{1+\phi}(\mu/\gamma)^n - \gamma\frac{\eta}{\eta+1}(\mu/\gamma)^{\eta+1}\right) + \sum_{k=2}^\infty \frac{u^{(k)}}{k!}\mu^k G_k$$

where $u^{(k)}$ denotes the kth derivative of u evaluated at ε = 0 and G_k denotes the kth moment of the distribution function G. The first derivative of expected utility with respect to μ is the sum of infinitely many terms, but only one of them involves the second moment of the distribution function G (G_2). That term is linear in G_2 with a negative coefficient because $u^{(2)}$ is negative, which means that the marginal benefit of μ falls with G_2. Thus the optimal μ falls with G_2, and the optimal (1-μ) increases.

There are a couple of reasons to think that the standard deviation s_ε of the random factor ε was greater after 2008 than it was before. One is that capital market and other aggregate events made it more difficult for market participants to distinguish bad outcomes that should be blamed on low effort from bad outcomes that were just unlucky. In this view, an aggregate shock does not necessarily change the total amount of idiosyncratic risk – the variation of labor productivity y around labor time and effort supplied – but rather changes its composition between the unobserved category ε and the observed category v. When the change is in the direction of more unobserved idiosyncratic risk, workers replace part of the full-information insurance lost by adding to
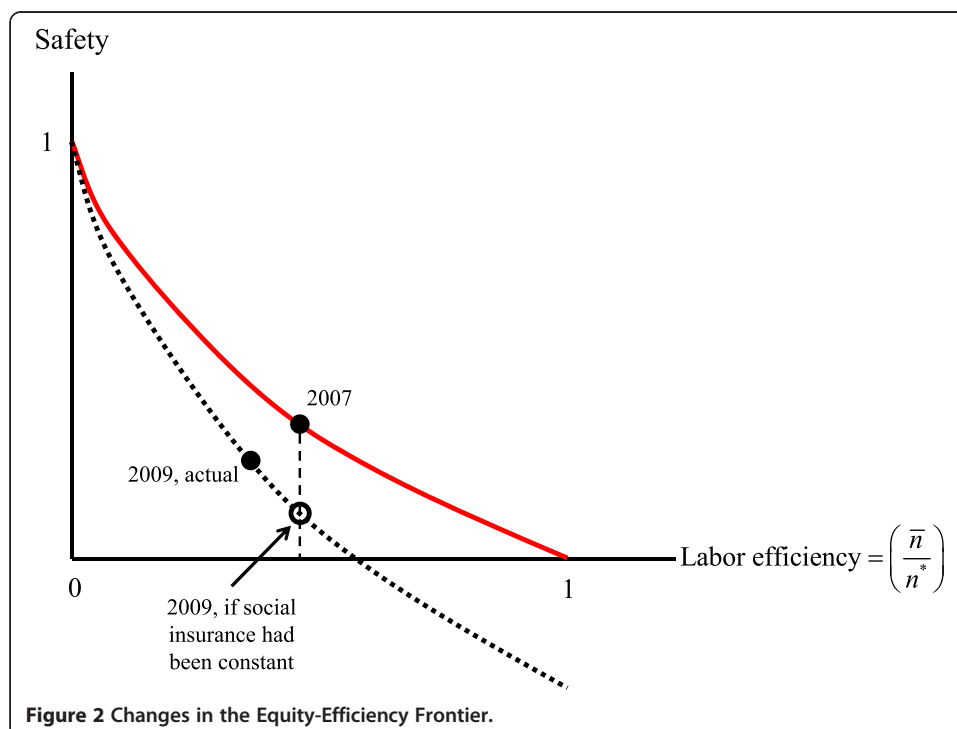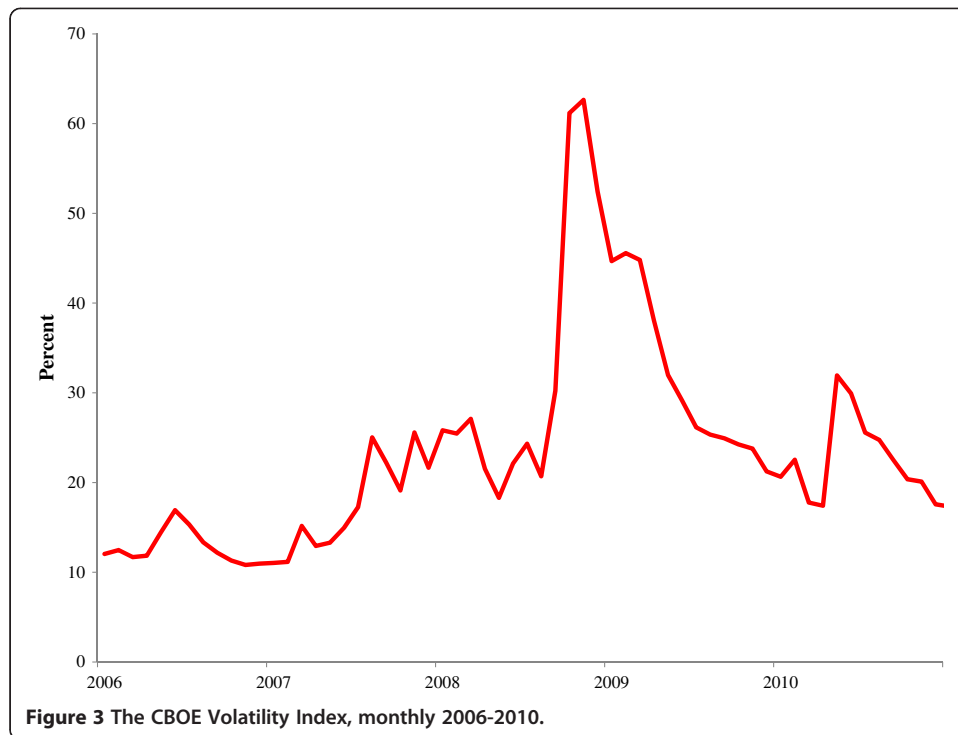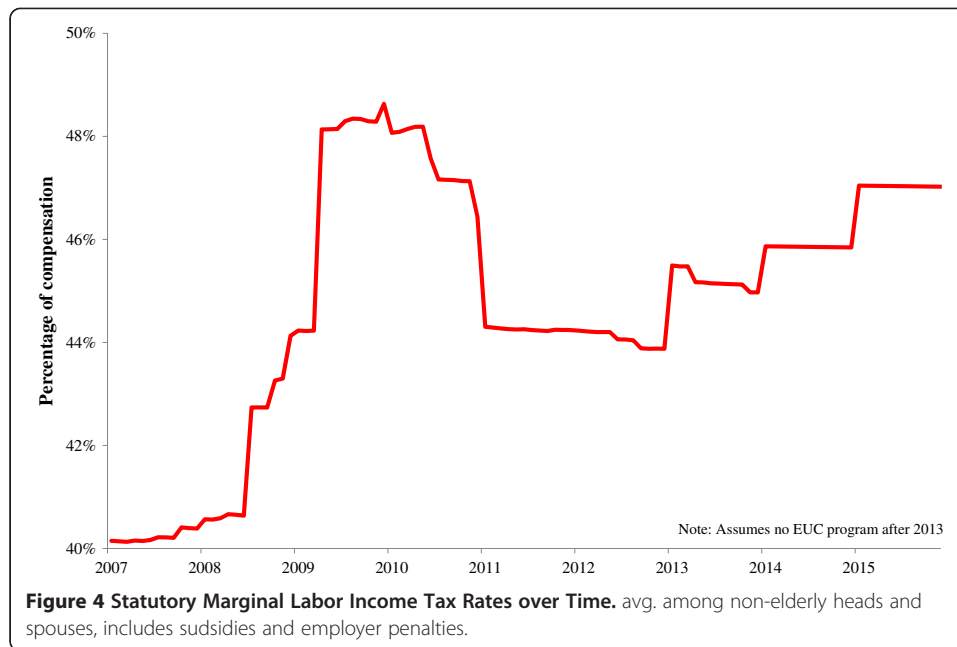

