

NBER WORKING PAPER SERIES

A NEW METHOD OF ESTIMATING RISK AVERSION

Raj Chetty

Working Paper 9988

<http://www.nber.org/papers/w9988>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

September 2003

I thank George Akerlof, John Campbell, David Card, Keith Chen, David Cutler, John Friedman, Ed Glaeser, Emir Kamenica, Louis Kaplow, Adam Looney, Greg Mankiw, Casey Mulligan, Mike Ostrovsky, Jesse Shapiro, and especially Martin Feldstein, Gary Chamberlain, Larry Katz, and Caroline Hoxby for helpful comments and discussions. Financial support from the National Science Foundation, National Bureau of Economic Research, and Harvard University is gratefully acknowledged. The views expressed herein are those of the authors and are not necessarily those of the National Bureau of Economic Research.

©2003 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 New Method of Estimating Risk Aversion

Raj Chetty

NBER Working Paper No. 9988

September 2003

JEL No. D8, G12, J22

ABSTRACT

This paper develops a method of estimating the coefficient of relative risk aversion (γ) from data on labor supply. The main result is that existing estimates of labor supply elasticities place a tight bound on γ , without any assumptions beyond those of expected utility theory. It is shown that the curvature of the utility function is directly related to the ratio of the income elasticity of labor supply to the wage elasticity, holding fixed the degree of complementarity between consumption and leisure. The degree of complementarity can in turn be inferred from data on consumption choices when employment is stochastic. Using a large set of existing estimates of wage and income elasticities, I find a mean estimate of $\gamma = 1$. I also give a calibration argument showing that a positive uncompensated wage elasticity, as found in most studies of labor supply, implies $\gamma < 1.25$. The estimate of γ changes by at most 0.25 over the range of plausible values for the complementarity parameter.

Raj Chetty

Department of Economics

University of California at Berkeley

521 Evans Hall, #3880

Berkeley, CA 94720

and NBER

chetty@econ.berkeley.edu

The curvature of utility over consumption, measured e.g. by the coefficient of relative risk aversion, is an important parameter for a broad set of economic problems. It is of obvious relevance in calibrating models of risk-taking behavior such as portfolio choice, insurance, and executive compensation. But the coefficient of relative risk aversion also enters into problems that require cardinal utility but do not explicitly involve uncertainty. It determines the deadweight cost of taxation and the benefits of redistribution, which are needed to identify optimal tax and welfare policies. In addition, if utility is time-separable, the curvature of utility is related to the path of consumption, savings, and labor supply over a lifetime.

Despite its importance, empirical estimates of the coefficient of relative risk aversion (γ) are limited and highly dispersed. One possible reason is the lack of field data that can be used to ascertain γ without making assumptions about agents' beliefs and planning horizons.¹ In the absence of compelling empirical evidence, most economists believe based on introspection that $\gamma \in (1, 5)$ while others contend that higher values of γ are reasonable.²

This paper proposes a new method of estimating risk aversion that uses widely available

¹Existing estimates of γ are highly dispersed. Estimates of γ from portfolio choice and equity premiums exceed 10 (Mehra and Prescott 1985, Kocherlakota 1996). Estimates of γ using data on insurance deductibles, premiums, and the fraction of insured assets range from 2 to 10 (Szpiro 1986, Dreze 1987). To the extent that agents are myopic or ill-informed about insurance or portfolio choice, these estimates of γ may be biased. A more recent literature uses experimental data and yields estimates of γ between 0 and 15 (see e.g. Wolf and Pohlman 1983, Barsky et. al. 1997, Metrick 1995). While these studies are informative, it is difficult to know whether individuals' behavior in experimental or hypothetical realms is reflective of their preferences in real life.

²See e.g. Kandel and Stambaugh (1991), who argue that $\gamma = 10$ does not seem unreasonable when one introspects about small gambles.

data on labor supply decisions. Surprisingly, existing estimates of labor supply elasticities place a tight bound on γ , without making any assumptions beyond expected utility maximization. The central result is that γ is directly related to the ratio of the income elasticity of labor supply to the price (substitution) elasticity of labor supply. To see why, recall that $\gamma \propto \frac{u_{cc}}{u_c}$ where u_c denotes the first derivative of utility with respect to consumption, and u_{cc} denotes the second derivative. An agent's labor supply response to a wage increase is directly related to u_c , the marginal utility of consumption: the larger the magnitude of u_c , the greater the benefit of an additional dollar of income, and the more the agent will work when w goes up. The labor supply response to an increase in income is related to how much the marginal utility of consumption changes as income changes, u_{cc} . If u_{cc} is large, the marginal utility of consumption falls sharply as income rises, so the agent will reduce labor supply significantly when his income rises. It follows that there is a connection between γ and the ratio of income and price elasticities.

Formalizing this connection requires an additional step. Labor supply data cannot be used in isolation to identify cardinal properties of the utility function because data on certainty behavior only identify utility functions up to a monotonic transformation.³ The cardinality of the utility function must be pinned down using information on risky decisions. To see how this can be done, observe that for a fixed degree of complementarity between

³In other words, identification of curvature requires a 1-1 map between observed choices and γ . In the general labor-leisure model, such a map does not exist because any monotonic transformation of utility generates the same labor supply choices.

consumption and leisure (u_{cl}), there is only one vN-M utility (up to affine transformations) that can be consistent with a given set of labor-leisure choices. Hence, given a value of u_{cl} , labor supply data can be used to estimate risk aversion. I show that u_{cl} can in turn be inferred from data on consumption choices when employment is stochastic, e.g. when individuals anticipate being laid off with some probability.⁴ Intuitively, the extent to which an agent chooses to smooth consumption across states in which labor supply differs (via an insurance policy or saving) reveals the degree of complementarity between consumption and leisure. Importantly, I find that the estimates of risk aversion are insensitive to the value of u_{cl} , implying that γ can be tightly bounded using data from only a simple static labor-leisure choice setting, where the assumption that agents maximize utility and “get things right” is likely to be most tenable. Unlike prior estimates of γ , which hinge on rational expectations and perfect foresight, the estimates of risk aversion using the method proposed here are more robust to myopia and overconfidence in planning ahead for risky events such as unemployment.

