

Raj Chetty: Summary of Research

My research has focused on theoretical and empirical issues in the design of tax and social insurance programs. My broad objective has been to analyze new models of economic behavior that better match empirical evidence (e.g. models of risk preferences or corporate behavior), and study the implications of these models for government policy. My work can be grouped into two categories: (1) risk preferences and social insurance and (2) behavioral responses to taxation.

Social Insurance. How much should the government spend on redistributive tax and social insurance programs such as unemployment insurance, disability insurance, welfare, and health insurance? The paper “**A General Formula for the Optimal Level of Social Insurance**” (Journal of Public Economics, 2006) provides one method of answering this question. Under some weak regularity conditions, I show that three parameters determine the optimal level of unemployment benefits in a dynamic model with arbitrary choice variables and constraints: (1) the consumption drop associated with unemployment, (2) the elasticity of unemployment durations with respect to benefits, and (3) the coefficient of relative risk aversion (γ). The formula for optimal benefits is a simple, empirically implementable rule that is robust to the underlying structure of the model of household behavior. This rule can be easily adapted to analyze the optimal design of other social insurance programs (e.g. health, disability).

While there are several estimates of the consumption drop and elasticity of duration in the literature, the value of the coefficient of relative risk aversion (γ) remains disputed. In “**A New Method of Estimating Risk Aversion**” (American Economic Review, 2006), I show that existing evidence on labor supply behavior and consumption-labor complementarity together place a tight upper bound on γ in the standard expected utility model. Using labor supply elasticity estimates from a large set of existing studies of labor supply, I find a mean estimate of $\gamma = 1$. I then show that generating $\gamma > 2$ would require that wage increases cause sharper reductions in labor supply than estimated in any of the studies. Intuitively, if the marginal utility of consumption diminishes rapidly with income (i.e., if γ is high), individuals should start working less as wages rise. In practice, higher wages lead to increased labor supply, leading to the upper bound on γ . The bound of $\gamma < 2$ implies that pure diminishing marginal utility itself does not provide significant justification for social insurance in the canonical expected-utility framework.

An important assumption in deriving this bound is that individuals consume a single composite commodity. The existence of a composite commodity requires that individuals can adjust all consumption goods freely in response to shocks such as unemployment. In practice, however, consumption of goods such as housing and cars may not be adjusted in response to small shocks because of adjustment costs. In “**Consumption Commitments and Risk Preferences**” (Quarterly Journal of Economics, 2007), Adam Szeidl and I explore the implications of such commitments for risk preferences. We first establish using CEX and PSID data that the average household devotes 50-66% of its budget to commitment goods. We then characterize risk preferences in a two good expected utility model with commitments. Commitments amplify risk aversion with respect to moderate-stake shocks and create a motive to take large-payoff gambles. The model can thus help resolve two basic puzzles in expected utility theory: the discrepancy between moderate-stake and large-stake risk aversion and lottery playing by insurance buyers. We apply the model to analyze issues such as optimal social insurance. For instance, we give an example showing that amplified

moderate stake risk aversion due to commitments can result in a significantly higher optimal benefit rate for unemployment insurance than calibrations based on the conventional one-good model.

In the working paper **“Marriage, Housing, and Portfolio Choice: A Test of Grossman-Laroque,”** Szeidl and I test whether adjustment costs in consumption affect choice under uncertainty in practice using data on portfolio choice. A key testable prediction of the commitments model is that recent home buyers (who are in the middle of their (S,s) adjustment bands) should be more risk averse than individuals who have drifted near the edges of their (S,s) bands over time. We show using a set of control groups that the timing of changes in marital status can be used to isolate exogenous variation in home tenure conditional on age and wealth. Using the timing of marital events as instruments, we find that the average investor reallocates \$1,500 from safe assets to stocks per year spent in a house. The estimates imply that recent home purchasers are twice as risk averse as individuals who have lived in their house for ten years, indicating that adjustment costs in consumption have quantitatively important effects on risk preferences.

