Research Review

Featured Paper

Why Did So Many People Make So Many *Ex Post* Bad Decisions? The Causes of the Foreclosure Crisis

Christopher L. Foote, Kristopher S. Gerardi, and Paul S. Willen

FEDERAL RESERVE BANK OF BOSTON

RESEARCH REVIEW

ISSUE NO. 17 JANUARY 2012–JUNE 2012

Research Review provides an overview of recent work by economists and policy analysts of the research department of the Federal Reserve Bank of Boston.

Research Review is available on the web at:
http://www.bostonfed.org/economic/ResearchReview/index.htm

Earlier issues of *Research Review* in hard copy (through Issue #11) are available without charge. To order copies of back issues, please contact the Research Library:

Research Library—D
Federal Reserve Bank of Boston
600 Atlantic Avenue
Boston, MA 02210
Phone: 617.973.3397
Fax: 617.973.4221
E-mail: boston.library@bos.frb.org

Views expressed in *Research Review* are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of Boston or the Federal Reserve System. The authors appreciate receiving comments.

Executive Summaries in This Issue

Federal Reserve Bank of Boston
Research Department Papers Series

Public Policy Discussion Papers present research bearing on policy issues. These are generally written for policymakers, informed business people, academics, and the informed public. Many of these papers contain research intended for publication in professional journals.

Working Papers present statistical or technical research. These are generally written for economists and others with strong technical backgrounds and are intended for publication in professional journals.

Public Policy Briefs present analysis on topics of current interest concerning the economy. These briefs are written by Boston Fed research economists, based on briefing materials or presentations prepared by Boston Fed research staff for senior Bank executives or outside audiences.

Research Reports present research on economic and policy issues of concern to New England's state and local policy leaders and others engaged in developing and implementing public policy in the region. These reports are written by Boston Fed economists and policy analysts.

Research department papers are available online:
http://www.bostonfed.org/economic/respubs.htm

p-12-1

Effects of Credit Scores on Consumer Payment Choice

by Fumiko Hayashi and Joanna Stavins

abstract and full text: http://www.bostonfed.org/economic/ppdp/2012/ppdp1201.htm
e-mail: fumiko.hayashi@kc.frb.org, joanna.stavins@bos.frb.org

MOTIVATION FOR THE RESEARCH

Over the last decade, debit card use grew rapidly, and debit cards are now the most commonly used noncash payment method in the United States. According to the 2010 Federal Reserve Payments Study (FRPS), debit card transactions represented 35 percent of total noncash retail payments in 2009. In contrast, credit card use accounted for 20 percent of total noncash retail payments in 2009 (Federal Reserve System 2010).

The rapid growth of debit cards has stimulated several studies on consumer payment choice. Previous studies highlighted several important factors that influence consumer payment choice, such as individual consumer characteristics, transaction characteristics, payment method attributes, and the price of and/or rewards for using certain payment methods. Most of these studies did not include factors that would limit available payment methods to consumers, because very few datasets contain the information necessary to examine the effects of such factors.

The paper's main goal is to investigate the effects of credit scores on consumer payment choice, especially on debit card and credit card adoption and use. Anecdotally, a negative relationship between debit use and credit scores has been observed (Lightspeed 2009); however, it is not clear what influences this relationship.

RESEARCH APPROACH

The authors' primary data source is the 2008 and 2009 Survey of Consumer Payment Choice (SCPC), a consumer survey conducted by the Federal Reserve Bank of Boston. Because the SCPC provides only self-reported FICO credit scores, the authors also use an external credit score measure from data provided by Equifax. The Equifax credit score is closely correlated with a FICO score, a measure of creditworthiness developed by the Fair Isaac Corporation. While the authors cannot merge the SCPC and the Equifax data precisely, they extract the Equifax credit score for each finely decomposed socioeconomic group and compare that external measure with the SCPC data.

Before exploring the effects of credit scores on consumer payment choice, the authors investigate what individual characteristics affect each consumer's credit score. Although they do not have access to the full set of variables that constitute the FICO score, they do have data on several questions designed to gauge financial stress that feed into the FICO score. The authors use an ordered probit model, with the credit score "index," a variable they develop from the midpoint of each credit score range selected by survey respondents, as a dependent variable. In addition to the basic consumer demographic characteristics, such as race, age, income, and education level, the independent variables include other consumer characteristics such as household size, marital status, work status, and access to new technologies.