Figure 2 Changes in the Equity-Efficiency Frontier.

Figure 3 The CBOE Volatility Index, monthly 2006-2010.

their imperfect-information insurance, despite the latter's cost in terms of labor ineffi-ciency. Figure 2 illustrates the change that might have occurred between 2007 and 2009. As workers began to face more risk—the equity-efficiency frontier shifted down from solid to dotted—the safety net could have remained as it was in 2007, in which case labor market efficiency might have remained constant too. But then workers would have substantially less safety, as with the hollow circle in Figure 2. In that unfor-tunate situation, workers may prefer to recover some of the lost safety by reducing labor efficiency, as with the solid circle on the dotted 2009 equity-efficiency frontier.[9]

A related argument is that the variance of ε increased without any reduction in the variance of v. Still, workers desire to have more imperfect-information insurance, des-pite its cost in terms of labor inefficiency. Either way, less labor is the result of the safety net expansion in the sense that keeping the safety net constant would have re-sulted in a time path for labor that was above its actual path. But the safety net expan-sions were themselves a response to another shock—increased uncertainty—and keeping the safety net constant might not have been the best response to that add-itional uncertainty. A constant safety net would have lessened the recession as mea-sured by the amount of the drop in aggregate labor market activity, but the cost of doing so would have been too high in terms of the risks that workers would have borne without government help. In this sense, one might say the change in the equity-efficiency tradeoff mediated through an expanded safety net was a major factor con-tracting the labor market after 2007.

A greater degree of risk aversion induces an optimal movement along Figure 1's equity-efficiency frontier in the direction of more social insurance. A smaller value for the stigma parameter ϕ can also increase the optimal amount of social insurance, although ϕ has a wealth effect in the other direction.

Figure 4 Statutory Marginal Labor Income Tax Rates over Time. avg. among non-elderly heads and spouses, includes sudsidies and employer penalties.