Our first paper on commitments and risk preference studied a model with a single agent and a simple stochastic structure without persistent uncertainty. To characterize the effects of commitments on optimal policy, it is helpful to analyze a general dynamic model with heterogeneous agents, where aggregate consumption dynamics are smooth. This is the objective of our working paper **“Consumption Commitments: Neoclassical Foundations for Habit Formation.”** In this paper, Szeidl and I show that aggregating the commitments model over a heterogeneous population of agents implies dynamics identical to those of a representative consumer economy with habit formation utility. Hence, consumption commitments can be viewed as non-psychological micro-foundations for “habit.” Szeidl and I are currently working on characterizing the optimal path and level of social insurance benefits in the commitments and habit models. Intuitively, commitments create a force toward providing more insurance for short-term, moderate-stake shocks relative to long-term welfare programs. This force does not emerge in standard models of habit, where the marginal welfare cost of a shock rises monotonically with its size. Hence, while commitments and habit may be equivalent descriptions of a macroeconomy at a positive level, they could have very different policy implications.

In related work, I have explored the implications of my research on the optimal level of social insurance benefits for developing economies. Empirical studies such as Townsend (1994) have found that income shocks do not cause large consumption fluctuations in developing economies. A common view is that if consumption is smooth, the benefits of additional (social) insurance must be small. In **“Consumption Smoothing and the Welfare Consequences of Social Insurance in Developing Economies”** (Journal of Public Economics, 2006), Adam Looney and I show that this conclusion may be incorrect. Intuitively, there are two scenarios under which consumption paths can be smooth. The first is a situation where financial and insurance markets are good and risk aversion is low. In this case, social insurance is not very valuable, because it would simply crowd out private insurance. The second case is a situation where insurance markets are incomplete and risk aversion is very high. In this case, individuals may rely on very costly consumption-smoothing mechanisms (such as taking children out of school) in order to maintain consumption (e.g. because of subsistence constraints). If this is the scenario that describes developing economies, insurance could be quite valuable despite the smoothness of observed consumption.

Looney and I provide evidence that the latter scenario is empirically relevant in **“Income Risk and the Benefits of Social Insurance: Evidence from Indonesia and the United States”** (NBER East Asian Seminar 16, 2007). Using panel surveys of households in Indonesia

and the United States, we find that food consumption falls by approximately ten percent when individuals become unemployed in both countries. However, in contrast with households in the U.S., Indonesians use costly methods such as reducing human capital investment to smooth consumption. They also experience large reductions in consumption of staple goods such as rice. These results point to the view that many Indonesians undertake costly efforts to maintain smooth consumption paths because of high risk aversion.

The projects described above all focus on evaluating the consumption-smoothing benefits of social insurance. The tradeoff in providing such social insurance is that these programs also create efficiency costs by distorting economic incentives. For example, there is a large literature showing that increases in unemployment insurance (UI) benefits raise unemployment durations. This result has been interpreted as evidence of “moral hazard” – a behavioral response to distorted marginal incentives to search. In the paper **“Moral Hazard vs. Liquidity and Optimal Unemployment Insurance,”** (Journal of Political Economy, 2008) I question this interpretation by showing that unemployment benefits also raise durations through a “liquidity” effect for households who cannot smooth consumption perfectly. I document the empirical importance of the liquidity effect in two ways. First, state-level increases in unemployment benefits have larger effects on durations for households that are likely to be liquidity constrained (e.g., those with low assets). Second, lump-sum severance payments significantly increase durations among constrained households. Together, the empirical estimates imply that 60% of the effect of unemployment benefits on durations is due to the liquidity effect rather than moral hazard. To evaluate the welfare implications of this finding, I derive a new formula for the optimal level of unemployment benefits in terms of the ratio of the liquidity effect relative to moral hazard. Relative to earlier approaches to optimal benefits, this formula – which uses revealed preference to infer the value of insurance – has two advantages: (1) it requires data only on unemployment durations and (2) it does not rely on state-independence or a specific parametric form of utility. Implementing the formula using the empirical estimates implies that the optimal unemployment benefit level exceeds 50% of the previous wage. The revealed preference approach proposed in this paper could be helpful in understanding the welfare consequences of other types of insurance such as health insurance, an issue that I have begun to explore in more recent work.