FICO Score and Card Use

2008

Weighted Mean Share (Percent)

FICO Ranges

2009

Weighted Mean Share (Percent)

FICO Ranges
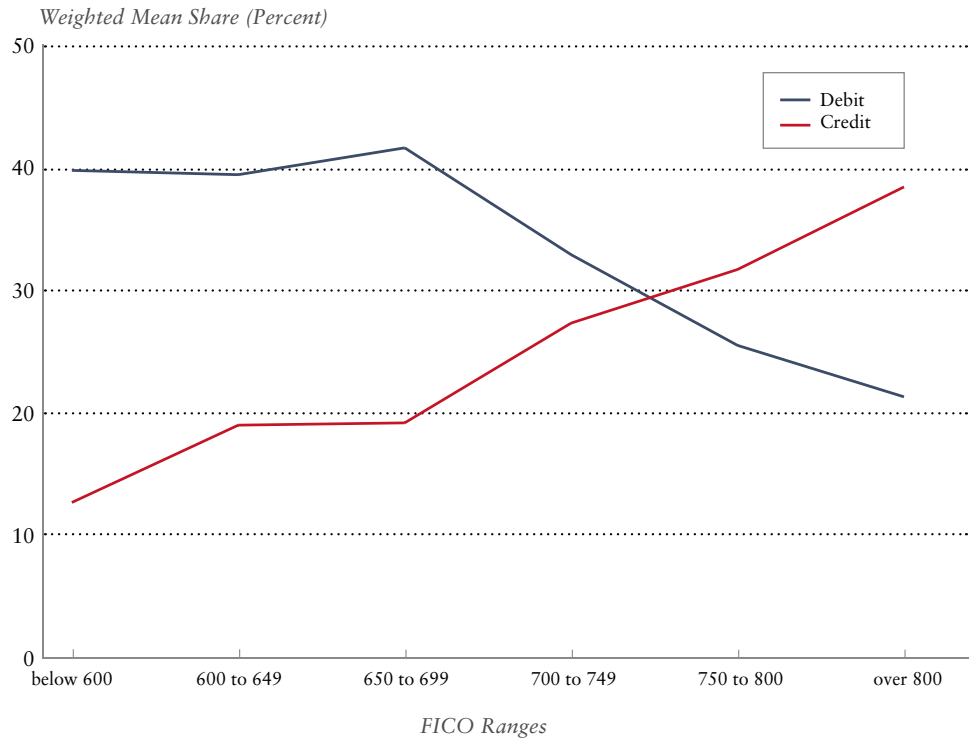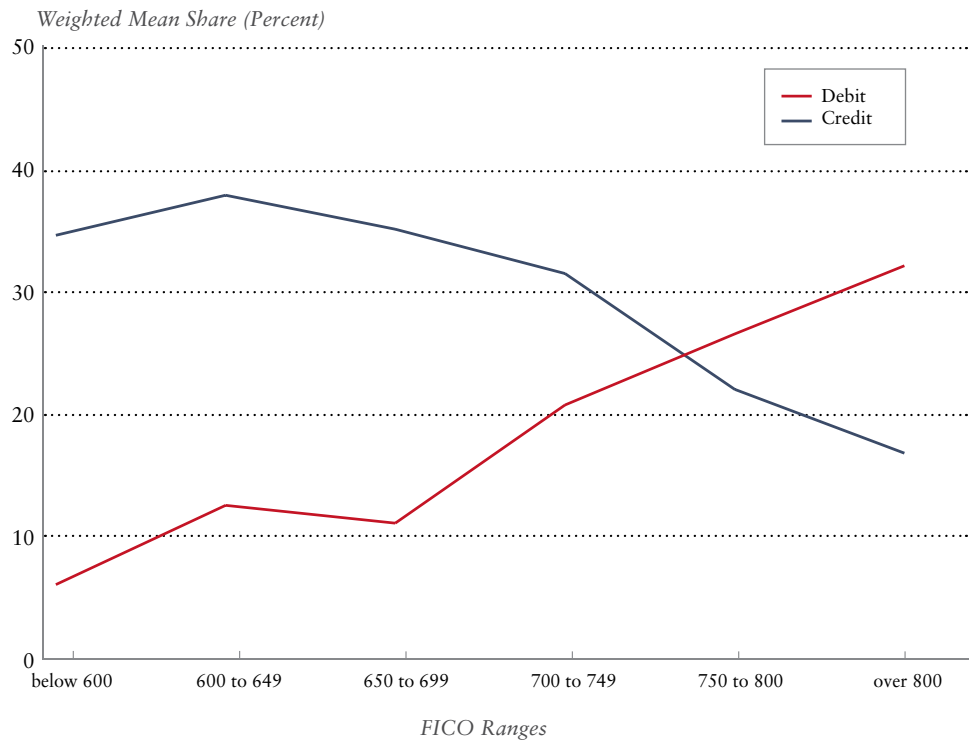
Source: 2009 Survey of Consumer Payment Choice (SCPC).