Quantitative estimates of risk and marginal tax rate changes since 2007

Economists have noted that uncertainty has increased a lot since 2007, and they have offered measures of the amount of added uncertainty (Bloom et al. 2011).[10] One commonly cited measure is the Chicago Board Options Exchange's Volatility Index, sometimes called the fear index or VIX, which is based on market forecasts of the volatility of an index of stock prices. The monthly version of the series is plotted in Figure 3, and has percentage units of thirty-day annualized rates of return (Chicago Board Options Exchange 2009). The index increased by a factor of two from the first half of 2007 to the second half, and then by at least another factor of two by the end of 2008. The index returned to its late-2007 level in early 2010 and again in late 2010.

Unlike the VIX index, the safety net did not quickly return to what it was before the recession. One possibility is that the transitory spike found with the VIX would not be found with measures of idiosyncratic risk like those featured in my model.[11] Equity options prices can also be used to measure firm-idiosyncratic risk (Stein and Stone 2012). Wage or earnings inequality can be indicators of household-idiosyncratic risk, and have increased significantly since 2007 (see Appendix 2), after increasing secularly for decades before that. Health spending has increased secularly as well, and it is sometimes argued that the added spending may be subjecting households to additional idiosyncratic financial risk (Seifert and Rukavina 2006), even though part of the rising spending is a symptom of welcome medical progress.

Even if the amount and composition of uncertainty remained constant, the degree of risk aversion could vary over time (perhaps in conjunction with aggregate risk indicators like the VIX). Perhaps changing risk aversion is a way to interpret President Roosevelt, but in any case rising risk aversion – as distinct from any changes in objective measures of risk – is commonly cited by finance economists as a reason for asset prices to fall (Campbell and Cochrane 1999). In my model, risk aversion is represented by the degree of concavity of the utility function u, with the amount of concavity increasing

the demand for social insurance. In this view, the economic expansion prior to 2007 might have been a time of low risk aversion in which people were especially willing to purchase risky assets and especially eager to forgo social insurance in order to have a more efficient labor market.

Figure 4 displays Mulligan's (2013b) measure of the statutory marginal labor income tax rate for household heads and spouses in the middle of the skill distribution, including the marginal tax rates implicit in the many federal and state anti-poverty programs. Its changes prior to 2012 derive primarily from changes in eligibility and benefit rules for the unemployment insurance program and related subsidies, and for the food stamp program. A jump at the start of 2013 reflects the expiration of the payroll tax cut and a jump at the beginning of 2014 reflects the expiration of Emergency Unemployment Compensation (EUC) and a collection of various implicit taxes (permanently) levied on employees by the Affordable Care Act.[12] The level of the marginal tax rate series (but not its changes) reflects actual rates of program participation, which is why the rates shown in Figure 4 are less than reported by studies that examine the rates faced by persons who participate in all programs for which they are eligible (Romich, Simmelink and Holt 2007).[13]

The marginal labor income tax rate began at about 40 percent and peaked in late 2009 at more than 48 percent. The marginal labor income tax rate is expected to reach almost that high in 2015, and remain elevated. In terms of the log self-reliance rate, that's a decrease of about 0.12 since 2007.

Holmstrom and Milgrom (1987, p. 323) give an example, consistent with the model above, with constant absolute risk aversion in which the optimal self-reliance rate (μ in my notation, α in theirs) has a closed form solution as a function of the standard deviation of the income risk being insured (s_ε in my notation), the coefficient of absolute risk aversion r, and a parameter governing the disutility of effort. Holding the effort-disutility parameter constant, the comparative statics for their model's optimal self-reliance rate are (my notation):

$$d \ln \mu = -(1-\mu)2(d \ln s_\varepsilon)-(1-\mu)(d \ln r) \qquad (6)$$

Appendix 1 contains the derivation of equation (6), which can be used to understand the size of the uncertainty or risk aversion impulse that would be needed to rationalize a safety net expansion like the one experienced since 2007.

Assuming for the moment that the optimal self-reliance rate μ was also about 0.56 (the average of late 2007 and late 2009), then the comparative static elasticities in equation (6) are about −0.88 and −0.44, respectively. Thus, an increase in the log of the standard deviation s_ε of about 0.14 would justify a reduction in the log self-reliance rate of 0.12. The same self-reliance rate reduction could also be justified by a 0.28 increase in the log of the coefficient of absolute risk aversion, or some combination of added variance and added risk aversion. Moreover, microeconomic wage data suggest that wage variability was high during the recession, with a log standard deviation change of about 0.07 (see Appendix 2). Thus, the Holmstrom and Milgrom (1987) approach does not require especially large increases in uncertainty or risk aversion in order to recommend an expansion of the safety net that would be about half of the actual safety net expansions, and more in combination.[14]

Regardless of whether it was optimal to reduce the log self-reliance rate by 0.12, it happened and that reduction is expected to significantly depress the labor market. Mulligan (2012) finds that marginal tax rate changes of the amounts shown in Figure 4, with a 0.4 wage elasticity of *aggregate* labor supply that is less than most reported in the literature (Chetty, et al. 2011), explain at least half of the actual reduction in per capita work hours between the end of 2007 and the end of 2009. With a larger (and thereby more reasonable) wage elasticity of aggregate labor supply, even more of the hours decline can be explained.