A weakness of the US-based evidence on liquidity vs. moral hazard in unemployment is that it relies on cross-sectional variation in lump sum income grants that may be endogenous. In addition, the US data does not permit a sharp characterization of the type of model that describes household behavior because of limitations in sample size and available variation. These issues motivated a paper (joint with David Card and Andrea Weber) entitled **“Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market”** (Quarterly Journal of Economics, 2007). This paper presents new tests of the permanent income hypothesis and other widely used dynamic models of household behavior using data from the labor market. We estimate the “excess sensitivity” of search behavior to cash-on-hand using sharp discontinuities in eligibility for severance pay and extended unemployment insurance (UI) benefits in Austria. Analyzing data for over one-half million job losers, we first obtain three empirical findings: (1) a lump-sum severance payment equal to two months of earnings reduces the job-finding rate by 8-12% on average; (2) an extension of the potential duration of UI benefits from 20 weeks to 30 weeks similarly lowers job-finding rates in the first 20 weeks of search by 6-10%; and (3) increases in the duration of search induced by the two programs have little or no effect on subsequent job quality, as measured by wages or the exit rate from the next job. We then use a simple search theoretic model to show how these estimates can be used to distinguish between several dynamic

models of household behavior, and develop a metric that can be used to calibrate such models to match our empirical findings. Our empirical findings are inconsistent with a simple permanent income model, as well as myopic “rule of thumb” behavior. The data point instead toward a model such as buffer-stock behavior (e.g. Deaton 1991), where forward looking agents have limited ability smooth income fluctuations. This results imply that temporary income support and tax rebate policies can have substantial economic effects, and provide a simple method of calibrating dynamic models when analyzing such policies in future work. For example, in the work I am doing with Szeidl on optimal social insurance with commitments, we plan to calibrate our model to match the moment estimated in the project with Card and Weber.

In related work using the Austrian data (“**The Spike at Benefit Exhaustion: Leaving the Unemployment System or Starting a New Job?**,” American Economic Review Papers and Proceedings, 2007) Card, Weber, and I revisit the literature on the “spike” in unemployment exit rates at the expiration of benefits. This sharp surge in the hazard rate is widely interpreted as evidence that recipients are waiting until their benefits run out to return to work, leading to a substantial efficiency loss. We find that the way unemployment spells are measured has a large effect on the magnitude of the spike at exhaustion, both in existing studies and in our Austrian data. Spikes are typically much smaller when spell length is measured by the time to next job than when it is defined by the time spent on the unemployment system. In Austria, the exit rate from registered unemployment rises by over 200% at the expiration of benefits, while the hazard rate of re-employment rises by only 20%. The difference between the two measures arises because many individuals drop off the unemployment registry once their benefits expire without returning to work. The modest spike in re-employment rates implies that fewer than 1% of jobless spells have an ending date that is manipulated to coincide with the end of UI benefits. Hence, the re-timing of job starts to coincide with benefit exhaustion is quantitatively a less important behavioral response to the provision of UI benefits than the smooth reduction in search effort that occurs throughout the spell.

Taxation. How do changes in tax policies affect corporate behavior? This question has traditionally been analyzed in neoclassical models of production. However, a large recent literature on “real options” has emphasized that models with non-convex adjustments costs best describe the lumpy dynamics of firm behavior empirically. In the paper “**Interest Rates, Irreversibility, and Backward-Bending Investment**” (Review of Economic Studies, 2007), I characterize the effects of exogenous changes in interest rates (induced e.g. by tax or monetary policies) in such a model. In particular, I study the effect of interest rates on investment in an environment where firms make irreversible investments with uncertain payoffs. In this setting, exogenous changes in the interest rate affect both the cost of capital and the cost of delaying investment to acquire information. These two forces combine to generate an aggregate investment demand curve that is a backward-bending function of the interest rate. At low rates, increasing the interest rate *raises* investment by increasing the cost of delay. The non-monotonic relationship between the interest rate and investment predicted by the model could help explain why empirical studies typically find little systematic relationship between the user cost of capital and investment.

