

NBER WORKING PAPER SERIES

A BOUND ON RISK AVERSION USING LABOR SUPPLY ELASTICITIES

Raj Chetty

Working Paper 12067

<http://www.nber.org/papers/w12067>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

February 2006

Department of Economics, University of California, Berkeley and NBER, 549 Evans Hall #3880, Berkeley CA 94709 USA (email: chetty@econ.berkeley.edu). This paper is based on the second chapter of my Ph.D. thesis at Harvard. I thank George Akerlof, John Campbell, David Card, Gary Chamberlain, David Cutler, Martin Feldstein, John Friedman, Ed Glaeser, Caroline Hoxby, Louis Kaplow, Larry Katz, Miles Kimball, Greg Mankiw, Richard Rogerson, anonymous referees, and numerous seminar participants for very helpful comments and discussions. Financial support from the National Science Foundation Graduate Fellowship, the National Bureau of Economic Research, and Harvard University is gratefully acknowledged. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Raj Chetty. 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.

A Bound on Risk Aversion Using Labor Supply Elasticities
Raj Chetty
NBER Working Paper No. 12066
February 2006
JEL No. D80, J20, J60

ABSTRACT

This paper shows that existing evidence on labor supply behavior places an upper bound on risk aversion in the expected utility model. I derive a formula for the coefficient of relative risk aversion (g) in terms of (1) the ratio of the income elasticity of labor supply to the wage elasticity and (2) the degree of complementarity between consumption and labor. I bound the degree of complementarity using data on consumption choices when labor supply varies randomly across states. Using labor supply elasticity estimates from thirty-three studies, I find a mean estimate of $g = 1$. I then show that generating $g > 2$ would require that wage increases cause sharper reductions in labor supply than estimated in any of the studies.

Raj Chetty
Department of Economics
UC- Berkeley
521 Evans Hall #3880
Berkeley, CA 94720
and NBER
chetty@econ.berkeley.edu

Expected utility is the canonical theory of choice under uncertainty in economics. In the expected utility model, risk aversion arises solely from the curvature of the utility function, typically measured by the coefficient of relative risk aversion (γ). This paper shows that evidence on the effects of wage changes on labor supply imposes a tight upper bound on the curvature of utility over wealth ($\gamma < 2$). Hence, the standard expected utility model cannot generate high levels of risk aversion without contradicting established facts about labor supply.

Labor supply behavior and risk aversion are tightly linked in the expected utility model because both are determined by the curvature of utility over consumption. To see the connection, consider the effect of a wage increase on labor supply in a static model where an agent maximizes utility over consumption and leisure. If the marginal utility of consumption diminishes quickly, the individual becomes satiated with goods as wages rise. A highly risk averse individual will therefore choose to consume more leisure (by reducing labor supply) as wages rise. More generally, a higher curvature of utility over consumption implies a lower uncompensated wage elasticity of labor supply.

The bound on risk aversion is obtained by combining this logic with empirical evidence on the wage elasticity. A well established finding of the labor supply literature is that wage increases do not cause sharp reductions in labor supply. This lower bound on the wage elasticity of labor supply places an upper bound on the curvature of utility over consumption and hence on risk aversion. The fact that individuals do not choose to reduce labor supply sharply when wages rise implies that their marginal utility of consumption does not diminish quickly, unless consumption and labor are very complementary.

If complementarity between consumption and labor is sufficiently strong, even highly risk averse individuals may choose not to reduce labor supply when wages rise because increased consumption makes work less painful. Therefore, bounding γ using labor supply elasticities requires that we first bound the degree of complementarity between consumption and labor. Such a bound can be obtained from evidence on consumption choices when agents

face uncertainty about labor supply. Intuitively, the extent to which an agent chooses to correlate consumption with labor across states where labor supply varies exogenously (e.g., because of job loss or disability) reveals the degree of complementarity. Combining the bound on complementarity with estimates of labor supply elasticities yields a bound on γ that does not rely on any assumptions beyond those inherent in expected utility theory.

I formalize the preceding logic in a dynamic lifecycle model with arbitrary non-separable utility over consumption and leisure. I derive a formula for γ in terms of the ratio of the income elasticity of labor supply to the substitution elasticity of labor supply along with the cardinal complementarity parameter. I bound the complementarity parameter using a set of estimates of the consumption drop associated with job loss and other exogenous shocks to labor supply. I then estimate γ using labor supply elasticity estimates from various types of microeconomic studies – e.g., structural lifecycle methods, natural experiments, and earned income responses – as well as macroeconomic observations such as the downward trend in labor supply over the past century. Using thirty-three sets of estimates of wage and income elasticities, the mean implied value of γ is 0.71, with a range of 0.15 to 1.78 in the additive utility case. At the upper bound for complementarity, the mean value of γ rises modestly, to 0.97.

I clarify why all the labor supply studies imply a low level of γ despite disagreement about the magnitudes of the elasticities using a calibration argument. I show that generating $\gamma > 2$ with a plausible level of complementarity requires an uncompensated wage elasticity of labor supply more negative than that estimated in any of the thirty-three studies.

The bound on risk aversion derived here contrasts with the much higher estimates of risk aversion obtained in studies of asset and insurance markets (e.g., Rajnish Mehra and Edward Prescott 1985, Narayana Kocherlakota 1996, Robert Barsky et. al. 1997, Alma Cohen and Liran Einav 2005, Justin Sydnor 2005). This paper therefore provides new evidence that the conventional expected utility model falls short of explaining choices under uncertainty in many domains. Importantly, the calibration argument here restricts risk preferences over

all risks, and not just the small gambles or low probability events that are the basis of many existing critiques (Chris Starmer 2000).

The paper proceeds as follows. Section I gives graphical intuition for the bounding argument, and derives a formula for risk aversion in terms of labor supply elasticities and complementarity between consumption and labor. Section II implements the formula using existing estimates of these parameters. Section III discusses how this paper is related to other recent calibration arguments for risk aversion and intertemporal substitution. Section IV concludes.

I Theory

Basic Setup. Consider a T period life-cycle model. Denote consumption in each period by c_t and labor supply by l_t . Let $U(c_1, \dots, c_T, l_1, \dots, l_T)$ denote utility over the consumption and labor streams. Let p_t denote the price of consumption in period t . Assume that U is smooth and that $U_{c_t} > 0, U_{l_t} < 0, u_{c_t c_t} < 0, u_{l_t l_t} < 0$. Let $w\theta_t$ denote the wage in period t and y unearned income (wealth) at time 0. In Thomas MaCurdy's (1981) terminology, a change in θ_t is a transitory wage change, while changes in w are permanent wage changes, i.e. shifts in the entire profile of wages over a lifetime.

The agent chooses a path of consumption and labor by solving

$$\begin{aligned} \max_{c_t, l_t} & U(c_1, \dots, c_T, l_1, \dots, l_T) \\ \text{s.t.} & p_1 c_1 + \dots + p_T c_T = y + w(\theta_1 l_1 + \dots + \theta_T l_T) \end{aligned}$$

It is convenient to rewrite this problem as a two-stage maximization:

$$\begin{aligned} \max_{c, l} & u(c, l) \text{ s.t. } c = y + wl & (1) \\ \text{where } & u(c, l) = \max_{c_t, l_t} U(c_1, \dots, c_T, l_1, \dots, l_T) \end{aligned}$$

$$\text{s.t. } p_1 c_1 + \dots + p_T c_T = c$$

$$\theta_1 l_1 + \dots + \theta_T l_T = l$$

In (1), c and l represent aggregates that capture total consumption and labor supply over the lifecycle. The function $u(c, l)$ is indirect utility over these two composite commodities. Our goal is to derive a bound for the coefficient of relative risk aversion of the indirect utility function $u(c, l)$, defined as follows:

$$\gamma(c, l) \equiv -\varepsilon_{u_c, c} = \frac{\partial u_c(c, l)}{\partial c} \frac{c}{u_c(c, l)} = -\frac{u_{cc}(c, l)}{u_c(c, l)} c \quad (2)$$

Note that γ is the curvature of utility over wealth – the parameter that determines risk preferences over immediately-resolved wealth gambles in an expected utility model – when total labor supply l is fixed. When l is variable, the curvature of utility over wealth is *strictly lower* than γ (see appendix A for a proof). Intuitively, if the agent can adjust labor supply, he has more flexibility to adjust to wealth shocks, and is less risk averse (Zvi Bodie et. al. 1992). A bound on γ therefore bounds risk aversion when l is endogenous as well.

