Research Statement

Erzo F.P. Luttmer
October 2009

Overview

Within my field of public economics, I focus my research on income redistribution and social insurance. The level of redistribution and social insurance available to individuals varies dramatically across time, across countries, and across localities and demographic groups within countries. My research papers seek to understand what drives the demand for redistribution and social insurance, and how the implementation of redistribution and social insurance programs can be improved. This work can be grouped into three related areas.

<u>1. The demand for income redistribution and social insurance.</u> Ten of my completed papers fall into this broad category [2, 5, 7, 9, 10, 11, 13, 15, 16, 17].[1] Understanding how much and why individuals value income redistribution and social insurance programs informs ongoing policy debates regarding the generosity and design of these programs. The fact that income redistribution and social insurance are not transacted in a marketplace has, however, presented a profound challenge to social scientists because the lack of a marketplace for redistribution and social insurance programs precludes the use of standard revealed preference techniques to study the demand for these programs. To overcome this challenge, I infer preferences for resource transfers to the unfortunate by relying on a variety of alternative methods. These methods include examining self-reported preferences, voting behavior, benefit take-up decisions, and charitable giving in randomized experiments.

From my papers that focus on redistribution, a common theme emerges; preferences for redistribution are not determined solely by financial self-interest but also by social effects [2, 7, 9, 10, 15, 16]. Social effects, sometimes called interdependent preferences, occur when individuals behave differently or have different outcomes because they are influenced directly by the behavior of others around them. These papers contribute not only to our understanding of the demand for redistribution but also to a broader empirical literature on the role of social effects. This literature has helped increase awareness among economists and policymakers of the limits to treating individuals as atomistically maximizing own self-interest.

With respect to social insurance, my papers examine circumstances that affect the benefits that people anticipate from social insurance [5, 11, 13, 17]. Social insurance is more beneficial if its payouts are received in situations when people value these payouts more. In the case of health insurance, a subject of fervent national debate, the event that triggers an insurance payout – a negative health shock – is also likely to directly affect how much people care about changes in their level of consumption. This recognition motivates my research on the effect of health on the marginal utility of consumption [11, 17], and the results of this research inform the policy debate on the optimal generosity of health insurance benefits. Similarly, social insurance is more valuable if individuals lack access to other channels of insurance, and this observation motivates my research examining the extent to which religious organizations insure individuals against economic shocks [5, 13].

[1] The numbers refer to the papers as numbered in my Curriculum Vitae.

2. The supply of social insurance and income redistribution programs. Six of my completed papers fall into this category [4, 6, 8, 12, 18, 20]. Just as the demand for social insurance and income redistribution reflects the benefits of these programs, one can view the marginal costs of these programs as determining their supply. Some of my research in this area is empirical and estimates the misallocation cost associated with income redistribution programs that use price controls [4, 8]. These misallocation costs had not been estimated before and these estimates allow for more informed policy decisions on the use of price controls as a means to redistribute income. My other papers in this area are theoretical in nature and show how certain design features can yield the same amount of redistribution with fewer economically costly distortions [6, 12, 18, 20]. This information is valuable for policymakers interested in the design or redesign of redistribution programs.

3. Cognitive limitations and the design of income redistribution and social insurance programs. I have only started doing research on cognitive limitations, also sometimes referred to as behavioral decision making, in the past three years and currently have four papers in this area [1, 3, 14, 21]. Income redistribution and social insurance programs often involve sophisticated incentive schedules that aim to minimize the economic distortions they impose. However, if incentive schedules are so complex that individuals have trouble understanding them, individuals will make suboptimal choices. Hence, it is important to take the cognitive limitations of individuals into account when designing public policy. Some of my research in this area examines individuals' response to and understanding of the complex implicit incentives in the Social Security system. Other research in this area examines cognition costs in other settings such as lending in peer-to-peer markets and decision making in elections.

My work to date uncovers several key findings: social effects influence the demand for redistribution, it is possible to achieve the same level of redistribution we currently observe but in less costly ways, and cognitive limitations affect the efficiency with which programs provide insurance and redistribution. Taken together, my research provides evidence that recognizing and addressing social effects and cognitive limitations will likely improve the design of programs that provide redistribution and social insurance.

I plan to continue working in each of these three areas. I am especially excited to pursue further research on social effects and behavioral decision making, subjects that I believe are relatively poorly understood in economics but are nonetheless very important for many areas of public policy, including areas unrelated to income redistribution and social insurance programs. Below, I describe each of my three research areas in more detail. The descriptions focus on completed papers, but I also mention current and future research projects.

1. The Demand for Income Redistribution and Social Insurance

My research on the demand for income redistribution and social insurance examines the factors that affect how much individuals value these programs. Below, I first describe my six papers on the demand for redistribution, explaining how and why these papers conclude that social effects are an important determinant of the demand for redistribution. I then discuss my four papers on social insurance. Although social effects may affect the demand for social insurance, these four papers focus on other determinants of the demand for social insurance.