The authors also include credit and debit card status, such as whether rewards are provided for using these payment methods, and whether the credit card(s) is (are) used for revolving credit (meaning the consumer carries a balance instead of paying the entire bill each month). Most importantly, they include variables that indicate current and past financial difficulties experienced by the consumers.

The authors model payment adoption and payment use by consumers of both debit cards and credit cards in order to test whether and how FICO scores affect adoption and use of debit and credit cards when controlling for other factors, such as consumer characteristics, payment method attributes, and the price and/or rewards associated with certain payment methods. They estimate adoption and use simultaneously using the Heckman (1976) selection model, which controls for potential selection bias in payment use. Their estimation technique is similar to that used in Schuh and Stavins (2010), but the analysis in this paper extends the previous paper in several ways. Much richer in information than the 2006 survey, the 2009 SCPC includes data on the respondents' FICO scores and nonadopters' perceptions of the various payment instruments. To the best of the authors' knowledge, this is the first paper to include FICO scores in the Heckman regressions of payment behavior.

The authors then investigate what a credit score implies for a consumer's payment choices. If a credit score significantly influences a consumer's access to credit, credit limits, or the cost of credit, then the negative relationship likely results from supply-side effects: consumers with lower credit scores cannot access credit by using a credit card, or accessing credit via a credit card may be too costly for them, and therefore they use their debit card instead. The SCPC dataset provides variables that are indicative of consumers' current credit conditions, such as credit card balances and financial difficulties, which help to disentangle supply- and demand-side effects.

Key Findings

Consumers with high credit scores were more likely to hold and to use a credit card and less likely to use a debit card.

- Results of the first-stage (adoption) regressions confirm that higher-scoring (lower-risk) consumers were more likely to hold a credit card and less likely to hold a debit card. Older consumers were less likely to adopt a debit card. Consumers with a college degree were more likely to adopt a credit card than consumers without such a degree. Convenience is a significant determinant of card adoption, and cost is significant in debit card adoption, both relative to the cost of credit cards and relative to all other payment methods. Bankruptcy has a negative and statistically significant effect on credit card adoption, while it has little effect on debit card adoption. The FICO score is statistically significant even after controlling for "bankruptcy (defaulted)" in the credit adoption regression.

- The second-stage (use) regressions indicate that there is a simultaneity bias of joint adoption and use decisions and that the two-step estimation used in this paper is more appropriate than an ordinary least squares regression. As in the adoption regressions, the coefficient on the FICO score is positive and statistically significant in the regression for credit card use, and negative and statistically significant in the debit card use regression, even when controlling for age, education, income, and other variables. Higher-scoring consumers were not only more likely to hold a credit card and less likely to hold a debit card, but conditional on their holding each card, they were also more likely to use a credit card for transactions, and less likely to use a debit card.

- Consistent with previous studies, consumers who got rewards for using their credit cards had a higher share of credit card transactions and a lower share of debit card transactions.

Consumers who get credit card rewards are likely to have higher FICO scores, but even holding the FICO score constant, receiving rewards affects payment use. The cost of holding and using credit cards affects the use of debit cards and vice versa.