I have also studied the effects of tax policies on firm behavior empirically. Emmanuel Saez and I examined the effects of the 2003 dividend tax cut on corporate behavior in the papers “**Dividend Taxes and Corporate Behavior: Evidence from the 2003 Dividend Tax Cut**” (Quarterly Journal of Economics, 2005) and “**The Effects of the 2003 Dividend Tax Cut on Corporate Behavior: Interpreting the Evidence**” (American Economic Review Papers and Proceedings, 2006). In these papers, we document a twenty percent increase in dividend amounts following

the 2003 tax cut along both the extensive and intensive margins. We find that the response to the tax cut was strongest in firms with strong principals whose tax incentives changed (those with large taxable institutional owners or independent directors with large share holdings), and in firms where agents had stronger incentives to respond (high share ownership and low options ownership among top executives). These results suggest that principal-agent issues play an important role in corporate responses to taxation.

The benchmark “old view” and “new view” models of dividend taxation do not fit these empirical findings, because they assume that firms simply maximize profits irrespective of ownership structure. In **“An Agency Theory of Dividend Taxation”**, Saez and I propose a simple alternative theory of dividend taxation in which managers and shareholders have conflicting interests, and show that it can explain the evidence. Using this agency model, we develop an empirically implementable formula for the efficiency cost of dividend taxation. The key determinant of the efficiency cost is the nature of private contracting. If the contract between shareholders and the manager is second-best efficient, deadweight burden follows the standard Harberger formula and is second-order (small) despite the pre-existing distortion of over-investment by the manager. If the contract is second-best inefficient – as is likely when firms are owned by diffuse shareholders because of incentives to free-ride when monitoring managers – dividend taxation generates a first-order (large) efficiency cost. An illustrative calibration of the formula using empirical estimates from the 2003 dividend tax reform in the U.S. suggests that the efficiency cost of raising the dividend tax rate could be close to the amount of revenue raised.

Saez and I have also explored the effects of the 2003 dividend tax cut on asset prices jointly with Joseph Rosenberg in **“The Effects of Taxes on Market Responses to Dividend Announcements and Payments: What Can we Learn from the 2003 Dividend Tax Cut?”** (in A. Auerbach and J. Slemrod, eds., 2007). This paper investigates the effects of capital gains and dividend tax rates on excess returns around announcements of dividend increases and ex-dividend days for U.S. corporations. Several studies have examined this issue and reached contradictory results using regression-based evidence in short windows around tax reforms. Using less parametric methods, we show that the relationship between tax rates and ex-day and announcement day premia is very sensitive to specification and sample period choices. Strong year-to-year fluctuations in the ex-day and announcement day premia greatly reduce statistical power, making it impossible to credibly detect even large tax effects. We conclude based on this meta-analysis that the important non-tax factors affecting these premia must be understood before credible progress can be made in evaluating the role of taxation in asset price responses.

I have also been interested in the taxation of individuals. A central assumption in existing models of taxation (and social insurance) is that individuals optimize fully relative to the incentives created by government policies. In **“Salience and Taxation: Theory and Evidence,”** (joint with Adam Looney and Kory Kroft), we present evidence that consumers under react to taxes that are not salient and characterize the welfare consequences of tax policies when agents make such optimization errors. The empirical evidence is based on two complementary strategies. First, we conducted an experiment at a grocery store posting tax inclusive prices for 750 products subject to sales tax for a three week period. Scanner data show that this intervention reduced demand for the treated products by 8 percent. Second, we find that state-level increases in excise taxes (which are included in posted prices) reduce alcohol consumption significantly more than increases in sales taxes (which are added at the register and are hence less salient). We develop simple, empirically implementable formulas for the incidence and efficiency costs of taxation that account for salience effects as well as other optimization errors. Contrary to conventional wisdom, the formulas imply

that the economic incidence of a tax depends on its statutory incidence and that a tax can create deadweight loss even if it induces no change in demand. Our approach to welfare analysis yields robust results because it does not require specification of a positive theory for why agents fail to optimize with respect to tax policies.