The Role of Social Effects in the Demand for Income Redistribution

My research on the demand for income redistribution finds that social effects are important determinants of preferences for income redistribution. Hence, these papers can also be viewed as contributing to the literature on social effects. Classical economists recognized the importance of social effects. For example, in *The Theory of Moral Sentiments* (1759) Adam Smith wrote extensively about how individuals' preferences and judgments are affected by others around them. While other social sciences, sociology in particular, view individuals' actions in the context of a social environment, modern mainstream economics tends to view individuals' behavior and outcomes as driven solely by self-interest and the market environment. This approach has been fruitful in many ways, and substantially eases the tasks of empirical and theoretical investigation. However, due to an array of findings in psychology, economics, and other disciplines, the limits of modeling individual behavior devoid of a social context are increasingly recognized. A better understanding of the role of social effects on individuals' outcomes and decision making is extremely valuable for policymakers, who need to anticipate how individuals will react to their policies.

The Effect of Social Comparisons on Subjective Well-Being

Do individuals experience a drop in well-being or utility when the standard of living of others around them increases? Modern mainstream economists by and large do not consider such social effects because they typically assume that utility is solely a function of one's own consumption. However, this assumption departs from an earlier tradition among classical economists, and from a broader view among other social scientists and the general public, that individuals care about their *relative* consumption in addition to their absolute level of consumption. The distinction between the absolute and relative formulations of utility has important implications for tax and expenditure policy, as numerous authors have shown. In particular, if utility depends on relative consumption, then one person's increase in consumption imposes a negative externality on others because it lowers the relative consumption of others. In this case, income redistribution programs that, as a side effect, discourage consumption are less distortionary than previously thought because they internalize the negative externality of consumption on others. The distinction between the absolute and relative formulations of utility is also pertinent to central policy concerns, such as the welfare implications of residential sorting by income, and to the debate about whether the poverty line should be absolute (a fixed consumption basket) or relative (a fraction of mean or median income).

My paper "**Neighbors as Negatives: Relative Earnings and Well-Being**," published in the *Quarterly Journal of Economics*, provides evidence that utility depends in part on relative position. Using panel data on individuals' self-reported happiness matched to Census data on local earnings, I find that, controlling for a host of individual characteristics including income, higher local earnings are associated with lower levels of happiness: an increase in neighbors' earnings (holding own income constant) and a similarly sized decrease in own income (holding neighbors' earnings constant) each lead to a reduction in happiness of about the same order of magnitude.

This paper makes three main contributions to the literature studying the effect of relative position on subjective well-being. First, it takes seriously the concern that living in an affluent area might affect one's *definition* of happiness even if it does not affect one's true or experienced well-being. To investigate this concern, I use other outcome measures that are less prone to definition changes in response to neighbors' earnings and conclude that such definition shifts do not likely drive the results. Second, the paper examines whether the inverse relationship between happiness and neighbors' earnings might be spurious due to omitted individual or local characteristics. The panel nature of the data enables me to run specifications with individual fixed effects while its detailed geographical information allows for the inclusion of local housing prices and state fixed effects. I also use a predicted measure of local earnings (based on the local industry occupation mix and national earnings by industry and occupation), which filters out many local earnings shocks caused by unobserved local factors that might simultaneously influence happiness. The results hold up under these specifications, reducing the concern that they are due to omitted variable bias. Third, the paper offers suggestive evidence concerning the mechanism mediating the negative relationship between neighbors' earnings and happiness. The results are stronger for people who socialize more with neighbors than for those who socialize more with friends outside the neighborhood. The paper's findings indicate that social preferences incorporating relative income concerns drive the negative association between neighbors' earnings and own well-being.

This research has stimulated considerable further study, at least as measured by citations. Currently, it is the most cited paper of all 40 papers published in the *Quarterly Journal of Economics* in 2005. This research also aroused significant interest in the popular press – it is described or mentioned in a dozen articles across various papers and magazines, including *The Wall Street Journal* and *The New Yorker*.

Network Effects and Welfare Cultures: The Role of Social Effects in Welfare Participation

The enrollment dynamics of welfare assistance is another realm in which social scientists have long suspected an important role of social effects. Welfare recipiency rates rose sharply in the late 1960s and early 1970s when benefits were made more generous, but rates remained roughly constant in the subsequent 20 years even as benefits became less generous. Recipiency rates have declined sharply since the mid 1990s, with the decline starting just before the Clinton Administration ended "welfare as we know it" with the 1996 welfare reform bill. Are these welfare recipiency rates solely determined by economic factors such as the generosity of benefits and labor market prospects, or did welfare recipiency rates stay high in the 1980s partly because stigma of welfare receipt had fallen? Similarly, was the rapid decline in welfare recipiency rates after the mid 1990s solely due to the changes in financial incentives caused by welfare reform, or did welfare recipients also respond to pressure from their peers who were now less likely to be on welfare?

It is impossible to convincingly answer these questions by examining time-series data on welfare use, benefit levels, and labor market conditions because many other unobserved aggregate fluctuations might affect welfare recipiency. Rather than using time-series data, Marianne Bertrand, Sendhil Mullainathan, and I estimate the effects of peer welfare use on own welfare use using individual-level data from the decennial Census. The resulting paper "**Network Effects and Welfare Cultures"** was published in the *Quarterly Journal of Economics*. The empirical challenge is to credibly distinguish any social or network effect of peers' welfare use from omitted variable bias or correlated effects, i.e., similar behavior caused by peers' facing the same unobserved determinants of welfare use. In order to do so, we exploit the fact that non-native English speakers are likely to have more contact with, and to be more influenced by, speakers of their own language in their area. Women from certain language groups have higher welfare recipiency rates than do women from other language groups. Thus, the average level of welfare recipiency in the own language group is a proxy for the attitudes of that group toward welfare and information from that group on how to navigate the welfare system. The attitudes of and the information from the own language group will have a greater influence on welfare recipiency if one has more contact with members of the own language group. We use the fraction of the population in the metropolitan area that belongs to the own language group as a proxy for contact availability of the own language group.

