

(JEP.Chetty)

Raj Chetty: 2013 Clark Medal Recipient

Martin Feldstein

Raj Chetty wrote to me during his first week as a freshman at Harvard College to ask for a job as a research assistant. Although my research assistants were usually much further along in their studies, the high school essay that he sent to me – a critical comment on Robert Fogel’s Time on the Cross – was so well done that I decided to interview him. A brief discussion convinced me that he was unusually bright and would be both productive as a research assistant and fun to work with.

After a few months working with Raj as a research assistant I realized that he was quite exceptional and should be investing in his own intellectual development rather than helping me with my current statistical research. I suggested a variety of things that he might read and we could discuss together. These included not only some papers in public economics but also risk theory and statistical decision theory.

Raj completed his Harvard B.A. in 3 years, graduating summa cum laude with a thesis on interest rates and business investment that was published in the *Review of Economic Studies* (Chetty, 2007). He went on to complete his PhD at Harvard in the next three years. He then went to Berkeley as an assistant professor in 2003. Harvard lured him back in 2009.

=====

Raj Chetty is eminently deserving of being awarded the John Bates Clark medal at the age of just 33. His research has transformed the field of public economics. His work is motivated by important public policy issues in the fields of taxation, social insurance, and public spending for education. He approaches his subjects with a creative redefinition of the problems that he studies and with new empirical methods based on experimental evidence or unprecedentedly large integrated sets of data. While his work is founded on basic microeconomics, he modifies this framework to take into account behavioral and institutional considerations. It is not surprising that the American Economic Association in its announcement of the Clark medal declared that “he has established himself in a few short years as arguably the best applied micro-economist of his generation.”

Chetty is a prolific scholar. It is difficult to summarize all of Chetty’s research or even to capture the details of his most significant papers. I have therefore chosen a

---

· Martin Feldstein is Professor of Economics at Harvard University. This essay will be published in the *Journal of Public Economics* in 2014 to note Raj Chetty’s receiving the John Bates Clark medal of the American Economic Association.

selection of Chetty's important papers dealing with taxation, social insurance and education that contributed to his selection as the winner of the John Bates Clark award. These examples from different aspects of public economics indicate Chetty's combination of selecting important issues, creatively extending existing theory, and applying novel empirical methods.

In several of these studies, Chetty works with co-authors. When I refer to Chetty as the author of each study it should be understood to mean Chetty and his coauthors. Full citations of the papers with the names of all the authors appear at the end of this paper.

### Taxation

Chetty's most widely cited paper, "Salience and Taxation: Theory and Evidence" (2009) is a good example of Raj's innovative style in conceptualizing a question and bringing novel evidence to bear. He begins by posing the question of whether consumers react to the taxes that are imposed on retail products in the way assumed by standard theory. He shows that they do not. He then explores the implications of the more realistic description of their behavior for the incidence and welfare loss of taxes.

His first clever strategy for assessing how consumers reacted to taxes on the products that they buy was to compare the reaction to state taxes on beer that are levied in some states as excise taxes and therefore are built into the price that consumers see on the shelf with the taxes on beer that are levied in other states as sales taxes that are collected at the cash register after the individual has already made his purchase decision. He found that the more "salient" price effect of the excise tax had a bigger impact on the quantity of beer purchased than a comparable size sales tax collected at the cash register.

Chetty confirmed this "salience" effect with by conducting a major experiment in a large grocery store. He persuaded the store's management to post tax-inclusive prices for several hundred randomly selected products for several weeks. Using scanner data, he found that the information about the higher prices reduced demand relative control products and nearby stores where the tax was only levied at the cash register.

This novel question and the imaginative source of evidence implies that consumers are making suboptimal purchase decisions and therefore that the tax imposes a higher utility loss to the extent that the consumers ignore the taxes. The consumer also bears more of the burden of the tax because there is less of a decline in the quantity purchased. Chetty then shows how to modify the traditional incidence and deadweight loss formulae to take into account the consumers' suboptimal behavior.

Another important Chetty paper that looks beyond the traditional model of individual behavior takes into account the individual's cost of adjusting to changed tax incentives ("Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," 2012) Many economists have noted that small changes in prices or taxes do not induce changes in behavior comparable to the effects of larger changes in prices or taxes. More formally, the elasticity of response appears to be greater in response to larger tax and wage changes. Chetty explains this apparent "threshold" effect by the cost that individuals face for adjusting to such changes.

He develops a formal model that calibrates these adjustment costs and uses this expanded model to reconcile differences between micro models and macro models of labor supply. Chetty then goes beyond this application to positive economics to show how traditional calculations of deadweight loss (excess burden) can be modified to take these adjustment costs into account.