The remainder of the paper is organized as follows. In the next section, I derive estimators for γ in the standard labor supply model, generalizing the model in steps to simplify the exposition. I begin by providing graphical intuition for the connection between labor supply and risk aversion in a setting where the marginal disutility of labor is constant. I then derive

⁴It is critical to have data on choices under *uncertainty*; barring additional assumptions, a cardinal value for u_{cl} cannot be inferred from the usual certainty settings in which we typically think about estimating the degree of complementarity between consumption and leisure (e.g. timing-of-work, household production).

an estimator for γ when utility is known to be additive in consumption and leisure ($u_{cl} = 0$). The next subsection considers the general case of arbitrary u_{cl} and derives an estimator for γ that uses both labor supply data and information on consumption choices when agents face unemployment risk. Finally, I discuss the implications of two assumptions made in all of the preceding derivations – that all agents have identical preferences and are able to make intensive labor supply choices. Not surprisingly, when risk preferences are heterogeneous, the estimator yields a measure of average risk aversion in the population. When agents can make only extensive labor supply choices, a corresponding formula for risk aversion based on income and price elasticities for labor force participation is obtained.

In Section 3, I implement these formulae using a large set of existing estimates of wage and income elasticities. I begin with the case in which $u_{cl} = 0$ and show that almost all studies of labor supply imply $\gamma \approx 1$ under this assumption. I then show that the estimate of γ changes by at most 0.25 for plausible perturbations in the degree of complementarity between consumption and leisure. The ability to estimate γ precisely from labor supply data may be surprising given the lack of consensus about the magnitudes of these elasticities. To clarify this point, I give a calibration argument showing that if the uncompensated labor supply curve is upward sloping (as almost all studies of labor supply find), $\gamma < 1.25$. The low estimate of γ is intuitive: since γ is related to the income effect divided by the price effect, as long as income effects are not too large relative to price effects – which we know they are not because the uncompensated supply curve is upward sloping – γ must be small.

1 Labor Supply and Risk Aversion

(a) Intuition

Consider an agent who has a vN-M utility function $u(c, l)$ over consumption and labor. This agent's expected utility from a gamble (\tilde{c}, \tilde{l}) is obtained by taking an expectation over the vN-M utility function:

$$U(\tilde{c}, \tilde{l}) = Eu(\tilde{c}, \tilde{l})$$

The goal of this paper is to identify the curvature of the underlying cardinal utility over outcomes, $u(c, l)$.⁵

This section illustrates the connection between risk aversion and labor supply graphically in the special case where the marginal disutility of labor is constant. Let c denote consumption, l labor supply, and w the wage. The agent's utility over consumption is $u(c)$; assume that $u_c > 0$ and $u_{cc} < 0$. Let ψ denote the marginal disutility of labor. I first derive an expression for compensated labor supply, $l^c(w)$, by solving an expenditure minimization problem:

$$\min c + w(1 - l) \text{ s.t. } u(c) - \psi l \geq \bar{v}$$

At an interior optimum, $l^c(w)$ satisfies the first order condition

$$u_c(u^{-1}(\bar{v} + \psi l^c)) = \frac{\psi}{w}$$

⁵I focus on the curvature of utility over consumption $\gamma = \frac{u_{cc}}{u_c} c$ here, but show below that deriving other curvature parameters of interest is straightforward once the value of γ is known.

This condition is intuitive: the individual chooses labor supply by equating the marginal utility of an extra dollar of consumption with the marginal disutility of working to earn that dollar. These choices are depicted by the intersections of the u_c and $\frac{\psi}{w}$ curves in Figure 1.

Now consider the effect of an increase in w from w_0 to w_1 on compensated labor supply. As shown in Figure 1, a change in w shifts the flat marginal disutility of labor curve downward. If the utility function is highly curved (case A), the marginal utility of consumption (u_c) falls quickly as labor supply and income rise. Consequently, the increase in w leads to a small increase in l_A^c . When the utility function is not very curved (case B), marginal utility declines slowly as a function of wealth and the same Δw leads to a larger increase in l_B^c . Figure 1 therefore illustrates that the compensated wage (price) elasticity of labor supply, $\varepsilon_{l,w}^c$ is inversely related to the curvature of utility over consumption.

The intuition for this relationship is as follows. Following a compensated wage increase, agents increase their labor supply up to the point where the marginal utility of an additional dollar is offset by the marginal disutility of the additional work necessary to earn that dollar. If utility is very curved, this condition is met by a small increase in labor supply. If utility is not very curved, the agent needs to increase l much more before his marginal utility of money falls sufficiently to equal the new $\frac{\psi}{w_1}$.