We find that living in an area with more people of the own language group increases welfare use more for women from a high welfare-using language group than for those from a low welfare-using language group. We take this as evidence that network effects are an important determinant of welfare use. The advantage of our identification strategy is that we can control for the most plausible omitted variables: we include a separate effect for each metropolitan area and for each language group, and we directly control for contact availability. To measure the importance of network effects, we calculate how much network effects magnify the response to changes in welfare policy. We find that if 100 people join the welfare rolls as a direct consequence of a more generous welfare policy, an additional 20 will join because of the network effect that the original 100 persons will exert on them.

How Culture Shapes Preferences for Income Redistribution

Whereas the "Network Effects" paper highlights the role of culture in the decision to participate in welfare, my research with Monica Singhal examines the effect of culture on preferences for redistribution. Though scholars have long believed that culture is an important determinant of preferences for redistribution, it has been hard to convincingly demonstrate this influence because of the difficulty of separating the effect of culture from the effect of the economic and institutional environment ("context"). In our paper "**Culture, Context, and the Taste for Redistribution**," Monica Singhal and I separate the effect of culture from the effect of context by examining the preferences of immigrants in the 32 countries of the European Social Survey. We relate immigrants' preferences for redistribution to the average preference in their birth countries, controlling extensively for individual characteristics and country-of-residence fixed effects. We find a strong positive relationship. This cultural effect is larger for non-voters, those with shorter tenure in the country of residence, and those who move to countries with a large number of immigrants from their own birth countries. Immigrants from countries with a higher preference for redistribution are also more likely to vote for a more pro-redistribution political party. The effect of culture also persists strongly into the second generation, confirming that cultural effects are passed on within immigrant groups. This research was recently featured in an article in *The Economist*.

Social Effects and Support for Welfare Spending

Differences in culture, however, may not be the only reason why the generosity of welfare benefits varies tremendously across international and state lines. The welfare system in most western European countries is considerably more generous than the U.S. welfare system. Even within the U.S., welfare benefit levels vary considerably, with cash benefit levels often differing by a factor of three across states. Moreover, there is a striking correlation between the generosity of redistribution and the demographic homogeneity of the population. Compared to many western European countries, the United States has a racially, ethnically, and culturally heterogeneous population, and engages in relatively little redistribution. Furthermore, the southern states in the U.S. tend to have both the highest racial heterogeneity and the lowest welfare benefit levels. My paper "**Group Loyalty and the Taste for Redistribution**," published in the *Journal of Political Economy*, examines whether population homogeneity causally affects welfare generosity and, if so, whether social effects drive this relationship. Since there are many unobserved variables that could confound the relationship between population homogeneity and generosity of redistribution at the national or state level, I explore how individuals' self-reported support for welfare spending depends on the level of welfare recipiency and the racial composition of welfare recipients in their neighborhoods. Because it is conceivable that individuals who dislike welfare might disproportionally choose to live in neighborhoods with low levels of welfare recipiency (causing a spurious correlation between support for welfare and local welfare recipiency), I use each metropolitan area's pattern of population location to predict for each individual the expected local racial composition and level of welfare recipiency. For example, if I know that an individual is white with a certain income level living in Boston in 1980, I use Census information to determine in which Boston Census tracts this individual likely resided, and then take a weighted average of welfare recipiency in those Census tracts.

I find that both black and non-black individuals reduce their support for welfare spending as the rate of welfare recipiency in their neighborhood increases. However, black individuals reduce their support for welfare spending most in response to non-black welfare recipients, while non-black individuals decrease their support for welfare spending most in response to black welfare recipients. This finding is evidence of "racial group loyalty": a willingness to be more generous to members of one's own racial group than towards members of other racial groups. One might wonder whether this result is truly a social effect. Perhaps non-black individuals living in areas with many non-black welfare recipients are more supportive of welfare because they themselves are more likely to be or become a welfare recipient. However, this does not seem to be the case. The estimate of group loyalty is stronger among richer persons, i.e., those who are least likely to use or have used welfare themselves. My estimates of group loyalty explain about 30% of the variation in benefit levels across states, suggesting that social effects are important determinants of welfare generosity.

Using Randomized Experiments to Revisit the Role of Race in Support for Redistribution

The issue of racial bias also figured prominently in the debate about the adequacy and speed of the response to Hurricane Katrina victims. As we could not think of a credible way to estimate the direct effect of the race of Katrina victims on the government response, Christina Fong and I designed a randomized experiment that investigates how the race of Katrina victims affects charitable giving to Katrina victims. The results of this experiment are described in the paper "**What Determines Giving to Hurricane Katrina Victims? Experimental Evidence on Racial Group Loyalty**," published in the *American Economic Journal: Applied Economics*. We experimentally manipulated individuals' perceptions of Katrina victims from a small city (Slidell, LA or Biloxi, MS) using a nationally representative sample of about 1100 Americans. We then asked each respondent to split \$100 between him or herself and the local chapter of

Habitat for Humanity in the small city in question. The respondent's choice was implemented randomly 10% of the time.