Bounding Risk Aversion: Graphical Example. The main result follows from the comparative statics implied by the agent’s first order condition for l . At an interior optimum, the marginal benefit of working an extra hour equals the marginal cost:

$$w u_c(y + wl, l) = -u_l(y + wl, l) \quad (3)$$

Figure 1 illustrates the calibration argument using this first order condition. It plots the marginal consumption utility of working an extra hour, $w u_c(y + wl, l)$ and the marginal disutility of working that hour, $-u_l(y + wl, l)$. The initial level of labor supply, l_0 , is determined by the intersection of these two curves at the initial wage w_0 . For simplicity, the figure is drawn for a case where the agent has no unearned income ($y = 0$).

Suppose first that the agent has additive utility over c and l ($u_{cl} = 0$). Consider the

effect of raising w by 1 percent on l . This change has two effects on the wu_c curve, which correspond to a substitution and income effect on labor supply. The substitution effect is that the number multiplying u_c rises by 1 percent, shifting the wu_c curve *upward* by 1 percent. The 1 percent increase in w also increases consumption (wl) at any given level of l by 1 percent. A 1 percent increase in consumption lowers u_c by $\varepsilon_{u_c,c} = \gamma$, so the 1 percent wage increase shifts the wu_c curve *downward* by γ percent via the income effect. The total shift in the wu_c curve is thus $(1 - \gamma)$ percent. This expression shows that higher γ makes the wage elasticity of labor supply more negative by magnifying the income effect. Intuitively, when γ is high, the marginal benefit of consumption falls quickly as the wage rises. This strengthens the incentive to consume more leisure (by reducing l) when w rises.

Since changes in w do not affect the $-u_l$ curve when $u_{cl} = 0$, it follows that

$$\partial l / \partial w > 0 \Leftrightarrow \gamma < 1$$

when $y = 0$. This result is the simplest version of the bound on risk aversion imposed by labor supply behavior. The remainder of the paper generalizes this bound to allow for positive unearned income ($y > 0$), a potentially negative wage elasticity of labor supply, and complementarity between c and l . These factors loosen the bound on γ slightly (to $\gamma < 2$), but the basic logic of the calibration argument is the same: If upward shifts in the wage profile do not cause sharp reductions in lifetime labor supply, γ must be small.

Complementarity between c and l causes shifts in the $-u_l$ curve in Figure 1 as w rises. If $u_{cl} > 0$, the $-u_l$ curve shifts outward when w rises and l rises more than it would if $u_{cl} = 0$. Consequently, the value of γ estimated from labor supply elasticities under the assumption that $u_{cl} = 0$ understates the true γ if $u_{cl} > 0$. This issue is addressed below using empirical evidence from studies of consumption smoothing to place bounds on the magnitude of u_{cl} . Given these bounds, the range of possible shifts in the $-u_l$ curve is narrow, as illustrated by the shaded region in Figure 1. The bound on γ is thus loosened modestly when plausible

levels of complementarity are permitted.

An Estimator for γ . To generalize the example in Figure 1, I derive a formula for γ in terms of labor supply elasticities. Implicitly differentiate (3) to obtain:

$$\begin{aligned}\frac{\partial l}{\partial y} &= -\frac{wu_{cc} + u_{cl}}{w^2u_{cc} + u_{ll} + 2wu_{cl}} \\ \frac{\partial l}{\partial w} &= -\frac{u_c + wlu_{cc} + lu_{cl}}{w^2u_{cc} + u_{ll} + 2wu_{cl}}\end{aligned}\tag{4}$$

Using the Slutsky decomposition for compensated labor supply ($\frac{\partial l^c}{\partial w}$)

$$\frac{\partial l^c}{\partial w} = \frac{\partial l}{\partial w} - l \frac{\partial l}{\partial y}\tag{5}$$

it follows that the ratio of the income effect to the substitution effect is given by

$$\frac{\partial l / \partial y}{\partial l^c / \partial w} = \frac{wu_{cc} + u_{cl}}{u_c}\tag{6}$$

Let $\varepsilon_{l,y} = \frac{\partial l}{\partial y} \frac{y}{l}$ denote the income elasticity of labor supply, $\varepsilon_{l,w}^c = \frac{\partial l^c}{\partial w} \frac{w}{l}$ the compensated wage elasticity of labor supply, and $\varepsilon_{u_c,l} = \frac{u_{cl}l}{u_c}$ the elasticity of the marginal utility of consumption with respect to labor. Some algebraic rearrangement gives

$$\gamma(y + wl) = -\frac{y + wl}{w} \frac{\partial l / \partial y}{\partial l^c / \partial w} = -(1 + \frac{wl}{y}) \frac{\varepsilon_{l,y}}{\varepsilon_{l,w}^c}(y, w) + (1 + \frac{y}{wl}) \varepsilon_{u_c,l}.\tag{7}$$

This equation shows that γ is determined by the ratio of the income elasticity of labor supply to the substitution elasticity of labor supply, with an adjustment for complementarity between c and l .¹ This is because the income effect is proportional to u_{cc} (how much the marginal consumption utility from working falls when y is raised) while the substitution

¹Note that (7) remains well defined when $y = 0$. In that case, the first term in (7) equals $-\frac{\partial lw}{\partial y} / \varepsilon_{l,w}^c$. The $\frac{\partial lw}{\partial y}$ term is the propensity to earn out of unearned income (in dollars rather than a percentage, which would be undefined).

effect is proportional to u_c (how much the marginal consumption utility from working rises when w is raised). For example, when utility is linear in c , there are no income effects in labor supply, and $\gamma = 0$. Note that the formula in for γ in (7) does not rely on any functional form assumptions; hence, the bounds derived below apply to any utility function.

Cardinality and Complementarity. It may be surprising that a unique value for γ can be identified from labor-leisure choices. Since non-linear monotonic transformations of $u(c, l)$ do not affect the choice of l , are there not infinitely many values of γ that could be associated with a given set of labor supply data? The reason that γ is identified in (7) is that any non-linear transformation of u would change the value of $\varepsilon_{u_c, l}$. For example, non-linear transformations of an additive u (with $u_{cl} = 0$) destroy additivity. Labor supply data are thus sufficient to identify γ *conditional* on the value $\varepsilon_{u_c, l}$, which pins down the cardinal normalization of u .

Since the cardinal complementarity parameter $\varepsilon_{u_c, l}$ is unknown, it must be estimated from choices under uncertainty. A natural method of estimating $\varepsilon_{u_c, l}$ is to examine the consumption choices of individuals who face exogenous variation in labor supply across states, e.g. due to a shock such as job displacement. Intuitively, if agents choose to consume a lot more in states where labor supply is high, c and l must be highly complementary; if in contrast labor supply fluctuations are not correlated with consumption changes, c and l must not be very complementary.

To obtain an estimate of $\varepsilon_{u_c, l}$ based on this logic, consider a setting with two states where agents work for l^1 hours in state 1 (which occurs with probability p) and l^2 hours in state 2 (probability $1 - p$). Assume that preferences are state-independent, i.e. the utility function in the two states is the same. Let w^s denote the wage in state s . Suppose the agent can trade consumption at an actuarially fair rate between the two states using an insurance policy. We will see below that if perfect insurance of this form is unavailable, the exercise below provides an upper bound for $\varepsilon_{u_c, l}$ and thereby an upper bound for γ .

Conditional on (l^1, l^2) , the agent chooses a consumption allocation (c^1, c^2) to maximize

expected utility:

$$\begin{aligned} & \max_{c^1, c^2} pu(c^1, l^1) + (1 - p)u(c^2, l^2) \\ \text{s.t.} \quad & pc^1 + (1 - p)c^2 = pw_1l^1 + (1 - p)w_2l^2 \end{aligned}$$

At the optimal (c^1, c^2) , marginal utilities are equated across the states:

$$u_c(c^1, l^1) = u_c(c^2, l^2)$$

The remainder of this section exploits this condition to link the $\varepsilon_{u_c, l}$ parameter of interest to a magnitude that can be empirically estimated. Let $\Delta c = c^2 - c^1$ and $\Delta l = l^2 - l^1$ denote the change in consumption and labor across the two states. A first-order Taylor expansion of u_c around c^1 gives:

$$u_c(c^2, l^2) = u_c(c^1, l^1) + u_{cc}(c^1, l^1)\Delta c + u_{cl}(c^1, l^1)\Delta l + R$$

where R , the remainder, must satisfy $\lim_{\Delta l \rightarrow 0} R = 0$. Therefore, in the optimal allocation,

$$\begin{aligned} -u_{cc}\Delta c &= u_{cl}\Delta l + R \\ \implies \gamma \frac{\Delta c}{c^1} &= \varepsilon_{u_c, l} \frac{\Delta l}{l^1} + \frac{R}{u_c(c^1, l^1)} \\ \implies \varepsilon_{u_c, l} &= \lim_{\Delta l \rightarrow 0} \gamma \frac{\Delta c}{c^1} / \frac{\Delta l}{l^1} \end{aligned} \tag{8}$$

Equation (8) shows that $\varepsilon_{u_c, l}$ is proportional to $\frac{\Delta c}{c} / \frac{\Delta l}{l}$, the percentage drop in consumption associated with a 1 percent difference in labor supply across states. This expression reflects the intuition described above: If the consumption change across states where labor supply differs is small, $\varepsilon_{u_c, l}$ must be small. The curvature of utility (γ) is also relevant because it determines the cost of consumption fluctuations in the expected utility model. The limit $\Delta l \rightarrow 0$ is necessary because $\varepsilon_{u_c, l}$ can be identified at a given point (c^1, l^1) without functional

form assumptions only by observing the effect of small variations in l on c .