The preceding argument relies on the assumption that disutility of labor, ψ , does not vary with l . When it does, the curvature of $\psi(l)$ is confounded with the curvature of u and $\varepsilon_{l,w}^c$ is no longer sufficient to recover γ . In this case, the elasticity of labor supply with respect to




Figure 1: Recovering γ from Labor Supply

unearned income, $\varepsilon_{l,y}$ is needed to separate the two curvature parameters. Abstractly, the income and compensated wage elasticities are both functions of the two curvatures. One can therefore back out γ and the curvature of ψ by solving a system of two equations and two unknowns, while holding fixed the degree of complementarity between consumption and leisure.

The next section derives the relationship between γ and labor supply elasticities formally in the more general case where the marginal disutility of labor varies with l . For simplicity of exposition, I continue to focus on the case in which utility is additive in consumption and labor ($u_{cl} = 0$).

(b) Base Case: Additive Utility

It is convenient to redefine $u(c, l)$ as the agent's utility over both consumption and labor.

Assume that $u_c > 0, u_l < 0, u_{cc} < 0, u_{ll} < 0$. This section also assumes additivity: $u_{cl} = 0$.⁶

An agent with wage w and unearned income y chooses (Marshallian) labor supply l by solving

$$\max_l u(y + wl, l)$$

At an interior optimum, l satisfies the first order condition

$$wu_c(y + wl, l) = -u_l(y + wl, l) \tag{1}$$

Consider the effects of increasing w and y on l :

$$\begin{aligned} \frac{\partial l}{\partial y} &= -\frac{wu_{cc}}{w^2u_{cc} + u_{ll}} \\ \frac{\partial l}{\partial w} &= -\frac{u_c + wl u_{cc}}{w^2u_{cc} + u_{ll}} \end{aligned}$$

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

$$\frac{\partial l^c}{\partial w} = \frac{\partial l}{\partial w} - l \frac{\partial l}{\partial y} \tag{2}$$

⁶Note that this restriction is stronger than assuming that the utility function permits an additively separable representation. For example, Cobb-Douglas utility is “additively separable” but does *not* satisfy the additivity restriction (however, the log of a Cobb-Douglas utility does).

it follows that

$$\frac{\partial l / \partial y}{\partial l^c / \partial w} = \frac{w u_{cc}(y + wl, l)}{u_c(y + wl, l)} \quad (3)$$

The definition of the coefficient of relative risk aversion at consumption c is

$$\gamma(c) \equiv -\varepsilon_{u_c, c} = \frac{\partial u_c(c)}{\partial c} \frac{c}{u_c(c)} = -\frac{u_{cc}(c)}{u_c(c)} c \quad (4)$$

which implies, using (3), that

$$\gamma(y + wl) = -\frac{y + wl}{w} \frac{\partial l / \partial y}{\partial l^c / \partial w} = -\left(1 + \frac{wl}{y}\right) \frac{\varepsilon_{l, y}}{\varepsilon_{l^c, w}}(y, w, l) \quad (5)$$

where $\varepsilon_{l, y}$ denotes the income elasticity, $\varepsilon_{l^c, w}$ the compensated wage (price) elasticity, and $\frac{wl}{y}$ the ratio of earned to unearned income.⁷

As in the graphical example, the coefficient of relative risk aversion is inversely related to the price elasticity. In addition, when disutility of labor is not constant, γ is directly related to the magnitude of the income elasticity. To see why the ratio of these elasticities determines risk aversion, recall that $\gamma \propto \frac{u_{cc}}{u_c}$ where u_c denotes the first derivative of utility with respect to consumption, and u_{cc} denotes the second derivative. An agent's labor supply response to a wage increase is directly related to u_c : the larger the magnitude of u_c , the greater the benefit of an additional dollar of income, and the more the agent will work

⁷Note that this estimator does not blow up as $y \rightarrow 0$, because $\varepsilon_{l, y}$ has y in the numerator; hence, this formula can be applied to the typical household in the U.S. that has negligible unearned income.

when w goes up. The labor supply response to an increase in income is related to how much the marginal utility of consumption changes as income changes, u_{cc} . A large income effect implies that the agent is willing to increase effort significantly in order to recoup lost income, which means that the marginal utility of consumption rises quickly as income falls, i.e. the magnitude of u_{cc} is large. The connection between γ and the ratio of income and price elasticities follows from these two observations.

The reader may be puzzled that we can identify a unique value for γ by observing only labor supply. 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 observed labor supply behavior? While this is true in general, the key is to observe that any non-linear transformation of u will change the value of u_{cl} . However, (5) was derived under the assumption that $u(c, l)$ is additive, i.e. $u_{cl} = 0$.

It should be emphasized that γ is the curvature of utility over *consumption*.⁸ When labor supply is not fixed, γ will not necessarily equal the curvature of utility over wealth. However, once we know γ , we have a complete map of the vN-M utility function, and can calculate any curvature of interest. For example, defining indirect utility over unearned income as

$$v(y) = u(y + wl(y)) - \psi(y)$$

⁸This is the parameter estimated by most studies of choice under uncertainty, insofar as labor supply is omitted from these analyses. More importantly, it is the relevant parameter for most models of risky behavior (e.g. portfolio choice and optimal social insurance).

it is shown in the appendix that the curvature of utility over unearned income is

$$-\frac{v_{yy}}{v_y}y = \gamma \frac{y + wl\varepsilon_{l,y}}{y + wl} < \gamma$$

Finally, one may worry that adjustment costs which prevent agents from reoptimizing fully in response to perturbations in w and y will affect the estimates of γ . This is not true because the method is biased only by factors that affect the price and income elasticities *differently*. Status-quo biases that make some individuals reluctant to change their labor supply or institutional constraints that make small adjustments in labor supply difficult do not affect the estimate.⁹

The preceding derivation hinges on the assumption that utility is additive ($u_{cl} = 0$), which turns out to be inconsequential for the estimates of risk aversion from labor supply. The next section relaxes this assumption.