We find highly significant effects of the race, income, and worthiness treatments on beliefs about the race, income, and worthiness of the recipients of aid from Habitat for Humanity in each city. In other words, the experiment succeeded in altering respondents' perceptions of Katrina victims. This study allows us to measure social effects because narrowly self-interested individuals would not condition their giving on perceived characteristics of recipients. We find that only the income treatment has a significant effect on giving: respondents are more generous when Katrina victims are perceived to be poorer. Perhaps surprisingly, the perceived race and worthiness of recipients do not significantly affect giving on average by either black or non-black respondents. Thus, our study suggests that to the extent that the private responses to Hurricane Katrina were slow and inadequate, race is unlikely the explanation. Although race does not affect giving on average, respondents who report feeling "close" to their racial or ethnic group, i.e., those who self-identify with their racial or ethnic group, give substantially more when victims are of the same race rather than of another race. However, respondents who do not feel close to their group give less when victims are of the same race. Hence, we find evidence of racial group loyalty only in subpopulations of respondents defined by their subjective identification with their racial or ethnic group.

Because the "Katrina" paper examines preferences for redistribution toward a particular population, namely individuals hit by a natural disaster, it seemed desirable to run a similar experiment that examines preferences for redistribution to welfare recipients and poor people more generally. This randomized experiment is described in my joint paper with Christina Fong "**Do Race and Fairness Matter in Generosity? Evidence from a Nationally Representative Charity Experiment**," which is a revise-and-resubmit at the *Journal of Public Economics*. Approximately 1000 respondents, who are representative of the U.S. adult population, participated in this experiment. The respondents viewed a short audiovisual presentation about one of two local charities that serve the poor in Tuscaloosa, AL. We used this presentation to randomly manipulate the perceived race and worthiness of the charity recipients. We then asked each respondent to split $100 between him or herself and the charity, and implemented this choice for 10% of respondents. The experiment yields three main findings. First, we find significant racial bias in perceptions of worthiness: respondents rate recipients of their own racial group as more worthy than recipients of other racial groups. Second, respondents give significantly more when the recipients are described as more worthy. These findings may lead one to expect that respondents would also give more generously when shown pictures of recipients belonging to their own racial group. However, our third result shows that this is not the case; despite our successfully manipulating perceptions of race and respondents' rating their own racial group as more deserving, giving does not respond significantly to recipient race. Thus, although our respondents do seem to rate ingroup members as more worthy, they appear to overcome this bias when it comes to giving. These findings complement those of the Katrina experiment. In the Katrina experiment, recipients were generally seen as worthy (because their need arose due to a natural disaster) but varied widely in their incomes, allowing us to detect effects of recipient income on giving, but not effects of perceived worthiness on giving. By contrast, in the Tuscaloosa experiment, recipients were all poor, but the perceived reasons for their poverty varied. As a result, we could estimate the effects of the perceived worthiness of the recipients on giving, but not the effect of recipient income on giving. The Tuscaloosa experiment confirms our finding of the Katrina experiment that on average there is no significant racial bias in giving. This result is not necessarily inconsistent with the finding of racial bias in my "racial group loyalty" paper, which is based mainly on data from the 1980s. The role of racial group loyalty may have diminished in the last two decades to the point that we cannot detect it in our experiments, which were run in 2006.

The Demand for Social Insurance

My papers on social insurance examine the factors that affect the benefits of social insurance. The fundamental reason why insurance is valuable is that people pay into insurance when their marginal utility of consumption is relatively low and that they receive insurance payouts when their marginal utility is relatively high. Economists typically assume that the event that triggers an insurance payout increases marginal utility only because the event causes a negative wealth or income shock to the individual. While this may be reasonable in many cases, this assumption is problematic in the case of health insurance because a negative health shock is likely to have a direct impact on the marginal utility of consumption even if all medical expenses are fully reimbursed. Yet, there is little evidence on whether individuals value additional consumption more in good health or in poor health.

The optimal design of health insurance, however, depends crucially on the direct effect of health on the marginal utility of consumption; the optimal fraction of medical expenses reimbursed by insurance is higher if individuals value consumption more in poor health. As Amy Finkelstein, Matthew Notowidigdo, and I explain in our paper "**Approaches to Estimating the Health State Dependence of the Utility Function**," published in the *American Economic Review, Papers and Proceedings*, the various methods that have been used to answer this question all face considerable empirical challenges, and the resulting evidence on state dependence is mixed. In our paper "**What Good Is Wealth without Health? The Effect of Health on the Marginal Utility of Consumption**," a revise-and-resubmit at the *Quarterly Journal of Economics*, we develop a simple model in which the impact of health on the marginal utility of consumption can be estimated from data on permanent income, health, and utility proxies. We estimate the model using the Health and Retirement Study's panel data on the elderly and near elderly, and proxy for utility with measures of subjective well-being. We find robust evidence that the marginal utility of consumption declines as health deteriorates. Our central estimate is that an increase in the number of chronic diseases by one standard deviation (or 0.63 diseases) is associated with an 11 percent decline in the marginal utility of consumption relative to this marginal utility when the individual has no chronic diseases. Some simple, illustrative calculations suggest that state dependence of this magnitude substantially reduces the optimal fraction of health expenditure covered by insurance.