Importantly, in the more realistic case where insurance markets are incomplete, consumption will fall beyond the optimal amount when labor supply is low. Hence, imperfections in insurance markets will make the observed consumption drop overstate the true complementarity-related consumption drop and consequently *overstate* the true values of $\varepsilon_{u_c, l}$ and γ .

Using (8) and (7), we can solve for γ to obtain an estimator for risk aversion in terms of magnitudes that can be empirically estimated:

$$\gamma = \left(1 + \frac{wl}{y}\right) \frac{-\varepsilon_{l,y}}{\varepsilon_{l^c,w}} / \left(1 - \left(1 + \frac{y}{wl}\right) \left[\lim_{\Delta l \rightarrow 0} \frac{\Delta c}{c} / \frac{\Delta l}{l}\right]\right) \quad (9)$$

Extensive Margin. The best established effects of wage changes are on the participation margin, perhaps because fixed costs of participation and institutional restrictions limit hours choices (see e.g. Joseph Altonji and Christina Paxson 1991). Estimates of participation elasticities can also be used to infer γ . Let θ denote the fraction of agents who work, $\varepsilon_{\theta,y}$ the income elasticity of participation, and $\varepsilon_{\theta,w}$ the wage elasticity of participation. Let $\frac{\Delta c}{c}$ denote the difference in consumption when working and not working chosen by the agent in an experiment involving uncertain labor supply analogous to the complementarity exercise described above. Under a constant- γ approximation of $u(c, l)$, a formula similar to (9) is obtained for γ :²

$$\gamma = \frac{\log\left[1 - \frac{\varepsilon_{\theta,y} w}{\varepsilon_{\theta,w} y}\right]}{\log\left[\left(1 - \frac{\Delta c}{c}\right)\left(1 + \frac{w}{y}\right)\right]} \quad (10)$$

²Details are given in Appendix C.

II Empirical Implementation

II.A Estimates of Complementarity

Equation (9) shows that an upper bound on $\frac{\Delta c}{c} / \frac{\Delta l}{l}$ is required to obtain an upper bound on γ . A bound on complementarity would ideally be derived from the consumption choices of agents who face small, permanent exogenous shocks to labor supply.³ The most obvious empirical analogs to this experiment are estimates of the consumption change associated with shocks such as job loss or disability. John Cochrane (1991) and Jonathan Gruber (1997, 1998) find that job loss causes a consumption drop of less than 10 percent. In subsequent work, Martin Browning and Thomas Crossley (2001) and Hans Bloemen and Elena Stancanelli (2005) show that consumption does not fall at all for individuals with positive liquid wealth prior to job loss. In addition, these studies find that higher unemployment benefits are associated with smaller consumption drops, and that with full insurance, there would be no drop at all. These results imply that most of the observed 10 percent consumption drop is due to imperfect insurance markets rather than complementarity between consumption and labor.

There are two concerns in connecting the 10 percent bound to the actual $\frac{\Delta c}{c} / \frac{\Delta l}{l}$ parameter of interest. First, the studies of job loss examine large fluctuations in l and therefore may not provide a good estimate of $\lim_{\Delta l \rightarrow 0} \frac{\Delta c}{c} / \frac{\Delta l}{l}$ if complementarity is much greater for small fluctuations in l than large ones. This concern is unlikely to be a serious problem in practice. Studies that examine smaller fluctuations in hours than full unemployment (e.g., Browning et. al. 1985) find estimates of $\frac{\Delta c}{c} / \frac{\Delta l}{l}$ that are of the same magnitude as those reported by studies of larger fluctuations in l . Moreover, most of the changes in labor supply resulting from changes in wages and unearned income tend to be large and discrete as well (e.g., from 20 to 40 hours). The range of Δl over which complementarity is estimated is therefore

³The shocks must be “exogenous” in the sense that they are involuntary changes in labor supply, as opposed to preference shocks that endogenously induce labor supply changes.

similar to the range over which the labor supply elasticities themselves are estimated. As equation (10) for the extensive margin case shows, if only discrete changes in labor supply are feasible, it is preferable to have estimates of the consumption drop when l fluctuates over a similar set of discrete values.⁴

The second concern, which is deeper, is that studies of job loss examine temporary fluctuations in labor (variation in l_t for a given period t) and not permanent fluctuations (variation in l). In the notation of the model, these studies estimate $\frac{\Delta c_t}{c_t} / \frac{\Delta l_t}{l_t}$ for a single period t rather than the desired value $\frac{\Delta c}{c} / \frac{\Delta l}{l}$ that reflects changes in lifetime aggregates. When utility is time non-separable, these two values need not be equal. The ratio of $\frac{\Delta c}{c} / \frac{\Delta l}{l}$ to $\frac{\Delta c_t}{c_t} / \frac{\Delta l_t}{l_t}$ is determined by the degree of cross-period complementarity in consumption.⁵ Intuitively, if consumption is complementary across periods (as in habit formation models), agents will be more reluctant to cut consumption in response to transitory fluctuations in labor than permanent ones. Durability of consumption and adjustment costs could further attenuate the short-run response.

To gauge the difference between short-run and long-run complementarity, I use evidence on consumption responses to long-term labor supply changes induced by disability or retirement. Cochrane (1991) finds that long-term disabilities cause a 11 percent drop in food consumption in the year that the shock occurs. Melvin Stephens (2001) shows that in the five years after disability occurs, consumption does not trend downward significantly, and is at most 10 percent lower than the pre-disability level. These results suggest that long-run complementarity ($\frac{\Delta c}{c} / \frac{\Delta l}{l}$) is not much greater than short-run ($\frac{\Delta c_t}{c_t} / \frac{\Delta l_t}{l_t}$) complementarity. If it were, there would be either a large immediate drop in consumption or a sharp downward trend in consumption in the years after disability.

⁴Relatedly, the estimates of γ based on participation elasticities – which require estimates of $\frac{\Delta c}{c}$ from fluctuations in labor force participation – yield very similar estimates of γ (see Table 1). This suggests that discreteness is unlikely to be an important source of bias here.

⁵See Appendix D for a formal derivation relating the two parameters. Karen Dynan (2001) finds no complementarity in consumption across periods in microdata, but studies using macro data find evidence of habit.

In related work, Paul Gertler and Gruber (2002) find that long-term health shocks leading to job loss are associated with less than a 20 percent reduction in non-health consumption (which includes durables) in Indonesia. Gertler and Gruber test whether incomplete insurance or complementarity between c and l is responsible for this drop in several ways. For instance, they show that the consumption drop is small in families where the person experiencing the shock is not the sole earner (because other household members help to smooth consumption). They conclude from this and other evidence that the complementarity-related portion of the 20 percent drop is close to zero.

One concern with the disability-based evidence is that the assumption of state-independent preferences may not hold for health shocks.⁶ Studies of retirement provide additional evidence on complementarity that helps mitigate such concerns. Mark Aguiar and Erik Hurst (2005) use detailed data on expenditures to show that expenditure drops at retirement by less than 15 percent.⁷ Douglas Bernheim et. al. (2001) show that there is no downward trend in expenditures in the years after retirement. These findings are also consistent with the claim that $\frac{\Delta c}{c} / \frac{\Delta l}{l}$ is not much larger than $\frac{\Delta c_t}{c_t} / \frac{\Delta l_t}{l_t}$.

In summary, evidence on the effect of job loss on consumption implies $\frac{\Delta c_t}{c_t} / \frac{\Delta l_t}{l_t} < 0.1$. An examination of the differences between this estimate and the long-run complementarity parameter of interest suggests a bound of $\frac{\Delta c}{c} / \frac{\Delta l}{l} < 0.15$.

II.B Labor Supply Elasticities

This section describes a set of elasticity estimates from studies of labor supply and reports the γ implied by each study. There is a controversial debate about which empirical methods yield the most reliable estimates of labor supply elasticities. I show that irrespective of the

⁶For example, Cochrane (1991, p974) notes that “sick people might lose their appetites” and therefore consume less. Insofar as health shocks reduce the taste for non-health consumption, the consumption drops associated with disability overstate the true level of complementarity between c and l .

⁷In Aguiar and Hurst’s time input model, the bound derived in this paper is a bound on the curvature of utility over expenditure, holding labor supply fixed. This remains an upper bound on curvature of utility over wealth, following the derivation in Appendix A.

method used to estimate the elasticities, the implied value of γ is always low.