Another impressive application of expanding the basic model was Chetty's analysis of the implications of taxpayer evasion of taxes. In a basic model without tax evasion, higher marginal income tax rates create a deadweight loss in three ways: reduced labor supply (broadly defined to include lower participation and hours, less effort, the choice of lower productivity jobs, etc.); a shift from taxable cash to untaxed fringe benefits, nice offices and other favorable working conditions; and a shift of consumer spending to categories that are tax deductible (mortgage interest, local government taxes, etc.). The resulting deadweight loss can be calculated using the effect of the higher marginal tax rate on total taxable income because all three effects are the results of responses to the same marginal tax rate, therefore implying that the taxable income is a form of Hicksian composite good ("Tax Avoidance and the Deadweight Loss of the Income Tax," Feldstein, 1999).

Chetty extended this analysis to allow for the realistic possibility that taxpayers may modify their reporting in order to avoid paying their full tax liabilities and that taxpayers may incorrectly estimate the probability of getting caught for this tax evasion. In this context, Chetty shows that the taxable income elasticity is no longer a sufficient statistic with which to calculate the deadweight loss ("Is the Taxable Income Elasticity Sufficient to Calculate the Deadweight Loss? The Implications of Evasion and Avoidance," 2009)

He then goes on to show how the deadweight loss formula based on the taxable income elasticity can be modified to derive a measure of deadweight loss as a weighted average of the elasticity of taxable income and the elasticity of total earnings. The relative weights placed on these two components depends on the extent to which changes in taxable incomes are driven by responses with real resource costs (such as changes in the form of compensation and of tax-deductible spending) versus responses that have private costs to the individual but not social cost (such as under-reporting taxable income.)

Tax rules in the United States and other countries discriminate against saving by taxing the return to capital. To offset this, the United States and other countries have special rules for reducing the tax on the return to saving relative to other forms of income. In the United States, this includes the 401k employer plans (in which employees have the opportunity to exclude part of their earnings if employers deposit those funds in long-term investment accounts) and the Individual Retirement Accounts in which individuals can choose to exclude part of their earnings by depositing those funds in similar investment accounts.

Although evidence shows that employees with Individual Retirement Accounts do increase their saving in that form, it is not clear whether this represents a net increase in the employees' total saving or is either a transfer from other forms of saving or is offset by increased mortgage borrowing. This cannot be resolved based on U.S. experience because of limits on the data on total individual assets and liabilities.

To resolve this issue, Chetty reached out to data from Denmark where there are complete records on the savings, assets and liabilities of everyone in the country ("Active vs. Passive Decisions and Crowd-out in Retirement Savings Accounts: Evidence from Denmark," 2012). The Danish tax system has both of the same types of saving incentive accounts as the United States and changed the rules of these programs during the sample period of Chetty's data.

In his analysis, Chetty distinguished between "passive savers" who do not respond to any of the changes in the tax rules applicable to saving and "active savers" who do respond. Employer contributions to 401k-type plans do raise the total saving of passive savers in Denmark but changes in the incentives in IRA-type plans do not alter the savings of this "passive saver" group. Active savers do respond to saving incentives in IRA-type plans but also offset their increased saving in these accounts by reducing their net saving in other accounts. This evidence suggests that IRA type saving incentives in Denmark do not increase national saving.

Chetty's analysis in this study focused on the implications for overall national saving and not on the welfare effect of changes in the effective marginal tax rate on saving through IRA-type plans. Taxes on the return to saving create a deadweight loss by altering the relative price of current and future consumption. Since saving is only the current outlay to purchase future consumption, a tax on saving creates a deadweight loss even if it does not alter the amount of saving (Feldstein TITLE). Chetty's analysis shows that the active savers do not increase their total saving in response to more favorable IRA-type rules but shift saving to a form that delivers a higher net of tax return and therefore higher future consumption. In doing so, the more favorable IRA rules reduce the deadweight loss associated with the existing taxation of saving.

The three studies that I have summarized all deal with the effects of taxes on the behavior of individuals. Chetty shifts his focus to the behavior of firms in his paper on the response of firms to the reduced rate of tax on dividends (“Dividend and Corporate Taxation in an Agency Model of the Firm”, 2010). This paper begins by confirming that the reduction in the rate of personal income tax on dividends in 2003 led to higher dividend payouts. He then shows that the increased dividend payout is greatest in firms where senior managers and the board of directors have substantial share ownership.

Chetty then goes on to develop a model in which shareholders and managers have conflicting interest over the desirability of dividends versus retained earnings. This agency model is then used to calculate the deadweight burden caused when higher taxes on dividends cause an increase in retained earnings.