Just as state dependence of the utility function is an important determinant of the benefit of social insurance, so is the extent to which individuals can insure themselves against adverse events through other means. In "**Insuring Consumption and Happiness through Religious Organizations**," Rajeev Dehejia, Thomas DeLeire, and I examine the degree to which religious organizations provide informal insurance in the United States. In this paper, published in the *Journal of Public Economics*, we examine how individuals' consumption and self-reported happiness respond to income shocks. If an individual is fully insured, consumption and happiness should not respond to own income shocks. Thus, the degree to which consumption and happiness are insensitive to movements in own income is a measure of effective insurance against income shocks. Using this measure, we find that individuals who are members of religious organizations are better insured against income shocks. This finding cannot be explained by differences in observable characteristics between members and non-members of religious organizations, which suggests that it is the religious organization itself or members of the congregation that provide the additional insurance to fellow members. These results help explain why the demand for government-provided social insurance tends to be lower among religious individuals and in more religious countries. This paper received considerable attention in the popular media including articles in the *New York Times* and the *Boston Globe*.

A related paper with the same coauthors along with Joshua Mitchell finds that religious organizations can help mitigate the long-term effects of growing up in a disadvantaged environment. This paper is titled "**The Role of Religious and Social Organizations in the Lives of Disadvantaged Youth**," and is published in the conference volume *An Economic Perspective on the Problems of Disadvantaged Youth*.

In a new project, "**Insurance Benefits from Progressive Taxes and Transfers**," Hilary Hoynes and I will examine the insurance value of state tax and transfer systems. This project is motivated by the observation that mobility across state lines hinders states' ability to redistribute income because states with more progressive tax and transfer systems will disproportionally attract low-income individuals and repel high-income individuals. However, state spending on transfer programs seems not to have declined despite the increased residential mobility of citizens. In this project, we are exploring a possible explanation for this puzzle: states' tax and transfer systems offer insurance against economic shocks, and high-income individuals also value this insurance. We have developed a methodology to measure the insurance value of state tax and transfer systems and are implementing this methodology using the Panel Survey on Income Dynamics. Our results will describe the importance of the insurance value across demographic groups, across states, and over time.

2. The Supply Side: Costs of Income Redistribution and Social Insurance Programs

What are the efficiency costs of income redistribution and social insurance programs? I have conducted a range of empirical projects to help answer this question. I look specifically at redistribution through price controls in housing and labor markets, and at the effects of the generosity of disability insurance on moral hazard. My theoretical work in this area examines how income redistribution programs can be designed to minimize costly distortions of behavior for any amount of redistribution.

How Redistribution through Price Controls Can Lead to Unintended Efficiency Costs

The paper "**The Misallocation of Housing under Rent Control**," written with Edward Glaeser and published as a lead article in the *American Economic Review*, examines how a redistribution program, rent control in New York City, generates unintended efficiency costs. While most research on rent control has focused on the negative effect of rent control on the supply and quality of housing, we examine how rent control interferes with how the existing stock of renters is allocated to the existing stock of apartments. A central theorem of economics is that (under suitable conditions) the price mechanism will result in a Pareto efficient allocation of goods: goods end up with consumers such that it is impossible to make someone better off without making others worse off by allocating goods to consumers in a different way. Loosely speaking, the price mechanism ensures that each good goes to the consumer who values that good most. Economists have long suspected that when the price mechanism is handcuffed, as it is under rent control, goods will be misallocated.

Our paper demonstrates empirically how severely rent control misallocates the existing stock of housing to the existing renters. To infer misallocation, one needs to find a pair of renters who would both be better off if they switched apartments (possibly with side payments).[2] However, we only observe a limited set of apartment characteristics in the Census data we use. Instead of observing each renter's preferences, we only observe some of their demographic characteristics. In light of these observations, it may seem impossible to credibly infer if a pair of renters would be better off switching apartments. We overcome this challenge by comparing population distributions rather than individual renters. A clarifying example follows. All else equal, one would expect a single person to value a large apartment less than the value placed on the same apartment by a family with children. Thus in a free rental market, singles would generally rent cheaper and smaller apartments while families would rent larger and more expensive apartments, though some singles with an idiosyncratic taste for large apartments could rent larger apartments than some families who don't happen to care for space. We measure the degree to which the demand for apartment size of these two different types of renters overlaps in cities without rent control, and assume that this overlap in preferences remains the same in New York. Thus, if elsewhere the median single person prefers an apartment that is the same size as the one rented by the 30th percentile of families, we assume that this is also the case in New York. Following this example, if we observe that in New York, the median single person lives in an apartment that is the same size as the one rented by the 60th percentile of families, we infer misallocation. In the paper, we develop a sophisticated statistical methodology based on this insight to estimate the misallocation of rental apartments in New York. We find that approximately 20% of renters in New York City would need to move in order to achieve an efficient allocation of housing.