Labor supply studies can be broadly classified into four categories: (1) The “static” approach estimates reduced-form labor supply responses to events such as tax changes, cross-sectional differences, or lottery winnings. Richard Blundell and MaCurdy (1999) show that these static estimates can be interpreted as labor supply responses to the permanent changes in wages and unearned income of interest when an appropriate set of controls for age and cohort are included. (2) The “life cycle” or “structural” literature, pioneered by MaCurdy (1981), explicitly models dynamic labor supply and consumption choices and backs out estimates of labor supply responses to permanent shifts in wage profiles and unearned income from life cycle variation in wages in a panel dataset. These estimates correspond more directly to the permanent wage-elasticities (e.g., $\varepsilon_{l,w}^c$) of interest, but identification of these models is often difficult because of the lack of exogenous shifts in wage profiles. Recent studies that combine the benefits of exogenous variation used in the static studies with the structural lifecycle approach give perhaps the most credible microeconomic estimates of long-run wage elasticities (Blundell et. al. 1998). (3) A more recent “earned income” literature, starting with Martin Feldstein (1995, 1999), examines the effect of tax reforms on total earned income as a means of capturing other margins of labor supply beyond hours (e.g., effort or job-related training). Estimates from this literature can be used to estimate γ by replacing the elasticity ratio $\frac{\varepsilon_{l,y}}{\varepsilon_{l,w}^c}$ used in (7) with $\frac{\varepsilon_{LI,y}}{\varepsilon_{LI^c,1-\tau}}$, where LI is labor income and $1-\tau$ the net-of-tax rate. (4) Finally, long-run macroeconomic trends and cross-country comparisons can be used to make inferences about long-run labor supply elasticities, potentially overcoming the institutional rigidities and some of the omitted variable biases that may affect the microeconomic studies.⁸

Table 1 presents a set of income and substitution elasticities from studies using each of these methods. The first two sets of estimates (hours and participation elasticities) are from studies that use the traditional static and lifecycle approaches. The third section shows es-

⁸The elasticities from the micro-level studies should yield consistent estimates of γ even if there are frictions which prevent agents from reoptimizing fully in the short-run. These frictions presumably attenuate both $\varepsilon_{l,y}$ and $\varepsilon_{l,w}^c$, leaving the ratio of the two elasticities unaffected.

timates from studies of earned income responses, and the fourth shows the macroeconomic evidence. The macro estimates are constructed using a lower bound on the uncompensated wage elasticity based on the secular downward trend in hours over the past century (documented e.g. by Casey Mulligan 2002) combined with estimates of substitution elasticities from other studies (see Appendix B for details).

To obtain a broad sense of the values of γ consistent with labor supply evidence, the table includes elasticity estimates for a wide range of groups, such as prime age males, married women, retired individuals, and low income families. Estimates of γ are computed at the mean values of y , w , and l in each study. Note that the mean values of $\frac{y}{wl}$ vary widely across the studies. For example, married women’s unearned income equals at least their husband’s income, which is generally larger than their own earned income.

Column (6) of Table 1 reports estimates of γ for the additive utility case. The overall (unweighted) mean estimate of γ across the 33 sets of elasticity estimates is $\gamma = 0.71$. Only 3 studies imply a value of γ above 1.25 when $u_{cl} = 0$.⁹ The macroeconomic evidence suggests slightly higher values of risk aversion than the microeconomic studies because the downward trend in labor supply over time implies a significantly larger income effect than substitution effect. The estimates from Blundell et. al.’s (1998) study, which perhaps addresses the central identification concerns in estimating labor supply elasticities most cleanly, yield $\gamma = 0.93$. Column (7) of Table 1 reports estimates of γ that account for complementarity consistent with the bound of $\frac{\Delta c}{c} / \frac{\Delta l}{l} = 0.15$. This adjustment increases the average estimate of γ to 0.97.

⁹John Pencavel (1986), Blundell and MaCurdy (1999), and Gruber and Emmanuel Saez (2002) summarize more than sixty other microeconomic studies that span various methodologies, nearly all of which imply $\gamma < 1.25$ as well.

II.C A Calibration Argument

The similarity of the estimates of γ across the labor supply studies despite their differences in methodology, definitions of labor supply, and sample composition may be surprising. This section provides a calibration argument that explains the consensus on γ . Intuitively, the consensus emerges from the uniform finding that $\varepsilon_{l,w}$ is not very negative, which implies that the income elasticity cannot be large relative to the substitution elasticity. This places an upper bound on γ because it depends on the ratio of these two elasticities.

To formalize this argument, consider first the common benchmark of an upward-sloping labor supply curve (Prescott (1986), Robert Hall and John Taylor (1991)).¹⁰ Using the Slutsky equation and (9), it follows that

$$\varepsilon_{l,w} \geq 0 \iff \gamma < 1 + \frac{y}{wl}$$

with additive utility. In the aggregate, $\frac{y}{wl}$ equals the ratio of capital income to labor income, which is $\frac{1}{2}$ in the U.S. Hence, with additive utility, $\varepsilon_{l,w} \geq 0$ implies $\gamma \leq 1.5$ for a representative agent. The skewed distribution of wealth implies that $\frac{y}{wl} < \frac{1}{2}$ for most households, implying that the bound on γ is tighter for many households. Note that if $y = 0$, $\gamma < 1$, consistent with Figure 1.

Table 2 generalizes this calibration result by showing the implied value of γ for several other cases, including cases where $\varepsilon_{l,w} < 0$ and cases with complementarity. Each column considers a different value for the ratio of the income effect of a 1 percent wage increase to the substitution effect, defined as $I/\varepsilon_{l,w}^c = -\frac{lw}{y}\varepsilon_{l,y}/\varepsilon_{l,w}^c$.¹¹ Each row represents a different value of the degree of complementarity. The table reports the implied γ in each cell assuming

¹⁰In a recent survey of 134 labor and public economists at 40 leading research institutions, Victor Fuchs, Alan Krueger, and James Poterba (1998) found that the vast majority of these experts believe that the best estimate of the uncompensated wage elasticity is weakly positive.

¹¹The Slutsky decomposition for a wage increase is $\varepsilon_{l,w} = \varepsilon_{l,w}^c + \frac{lw}{y}\varepsilon_{l,y}$, where the first term on the right hand side is the substitution effect and the second is the income effect. Hence $I = -\frac{lw}{y}\varepsilon_{l,y}$ corresponds to the (absolute value of) the income effect of a wage increase.

$\frac{y}{wl} = \frac{1}{2}$ (see Appendix B for details). For instance, the benchmark case of $\varepsilon_{l,w} = 0$ implies $I/\varepsilon_{l,w}^c = 1$ (income and substitution effects cancel exactly). With no complementarity this yields $\gamma = 1.5$, consistent with the derivation above.

The calibrations show that γ does not rise much if the labor supply curve is downward sloping to the extent suggested by the macroeconomic evidence in part D of Table 1. The macro evidence, which yields the most negative estimates of $\varepsilon_{l,w}$ of all the studies, implies $I/\varepsilon_{l,w}^c$ less than $\frac{4}{3}$ (see Appendix B). At this value, γ rises to 2. The calibrations also show that γ is not very sensitive to the degree of complementarity. With $I/\varepsilon_{l,w}^c = 1$ and the upper bound complementarity value of $\frac{\Delta c}{c}/\frac{\Delta l}{l} = 0.15$, γ rises to 1.94. The bottom line is that generating γ significantly greater than 2 would require complementarity and labor supply patterns that contradict evidence to date sharply.

III Discussion

A few recent papers have also conducted “internal consistency checks” of standard models of consumption behavior. Most relevant is Susanto Basu and Miles Kimball [BK] (2002), who build on Robert King et. al. (1988). BK show that reconciling low estimates of the elasticity of intertemporal substitution (EIS) with $\varepsilon_{l,w} \geq 0$ requires either strong complementarity between consumption and labor or time non-separable utility. To see how our results are related, consider the case where utility is additive over c and l . Here, the BK result is that *time* separability is inconsistent with $\varepsilon_{l,w} > 0$ and low EIS. In contrast, this paper shows that *state* separability (expected utility theory) is inconsistent with $\varepsilon_{l,w} > 0$ and high γ . The two results thus address two aspects of preferences – intertemporal substitution and risk aversion – that are empirically and intuitively distinct (Hall (1988), Philippe Weil (1990), Larry Epstein and Stanley Zin (1991)). While the BK result leaves γ unidentified, the bound in this paper leaves the EIS unrestricted because $U(\cdot)$ is permitted to be an arbitrary time

non-separable function.¹² Similarly, while habit formation (which drops time separability) can resolve the BK bound on the EIS, it does not relax the bound on risk aversion.