Wages and wedges in the labor market

Labor wedges are another way to categorize various effects of uncertainty and other factors on labor market outcomes, and are especially helpful for detecting and evaluating theories of labor market fluctuations that emphasize time-varying uncertainty, risk aversion, or taxation. Labor wedges are constructed solely from measures of market outcomes – productivity, wages, consumption, work hours, and related variables – in order to represent changes in the gap between the marginal productivity of labor and households' marginal rate of substitution (*MRS*, sometimes described as the reservation wage of the marginal worker). The overall wedge can be decomposed into the sum of two wedges: an "employer wedge" or productivity-wage wedge that measures the gap between labor productivity and real wages and an "employee wedge" that measures the gap between real wages and *MRS*. Equation (7) displays these relationships in log-change form:

$$d\ln(\text{labor productivity}) = d\ln(\text{employer wedge}) + d\ln(\text{real wage})$$
$$d\ln(\text{real wage}) = d\ln(\text{employee wedge}) + d\ln MRS \tag{7}$$

where the d operator denotes changes over time. Parkin (1988), Hall (1997), and Chari et al. (2007) interpret labor wedges as an indicator of economists' misspecifications of labor market phenomena (that is, important forces left out of the market-clearing model), but Mulligan (2002), Mulligan (2005) and Galí et al. (2007) argue that labor wedges are market participants' revealing – through their behavior – the taxes and other market distortions that they face.
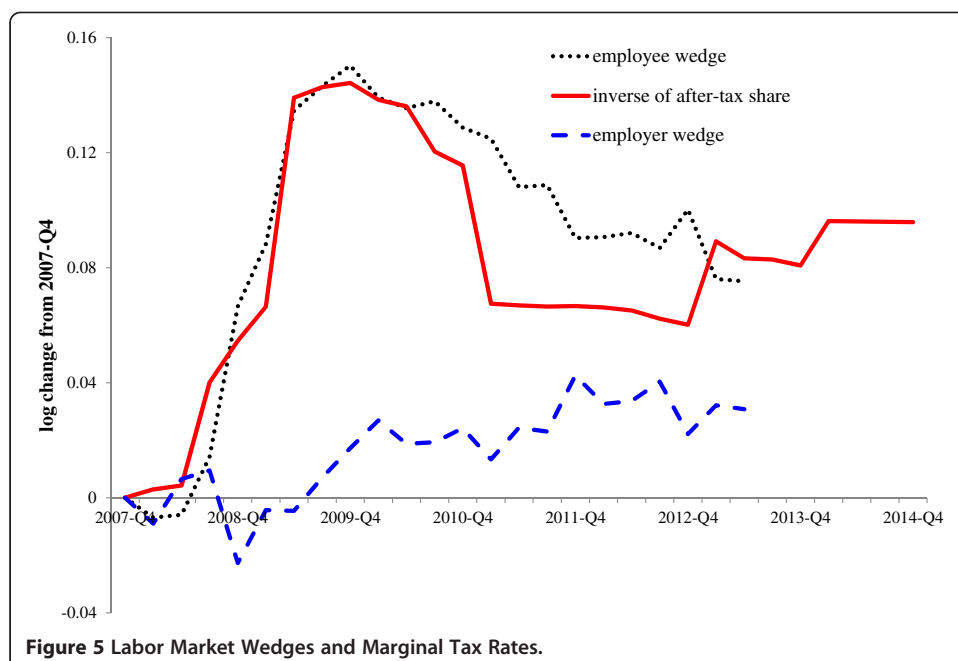
Theories of uncertainty and labor fluctuations have different predictions for the labor wedge. In Bloom, et al. (2011, see especially their Equation 12), uncertainty does not create any employee wedge. Uncertainty reduces labor in their model by reducing real wages (both wages paid by employers and wages received by employees) and moving households along their labor supply curves. Uncertainty does create an employer wedge in their model, because hiring labor up to the point where wage equals marginal product ignores the option value of delaying hires in the face of uncertainty. Uncertainty also creates an employer wedge in Fernandez-Villaverde, et al. (2012)'s model: by increasing the degree to which employers mark up their product prices over marginal cost. In this way, both models predict that labor productivity grows more than compensation per hour whenever the degree of uncertainty increases.
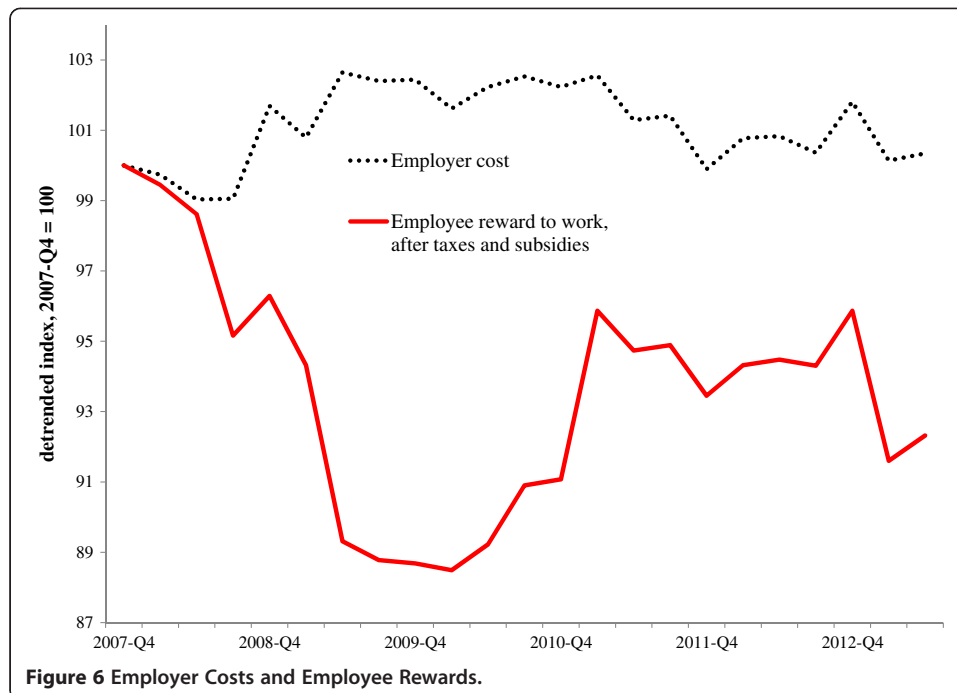
Uncertainty and risk aversion create an employee wedge in my model to the extent that they motivate redistribution and at least part of the associated marginal tax rates fall on employees. If employees perceive a new cost of employment, such as a new payroll tax on employees, added child-care costs, or new subsidies for the unemployed, then the