Much as rent control misallocates apartments, the minimum wage has the potential to misallocate jobs. Specifically, if the minimum wage leads to an excess of workers interested in a limited number of minimum wage jobs, do these jobs go to those workers who value these jobs the most instead of to those who would rather not work if the wage had been slightly lower? I examine this question in "**Does the Minimum Wage Cause Inefficient Rationing?**," published in the *B.E. Journal of Economic Analysis & Policy*. The 1990-1991 Federal minimum wage hike had a significant effect on states with low state minimum wages but little or no effect on states that already had a high state minimum wage. I exploit this cross-state variation to estimate whether, in states where the minimum wage became more

[2] Of course, there could be more complicated switches. A slightly more complicated switch would have A get B's apartment, B get C's apartment, and C get A's apartment. In the paper, we allow for arbitrarily complicated switches.

binding, a higher fraction of minimum wage workers had characteristics that indicate that these workers really value a job. The results show no evidence of inefficient job rationing due to the minimum wage. If anything, a more binding minimum wage causes the jobs to go to those who value them most.

Some Evidence on the Efficiency Cost of Generous Disability Insurance

An unintended cost of social insurance programs is that individuals distort their behavior in response to the incentives provided by these programs. Although the existence of this phenomenon, also known as moral hazard, is well-known, how important moral hazard is in the context of disability insurance is subject to debate. In a new research project, "**Moral Hazard in Dutch Disability Receipt: Evidence from an Age Discontinuity in Disability Eligibility Reviews**," Lex Borghans, Anne Gielen, and I are assessing the extent of moral hazard in disability insurance by exploiting the discontinuity introduced by the 1993 Dutch Disability Insurance Reform Act. The act implemented disability eligibility reviews based on stricter standards for all recipients younger than 45 years on the date the act was passed. Using administrative data on the universe of disability recipients in the Netherlands, we are implementing a regression discontinuity design to estimate moral hazard in disability receipt.

Can Alternative Tax Systems Reduce the Efficiency Cost of Income Taxation?

In addition to my empirical work on the efficiency costs of social insurance and redistribution programs, I have also written two theoretical papers on this topic. In the paper "**Can Income Redistribution Be Privatized?**," I examine whether competition in "the production of income redistribution" can lead to efficiency gains, just as competition in the production of ordinary goods is generally believed to produce more efficient outcomes than can be achieved through central planning. The income tax and transfer system leads to efficiency losses because it drives a wedge between the private and social return to work. Are these efficiency losses an unavoidable consequence of income redistribution? Seminal work by James Mirrlees has answered this question: redistribution is costly because individuals have private information about their levels of ability. If the government could observe ability levels, it could redistribute without any efficiency cost by using taxes based on ability but not on income. However, in a world with private information, knowledge about ability levels can be extracted only by distorting individuals' behavior. Such distortions make redistribution costly. Mirrlees derived the optimal non-linear income tax schedule that extracts private information about abilities and implements the desired redistribution with the least distortions.

There may be other instruments than the optimal income tax schedule to extract private information about ability. Ignoring such instruments, such as the type of work assignment, leads to avoidable inefficiencies in the tax schedule. I explore whether there are ways in which these avoidable efficiency losses can be reduced or eliminated. Ronald Coase argued that when the social return to an action differs from the private return due to an externality, the efficient action will still be taken as long as the person choosing the action can costlessly bargain with those affected by the externality. My paper shows that Coase's insight carries over to income taxation. In the case of taxation, the tax creates a wedge between the private and social return to work, i.e., it creates a fiscal externality. If the party receiving that externality can bargain with the individual paying the tax, they will reach the second-best efficient outcome, fully exploiting efficiency gains from any observable components of the person's ability and from the use of alternative instruments for discriminating between individuals of different abilities. In theory, I show, we could set up a decentralized system for income redistribution that creates incentives for private bargaining to eliminate the avoidable efficiency losses associated with income taxation.

The income tax schedule is one of many incentive schedules faced by individuals. Incentive schedules usually balance two competing objectives. For example, income tax schedules are the result of a tradeoff between achieving a fairer income distribution and the efficiency cost of distorting work incentives. In the paper, "**Schedule Selection by Agents: from Phone Plans to Tax Tables**," Richard Zeckhauser and I prove that one can generally achieve the same objective (say, the amount of income redistribution) at a lower efficiency cost by requiring agents with private information to select from a menu of incentive schedules. This efficiency gain will be achieved if, and only if, agents will receive further private information after selecting the incentive schedule but before taking the action that determines where on the incentive schedule they end up. This information structure is relevant in many applications. For example, such a structure occurs in the private sector when an individual chooses from a menu consisting of two cell phone plans, one with a low monthly fee and a high rate per minute and one with a high monthly fee and a low rate per minute. The individual might know that she generally will be a heavy cell phone user (her initial private information), but before she chooses her phone plan, she is unlikely to know whether some lengthy calls will be desirable that month (her subsequent private information). We develop the theory underlying optimal menus of non-linear schedules and prove that a menu of schedules exists that offers a strict first-order interim Pareto improvement over the optimal single non-linear schedule.

The possibility of offering a menu of schedules warrants more consideration by policymakers because schedule selection may also offer substantial welfare gains in practice. A simulation based on U.S. earnings data shows that moving to a tax system that allows individuals to choose their tax schedule increases social welfare by the same amount as would occur from a 4.0 percent windfall gain in the government budget (or about $600 per filer per year). The resulting reduction in distortions accounts for about two thirds of the increase in social welfare while the remainder comes from an increase in redistribution.