Matthew Rabin (1999) and Louis Kaplow (2005) also give calibration results for risk preferences in an expected utility model. Rabin shows that expected utility cannot generate a reasonably high level of moderate-stakes risk aversion without creating unreasonably high large-stakes risk aversion. Kaplow shows that estimates of the income elasticity of the value of a statistical life bound γ because the rich would pay much more to save their lives if the marginal utility of non-health consumption fell quickly with wealth. Each of these calibration arguments illuminates the restrictions inherent in expected utility theory in a different way.

IV Conclusion

A large literature on labor supply has found that the uncompensated wage elasticity of labor supply is not very negative. This observation places a bound on the rate at which the marginal utility of consumption diminishes, and thus bounds risk aversion in an expected utility model. The central estimate of the coefficient of relative risk aversion implied by labor supply studies is 1 (log utility) and an upper bound is 2, accounting for substantial complementarity between consumption and labor. The intuition for this tight bound is simple: If the marginal utility of wealth diminishes rapidly, why don't people choose to work much less when their wages rise?

This result implies that diminishing marginal utility of wealth plays a secondary role in generating the high levels of risk aversion estimated in many studies of choice under uncertainty. An additional, quantitatively powerful source of risk aversion must be identified

¹²Another way to see this point is to consider Kreps-Porteus utility. When the only risk at issue is an immediately resolved one, the Kreps-Porteus specification is a special case of the general time non-separable class of utility functions analyzed above. Consequently, the arguments above bound risk aversion over immediately-resolved wealth gambles for a Kreps-Porteus utility, but do not pin down the EIS.

to explain observed behavior in these cases.¹³ Testing alternative models of risk preferences under the constraints on curvature imposed by labor supply behavior would be an interesting direction for further research. More generally, examining how one domain of behavior (such as labor supply) disciplines the conclusions drawn in another domain (such as choice under uncertainty) could be a useful method of developing unified, internally consistent theories of economic behavior.

¹³Recent examples of theories that introduce additional sources of risk aversion beyond diminishing marginal utility include Botond Koszegi and Rabin's (2005) model of reference-dependent risk preferences and Raj Chetty's (2004) model of consumption commitments and local risk aversion.

References

- Altonji, Joseph G. and Paxson, Christina H.** “Labor Supply, Hours Constraints, and Job Mobility.” *Journal of Human Resources*, 1992, 27(2), pp. 256-278.
- Aguiar, Mark and Hurst, Erik.** “Consumption vs. Expenditure.” *Journal of Political Economy*, 2005, 113(5), pp. 919-948.
- Auten, Gerald and Carroll, Robert.** “The Effect of Income Taxes on Household Behavior.” *Review of Economics and Statistics*, 1999, 81, pp. 681-693.
- Barsky, Robert B.; Juster, F. Thomas; Kimball, Miles S. and Shapiro, Matthew D.** “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study.” *Quarterly Journal of Economics*, 1997, 112, pp. 537-580.
- Bernheim, B. Douglas; Skinner, Johnathan and Weinberg, Steven.** “What Accounts for the Variation in Retirement Wealth Among U.S. Households?” *American Economic Review*, 2001, 91(4), pp. 832-857.
- Basu, Susanto and Kimball, Miles.** “Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption.” University of Michigan mimeo, 2002.
- Blau, Francine and Kahn, Lawrence.** “Changes in the Labor Supply Behavior of Married Women: 1980-2000.” NBER Working Paper 11230, 2005.
- Bloemen, Hans and Stancanelli, Elena.** “Financial Wealth, Consumption Smoothing and Income Shocks Arising from Job Loss,” *Economica*, 2005, 72(3), pp. 431-452.
- Blundell, Richard; Duncan, Alan and Meghir, Costas.** “Estimating Labor Supply Responses Using Tax Reforms.” *Econometrica*, 1998, 66(7), pp. 827-862.
- Blundell, Richard and MaCurdy, Thomas.** “Labor Supply: A Review of Alternative Approaches,” in Ashenfelter, Orley and David Card, eds., *Handbook of Labor Economics*. Vol. 3. Amsterdam: North-Holland, 1999.
- Bodie, Zvi; Merton, Robert and Samuelson, William.** “Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model.” *Journal of Economic Dynamics and Control*, 1992, 16, pp. 427-450.
- Browning, Martin and Crossley, Timothy.** “Unemployment Insurance Benefit Levels and Consumption Changes.” *Journal of Public Economics*, 2001, 80, pp. 1-23.
- Browning, Martin; Deaton, Angus and Irish, Margaret.** “A Profitable Approach to Labor Supply and Commodity Demands Over the Life Cycle.” *Econometrica*, 1985, 53, pp. 503-44.
- Browning, Martin; Hansen, Lars Peter and Heckman, James.** “Micro Data and General Equilibrium Models.” in John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Vol. 1A, Amsterdam: North Holland, 1999.
- Chetty, Raj.** “Consumption Commitments, Unemployment Durations, and Local Risk Aversion.” NBER Working Paper 10211, 2004.
- Cochrane, John H.** “A Simple Test of Consumption Insurance.” *Journal of Political Economy*, 1991, 99, pp. 957-76.

- Cohen, Alma and Einav, Liran.** “Estimating Risk Preferences from Deductible Choice.” NBER Working Paper 11461, 2005.
- Davis, Steven and Henrekson, Magnus.** “Tax Effects on Work Activity, Industry Mix and Shadow Economy Size: Evidence from Rich-Country Comparisons.” NBER Working Paper 10509, 2004.
- Dynan, Karen.** “Habit Formation in Consumer Preferences: Evidence from Panel Data.” *American Economic Review*, 2000, 90(6), pp. 391-406.
- Eissa, Nada and Hoynes, Hillary.** “The Earned Income Tax Credit and the Labor Supply of Married Couples.” NBER Working Paper 6856, 1998.
- Epstein, Larry and Zin, Stanley.** “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis.” *Journal of Political Economy*, 1991, 99, pp. 263-286.
- Feldstein, Martin.** “The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act.” *Journal of Political Economy*, 1995, 103, pp. 551-572.
- Feldstein, Martin.** “Tax Avoidance and the Deadweight Loss of the Income Tax.” *Review of Economics and Statistics*, 1999, 81, pp. 674-680.
- Friedberg, Leoria.** “The Labor Supply Effects of the Social Security Earnings Test.” *Review of Economics and Statistics*, 2000, 82, pp. 48-63.
- Fuchs, Victor; Krueger, Alan B. and Poterba, James M.** “Economists’ Views about Parameters, Values, and Policies: Survey Results in Labor and Public Economics.” *Journal of Economic Literature*, 1998, 36, pp. 1387-1425.
- Gertler, Paul and Gruber, Jonathan.** “Insuring Consumption Against Illness.” *American Economic Review*, 2002, 3, pp. 51-70.
- Gruber, Jonathan.** “The Consumption Smoothing Benefits of Unemployment Insurance.” *American Economic Review*, 1997, 87, pp. 192-205.
- Gruber, Jonathan.** “Unemployment Insurance, Consumption Smoothing, and Private Insurance: Evidence from the PSID and CEX.” *Research in Employment Policy*, 1998, 1, pp. 3-31.
- Gruber, Jonathan and Saez, Emmanuel.** “The Elasticity of Taxable Income: Evidence and Implications.” *Journal of Public Economics*, 2002, 84(1), pp 1-32.
- Hall, Robert.** “Intertemporal Substitution in Consumption.” *Journal of Political Economy*, 1988, 96, pp. 339-357.
- Hall, Robert E., and Taylor, John B.** *Macroeconomics: Theory, Performance, and Policy*. 3rd ed. New York: Norton, 1991.
- Imbens, Guido; Rubin, Donald B. and Sacerdote, Bruce I.** “Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players.” *American Economic Review*, 2001, 91, pp. 778-794.
- Kaplow, Louis.** “The Value of a Statistical Life and the Coefficient of Relative Risk Aversion.” *The Journal of Risk and Uncertainty*, 2005, 31(1), pp. 23-34.
- King, Robert; Plosser, Charles and Rebelo, Sergio.** “Production, Growth and

Business Cycles I. The Basic Neoclassical Model.” *Journal of Monetary Economics*, 1988, 21, pp. 309-341.

Kocherlakota, Narayana. “The Equity Premium: It’s Still a Puzzle.” *Journal of Economic Literature*, 1996, 24, pp. 42-71.

Koszegi, Botond and Rabin, Matthew. “Reference-Dependent Risk Attitudes.” UC-Berkeley mimeo, 2005.

MaCurdy, Thomas. “An Empirical Model of Labor Supply in a Life-Cycle Setting.” *The Journal of Political Economy*, 1981, 89, pp. 1059-1085.