### Social Insurance

Chetty made a major contribution to the study of unemployment insurance in his paper “Moral Hazard vs. Liquidity and Optimal Unemployment Insurance “(2008). There is substantial evidence that higher unemployment insurance benefits lead to longer periods of unemployment. Chetty begins by decomposing the reasons for the increased duration into liquidity effects and moral hazard effects. He finds that increases in benefits have much bigger effects for households that are liquidity constrained. He also shows that lump sum unemployment benefits have bigger effects on liquidity-constrained households.

The optimal level of unemployment insurance benefits depends on balancing the insurance protection that benefits provide and the loss of output caused by the moral hazard of the induced increase in unemployment. The protection depends on the unemployed individual’s risk aversion. Earlier studies used the measures of risk aversion derived from portfolio decisions, which are inherently measures of long-term risk aversion. Chetty notes that the unemployed individual at the time of his unemployment typically has substantial fixed commitments – e.g., rent or mortgage payments – that cannot be changed in the short run. This in turn means that any loss of income associated with unemployment represents a larger fraction of the individual’s uncommitted income, implying that the relevant risk aversion parameter is larger than it would be when considering long-term unconstrained decisions.

Chetty then derives a formula for the optimal level of unemployment benefits that depends on the reduced form elasticity with respect to liquidity and the moral hazard elasticity. Because the implied level of risk aversion is higher than in other previous studies, Chetty concludes that the optimal level of unemployment benefits is substantially higher than previous studies found and also higher than the benefit levels that are typical in the United States.

The Chetty calculation of the optimal level of unemployment benefits assumes a structure of the unemployment insurance system that provides government financed benefits proportional to previous wages. My own research on unemployment insurance led me to favor an alternative approach of mandatory unemployment insurance saving accounts that provide liquidity when an individual is unemployed based on previous mandatory saving by that individual (Feldstein, title, year). Chetty's evidence on the importance of liquidity for the unemployed reinforces the desirability of such a structure that provides liquidity with relatively little risk of moral hazard because the unemployed individual is generally using his own previously accumulated funds.

While unemployment insurance is designed to provide cash benefits to individuals with a short-term loss of income, other programs like the Earned Income Tax Credit (EITC) are intended to provide cash to individuals whose income is low for a sustained period of time. The structure of the EITC makes benefits a function of the individual's earned income. If individuals understand the rules, they can adjust their income or the amount of income that they report, to maximize their benefit.

Chetty explores the extent to which individuals do this by using a remarkable set of administrative tax return data on all individuals who filed for EITC in the fourteen years from 1996 to 2009 ("Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of EITC on Earnings," 2012). Because of the enormous size of the data set, Chetty can study the extent to which individuals bunch their earnings at the exact income level that maximized their EITC payments.

More specifically, Chetty begins by focusing on self-employed individuals since that group has the greatest control over their reported earnings. He finds that the degree of "sharp bunching" at the benefit maximizing level of income differs substantially among geographic areas. Chetty infers that differences among areas are not random but reflects knowledge gained in the area since self-employed individuals who move to an area will increase their benefit if the extent of "sharp bunching" is higher in their new area than it was in the area from which they came.

Chetty then uses this area measure of "sharp bunching" to assess the extent to which employed individuals (in addition to the self employed) are also able to report W-2 wage incomes that come close to maximizing their benefits. He finds that they do have this ability to learn from the knowledge in their neighborhood to optimize in this way.

### Education

Spending on education is one of the most substantial and important activities of the government. State and local governments spend nearly \$1 trillion on education and local governments alone spend more than \$600 billion on primary and secondary education. Economists are interested in ways to make that educational spending

more productive and in understanding the impact of educational quality on later life outcomes.

To study these important issues, Chetty amassed and linked amazing sets of data so that he could trace the success of students from primary school to their college attendance and their earnings at age 27, the most recent year of the tax data that he could link to the primary school records.

The first of these education studies evaluated the impact of children's early primary education on their important outcomes later in life, including college attendance and incomes at age 27 ("How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project Star," 2010) Project STAR was a careful experiment in Tennessee that randomly assigned students in grades K through 3 to different classrooms with different teachers.

Chetty's first dramatic finding was that kindergarten test scores are highly correlated with adult outcomes including college attendance and earnings at age 27. His analysis also showed that students who were randomly assigned to small classes as well as students with more experienced teachers had higher earnings later in life. There is also a contagion effect in that some classes in grades K through 3 produced better lifetime performance for the class as a whole.

In a second major education study, Chetty studied whether teachers who increase their pupils test scores actually help them with achievements later in life. This is another tour de force of combining massive independent data sets to be able to answer important questions. There is no fancy economic theory but rather thoroughly convincing evidence based on bringing large amounts of relevant data to bear on the question.