(c) Complementarity Between Labor and Consumption

When $u_{cl} \neq 0$, (3) becomes

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

⁹To see this formally, note that an adjustment cost or status quo bias can be modeled as a cost $k(l, l_0)$ of changing labor supply to l from l_0 . Since we make no assumptions about the way in which labor supply l enters $u(c, l)$, (5) still obtains. The reason is that curvature is identified from the ratio of income and price effects, and $k > 0$ attenuates both effects.

which, after some rearrangement, implies that

$$\gamma(y + wl) = \left(1 + \frac{wl}{y}\right) \frac{-\varepsilon_{l,y}}{\varepsilon_{l^c,w}} + \left(1 + \frac{y}{wl}\right) \varepsilon_{u_c,l} \quad (7)$$

where $\varepsilon_{u_c,l}$ denotes the elasticity of the marginal utility of consumption with respect to labor supply. Note that labor supply data is sufficient to identify γ given *any* value of $\varepsilon_{u_c,l}$ because no non-linear transformation of u will leave $\varepsilon_{u_c,l}$ unchanged.

It remains to estimate $\varepsilon_{u_c,l}$ from choices under *uncertainty*. One way to do so is by observing data on the consumption choices of individuals who face uncertainty in labor supply. For example, if an agent chooses to keep consumption fairly constant across states in which he is employed and unemployed, $\varepsilon_{u_c,l}$ must be small.

To derive the relationship between consumption choices when employment is stochastic and $\varepsilon_{u_c,l}$ formally, consider a world with two states. Agents supply l_1 units of labor in state 1 and l_2 units of labor in state 2. Suppose that the agent can trade consumption fairly between the two states by purchasing state-contingent commodities (e.g. using an insurance policy). He chooses consumption in the two states by maximizing expected utility

$$\begin{aligned} & \max_{c_1, c_2} pu(c_1, l_1) + (1 - p)u(c_2, l_2) \\ \text{s.t.} \quad & pc_1 + (1 - p)c_2 = pwl_1 + (1 - p)wl_2 \equiv \bar{W} \end{aligned}$$

The agent's first-order condition for consumption is obtained by equating marginal utilities

across the two states:

$$u_c(c_1, l_1) = u_c(\bar{W} - c_1, l_2)$$

Now, suppose we observe data from the following experiment. Assume that the agent starts out supplying a constant $l_1 = l_2 = \bar{l}$ units of labor in each state. Suppose there is a balanced-budget change in labor supply, increasing state 1 labor supply by δ_1 units while decreasing state 2 labor supply by δ_2 units to keep expected earnings fixed at \bar{W} . We can think of this as a decision to increase work effort in state 1 to compensate for (partial) unemployment in state 2.

Differentiating the first order condition with respect to l_1 while holding earnings fixed at \bar{W} yields the following identity:

$$\varepsilon_{u_c, l} = \gamma \varepsilon_{c_1, l_1}$$

Here, ε_{c_1, l_1} denotes the elasticity of consumption with respect to labor supply in state 1 (while labor supply in state 2 changes so that total income remains constant). Plugging this expression into (7) and solving gives an estimator for risk aversion in terms of ε_{c_1, l_1} :

$$\gamma = \left(1 + \frac{wl}{y}\right) \frac{-\varepsilon_{l, y}}{\varepsilon_{l^c, w}} / \left(1 - \left(1 + \frac{y}{wl}\right) \varepsilon_{c_1, l_1}\right) \quad (8)$$

This formula reduces to (5) when utility is additive in labor and consumption ($\varepsilon_{c_1, l_1} = 0$). When consumption and labor are complements, $\varepsilon_{c_1, l_1} > 0$, and the actual γ is higher relative

to the estimate obtained when additive utility over c and l is assumed.¹⁰

While one can introspect about the magnitude of ε_{c_1, l_1} , it is helpful to map existing estimates of the consumption drop during unemployment into values of ε_{c_1, l_1} . To do so, normalize the agent's labor supply to 1 when working and 0 when unemployed. Since we are using a discrete change in labor supply to infer an elasticity, we must choose a functional form for consumption in terms of labor supply. For simplicity, I assume a linear form:

$$c = a + bl$$

In this case, $\varepsilon_{c_1, l_1}(l_1 = 1) = \frac{b}{a+b}$, which is precisely the percentage drop in consumption from the employed to unemployed state. Estimates of the consumption drop are small: for example, Gruber (1997) estimates $\frac{b}{a+b} = 0.068$ using data from the PSID.¹¹ It should be stressed that this estimate gives an *upper* bound for the true ε_{c_1, l_1} if the agent does not smooth consumption across states to the optimal level because of the inadequacy of available insurance policies, myopia, or overconfidence about the probability of unemployment. Hence, the empirical evidence suggests that the departure from additivity ($u_{cl} = 0$), if any, is not large. Since u_{cl} is the only parameter directly estimated from choices under uncertainty,

¹⁰The sign of ε_{c_1, l_1} is theoretically ambiguous. If consumption requires time, as in Becker (1965), leisure and consumption are complements ($\varepsilon_{c_1, l_1} < 0$). On the other hand, work-related expenses can make labor and consumption complementary ($\varepsilon_{c_1, l_1} > 0$). I show below that regardless of the sign of ε_{c_1, l_1} , the estimates of γ are not very sensitive to its magnitude.