MaCurdy, Thomas; Green, David and Paarsch, Harry. “Assessing Empirical Approaches for Analyzing Taxes and Labor Supply.” *Journal of Human Resources*, 1990, 25, pp. 415-90.

Mehra, Rajnish and Prescott, Edward C. “The Equity Premium: A Puzzle.” *Journal of Monetary Economics*, 1985, 15, pp. 145-161.

Mulligan, Casey. “A Century of Labor-Leisure Distortions.” NBER Working Paper 8774, 2002.

Pencavel, John. “Labor Supply of Men: A Survey.” in Ashenfelter, Orley and Richard Layard, eds., *Handbook of Labor Economics*. Vol. 1. Amsterdam: North Holland, 1986.

Prescott, Edward. 1986. “Theory Ahead of Business Cycle Measurement.” *Federal Reserve Bank of Minneapolis Quarterly Review*, 1986, 10(3), pp. 9-22.

Prescott, Edward. “Why Do Americans Work So Much More Than Europeans?” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2004, 28(1).

Rabin, Matthew. “Risk Aversion and Expected-Utility Theory: A Calibration Theorem.” *Econometrica*, 2000, 68, pp. 1281-1292.

Starmer, Chris. “Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice Under Risk.” *Journal of Economic Literature*, 2000, 38, pp. 332-382.

Stephens, Melvin. “The Long-Run Consumption Effects of Earnings Shocks.” *The Review of Economics and Statistics*, 2001, 83(1), pp. 28-36.

Sydnor, Justin. “Sweating the Small Stuff: The Demand for Low Deductibles in Homeowners Insurance.” UC-Berkeley mimeo, 2005.

Weil, Philippe. “Non-Expected Utility in Macroeconomics.” *Quarterly Journal of Economics*, 1990, 29-42.

Appendix A: Curvature of utility over wealth

Define indirect utility over wealth when l is endogenous as

$$v(y) = u(y + wl(y), l(y))$$

Since the envelope condition requires

$$v_y(y) = u_c(c(y), l(y))$$

it follows that

$$v_{yy} = u_{cc} \frac{\partial c}{\partial y} + u_{cl} \frac{\partial l}{\partial y}$$

Recall the expression for $\partial l / \partial y$ in (4):

$$\frac{\partial l}{\partial y} = K(wu_{cc} + u_{cl}) \quad (11)$$

where $K = -\frac{1}{w^2u_{cc} + 2wu_{cl} + u_{ll}}$. Equation (5) implies that $\frac{\partial l^c}{\partial y} = -\frac{u_c}{w^2u_{cc} + 2wu_{cl} + u_{ll}}$. Utility maximization requires $\frac{\partial l^c}{\partial w} > 0$, implying that $K > 0$.

Recognizing that $\partial c / \partial y = 1 + w \partial l / \partial y$, it follows that

$$v_{yy} = u_{cc} + u_{cc}w \frac{\partial l}{\partial y} + u_{cl} \frac{\partial l}{\partial y}$$

Now plug in using (11) for $\partial l / \partial y$ in the preceding expression to obtain

$$\begin{aligned} v_{yy} &= u_{cc} + K[w^2u_{cc}^2 + wu_{cc}u_{cl} + u_{cl}^2] \\ &= u_{cc} + K[wu_{cc} + u_{cl}]^2 \end{aligned}$$

It follows that $v_{yy} > u_{cc}$ which implies

$$\gamma^y = \frac{-v_{yy}}{v_y} y < \frac{-u_{cc}}{v_y} y = \frac{-u_{cc}}{u_c} c \frac{y}{c} = \gamma \frac{y}{c} = \gamma \frac{y}{y + wl} < \gamma$$

This proves that $\gamma^y < \gamma$, i.e. that the curvature of utility over wealth is lower when l is endogenous.

Appendix B: Construction of Tables 1 and 2

Notes on Table 1: In part A of Table 1, the first two rows assume $\frac{y}{wl} = \frac{1}{2}$ because MaCurdy (1981) does not report the mean ratio of unearned to earned income in his sample and the Blundell and MaCurdy (1999) elasticity estimates are an average across several different studies, some of which do not report $\frac{y}{wl}$. All other rows in part A use the mean reported values of y and wl in conjunction with the elasticity estimates reported in that study. In part B, I use the CRRA approximation used to derive equation (10) to estimate γ with the reported extensive-margin elasticities. In part C, I use the Imbens. et. al. income elasticity estimate in conjunction with the compensated wage elasticity estimates from the other studies with $\frac{y}{wl} = \frac{1}{2}$. The compensated wage elasticity estimates in the earned income literature are the elasticity of earned income with respect to the net of tax rate.

In part D, for the Blau and Kahn (2005) study, I take the average of the three sets of substitution elasticities reported for three different periods. The income elasticity is defined as the elasticity of women's hours with respect to husband's wages and computed in corresponding fashion. I estimate γ using the mean value of y and wl reported by Blau and Kahn for their sample.

For the remaining two studies in part D, I first estimate the uncompensated wage elasticity $\varepsilon_{l,w}$ from Mulligan (2002), who reports a 25 percent drop in aggregate hours over the 20th century while real hourly wages rose by roughly a factor of 8. This implies $\varepsilon_{l,w} \approx -0.035$. To account for the possibility that labor supply might be less arduous than it was 100 years ago (e.g. individuals get more breaks today), I double this value to obtain $\varepsilon_{l,w} = -0.07$. Note that placing a lower bound on $\varepsilon_{l,w}$ leads to an upper bound on γ given an estimate of $\varepsilon_{l,w}^c$. Estimates of the compensated wage elasticity are obtained from other studies that compare trends or levels across countries with varying tax and transfer regimes (Prescott 2004, Davis and Henrekson 2004). These tax responses can be interpreted as compensated wage elasticities of aggregate labor supply since non-transfer government expenditure can be viewed as unearned income in the aggregate. Income elasticities are then computed for each study using the Slutsky equation under the assumption $\varepsilon_{l,w} = -0.07$ with $\frac{y}{wl} = \frac{1}{2}$. Finally, I compute γ using the resulting compensated wage and income elasticities with $\frac{y}{wl} = \frac{1}{2}$.

The overall mean estimates of γ are unweighted means of the values reported in each study. In computing the mean, the Blundell and MaCurdy (1999) values are given a weight of 20 since this line represents an average of twenty different studies.

Notes on Table 2: The formula used for the calibrations reported in Table 2 is derived as follows. Rewrite the Slutsky equation given in (5) in terms of elasticities:

$$\varepsilon_{l^c,w} = \varepsilon_{l,w} - \frac{lw}{y} \varepsilon_{l,y}$$

Let $I \equiv -\frac{lw}{y} \varepsilon_{l,y} = -\frac{\partial wl}{y}$ denote the income effect of a wage increase. Then we can write γ

in terms of I as:

$$\gamma = (1 + \frac{y}{wl}) \frac{I}{\varepsilon_{l,w}^c} / (1 - (1 + \frac{y}{wl} \frac{\Delta c}{c} / \frac{\Delta l}{l})) \quad (12)$$

The values reported in the table are computed using this formula with $\frac{y}{wl} = \frac{1}{2}$.

To derive the bound of $\frac{I}{\varepsilon_{l,w}^c} < \frac{4}{3}$ implied by the macro trend evidence described in the text, note that $\varepsilon_{l,w} = -0.07$ is a lower bound on the uncompensated wage elasticity for reasons described above. Given this parameter, it is necessary to place a lower bound on $\varepsilon_{l,w}^c$ to obtain an upper bound on $\frac{I}{\varepsilon_{l,w}^c}$ and γ . Most studies find $\varepsilon_{l,w}^c$ above 0.2, with the macroeconomic evidence suggesting larger values. With $\varepsilon_{l,w} = -0.07$ and $\varepsilon_{l,w}^c = 0.2$, the Slutsky equation implies that $\frac{I}{\varepsilon_{l,w}^c} \approx \frac{4}{3}$.