households' *MRS* falls more, or rises less, than wages do. In theory, the percentage increase in the employee wedge is equal to the percentage reduction in the post-fisc share of wages accruing to employees as a consequence of their work decision.

Uncertainty and risk aversion create an employer wedge in my model to the extent that they motivate redistribution and at least part of the associated marginal tax rates fall on employers. In theory, an employer wedge is created when employers perceive a new cost of employment in addition to employee compensation, such as the fines levied by the Affordable Care Act.[15] These taxes create a productivity-wage wedge of the same percentage as the tax itself. Anticipated severance costs, and other costs of downsizing, create the same sort of wedge between contemporaneous wages and productivity because they are employment costs over and above the contemporaneous wage payments, and this employer wedge would grow as uncertainty increased the probability of downsizing events.

Figure 5 displays the quarterly employer (dashed) and employee (dotted) wedges from 2007 through mid-2013, applying the formula (7) measuring productivity as real GDP per manhour, the real wage as real employee compensation per manhour, and the log *MRS* change as the log change in real consumption per capita plus the log change in hours worked per adult.[16] The employer wedge is relatively constant through mid-2009, after which time its log increases almost 0.05 from early 2009 to the end of 2011 when it is 0.04 greater than it was when the recession began. The most recent data (the first half of 2013) suggest that the employer wedge remains more than 0.03 above its pre-recession value. The change in the employer wedge over a short time frame is significant by historical standards, and roughly of the magnitude of a three percentage point increase in the employer payroll tax. It is consistent with the hypothesis that *something* reduced the demand for labor (at a given level of real compensation per hour) after mid-2009. That "something" could have been the effects of uncertainty

Figure 5 Labor Market Wedges and Marginal Tax Rates.

Figure 6 Employer Costs and Employee Rewards.

working through labor demand channels, as in the Bloom, et al. (2011) and Fernandez-Villaverde, et al. (2012) models.[17]

However, the change in the employer wedge is small in comparison to changes in the employee wedge, which looks more like a change in the employee payroll tax exceeding ten percentage points. The fact that most of the labor wedge during this period appears on the employee side is one of the reasons why it is worth considering the idea that uncertainty affects the labor market through labor supply a lot more than it works through labor demand.

Note that both wedges have log changes that are in the same units, which are essentially units of a tax rate. For example, to a first order approximation a two percent payroll tax would increase the log employer wedge by 0.02 if levied on employer (and not counted in employee compensation for the purposes of measuring the real wage) and increase the log employee wedge by 0.02 if levied on employee. That does not mean that the wedges *are* taxes: the wedges are constructed from behavior, not tax information. But if Mulligan (2002), Mulligan (2005) and Galí et al. (2007) are right, the wedges are market participants' revealing, through behavior, the direction and magnitude of the taxes that they face. Their theory can be tested by comparing measured behavior as summarized by the employee wedge to a measure of marginal tax rates falling on employees, such as the solid series in the Figure.[19] In theory, the solid and dotted series should coincide, at least if the marginal tax rate and other series were not measured with error and the proper marginal rate of substitution function were used.[18] Although the two series are constructed from completely different data sources, they coincide closely through 2010 both in terms of direction and magnitude. Both the after-tax share and the employee wedge fall during 2011, but the wedge falls less. The two series had come back together by early 2013. The congruence of the solid and dotted series are a second reason why it is worth considering the idea that uncertainty affects the labor market through redistribution and labor supply at least as much as it works through labor demand.