¹¹Similarly, Browning and Crossley (2001) use data from the Canadian Out of Employment Panel to show that the consumption drop is not statistically distinguishable from zero for households that have positive liquid assets before their unemployment spell.

it follows that the labor supply method of estimating risk aversion relies primarily on the assumption that agents make static labor-leisure choices by maximizing utility. The estimates of risk aversion in this paper are therefore more likely to be robust to myopia or ill-informed planning for risky events than existing estimates from domains such as insurance purchases or portfolio choice.

That concludes the analysis of risk aversion and labor supply in the standard labor-leisure choice model. The final two subsections explore the implications of relaxing two assumptions of this model: (1) homogeneous preferences and (2) the ability to make intensive labor supply decisions.

(d) Extension 1: Unobservable Heterogeneity in Preferences

How should the estimators (5) and (8) be interpreted in a world with unobserved heterogeneity in preferences? To answer this question, suppose there are N types of agents, who have utilities $u^1(c, l), u^2(c, l), \dots, u^n(c, l)$. To simplify the exposition, assume that all the utilities are additive in c and l . Let α_i denote the fraction of type i agents.

Let us define a utility function $\bar{u}(c)$ for a representative agent by taking a weighted average of the individual utilities as follows:

$$\bar{u}(c, l) = \sum_{i=1}^J \omega_i u^i(c, l)$$

where the weights

$$\omega_i = \frac{\alpha_i / (w^2 u_{cc}^i + u_{ll}^i)}{\sum_{i=1}^J \alpha_i / (w^2 u_{cc}^i + u_{ll}^i)}$$

for $i = 1, \dots, J$ sum to 1. Note that in general, the representative agent's utility will differ depending on the values of (y, w, l) .

It is shown in the appendix that in a heterogeneous population, the estimate $\hat{\gamma}(y + wl)$ equals the coefficient of relative risk aversion ($\gamma^{\bar{u}}$) for the representative agent with utility \bar{u} .¹²

$$\gamma^{\bar{u}} = -(y + wl) \frac{\bar{u}_{cc}}{\bar{u}_c} = \hat{\gamma}$$

Hence, in a heterogeneous economy, the estimator $\hat{\gamma}(y + wl)$ can be interpreted as a weighted average of the risk aversions of agents in the marginal group at (y, w, l) .¹³ The marginal group consists of those who change their labor supply in response to perturbations in w and y . The preferences of those who are constrained and cannot make such changes could be quite different.

(e) Extension 2: Extensive Labor Supply Decisions

Many individuals are unable to choose the number of hours they work. They face the narrower choice of either working for a fixed number of hours or not working at all.¹⁴

¹²In general, $\gamma^{\bar{u}}$ is *not* equal to the weighted average of each type's γ because coefficients of relative risk aversion do not aggregate linearly.

¹³Note that we can observe many different marginal groups by estimating $\hat{\gamma}$ at various levels of (y, w, l) .

¹⁴The purpose of this section is to show that inferences about curvature can be made even when changes in labor supply are lumpy. The nature of the lumpiness itself (e.g. whether there are restrictions on hours worked in a week or weeks worked in a year) is not important.

To model these extensive labor supply decisions, assume that the agent makes a binary decision to work and supply 1 unit of labor or not work. As above, let y denote unearned income and w the 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

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

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

$$\theta(y, w) = \int_0^{\hat{\psi}(y, w)} f(\psi) d\psi \tag{9}$$

It follows that

$$-\frac{\partial \theta / \partial y}{\partial \theta / \partial w} = \frac{u_c(y) - u_c(y + w)}{u_c(y + w)} \tag{10}$$

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. In a world 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:¹⁵

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

Under this assumption, (10) implies

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

Solving for γ yields

$$\gamma = \frac{\log[1 - \frac{\varepsilon_{\theta,y} w}{\varepsilon_{\theta,w} y}]}{\log[1 + \frac{w}{y}]} \quad (11)$$

Finally, a model of unemployment analogous to that above can be used to derive an

¹⁵If $\gamma(c)$ 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 $c \in [y, y + w]$.

estimator for γ when utility is not additive:

$$\gamma = \frac{\log[1 - \frac{\varepsilon_{\theta,y} w}{\varepsilon_{\theta,w} y}]}{\log[(1 - \frac{\Delta c}{c})(1 + \frac{w}{y})]} \quad (12)$$

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

This concludes the theoretical portion of the paper. The next section implements the formulae derived above using existing estimates of labor supply elasticities.

2 Empirical Implementation

A large number of studies have estimated “static” wage and income elasticities using exogenous variation in unearned income and wages such as tax changes, cross-sectional differences, or lottery winnings.¹⁶ The traditional literature, summarized by Pencavel (1986) and Blundell and MaCurdy (1999), defines labor supply as hours worked or work participation. More recently, Feldstein (1995) emphasizes that hours worked is only one component of labor supply and that other margins such as effort or training might adjust as well. When a multi-dimensional definition of labor supply is incorporated into the model above, (5) still obtains except that the elasticity ratio $\frac{\varepsilon_{l,y}}{\varepsilon_{l^c,w}}$ is replaced by $\frac{\varepsilon_{LI,y}}{\varepsilon_{LI^c,1-\tau}}$, where LI is labor income and $1 - \tau$ the net-of-tax rate. This result follows the lines of Feldstein (1999), who shows

¹⁶Of course, these estimates are obtained from a world in which consumers make dynamic choices. However, a wide variety of empirical strategies have been designed to identify precisely the relevant elasticities for a static model. Moreover, the calibration arguments given below suggest that even if the estimates are not very precise, the main result of the paper – that γ is quite low – still goes through.

that the elasticity of taxable labor income with respect to the net-of-tax rate captures all margins on which taxable income can be adjusted. The reason is that the relative prices of each mechanism of adjustment remain fixed when tax rates change.