How Much Protection Against Natural Disasters Should the Government Provide?

Natural disasters represent a significant risk faced by citizens. What are the appropriate roles for private and public investment in protecting against such risks? I have written two theoretical papers, both joint with Carolyn Kousky and Richard Zeckhauser, that address this question. This research shows that because private (housing) capital is more likely to locate in better-protected areas, the social value of additional protection will increase with the level of protection provided. The resulting outcome may be an inefficient equilibrium: few people locate to an at-risk area because the level of protection is low, and given the low number of people it seems not worthwhile to increase the level of protection. The paper "**Private Investment and Government Protection**," published in the *Journal of Risk and Uncertainty*, explains that rather than taking the number of people as given, as is common in cost-benefit analyses, policymakers should take the private response to protection into consideration. They may then find that it is worthwhile to increase the level of protection precisely because the amount of private capital benefiting from protection will be greater at the higher level of protection. A second paper, "**Permits to Elicit Information**," provides a mechanism to solve this coordination problem and is published in the *Proceedings of the National Tax Association*.

3. Cognition Costs and the Design of Public Policy

My third research area focuses on the interaction between the cognitive limitations of individuals and the design of public policy. In particular, income redistribution and social insurance programs often involve sophisticated incentive schedules that aim to minimize the economic distortions of these programs. However, if incentive schedules are so complex that individuals have trouble understanding them, individuals will make suboptimal choices. Hence, it is important to take the cognitive limitations of individuals into account in the design of public policy. I have only started research in this area in the last three years, but I am excited to expand my research in this direction in the future, particularly given the increasing attention devoted to decision-making research at the Kennedy School and in the policy arena.

How Cognitive Limitations of Voters Can Affect Election Outcomes

One pertinent application of cognition costs is in the realm of ballot design. If voters are fully rational and have negligible cognition costs, then ballot layout should not affect voting behavior. In the paper "**Who Misvotes? The Effect of Differential Cognition Costs on Election Outcomes**," published in the *American Economic Journal: Economic Policy*, Kelly Shue and I use the quasi-random variation in layouts of 157 types of ballots used in the 2003 California Recall Election to study how different ballot designs led to systematic deviations in vote shares of candidates. We refer to votes for a candidate that are solely attributable to that candidate's favorable position on the ballot as misvotes because differences in ballot layout should not affect the vote of a fully rational voter. Studying misvotes is important for two reasons. First, and at a more fundamental level, misvotes provide evidence on human cognition: how does ballot design alter voting and which types of voters are affected most strongly? Second, misvotes can change election outcomes. Although the literature has predominantly focused on which candidate gains from misvotes, what often matters more for the outcome of an election is which of the two frontrunners loses more votes to misvotes. This point is perhaps best illustrated by the unusually large number of votes that Patrick Buchanan received on the "butterfly ballot" of Palm Beach County, Florida, during the 2000 U.S. Presidential Election. While even Buchanan himself acknowledges that probably most of the approximately 2000 extra votes that he received in Palm Beach County were due to voter mistakes, the key question is whether these 2000 misvotes came sufficiently more from would-be Al Gore voters than from would-be George Bush voters to overturn the 537 vote margin of victory in that state for George Bush.

We find that minor candidates' voteshares almost double when their names are adjacent to the names of major candidates. All else equal, voteshare gains are larger in precincts with higher percentages of poorly educated, low-income, or third-party voters. Evidence suggests that, if the Republican frontrunner Arnold Schwarzenegger and Democratic frontrunner Cruz Bustamante had been in a tie, ballot position effects would have disproportionally favored

Schwarzenegger by taking fewer votes away from him. Although in the case of the California Recall election, the ballot position effects were much smaller than Schwarzenegger's margin of victory, they can easily influence the outcome in close elections. We also explore which voting technology platforms and brands mitigate misvoting, and show that simple modifications to ballot design can significantly reduce misvoting.

Do Cognitive Limitations Affect People's Responses to the Social Security Program?

Cognitive limitations do not only cause misvotes in elections, but they can also affect how individuals respond to complicated social insurance programs. In particular, the Social Security benefit rules provide implicit labor supply incentives because future benefits generally increase with the amount of Social Security taxes paid. However, given the considerable complexity of the Social Security benefit rules, some economists have wondered whether individuals recognize this tax-benefit link and have argued that investment-based Social Security reform will improve economic efficiency by strengthening the perceived link between retirement contributions and retirement benefits. In the paper "**Labor Supply Responses to Marginal Social Security Benefits: Evidence from Discontinuities**," forthcoming in the *Journal of Public Economics*, Jeffrey Liebman, David Seif, and I estimate the effects of the marginal Social Security benefits that accrue with additional earnings on three measures of labor supply: retirement age, hours worked, and labor earnings. We develop a new approach to identifying these incentive effects by exploiting five provisions in the Social Security benefit rules that generate discontinuities in marginal benefits. An example of a Social Security rule that creates a discontinuity in incentives is that a couple has the choice of claiming either 150% of the benefits of one spouse or of both spouses claiming 100% of their own benefits. Thus, we can compare two similar individuals, except that it is optimal for only one of them to claim spousal benefits because one individual has a spouse that earns slightly more than the spouse of the other individual. The individual who claims spousal benefits receives no work incentive from Social Security because her benefits do not increase with additional work, whereas the individual claiming on her own record does have positive work incentives. We can then compare the labor supply of these individuals to estimate the effect of Social Security incentives on labor supply.