Chetty starts with data on 2.5 million pupils in grades 3 through 8 linked to tax records of their parents and of themselves as adults. Also integrated into the linked data are the identities of the colleges that they later attend, their incomes, and even the places where they live as adults.

The teachers of these 2.5 million pupils are evaluated based on their "value added," i.e., the increase in test scores of these pupils in primary school. Using school district data for each pupil and taking into account each student's prior test scores, Chetty studies the effect of changing teacher assignments to show that there is little or no statistical bias in assessing the value added of individual teachers. ("Measuring the Impact of Teachers I: Developing Unbiased Estimates of Teacher Value Added," 2013)

The evidence also shows powerful effects of the high value added teachers. Those pupils who had high value added teachers are more likely to attend college, attend higher ranked colleges, live in higher SES neighborhoods and earn higher salaries as

adults. (“Measuring the Impact of Teachers II: Teacher Value Added and Student Outcomes in Adulthood,” 2013)

Perhaps the most amazing finding in this study is that replacing a teacher in the bottom five percent of the value-added distribution with an average teacher would “increase the present value of students’ lifetime income by more than \$250,000 for the average classroom in our sample.”

In short, test-score effects identify teachers that can consistently improve student performance and those improved test scores have very substantial positive effects of the students’ lifetime earnings.

### Chetty as a Teacher

Although the Clark medal is awarded on the basis of the economist’s research contribution, Chetty has made and is continuing to make a major contribution to the teaching of public economics.

As I noted earlier, several of the papers that I summarized were produced by Chetty working with small teams of researchers, some of whom were then very bright graduate students. Their names appear in the list of references at the end of this paper. Their experience creates a group of researchers who will at least aspire to follow Chetty’s methods and approach.

When Chetty came to Harvard as a professor he organized the Lab for Applied Economic Projects to encourage active collaboration among graduate students and faculty of the economics department as well as researchers in other Harvard faculties and visitors from other universities.

Chetty has also accepted the responsibility to be the editor of the Journal of Public Economics and the co-director (with Amy Finkelstein) of the Public Economics program of the National Bureau of Economic Research.

At Harvard, Chetty not only teaches public economics to both graduate students and undergraduates but has also produced very detailed lecture notes and videos that are available on line, certainly the best modern “textbook” of public economics (<http://www.rajchetty.com/index.php/lecture-videos>)

Cambridge, MA  
December 2013



## References

Raj Chetty, "Interest Rates, Irreversibility, and Backward-Bending Investment," Review of Economic Studies 74(1) 67-91, 2007

Raj Chetty, Adam Looney and Kory Kroft, "Salience and Taxation: Theory and Evidence," American Economic Review 99(4) 1145-1177 September 2009

Raj Chetty, "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," Econometrica 80(3) 969-1018, 2012

Raj Chetty, "Is the Taxable Income Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance," American Economic Journal: Economic Policy 1(2) 31-52, 2009

Raj Chetty, John Friedman, Soren Leth-Petersen, Torben Nielsen and Tore Olsen "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," Quarterly Journal of Economics, forthcoming

Raj Chetty and Emmanuel Saez, "Dividend and Corporate Taxation in an Agency Model of the Firm," American Economic Journal: Economic Policy 2(3) 1-31, 2010

Raj Chetty, "Moral Hazard vs. Liquidity and Optimal Unemployment Insurance," Journal of Political Economy, 116 (2), 173-234, 2008

Raj Chetty, John Friedman and Emmanuel Saez, "Using Differences in Knowledge Across Neighborhoods to Uncover the Impact of the EITC on Earnings," American Economic Review, forthcoming

Raj Chetty, John Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Schanzenbach, and Danny Yagan, "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR," Quarterly Journal of Economics, 126 (4) 1593-1660, 2011

Raj Chetty, John Friedman and Jonah Rockoff, "Measuring the Impact of Teachers I: Developing Unbiased Estimates of Teacher Value Added," American Economic Review, forthcoming

Raj Chetty, John Friedman and Jonah Rockoff, "Measuring the Impact of Teachers II: Teacher Value Added and Student Outcomes in Adulthood," American Economic Review, forthcoming

Raj Chetty, "Public Economics: Economics 2450a",  
<http://www.rajchetty.com/index.php/lecture-videos>

Martin Feldstein, "Tax Avoidance and the Deadweight Loss of the Income Tax," Review of Economics and Statistics, 81(4) 674-80. November 1999

Martin Feldstein, "The Welfare Cost of Capital Income Taxation," Journal of Political Economy 86(2) S29-51, 1978

Martin Feldstein and Daniel Altman, "Unemployment Insurance Saving Accounts," Tax Policy and the Economy, (21) 2006

END