These findings that the employer wedge changes less than the employee wedge, whose time series is well approximated by the series for employee marginal tax rates, are reflected in the wage data. Figure 6 displays indices of real wages relative to a 0.5 percent per year trend, measured on both sides of the labor market. The solid series is the index of employees' financial reward to work after taxes and subsidies. It falls sharply and remains depressed.

In contrast, the dotted employer cost series rises somewhat during the first two years. It is stable for the next two years, and falls slightly during 2011. By 2013, the employer cost index is roughly the same as it was before the recession began. Thus, on the scale of changes in the reward to work, employer costs hardly fell, if at all, which contradicts the hypothesis that the labor market was primarily depressed by heightened uncertainty working through demand-depressing channels.

Conclusions

Safety net programs face a well-known equity-efficiency tradeoff: providing more resources for the poor can raise their living standards, but it also gives them less incentive to raise their own living standards. Most societies somehow balance the tradeoff by permitting outcomes to vary across people—less-than-perfect equity—yet still having some amount of redistribution.

The recent housing crash, financial crisis, or even political events may have altered the nature of this tradeoff or altered society's willingness to tolerate labor market inefficiency in order to have more equity. The tradeoff would be altered by a deterioration of information about labor market outcomes that would normally help distinguish luck from the consequences of worker effort. It could also be altered by an increase in non-diversifiable risk. An increase in risk aversion, or possibly a reduction in welfare stigma, induces an optimal movement along that tradeoff in the direction of more social insurance. Under any of these scenarios, it would be no surprise that the social safety net expanded and labor market activity fell as much as they did, and uncertainty in one way or another would be the ultimate cause.

A previous literature has emphasized effects of uncertainty on the labor market through labor demand, in the sense that an employer's desired hiring depends not only on what employees will be paid in the short term but also on the degree of uncertainty about their future pay or other future costs of doing business. To the degree that wages adjust over time to reflect supply and demand, more uncertainty should in this view cause real wages to grow less than productivity, which did not happen in the two or three years prior to mid-2009, but did happen after that. The growth of the employer wedge after mid-2009 amounted to about a 3–5 percent increase in the marginal payroll tax rate on employers, and is large by historical standards.

However, much of what happened in the labor market happened on the employee side. The marginal tax rate on employees, most of it implicit in safety net program spending rules, increased three or four times more than the employer wedge did. To the degree that fear and uncertainty about the economy motivated the new safety net program rules, they affected the labor market more through labor supply than they did through labor demand.

This paper examines the demand for social insurance in the context of an optimal social insurance model. I do not interpret the model results as proving that the actual safety net expansions of recent years were optimal. Indeed, the actual expansions have several features that depart from economically efficient program design, such as subjecting parts of the population to 100 percent tax rates (Mulligan 2013a), subjecting households to additional policy randomness (Baker et al. 2011), and allowing households to participate in multiple uncoordinated insurance schemes (Mulligan 2012). Future research should look at these weaknesses of the actual safety net expansions, and reach conclusions about possible policy improvements. Nevertheless, the model shows why an event that increases uncertainty may lead through the political process to an expanded safety net, despite the fact that safety net expansions contract the labor market.

Endnotes

[1] 60 leading economists were surveyed by *Foreign Policy* in 2011 and asked to "name the single most important reason U.S. jobs have not returned more quickly", and the number one response was "uncertainty" (Foreign Policy 2012). See also Bloom et al. (2013), Becker et al. (2010), Galston (2013), Kasperowicz (2011).

[2] See for example, Bloom et al. (2011) or Stokey (2013). The aforementioned survey of economists also cites "uncertainty", "inadequate demand" and "investment drought" as possible number-one reasons for the depression in the labor market.

[3] See especially Bloom et al.'s (2011) Equation 12. The body of this paper further discusses the labor wedge. More work is also needed to examine whether uncertainty *increases* investment, and labor more directly, through precautionary savings motives (Carroll 1997).

[4] Here I assume a wage elasticity of aggregate labor demand of −3.3 (as it would be with an aggregate Cobb-Douglas production function with labor share 0.7) and a wage elasticity of aggregate labor supply of one. The labor quantity impact would be even less if the wage elasticity of labor supply were less than one, or investment industries were more capital intensive than the rest of the economy.

[5] In principle, redistribution can also create a labor wedge on the demand side of the market by creating additional employer costs associated with employment (more on this below). Another difference between my approach and the investment-driven models is that the former links uncertainty and labor with public policy whereas the latter links them with employer willingness to invest.

[6] See also Mulligan (2012) and Bertola (2013), who each use a special case of Holmstrom and Milgrom (1987) to qualitatively model changes in labor market policies after "crises".

[7] Both safety and efficiency measures vary between zero and 1. Safety is $(1 + s_e)^{-1}$ and efficiency is 1 when the self-reliance rate is 1 (no sharing of the imperfect information risks). Safety is 1 and efficiency is zero when all of the risks are shared.

[8] Equation (5) is derived from equation (3) by setting individual labor at its privately optimal amount, and imposing the aggregate consistency condition $\bar{n} = n$.