- All respondents rate credit cards as more costly than debit cards. Among credit card reward recipients, credit cards get progressively better (relatively less costly) as FICO scores increase. In other words, individuals with lower FICO scores assess credits cards as more costly than do higher-scoring individuals. The pattern is not as clear among consumers who do not receive rewards for using a given payment method, and this suggests that some of the difference in perceived cost among people with low FICO scores and high FICO scores may arise from differences in rewards received, rather than from differences in fees or interest rates paid on credit card debt.

- Using the Equifax data to obtain each consumer's total credit limit, summed over all his or her credit cards, as well as the average credit limit per card, the authors find a positive correlation for both 2008 and 2009 between a consumer's credit limit and credit score and an even stronger positive correlation between the average credit limit per card and credit score, indicating that consumers with a lower credit score were provided lower credit limits than those with a higher credit score. Because the Equifax data do not include information on consumers' income or net worth, it is not observable whether consumers with lower credit scores are provided lower credit limits relative to their income or net worth. Nevertheless, based on these results, one cannot reject the possibility of supply-side effects on consumer payment choice—meaning that more frequent use of debit cards among low-scoring consumers may be a result of their having lower credit limits.

- Results from the Equifax data also indicate a negative correlation between credit utilization (percent of credit limit used) and credit score that is even stronger than the correlation between credit limit and credit score: low-scoring consumers have much higher credit utilization rates than those with higher scores. The causality may run the other way: high credit card utilization rates may cause low scores. Nevertheless, the finding could imply credit limitations for consumers with a lower credit score—due to credit limits, greater liquidity needs in the past, or both.

- Another finding from the Equifax data is that the percentage of the credit limit that is revolved—known as credit card debt—is also negatively correlated with the credit score; in other words, low-scoring consumers carry more credit card debt. However, as the SCPC data show, the relationship between credit score and credit card debt is not monotonic: both the probability of revolving and the amount of debt carried on credit cards drops only above a FICO score of 750. For consumers with FICO scores below 750, there is no clear relationship between revolving and credit scores. While the higher rates of adoption and use of debit cards among consumers with lower FICO scores could be caused by behavioral factors—they may turn to debit cards as a self-restraining tool to help them lower their debt (Sprenger and Stavins 2010), one cannot reject the possibility that the relationship is caused by supply-side credit constraints.

- Even when controlling for demographic and financial variables, consumers who had lost their job in the previous 12 months had a higher share of debit card transactions relative to the rest of the sample, and there was no significant effect on the use of credit cards. Instead of relying more heavily on their credit cards, recently laid-off workers used debit cards more frequently. This may imply that consumers do not necessarily increase the amount

of their credit card balances that they revolve due to a demand shock, such as a job loss. The payment behavior of those consumers who recently lost their jobs could reflect either recent job losers' avoidance of the possibility of going into debt or their expectation that their credit limits would be lowered as a result of their job loss, resulting in an increased likelihood of taking on more debt unless they change their payment behavior.

- Regional differences could be associated with supply-side-related variation in the terms of banking or credit, such as interest rates on deposit accounts or on credit card loans, and these differences could affect consumer payment behavior. Merchant acceptance of credit and debit cards may also vary by region, which likely limits the payment options available to consumers. Although it is possible that consumer preferences for payment methods vary by region, the regional differences likely underscore the importance of supply-side factors and network effects. Results from the authors' testing whether the effect of credit scores on payment behavior disappears with regional or state fixed effects are inconclusive.

IMPLICATIONS

The authors' results suggest that there is a negative relationship between debit card use and credit score, and a positive relationship between credit card use and credit score, even after controlling for various consumer characteristics, payment method attributes, and rewards on payment cards. A new rule, effective on October 1, 2011, reduced the interchange fees for transactions charged on the debit cards that are issued by large financial institutions. Some large financial institutions reacted to this rule by announcing higher debit card fees to recover their lost interchange fee revenues. Because consumers with low credit scores, such as the FICO score, are the ones who use debit cards more intensively, they are likely to be especially adversely affected if their banks introduce debit card fees. Based on the authors' data, younger, less educated, and lower-income consumers are more likely than other demographic groups to be affected by higher debit card fees, especially if their access to alternative payment methods is limited.