In a followup project, I am investigating the effects of salience and information in income taxation (joint with Emmanuel Saez). In collaboration with H&R Block, we implemented an experiment involving 40,000 Earned Income Tax Credit (EITC) claimants in 140 H&R Block offices in Chicago from January-April 2007. We are studying whether providing simple information on the structure of the EITC to eligible tax filers during tax preparation affects labor supply. For example, tax preparers will tell an individual who is earning \$10 per hour and is in the phase-in region of the EITC that she should actually think of her wage as \$14 per hour. We will then test whether this simple information about marginal tax rates affects individuals' earnings in the subsequent year (relative to a control group). If we find that information matters, we plan to study theoretically the optimal design of income tax policies when not all agents are fully informed about the incentives created by the tax code.

My most recent work focuses on understanding the long-run effects of income taxation. Existing micro-econometric studies of labor supply typically focus on identifying the effects of small changes in marginal income tax rates within a 2-3 year period. This approach could lead to attenuated estimates of the long-run effect of tax policies in an environment where agents have limited attention or face other frictions (e.g. costs in switching jobs or consumption plans). To evaluate whether these studies are likely to recover "fundamental" wage elasticities relevant for long-run tax policy, I am calibrating the individual utility cost (in dollar terms) of *ignoring* the tax changes studied in the literature using NBER's TaxSim calculator. For instance, we ask, "How much would an individual pay to learn that his marginal income tax rate has changed from 28 to 33%?" Based on the results in our earlier work, we expect that the cost of not reoptimizing in response to such tax changes is very small: if one is already optimizing relative to the original tax rate (28%), the welfare gain from re-optimizing relative to the 5% rate change is second-order. The implication is that existing studies may not be very informative about the long run effects of taxation, when individuals are likely to become aware of the general characteristics of the new tax regime (e.g. high taxes in Europe vs. lower taxes in the U.S.). We hope that this critique of the existing empirical literature will motivate researchers to find new methods of estimating the long-run effects of taxes and other policies.

To this end, I am currently working on a project which develops a new method of estimating the long-run effects of permanent changes in tax rates. The project (joint with John Friedman, Anders Fredriksson, and Luigi Pistaferri) starts from the observation that many of the differences between behavior in the short and long run arise from fixed costs (costs of switching jobs, consumption plans, acquiring information, etc.). All of these fixed costs induce (S,s) behavior: individuals are less responsive in the short-run because they adjust behavior only when the difference between their current decision and the optimal one exceeds a certain threshold. A key prediction of (S,s) models is that we can learn about the long run by comparing the *short run* effect of large and small tax changes. Large changes immediately move more people across the thresholds of the (S,s) band, exactly as would occur through idiosyncratic shocks in the long run. To formalize this intuition, we have derived an estimator for the long run effect of taxation that relies purely on short-run effects of tax changes of different sizes. We are implementing this strategy empirically using administrative tax records spanning 1976-2005 for the universe of individuals in Denmark. The quality and size of the Danish data, coupled with the substantial variation in tax rates over the past thirty years,

offers the promise of relatively precise long-run elasticity estimates.

One measure of “labor supply” used in the projects described above is reported taxable income, which has become a benchmark measure to calculate the deadweight loss of taxation since Feldstein (1999). In the working paper “**Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance,**” I reevaluate the taxable income elasticity as a measure of deadweight loss when individuals can shift income to evade or avoid taxes. In many situations, part of the cost of evasion and avoidance reflects a transfer to another agent in the economy. I show that in such situations, excess burden depends on a weighted average of the taxable income and total earned income elasticities, with the weight determined by the marginal resource cost of income shifting. This generalized formula implies that the efficiency cost of taxing high income individuals is not necessarily large despite evidence that their reported incomes are highly sensitive to tax rates. This result does not affect the use of the taxable income measure as a way to measure behavioral responses in the projects described above, but it suggests that additional research may be needed to fully understand the normative implications of the empirical estimates.