Table 1 presents a set of income and wage elasticities for three definitions of labor supply: (A) Hours worked, (B) Participation, and (C) Earned income.¹⁷ To get a sense of the plausible range for γ that is consistent with labor supply behavior, I include elasticity estimates for a wide range of groups, such as prime age males, married women, retired individuals, and low income families. In addition, I consider studies with many different approaches to estimating the elasticities, such as welfare reforms, tax changes, cross-sectional variation, and lottery payouts. The inclusion of a diverse set of studies yields a substantial amount of variation in the elasticities, ranging from 0.035 to 1.0 for the compensated wage elasticity and -0.08 to -0.3 for the income elasticity. In general, elasticity estimates for groups who are not as attached to the labor force (married women and older individuals) tend to be higher than the elasticity estimates for groups with greater labor force attachment (prime age men). Despite this variation in the estimates, the implied range for γ ends up being quite tight, largely because the *ratios* of the unearned income to wage elasticities are actually fairly stable across studies.

To estimate γ , I first consider the additive utility case and apply (5) and (11). Column

¹⁷In part (C) of the table, it is important to distinguish between *taxable* labor income and total labor income. In the calculation of γ , we are interested in changes in *total* earned income, irrespective of the form of compensation. The measure of income used in the table (AGI) may not capture forms of compensation that are not reported on tax returns such as office perks.

(6) of Table 1 reports estimates of γ at the average value of $\frac{y}{w}$ and l in each study.¹⁸ Rather than attempting to discern the coefficient of relative risk aversion at specific income levels, I group all the estimates together to identify a representative agent’s average γ across a broad range of incomes. The (unweighted) mean γ is 1.0 using hours, 0.29 using participation, and 1.19 using earned income elasticities. The overall mean estimate of $\gamma = 1.01$, which implies that a 10% increase in consumption reduces the marginal utility of consumption by 10%.

Despite the lack of a consensus among labor economists about the magnitudes of income and wage elasticities, every study except one implies $\gamma < 1.5$. This similarity of the estimates of γ despite the use of different methodologies, definitions of labor supply, and groups of the population is striking. The following calibration exercise sheds some light on the source of this consensus. Rewriting the Slutsky equation for compensated labor supply (2) in terms of elasticities yields

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

Most studies of labor supply find that uncompensated labor supply curves are weakly upward sloping ($\varepsilon_{l,w} \geq 0$). This places a lower bound on $\varepsilon_{l^c,w}$ of $-\frac{lw}{y}\varepsilon_{l,y}$ which implies using (5) that

$$\gamma < 1 + \frac{y}{wl}$$

¹⁸Following Hausman (1985), y is defined as “virtual” unearned income to account for the progressivity of the U.S. tax system. Since the earned income estimates combine different studies, they are evaluated at $\frac{y}{wl} = \frac{1}{4}$, which reflects the median value of unearned to earned income in the US (see below).

The ratio of unearned income to earned income varies across the population, but in the aggregate it equals the ratio of capital income to labor income, which is $\frac{1}{2}$ in the U.S. This places an upper bound of $\gamma = 1.5$ for a representative agent whose utility is an income-weighted average of individual utilities. Since capital income is highly concentrated, if we are interested in the curvature of an hours-weighted average of utilities, the relevant value of $\frac{y}{wl}$ is much lower; a reasonable estimate is $\frac{1}{4}$.¹⁹ In this case, $\varepsilon_{l,w} > 0$ implies $\gamma < 1.25$. In fact, to generate $\gamma > 2$ with an income elasticity of -0.1 , we must have $\varepsilon_{l,w} < -0.15$. Almost every existing study of labor supply finds $\varepsilon_{l,w} > -0.15$.²⁰ These calibration results are important because they show that the estimates of risk aversion derived in this paper are unaffected by the ongoing debate among labor economists about appropriate methods of estimating labor supply elasticities.²¹

Taking complementarity between labor and consumption into account does not change these results significantly. Column (7) of Table 1 reports estimates that account for the degree of complementarity implied by an unemployment consumption drop of $\frac{\Delta c}{c} = 10\%$. This adjustment increases the average estimate of γ to 1.24. As noted above, if some of this drop in consumption reflects inadequate consumption smoothing across states rather

¹⁹Tabulations by the US Census Bureau (1999, Table E) adjusted for the progressivity of the income tax indicate that $\frac{y}{wl} \approx \frac{1}{4}$ for the median family in the U.S., which has an income of approximately \$40,000.

²⁰Pencavel (1986), Blundell and MaCurdy (1999), and Gruber and Saez (2000) summarize more than sixty studies with an array of methodologies. All find uncompensated wage elasticities greater than -0.15 .

²¹For instance, one may be concerned that existing empirical strategies do not properly identify “static” elasticities because income and wage shocks may be perceived as transitory. However, the results of this paper still go through if one agrees that a permanent upward shift in an agent’s wage profile is unlikely to sharply reduce his lifetime labor supply.

than optimal behavior, the observed $\frac{\Delta c}{c}$ overstates the degree of complementarity between consumption and labor. In this case, the estimates in column (7) give an *upper* bound for γ . In fact, generating a mean estimate of $\gamma > 2$ requires $\varepsilon_{c_1, l_1} = \frac{\Delta c}{c} > 0.25$, which seems implausibly large. Therefore, the band for γ remains narrow even though our estimate of $\varepsilon_{u_c, l}$ may be imprecise and possibly biased by myopia or inadequate planning for the risk of job loss. The estimate of $\gamma < 1.25$ thus relies primarily on the assumption that agents properly maximize utility in making static labor-leisure choices, and is robust to deviations from optimal behavior in more complicated domains.