This paper tests various hypotheses concerning the determinants of the relationship between credit scores and payment behavior and finds support for supply-side factors related to credit constraints placed on consumers with low credit scores. The next phase of this research will focus on further separating demand-side from supply-side factors that influence the effect of credit scores on payment behavior.

p-12-2

Why Did So Many People Make So Many *Ex Post* Bad Decisions? The Causes of the Foreclosure Crisis

by Christopher L. Foote, Kristopher S. Gerardi, and Paul S. Willen

MOTIVATION FOR THE RESEARCH

Losses on U.S. residential real estate helped spark the largest financial crisis since the Great Depression. Why did so many actors, ranging from individual homebuyers to investment banks, make decisions that in hindsight turned out to be disastrous? A widely held

explanation contends that well-informed mortgage insiders used the securitization process to take advantage of less well-informed outsiders: for example, one variation of this view suggests that some adjustable-rate mortgages were intentionally designed to fail. Essentially, what might be called the insider/outsider explanation holds that distorted incentives and information were the root causes of the housing crisis. People made bad decisions because they were deceived by mortgage-industry insiders who had superior information about the quality of real estate investments.

The authors of this paper present 12 facts that refute this insider/outsider interpretation. They offer an alternative explanation that is based on overly optimistic forecasts about future U.S. house prices. The authors contend that optimistic price beliefs, not incentive or information problems in the pooling of loans into securities, best rationalize the real-time decisions made by borrowers, lenders, intermediaries, and investors during the pre-crisis years.

Research Approach
The authors rely on the historical record, their own previous work, and research conducted by others to construct their 12 facts. The authors then argue that the 12 facts are inconsistent with the insider/outsider view of the crisis, though these facts are quite consistent with an explanation based on optimistic house price expectations.

Key Findings
Fact 1: Resets of adjustable-rate mortgages did not cause the foreclosure crisis.
A popular explanation for why borrowers took out adjustable-rate mortgages (ARMs) they ultimately could not repay is that lenders misled them by giving them loans with terms that initially appeared affordable but proved otherwise. Yet if all the complex mortgage products had been replaced with fixed-rate loans, at most only 12 percent of the foreclosures that took place from 2007 through 2010 would have been averted. A broader examination of all the foreclosures that took place from 2007 through 2010 shows that fixed-rate mortgages accounted for 59 percent of all U.S. foreclosures during this period.

Fact 2: No mortgage was "designed to fail."
A wide array of nontraditional mortgage products was available during the housing boom, ranging from subprime mortgages given to borrowers with poor credit histories, option ARMs, reduced-documentation loans, and loans requiring no downpayment. Some critics claim that these products were "designed to fail," such that no reasonably informed borrower would willingly assume these loans. Yet the vast majority of all mortgages originated between 2000 and 2006 were successful for borrowers and lenders. The fact that the failure rates on all nontraditional mortgages rose at the same time suggests not an intentionally flawed design but rather that these products were not designed to withstand the stunning (and unprecedented) nationwide drop in house prices that began in 2006.

Fact 3: There was little innovation in mortgage markets in the 2000s.
Somewhat related to fact 2, a popular critique holds that the rise in nontraditional mortgages, particularly the growth of the payment-option ARM, meant that the housing boom was fueled in part by intense innovation in the mortgage market. While it is true that historically most mortgages prior to 1981 were fixed-rate loans, the emergence of nontraditional mortgages predates by some decades the housing boom of the 2000s. The payment-option ARM was invented in 1980 and approved for widespread use by the Federal Home Loan Bank Board and the Office of the Comptroller of the Currency in 1981. It is true that this

loan product was first used mostly in California and was almost exclusively held in bank portfolios, as the payment-option ARM generated floating-rate interest income and eliminated the lender's interest-rate risk. At the same time, it smoothed out payment fluctuations for borrowers. It was only in 2004 that option ARMs showed up in datasets of securitized mortgage-backed securities (MBS).