[9] A related argument is that the equity-efficiency tradeoff is fixed, but the optimal balance of equity and efficiency has changed since 2007 because people became more risk-averse, at least temporarily. See Holmstrom and Milgrom (1987, 323) for a risk-aversion comparative static. Also note that the Figure 2's horizontal axis does not cross

the vertical axis at zero, but rather at the level of safety that would have prevailed in 2007 without redistribution.

[10]As noted below, economists have not suggested of late that the labor market effects of uncertainty would work through marginal tax rates, but rather that uncertainty discourages firms from undertaking new investment projects or otherwise expanding their payrolls.

[11]The VIX could be an indicator of idiosyncratic risk if some of the volatility of average stock prices derives from redistribution shocks, but otherwise might be better interpreted as an aggregate risk indicator. Another possibility is that safety net expansions create a political constituency that allows them to outlast their original economic function.

[12]Beginning in 2015, the Affordable Care Act also levies penalties on *employers*, and these penalties are reflected in Figure 4's tax rate series. After Figure 4 was prepared for publication, some of the 2015 employer penalties were delayed until 2016.

[13]Formally, my index is a fixed-weighted index akin to the Marshall-Edgeworth index, with the weights determined by averaging participation rates in 2007 with rates in 2010. That one set of weights is used for every month's index.

[14]This approach is arguably too powerful because it predicts significant increases in the amount of redistribution when they do not occur. See, for example, the Meltzer and Richard (1981) model of inequality and redistribution (which says that inequality should lead to more redistribution for many of the reasons that uncertainty leads to more redistribution in my model) and Peltzman (1980) and Perotti (1996) on Meltzer-Richard's incongruence with the time series and cross-country evidence.

[15]As of the time of writing, I do not know whether national accountants will classify the new employer penalties as business taxes (like they do for corporate tax payments) or employee compensation (like they do for employer payroll taxes).

[16]This marginal rate of substitution function corresponds to a representative agent's utility function of the form $u(c, n, P, A) = P \ln(c/P) - \gamma \frac{\eta}{\eta+1} A \left(\frac{n}{A}\right)^{\eta/(\eta+1)}$, where c is aggregate real consumption, n is aggregate hours worked, P is the total population, and A is the age-adjusted adult population (a slight age adjustment reflects the increased propensity during these years of the adult population to be near or beyond retirement ages; see Mulligan 2012 for details). The constant η is usually interpreted as the Frisch elasticity of aggregate labor supply with respect to after-tax real wages, and is taken to be one for the purposes of constructing Figure 5's employee wedge. The employee wedge is similar with η equal to say, 0.75, but would increase significantly more if η were taken to be less than 0.5.

[17]A growing employer wedge could also reflect increases in the amount of social insurance via employer-related instruments, such as the employer fines used in the new Affordable Care Act to finance part of its health insurance subsidies.

[18]Specifically, Figure 5's solid series is minus one times the log change in the after-tax share formed from Figure 4's marginal tax rate series between 2007 and 2014.

[19]Among other things, for the purpose of comparing the inverse of the after-tax share with the employee wedge, the ACA's employer penalty (beginning in 2015) should be treated consistently between the two series, as I have done in Figure 5 because that figure does not show a tax measured beyond 2014-Q4.

Appendix 1: ARA closed-form solution

Following Holmstrom and Milgrom (1987, 323), assume a normal distribution for ε and let the utility function be of the ARA form $u(x) = -e^{-rx}$, let the wage elasticity of effort supply be one, and let the stigma cost ϕ be zero. Equation (5) becomes

$$\int u\left(\mu\varepsilon + \frac{1+\phi\mu}{1+\phi}(\mu/\gamma)^{\eta} - \gamma\frac{\eta}{\eta+1}(\mu/\gamma)^{\eta+1}\right)dG(\varepsilon) = -e^{r/\gamma\left[\mu^2\gamma rs_\varepsilon^2/2 - \mu + \mu^2/2\right]} \tag{8}$$

where s_ε is the standard deviation of the idiosyncratic private information risk. Maximizing with respect to the self-reliance rate μ,

$$\mu = \left(1 + rs_\varepsilon^2\gamma\right)^{-1} \tag{9}$$

Holding constant the effort disutility parameter γ, the comparative statics are the same as shown in the main text's equation (6).

Appendix 2: idiosyncratic risk as residual wage inequality

Two key parameters in the model are changes over time in the risk aversion coefficient and in the amount of idiosyncratic private information risk. Both are difficult to measure directly. Outcome variability can be measured, but it's difficult to know how much outcome variability is the consequence of effort, or the consequence of anticipated factors, or the consequence of shocks whose origins are publicly known and thereby insured without moral hazard. Advancing risk measurement is beyond the scope of this paper, but this appendix offers a crude measure that might indicate the general direction of idiosyncratic private information risk changes since 2007 and perhaps roughly indicate their magnitude.