3 Conclusion

This paper has shown that labor supply elasticities can be used to place a tight bound on risk aversion in an expected utility framework without making assumptions about the probability distributions that govern gambles or agents' beliefs in risky situations. The key result is that risk aversion is directly related to the ratio of the income elasticity of labor supply to the price elasticity, holding fixed the degree of complementarity between consumption and leisure. The degree of complementarity can in turn be estimated from consumption choices when labor supply is stochastic.

Using a broad array of existing estimates of income and wage elasticities, I find an (unweighted) mean estimate of $\gamma \approx 1$ and show that the estimates of almost every study of labor supply imply $\gamma < 1.25$. The estimate of γ changes by at most 0.25 over the

range of plausible values for the complementarity parameter. Consequently, we are able to bound γ quite tightly despite the lack of consensus about the magnitudes of labor supply elasticities and the limited amount of evidence on the complementarity parameter. The intuition underlying the low coefficient of relative risk aversion is clear when one thinks about the labor supply behavior we would observe if γ were high. If γ were high, the marginal utility of money would deteriorate quickly with income, and labor supply would diminish significantly as income rises. But existing empirical evidence strongly rejects this hypothesis.

Two caveats deserve mention: First, this estimate applies only to the range of incomes over which the relevant elasticities are estimated (roughly \$10,000-\$200,000). Second, the estimate only reflects the curvature of the utility function for agents who are able to freely adjust their labor supply and consumption decisions in response to wage and income changes. Many people do not have such flexibility because of prior commitments, and it would be useful to identify the curvature of their utilities over wealth.²²

²²Chetty (2002) uses data on unemployment durations to show that the coefficient of relative risk aversion is an order of magnitude larger when individuals are constrained by prior commitments.

References

- Auten, Gerald and Robert Carroll, "The Effect of Income Taxes on Household Behavior," *Review of Economics and Statistics*, 81 (1999), 681-693.
- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics*, CXII (1997), 537-580.
- Becker, Gary, "A Theory of the Allocation of Time," *Economic Journal*, 75 (1965), 493-517.
- Blundell, Richard and Thomas MaCurdy, "Labor Supply: A Review of Alternative Approaches," in Ashenfelter, Orley and David Card, eds., *Handbook of Labor Economics* 3 (Amsterdam: North-Holland, 1999).
- Browning, Martin and Timothy Crossley, "Unemployment Insurance Benefit Levels and Consumption Changes," *Journal of Public Economics*, 80 (2001), 1-23.
- Chetty, Raj. "Consumption Commitments, Unemployment Durations, and Local Risk Aversion," Harvard University Mimeograph (2002).
- Dreze, Jacques, *Essays on Economic Decisions Under Uncertainty* (Cambridge: Cambridge Univ. Press, 1987).
- Eissa, Nada and Hoynes, Hillary, "The Earned Income Tax Credit and the Labor Supply of Married Couples," NBER Working Paper 6856 (1998).
- Feldstein, Martin, "The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act," *Journal of Political Economy*, 103 (1995), 551-572.
- Feldstein, Martin, "Tax Avoidance and the Deadweight Loss of the Income Tax," *Review of Economics and Statistics*, 81 (1999), 674-680.
- Friedberg, Leoria, "The Labor Supply Effects of the Social Security Earnings Test," *Review of Economics and Statistics*, 82 (2000), 48-63.
- Gertner, Robert, "Game Shows and Economic Behavior: Risk-Taking on 'Card Sharks'," *Quarterly Journal of Economics*, 108 (1993), 507-22.
- Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 87 (1997), 192-205.
- Gruber, Jonathan and Emmanuel Saez, "The Elasticity of Taxable Income: Evidence and Implications," NBER Working Paper 7512 (2000).
- Hausman, Jerry, "Taxes and Labor Supply," in Feldstein, Martin and Alan Auerbach, eds., *Handbook of Public Economics* 1 (Amsterdam: North Holland, 1985).
- Imbens, Guido, Donald B. Rubin, and Bruce I. Sacerdote, "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players," *American Economic Review*, 91 (2001), 778-794.
- Kahneman, Daniel, Peter P. Wakker, and Rakesh Sarin, "Back to Bentham? Explorations of Experienced Utility," *Quarterly Journal of Economics*, CXII (1997), 375-405.