Appendix C: Extensive Margin

Suppose that the agent makes a binary decision to work and supply 1 unit of labor or not to work at all. Let y denote unearned income and w the additional income earned by working. Returning temporarily to additive utility over consumption and leisure, redefine $u(c)$ as the utility from consumption. Let ψ denote disutility of supplying 1 unit of labor. The agent chooses labor supply by solving

$$\max_{l \in \{0,1\}} u(y + wl) - \psi l$$

He works if his disutility of labor is less than the utility of an additional w units of consumption, i.e. if

$$\psi < \hat{\psi}(y, w) \equiv u(y + w) - u(y)$$

Suppose there is heterogeneity in disutility of labor in the economy given by a smooth density $f(\psi)$. Then the fraction of workers who participate in the labor force is

$$\theta(y, w) = \int_0^{\hat{\psi}(y, w)} f(\psi) d\psi \quad (13)$$

It follows that

$$-\frac{\partial \theta / \partial y}{\partial \theta / \partial w} = \frac{u_c(y) - u_c(y + w)}{u_c(y + w)} \quad (14)$$

This expression shows that the percent change in marginal utility of wealth from y to $y + w$ is equal to the ratio of the income and wage effects on labor supply. In the intensive labor supply model, we could compute $\gamma(c)$ at any level c without making any functional form assumptions because we could observe how marginal utility changes for small changes in income. With extensive labor supply decisions, we observe only the change in marginal utility between y and $y + w$. Consequently, we need to make a functional form assumption for $u(c)$ to translate the change in marginal utilities into a coefficient of relative risk aversion. I assume CRRA utility.¹⁴ Then

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

Under this assumption, (14) implies

$$-\frac{\partial \theta / \partial y}{\partial \theta / \partial w} = \frac{y^{-\gamma} - (y + w)^{-\gamma}}{(y + w)^{-\gamma}}$$

¹⁴If $\gamma(c)$ actually varies with c , this method yields the best constant- γ fit of the data, which can be loosely interpreted as the average $\gamma(c)$ in the region $c \in [y, y + w]$.

Solving for γ yields

$$\gamma = \frac{\log[1 - \frac{\varepsilon_{\theta,y} w}{\varepsilon_{\theta,w} y}]}{\log[1 + \frac{w}{y}]}$$

Finally, a model of exogenous extensive-margin labor supply shocks analogous to that in the intensive margin case can be used to derive an estimator for γ when utility is not additive:

$$\gamma = \frac{\log[1 - \frac{\varepsilon_{\theta,y} w}{\varepsilon_{\theta,w} y}]}{\log[(1 - \frac{\Delta c}{c})(1 + \frac{w}{y})]} \quad (15)$$

where $\frac{\Delta c}{c}$ denotes the consumption drop associated with job loss.

Appendix D: Short-Run vs. Long-Run Complementarity: $\frac{\Delta c_t}{c_t} / \frac{\Delta l_t}{l_t}$ and $\frac{\Delta c}{c} / \frac{\Delta l}{l}$

This appendix establishes a connection between the consumption drop observed from a transitory labor supply shock and the corresponding drop that would result from a permanent labor supply shock. When the short-run drop understates the size of the long-run drop, I show that the ratio of these two values can be bounded by a parameter that measures the strength of cross-period complementarities in utility.

Let $\frac{\Delta c_t/c_t}{\Delta l_t/l_t}|_{LR}$, denote the consumption change in period t for a permanent labor supply shock that raises labor supply proportionally in all periods by $\Delta l_t/l_t = \Delta l/l$. Let $\frac{\Delta c_t/c_t}{\Delta l_t/l_t}|_{SR}$ denote the response of consumption in period t to a transitory labor supply fluctuation that shifts l_t alone. The analysis below requires the following assumption.

Assumption A1. Within-period complementarity across consumption and labor is stronger than cross-period complementarities: $|U_{c_t l_t}| > |U_{c_t l_{t+1}}|$.

This condition is intuitive: it requires that a transitory shock to l_t generates smaller Δc_s for $s \neq t$ than a permanent shock to l that changes $l_t \forall t$.

I begin by considering the effect of transitory variation labor supply on consumption. Assume without loss of generality that the agent faces uncertainty in l_1 , while all subsequent l_n are known with certainty. Denote the two states of the world with superscripts a and b and assume $l_1^a > l_1^b$. An agent maximizing utility over his lifecycle chooses his consumption allocation across the two states to satisfy:

$$U_{c_1}(c_1^a, c_2^a, \dots, c_n^a, l_1^a, l_2, \dots, l_n) = U_{c_1}(c_1^b, c_2^b, \dots, c_n^b, l_1^b, l_2, \dots, l_n)$$

Take a Taylor expansion to simplify the right hand side of this expression:

$$U_{c_1}^b = U_{c_1}^a + U_{c_1 c_1} \Delta c_1 + U_{c_1 l_1} \Delta l_1 + \sum_{n>1} U_{c_1 c_n} \Delta c_n + \sum_{n>1} U_{c_1 l_n} \Delta l_n$$

Here and subsequently we drop the remainder from this expression, which is appropriate if we ultimately take the limit as $\Delta l \rightarrow 0$ as in the text. Since $\Delta l_n = 0 \forall n > 1$, the optimality condition simplifies (taking the limit as $\Delta l_1 \rightarrow 0$) to:

$$-U_{c_1 c_1} \frac{\Delta c_1}{\Delta l_1}|_{SR} = U_{c_1 l_1} + \sum_{n>1} U_{c_1 c_n} \frac{\Delta c_n}{\Delta l_1}$$

Let $\gamma_{c_1} = -\frac{U_{c_1 c_1}}{U_{c_1}} c_1$, the curvature of utility over period 1 consumption. Algebraic rearrangement of the preceding condition gives

$$\gamma_{c_1} \frac{\Delta c_1/c_1}{\Delta l_1/l_1}|_{SR} = \varepsilon_{U_{c_1}, l_1} + \sum_{n>1} \varepsilon_{U_{c_t}, c_n} \frac{\Delta c_n/c_n}{\Delta l_1/l_1}|_{SR} \quad (16)$$

Now consider permanent variation in l , resulting in a constant proportional change $\frac{\Delta l_t}{l_t}$

for all t . A corresponding optimality condition for the consumption vector is:

$$U_{c_1}(c_1^a, c_2^a, \dots, c_n^a, l_1^a, l_2^a, \dots, l_n^a) = U_{c_1}(c_1^b, c_2^b, \dots, c_n^b, l_1^b, l_2^b, \dots, l_n^b)$$

Taking a Taylor expansion of RHS as above and simplifying gives the approximate requirement:

$$-U_{c_1 c_1} \frac{\Delta c_1}{\Delta l_1} |_{LR} = U_{c_1 l_1} + \sum_{n>1} U_{c_1 c_n} \frac{\Delta c_n}{\Delta l_1} |_{LR} + \sum_{n>1} U_{c_1 l_n} \frac{\Delta l_n}{\Delta l_1} |_{LR}$$

which implies

$$\gamma_{c_1} \frac{\Delta c_1/c_1}{\Delta l_1/l_1} |_{LR} = \varepsilon_{U_{c_1}, l_1} + \sum_{n>1} \varepsilon_{U_{c_1}, c_n} \frac{\Delta c_n/c_n}{\Delta l_1/l_1} |_{LR} + \sum_{n>1} \varepsilon_{U_{c_1}, l_n} \frac{\Delta l_n/l_n |_{LR}}{\Delta l_1/l_1} \quad (17)$$

Comparing (16) and (17), it follows that

$$\frac{\Delta c_1/c_1}{\Delta l/l} |_{LR} = \frac{\Delta c_1/c_1}{\Delta l_1/l_1} |_{SR} + \sum_{n>1} \frac{\varepsilon_{U_{c_1}, c_n}}{\gamma_{c_1}} \left(\frac{\Delta c_n/c_n}{\Delta l/l} |_{LR} - \frac{\Delta c_n/c_n}{\Delta l_1/l_1} |_{SR} \right) + \sum_{n>1} \frac{\varepsilon_{U_{c_1}, l_n}}{\gamma_{c_1}} \frac{\Delta l_n/l_n |_{LR}}{\Delta l/l} \quad (18)$$

Since we are interested in placing bounds on complementarity, I focus on the case where c_t and l_t are complements, i.e. $\frac{\Delta c_t/c_t}{\Delta l_t/l_t} |_{SR} > 0$. Equation (18) indicates that the short run consumption drop understates the long-run drop if the last two terms on the right hand side are positive. Under the assumption that cross-period complementarity is weaker than within-period complementarity, it follows that $\frac{\Delta c_n/c_n}{\Delta l_1/l_1} |_{SR} < \frac{\Delta c_n/c_n}{\Delta l/l} |_{LR}$. Therefore, complementarity of consumption and labor across periods ($\varepsilon_{U_{c_1}, c_n} > 0$ and $\varepsilon_{U_{c_1}, l_n} > 0$) will lead to $\frac{\Delta c_1/c_1}{\Delta l/l} |_{LR} > \frac{\Delta c_1/c_1}{\Delta l_1/l_1} |_{SR}$. Since this is the case of interest in deriving a bound on risk aversion, I focus on it below.

Note that

$$\frac{\Delta c/c}{\Delta l/l} = \frac{\sum c_t \left(\frac{\Delta c_t/c_t}{c_t} / \frac{\Delta l}{l} \right) |_{LR}}{\sum c_t}$$

equals a consumption-weighted average of the consumption change (across all periods) associated with the long-run labor supply shock $\Delta l/l$. Hence, $\frac{\Delta c/c}{\Delta l/l} \leq \max_t \frac{\Delta c_t/c_t}{c_t} / \frac{\Delta l}{l} |_{LR}$, i.e. the total consumption drop is lower than the largest drop in a given period. Without loss of generality, assume that the period with the largest drop is $t = 1$. Let

$$\delta_t = \frac{\Delta c_t/c_t}{\Delta l/l} |_{LR} - \frac{\Delta c_1/c_1}{\Delta l/l} |_{LR}$$

denote the deviation of the consumption change in period t from the period 1 change.