One measure is a residual from an hourly log wage regression on indicator variables for white; calendar year; month of year; the interaction of an age quartic with educational attainment and sex; the interaction of educational attainment, sex, and presence of children under eighteen; and all interacted with married (spouse present). The sample was full-time employed persons aged 26–64 in the monthly CPS Merged Outgoing Rotation Groups for 2000 through 2010. The standard deviation of these residuals averaged 0.512 for 2000–2006, as compared to 0.551 for 2009–10: a log change of 0.073.

References
Baker SR, Bloom N, Davis SJ (2011) Measuring Economic Policy Uncertainty. University of Chicago Booth School of Business (unpublished manuscript)
Becker GS, Davis SJ, Murphy KM (2010) Uncertainty and the Slow Recovery. Wall Street J: on-line edition
Bertola G (2013) Labor Market Policies and European Crises, EDHEC Business School
Bloom N, Floetotto M, Jaimovich N, Saporta-Eksten I, Terry S (2011) Really Uncertain Business Cycles. Stanford University (unpublished manuscript)
Bloom N, Kose MA, Terrones ME (2013) Held back by uncertainty. Finance Dev 50(no. 1):38–41
Campbell JY, Cochrane JH (1999) By force of habit: a consumption-based explanation of aggregate stock market behavior. J Polit Econ 107(2):205–251

Carroll CD (1997) Buffer-stock saving and the life cycle/permanent income hypothesis. Q J Econ 112(no. 1):1–55

Chari VV, Kehoe PJ, McGrattan ER (2007) Business Cycle Accounting. Econometrica 75(3):781–836

Chetty R, Guren A, Manoli D, Weber A (2011) Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. Am Econ Rev 101(2):1–6

Chicago Board Options Exchange (2009) The CBOE Volatility Index - VIX., http://www.cboe.com/micro/vix/vixwhite.pdf (accessed January 6, 2012)

Fernandez-Villaverde J, Guerron-Quintana PA, Kuester K, Rubio-Ramirez J (2012) Fiscal Volatility Shocks and Economic Activity. UPenn (unpublished manuscript)

Foreign Policy (2012) The FP Survey: The Economy., http://www.foreignpolicy.com/articles/2012/10/08/ the_fp_survey_the_economy (accessed January 30, 2014)

Galí J, Gertler M, Lopez-Salido JD (2007) Markups, gaps, and the welfare costs of business fluctuations. Rev Econ Stat 89:44–59

Galston W (2013) Policy uncertainty paralyzes the economy. Wall Street J: , on-line edition

Garen JE (1994) Executive compensation and principal-agent theory. J Polit Econ 102(no. 6):1175–1199

Hall RE (1997) Macroeconomic fluctuations and the allocation of time. J Labor Econ 15(no. 1, Part 2):S223–S250

Holmstrom B, Milgrom P (1987) Aggregation and linearity in the provision of intertemporal incentives. Econometrica 55(2):303–328

Kasperowicz P (2011) Senate GOP seeks regulatory 'time out., http://thehill.com/blogs/floor-action/senate/181127-senate-gop-seeks-regulatory-time-out (accessed February 1, 2014)

Meltzer AH, Richard SF (1981) A rational theory of the size of government. J Polit Econ 89(5):914–927

Mirrlees JA (1971) An exploration in the theory of optimum income taxation. Rev Econ Stud 38(no. 114):175–208

Mulligan CB (2002) A Century of Labor-Leisure Distortions, vol 8774, NBER working paper

Mulligan CB (2005) Public policies as specification errors. Rev Econ Dyn 8(no. 4):902–926

Mulligan CB (2012) The Redistribution Recession. Oxford University Press, New York, redistributionrecession.com

Mulligan CB (2013a) Average Marginal Tax Rates under the Affordable Care Act, vol 19365, NBER working paper

Mulligan CB (2013b) The ARRA: Some Unpleasant Welfare Arithmetic, vol 18591, NBER working paper

Okun AM (1975) Equality and Efficiency: The Big Tradeoff. Brookings Institution Press, Washington, DC

Parkin M (1988) A method for determining whether parameters in aggregative models are structural. Carn-Roch Conf Ser Public Policy 29(1):215–252

Peltzman S (1980) The growth of government. J Law Econ 23(2):209–287

Perotti R (1996) Growth, income distribution, and democracy: what the data say. J Econ Growth 1(2):149–187

Romich JL, Simmelink J, Holt S (2007) When working harder does not pay: low-income working families, tax liabilities, and benefit reductions. Fam Soc-J Contemp Soc Serv 88(3):418–426

Rosen S (1982) Authority, control, and the distribution of earnings. Bell J Econ 13:311–323

Seifert RW, Rukavina M (2006) Bankruptcy is the tip of a medical-debt iceberg. Health Aff 25(2):w89–w92

Stein LCD, Stone EC (2012) The Effect of Uncertainty on Investment, Hiring, and R&D: Causal Evidence from Equity Options. Arizona State University (unpublished manuscript)

Stokey NL (2013) Wait-and-See: Investment Options under Policy Uncertainty, vol 19630, NBER working paper

White House (2012) Franklin D. Roosevelt., http://www.whitehouse.gov/about/presidents/franklindroosevelt (accessed September 13, 2012)