- Kandel, Shmuel and Robert F. Stambaugh, "Asset Returns and Intertemporal Preferences," *Journal of Monetary Economics*, 27 (1991), 39-71.
- Kocherlakota, Narayana, "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature*, 24 (1996), 42-71.
- MaCurdy, Thomas, David Green, and Harry Paarsch, "Assessing Empirical Approaches for Analyzing Taxes and Labor Supply," *Journal of Human Resources*, 25 (1990), 415-90.
- Machina, Mark, "Choice Under Uncertainty: Problems Solved and Unsolved," *Journal of Economic Perspectives*, 1 (1987), 121-154.
- Mehra, Rajnish and Edward C. Prescott, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15 (1985), 145-161.
- Metrick, Andrew, "A Natural Experiment in Jeopardy!," *American Economic Review*, 85 (1995), 240-253.
- Pencavel, John, "Labor Supply of Men: A Survey" in Ashenfelter, Orley and Richard Layard, eds, *Handbook of Labor Economics 1* (Amsterdam: North Holland, 1986).
- Prelec, Drazen, "The Probability Weighting Function," *Econometrica*, 66 (1998), 497-527.
- Rabin, Matthew, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68 (2000), 1281-1292.
- Saez, Emmanuel, "The Effect of Marginal Tax Rates on Income: A Panel Study of 'Bracket Creep'," NBER Working Paper 7367 (1999).
- Szpiro, George, "Measuring Risk Aversion: An Alternative Approach," *The Review of Economics and Statistics*, 68 (1986), 156-159.
- Tversky, Amos and Craig Fox, "Weighting Risk and Uncertainty," *Psychological Review*, 102 (1994), 269-283.
- U.S. Census Bureau, Current Population Reports P60-209. Money Income in the United States (1999).
- Wolf, Charles and Larry Pohlman, "The Recovery of Risk Preferences from Actual Choices," *Econometrica*, 51 (1983), 843-850.
- Wu, George and Richard Gonzalez, "Curvature of the Probability Weighting Function," *Management Science*, 42 (1996), 1676-1690.

Appendix

(a) Curvature of utility over unearned income

The Envelope theorem implies that

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

and it follows that

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

Recognizing that $\partial c/\partial y = 1 + w\partial l/\partial y$, we obtain

$$\gamma^y = \gamma\varepsilon_{c,y}$$

where $\varepsilon_{c,y}$ denotes the income elasticity of consumption. Finally, observe that

$$\varepsilon_{c,y} = \frac{y + wl\varepsilon_{l,y}}{y + wl}$$

where $\varepsilon_{l,y}$ is the income elasticity of labor supply.

Note that $\varepsilon_{c,y} < 1$ implies $\gamma^y < \gamma^c$. If utility over consumption is not very curved ($\gamma = 1$), utility over wealth must be even less curved.

(b) Heterogeneity in Preferences

Assume that the first J types supply exactly l units of labor when they have unearned income y and wage w . The aggregate income and compensated wage effects at (y, w, l) are

$$\begin{aligned}\partial l/\partial y &= \sum_{i=1}^J \alpha_i \partial l^i/\partial y = \sum_{i=1}^J \alpha_i \frac{wu_{cc}^i}{w^2u_{cc}^i + u_{ll}^i} \\ \partial l^c/\partial w &= \sum_{i=1}^J \alpha_i \partial l^{c,i}/\partial w = \sum_{i=1}^J \alpha_i \frac{wu_c^i}{w^2u_{cc}^i + u_{ll}^i}\end{aligned}$$

which implies that the estimator in (5) is

$$\hat{\gamma}(y + wl) = -\frac{y + wl}{w} \frac{\partial l/\partial y}{\partial l^c/\partial w} = -\frac{y + wl}{w} \frac{\sum_{i=1}^J \alpha_i \frac{wu_{cc}^i}{w^2u_{cc}^i + u_{ll}^i}}{\sum_{i=1}^J \alpha_i \frac{u_c^i}{w^2u_{cc}^i + u_{ll}^i}}$$

It follows that

$$\gamma^{\bar{u}} = -(y + wl) \frac{\bar{u}_{cc}}{\bar{u}_c} = -(y + wl) \frac{\sum_{i=1}^J \omega_i u_{cc}^i}{\sum_{i=1}^J \omega_i u_c^i} = \hat{\gamma}$$

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

Study (1)	Sample (2)	Identification (3)	Income Elasticity (4)	Compensated Wage Elasticity ^a (5)	γ Additive (6)	γ $\Delta c/c=0.1$ (7)
<i>A. Hours</i>						
Blundell and MaCurdy (1999) ^b	Men	Various	-0.120	0.567	1.06	1.21
MaCurdy et al (1990)	Married Men	Cross Section	-0.010	0.035	1.47	1.68
Eissa and Hoynes (1998)	Married Men, Inc < 30K	EITC Expansions	-0.030	0.192	0.88	1.01
	Married Women, Inc < 30K	EITC Expansions	-0.040	0.088	0.64	0.98
Friedberg (2000)	Older Men (63-71)	Soc. Sec. Earnings Test	-0.297	0.545	0.93	1.23
Average					1.00	1.22
<i>B. Participation^c</i>						
Eissa and Hoynes (1998)	Married Men, Inc < 30K	EITC Expansions	-0.008	0.033	0.44	0.50
	Married Women, Inc < 30K	EITC Expansions	-0.038	0.288	0.15	0.45
Average					0.29	0.48
<i>C. Earned Income^d</i>						
Imbens et al (2001)	Lottery Players in MA	Lottery Winnings	-0.110			
Feldstein (1995)	Married, Inc > 30K	TRA 1986		1.040	0.53	0.60
Auten and Carroll (1997)	Single and Married, Inc>15K	TRA 1986		0.660	0.83	0.95
Saez (1999)	NBER Tax Panel 1979-81	Tax Bracket Creep		0.250	2.20	2.51
Average					1.19	1.36
Overall Average					1.01	1.24

^aIn part (C), this column gives the elasticity of earned income with respect to the net-of-tax rate

^bThis row uses an average of the 20 elasticities reported in Blundell and MaCurdy (1999) and assumes $y/wl=1/4$

^cParticipation elasticities assume CRRA utility

^dSince studies on earned income do not estimate income elasticities, I use the Imbens et. al. estimate in each case