We find that individuals approaching retirement (age 52 and older) respond to the Social Security tax-benefit link on the extensive margin of their labor supply decisions. We estimate that a 10 percent increase in the net-of-tax share reduces the two-year retirement hazard by a statistically significant 2.0 percentage points from a base rate of 15 percent. The evidence with regards to labor supply responses on the intensive margin is more mixed: we estimate that the elasticity of hours with respect to the net-of-tax share is 0.4 and statistically significant, but we do not find a statistically significant earnings elasticity.

To study individuals' understanding of the Social Security benefit rules more directly, Jeffrey Liebman and I designed a 30-minute survey that was fielded to a nationally representative sample of 2500 individuals between the ages of 50 and 70. A first investigation of the results of this survey is described in the paper "**The Perception of Social Security Incentives for Labor Supply and Retirement: The Median Voter Knows More Than You'd Think**," in which we find that the majority of respondents believe that their Social Security benefits increase with labor supply, i.e., that the Social Security benefit rules provide a positive work incentive. The magnitude of this perceived incentive varies across respondents, but people generally cite an incentive that is somewhat greater than the actual figure. We also surveyed people about their understanding of various provisions in the Social Security benefit rules. We find that some of these provisions (e.g., effects of delayed benefit claiming and rules on widow benefits) are relatively well understood while others (e.g., rules on spousal benefits, provisions on which years of earnings are taken into account) are poorly understood.

An ongoing part of this project is a field experiment in which we provided detailed information about certain incentives in the Social Security benefit rules to a random subsample of working respondents aged 60 to 65. The information intervention, which took place in February/March 2009, consisted of a 10-page brochure sent by mail and a 15-minute interactive web-based tutorial. One year from the informational intervention, we will survey the complete sample to test whether the intervention affected perceptions of the incentives, retirement intentions, and realized retirement decisions. We will continue to follow the respondents for the coming five years to monitor longer-term retirement responses.

The results of this research are of considerable interest for policymakers because they inform the debate on Social Security reform. The U.S. Senate Special Committee on Aging has invited me to provide written testimony for a hearing on Social Security held in June of 2009. In that testimony, I describe which aspects of Social Security benefits the public understands relatively well, and which aspects are poorly understood. I suggest two principal ways in which misperceptions of Social Security benefit rules can be reduced. First, I offer suggestions on how to improve information dissemination to the public about these rules. Second, I suggest ways in which the more complicated rules could be simplified without changing fundamental features of Social Security. These simplifications would make it easier to for the public to understand what determines their benefits.

How Sophisticated Are Lenders in Online Peer-to-Peer Credit Markets?

Whereas my research on Social Security examines the cognitive ability of individual people to understand their implicit incentives, in many other situations the collective ability of groups of individuals to make the correct inference matters. One such setting concerns online peer-to-peer credit markets, which have grown substantially over the past couple of years. In these markets, individuals lend directly to other individuals, without the use of a financial intermediary. These transactions are facilitated by a firm, such as Prosper.com, that maintains a website where prospective borrowers can post a listing specifying the size of the loan sought, the purpose of the loan, and any other information that the prospective borrower wishes to share (e.g., a picture or a narrative). In addition, the firm runs a credit report on the prospective borrower and posts some of the information of the credit report to the borrower's listing, though the firm does not reveal the exact credit score. Prospective lenders then bid on listings, specifying how much they are willing to lend and at which minimum interest rate. If sufficient bids arrive to fund the entire amount requested, the loan is made at the interest rate determined by the bidding process. A typical loan on Prosper.com is a three-year $7000 loan at a 15% interest rate funded by several dozen lenders. Lenders generally diversify by investing relatively small amounts ($50-$500) in many loans. The firm sponsoring the online marketplace receives a service fee for each loan. In return, it collects payments from the borrowers and distributes these payments to the lenders. It also sells any loans that default to a collection agency, distributing the proceeds to the relevant lenders.

The sustainability of these alternate credit markets crucially depends on the collective ability of lenders to infer borrowers' creditworthiness. In the project "**Screening in New Credit Markets: Can Individual Lenders Infer Borrower Creditworthiness in Peer-to-Peer Lending?**," Raj Iyer, Asim Khwaja, Kelly Shue, and I examine this ability using a methodology that takes advantage of lenders' not observing a borrower's exact credit score but only seeing an aggregate credit category in the particular market considered. We find that lenders are able to utilize available information to infer up to a third of the variation in creditworthiness that is captured by a borrower's credit score. This inference is economically significant and allows individuals to lend at a rate that is 200 basis points lower for borrowers with (unobserved to lenders) better credit scores within a category. While lenders infer the most from standard "hard" information that is traditionally used by banks for screening, they also learn from the subjective or "soft" information contained in pictures and personal descriptions. We conclude that lenders in peer-to-peer markets collectively prove apt at extracting information about creditworthiness from the loan listings. Peer-to-peer markets provide some advantages over traditional lending markets in extracting information from non-traditional sources.