Rewrite (18) as

$$\begin{aligned} \frac{\Delta c_1/c_1}{\Delta l/l} \Big|_{LR} \left(1 - \sum_{n>1} \frac{\varepsilon_{U_{c_1}, c_n}}{\gamma_{c_1}}\right) &= \frac{\Delta c_1/c_1}{\Delta l_1/l_1} \Big|_{SR} + \sum_{n>1} \frac{\varepsilon_{U_{c_1}, c_n}}{\gamma_{c_1}} \left(\delta_n - \frac{\Delta c_n/c_n}{\Delta l_1/l_1} \Big|_{SR}\right) \\ &+ \sum_{n>1} \frac{\varepsilon_{U_{c_1}, l_n}}{\gamma_{c_1}} \frac{\Delta l_n/l_n \Big|_{LR}}{\Delta l/l} \end{aligned}$$

When consumption is complementarity across periods, it can be shown that $\frac{\Delta c_n/c_n}{\Delta l_1/l_1} \Big|_{SR} > 0$. Intuitively, the agent will choose higher consumption at all times in the high transitory-labor state to maintain similar consumption streams across periods in the two states. Since $\delta_t \leq 0 \forall t$ by construction, it follows that

$$\begin{aligned} \frac{\Delta c/c}{\Delta l/l} \left(1 - \sum_{n>1} \frac{\varepsilon_{U_{c_1}, c_n}}{\gamma_{c_1}}\right) &\leq \frac{\Delta c_1/c_1}{\Delta l_1/l_1} \Big|_{SR} + \sum_{n>1} \frac{\varepsilon_{U_{c_1}, l_n}}{\gamma_{c_1}} \frac{\Delta l_n/l_n \Big|_{LR}}{\Delta l/l} \\ &\leq \max_t \left[\frac{\Delta c_t/c_t}{\Delta l_t/l_t} \Big|_{SR}\right] + \sum_{n>1} \frac{\varepsilon_{U_{c_1}, l_n}}{\gamma_{c_1}} \frac{\Delta l_n/l_n \Big|_{LR}}{\Delta l/l} \end{aligned}$$

where $\max_t \left[\frac{\Delta c_t/c_t}{\Delta l_t/l_t} \Big|_{SR}\right]$ is the largest observed short-run consumption drop over the agent's lifecycle. Since $\Delta l_n/l_n \Big|_{LR}$ equals $\Delta l/l$ by construction in all periods, rearrangement gives

$$\begin{aligned} \frac{\Delta c/c}{\Delta l/l} &\leq \frac{1}{\mu} \max_t \left[\frac{\Delta c_t/c_t}{\Delta l_t/l_t} \Big|_{SR}\right] \\ \mu &= \frac{\gamma_{c_1} \Delta c/c - \sum_{n>1} \varepsilon_{U_{c_1}, c_n} \Delta c/c - \sum_{n>1} \varepsilon_{U_{c_1}, l_n} \Delta l/l}{\gamma_{c_1} \Delta c/c} \end{aligned} \quad (19)$$

Equation (19) shows that the ratio of $\frac{\Delta c}{c} / \frac{\Delta l}{l}$ to $\frac{\Delta c_t}{c_t} / \frac{\Delta l_t}{l_t}$ is determined by μ , the degree of cross-period complementarity in consumption. The denominator of μ represents the effect of raising period 1 consumption by a percentage $\Delta c/c$ (and leaving all other consumption and labor levels fixed) on U_{c_1} (the marginal utility of consumption in period 1). The numerator of μ is the effect on U_{c_1} of increasing c in *all* periods by the same fixed percentage $\Delta c/c$ while increasing labor supply in all periods $n > 1$ by a fixed percentage $\Delta l/l$. Hence, μ represents how much changing c_n and l_n in all other periods besides period 1 dampens the change in U_{c_1} relative to a change in only c_1 . Note that if there are no cross-period complementarities in U , $\mu = 1$.

TABLE 1
Labor Supply Elasticities and Implied Coefficients of Relative Risk Aversion

Study (1)	Sample (2)	Identification (3)	Income Elasticity (4)	Compensated Wage Elasticity (5)	γ Additive (6)	γ $\Delta c/c=0.15$ (7)
<i>A. Hours</i>						
MaCurdy (1981)	Married Men	Panel	-0.020	0.130	0.46	0.60
Blundell and MaCurdy (1999)	Men	Various	-0.120	0.567	0.63	0.82
MaCurdy, Green, Paarsch (1990)	Married Men	Cross Section	-0.010	0.035	1.47	1.81
Eissa and Hoynes (1998)	Married Men, Inc < 30K	EITC Expansions	-0.030	0.192	0.88	1.08
	Married Women, Inc < 30K	EITC Expansions	-0.040	0.088	0.64	1.34
Friedberg (2000)	Older Men (63-71)	Soc. Sec. Earnings Test	-0.297	0.545	0.93	1.46
Blundell, Duncan, Meghir (1998)	Women, UK	Tax Reforms	-0.185	0.301	0.93	1.66
Average					0.69	0.94
<i>B. Participation</i>						
Eissa and Hoynes (1998)	Married Men, Inc < 30K	EITC Expansions	-0.008	0.033	0.44	0.48
	Married Women, Inc < 30K	EITC Expansions	-0.038	0.288	0.15	0.30
Average					0.29	0.39
<i>C. Earned Income</i>						
Imbens, Rubin, Sacerdote (2001)	Lottery Players in MA	Lottery Winnings	-0.110			
Feldstein (1995)	Married, Inc > 30K	TRA 1986		1.040	0.32	0.41
Auten and Carroll (1997)	Single and Married, Inc>15K	TRA 1986		0.660	0.50	0.65
Average					0.41	0.53
<i>D. Macroeconomic/Trend Evidence</i>						
Blau and Kahn (2005)	Women	Cohort Trends	-0.278	0.646	0.60	1.29
Davis and Henrekson (2004)	Europe/US aggregate stats	Cross-Section of countries	-0.251	0.432	1.74	2.25
Prescott (2004)	Europe/US aggregate stats	Cross-Country time series	-0.222	0.375	1.78	2.30
Average					1.37	1.95
Overall Average					0.71	0.97

NOTES -- All risk aversion estimates are computed at sample means of y and wl unless noted otherwise. In Part A, the Blundell and MaCurdy estimates are an unweighted average of the 20 elasticities reported in that study and assumes $y/wl=1/2$. In Part B, calculations of γ assume CRRA utility. In Part C, compensated wage elasticity column reports the elasticity of earned income with respect to the net-of-tax rate. For these studies, the Imbens et. al. estimate of the income elasticity is used to compute γ . In Part D, income elasticities for the Davis and Henrekson and Prescott studies are computed from estimates in Mulligan (2002). See Appendix B for further details on the construction of this table.

TABLE 2

Labor Supply, Complementarity, and Risk Aversion: Calibration Results

		Labor Supply Elasticity Ratio: $l/\varepsilon_{l,w}^c$				
		0.33	0.66	1.00	1.33	1.66
Complementarity $(\Delta c/c)/(\Delta l/l)$	0.00	0.50	0.99	1.50	2.00	2.49
	0.05	0.54	1.07	1.62	2.16	2.69
	0.10	0.58	1.16	1.76	2.35	2.93
	0.15	0.64	1.28	1.94	2.57	3.21
	0.20	0.71	1.41	2.14	2.85	3.56

NOTES -- This table shows the implied value of γ for various income/substitution elasticity ratios and consumption-labor complementarity levels. Values of γ are computed using equation (12) with $y/wl=1/2$. See Appendix B for additional details.




FIGURE 1

Risk Aversion and the Uncompensated Wage Elasticity of Labor Supply

NOTE—This figure illustrates the labor supply decision of an agent who has no unearned income ($y = 0$) at two wage levels (initial wage w_0 and new wage $w_1 > w_0$). The downward-sloping lines show the marginal consumption utility of working for an extra hour and the upward sloping lines show the marginal disutility of working that hour. The optimal level of labor supply is determined by the intersection of these curves. The effect of the wage increase on labor supply is shown for two cases under the assumption that $u_{cl} = 0$: (A) $\gamma < 1$, where the increase in w raises labor supply from l_0 to l_A ; and (B) $\gamma > 1$, where the same increase in w reduces labor supply from l_0 to l_B . If $u_{cl} \neq 0$, changes in w shift the marginal disutility of labor curve as shown in the shaded region, loosening the bound on γ .