Fact 4: Government policy toward the mortgage market did not change much from 1990 to 2005.
While many blame the foreclosure crisis on lax government regulation of mortgage markets, an influential minority contends that looser underwriting and downpayment requirements were instituted in the service of federal policies enacted in the 1990s to broaden homeownership. Yet there is no evidence to support the contention that government-led housing initiatives implemented over the last 20 years loosened mortgage lending standards and contributed to the foreclosure crisis. What can be described as massive government intervention in the mortgage market occurred in the 1940s, when the GI Bill allowed veterans to purchase homes with small or no downpayments, and obligated the federal government to take a first-loss position equal to 50 percent of these loans. The loan limits on Veteran's Administration (VA) loans were subsequently and repeatedly raised, and similar guarantees were later added to loans administered by the Federal Housing Administration (FHA). By the late 1960s the average downpayment on a VA mortgage was about 2 percent, and by no standard can a VA (or FHA) loan be considered a "niche" product. Compared to these earlier decades, recent data on loan-to-value (LTV) ratios suggests no major federal mortgage market intervention in the 1990s and 2000s. Granted, during the 2003–2006 housing boom, there was an increase in zero-downpayment financing, but even before the boom most U.S. borrowers got mortgages without having to post 20 percent downpayments.

Fact 5: The originate-to-distribute model was not new.
The 2010 Dodd-Frank financial reforms were partly motivated by the idea that the originate-to-distribute (OTD) model of mortgage lending was responsible for much of the financial crisis. In part, this idea rests on the increase in securitized instruments, such as mortgage-backed securities (MBSs). Yet while the individual players changed, the OTD model of lender-servicers' originating loans and then selling these portfolios to other institutions has been central to the functioning of the U.S. mortgage market since the immediate postwar period. The OTD model has evolved from what it was in the 1950s, a model where mortgage companies typically sold their loans to insurance companies, which kept them on portfolio as whole loans, but the institutional framework has largely remained intact. During the 1970s the OTD model was emulated by other financial institutions, most notably savings and loan associations, and the issuance of MBSs largely guaranteed by Ginnie Mae. The early 1980s saw Fannie Mae and Freddie Mac become dominant players in the U.S. mortgage market and the rise of the private-label securities market, which in the 2000s grew at the expense of the agency market. So while the actors in the OTD framework changed over time, the basic model of the delegated underwriting of loans had been in place since the early 1950s.

Fact 6: MBSs, CDOs (collateralized debt obligations), and other "complex financial products" had been widely used for decades.
Some understandable confusion exists between the OTD model and securitization, as securitization implies "originate-to-distribute" but elides the fact that the OTD model had existed for decades before mortgages were securitized in the 2000s. As noted in fact 5, the OTD model first featured the sale of whole loans to insurance companies, but by the 1980s the three federal housing agencies began to arrange and/or insure pass-through securities, whereby

investors could buy a pro-rated share of a pool of mortgages. In the early 1990s, CDOs were designed as a way for banks to sell the risk on pools of commercial loans. Over time, financial institutions realized this same instrument could be used for pools of risky tranches from securities, including private-label MBSs. In 2000, investment banks started to combine the lower-rated tranches of MBSs, typically subprime asset-backed securities (ABSs), with other forms of securitized debt to create a CDO known as the ABS CDO. The poor performance of this instrument in the early 2000s was widely blamed on the inclusion of nonmortgage assets, like tranches from car loans or credit cards, so the ABS CDO came to be dominated by tranches from subprime mortgages. Thus, the growth in securitized mortgages that took place by the mid-2000s was supported by an institutional and legal framework that had been in place since the early 1990s, and the idea that the boom in securitization served as an exogenous event that sparked the housing bust is not supported by the institutional history of the U.S. mortgage market.

Fact 7: Mortgage investors had lots of information.
The idea that mortgage industry experts knowingly withheld information about the securities they structured and sold is one of the pillars upon which the insider/outsider theory of the crisis rests. But the real story is that issuers supplied potential investors with a great deal of detailed information. Prospectuses included key credit-quality variables, such as LTV ratios, borrower credit scores, and loan documentation status. MBS issuers were careful to document the extent to which they did not verify a borrower's income and assets. Investors knowingly bought low doc/no doc loans, and all issuers provided monthly loan-level information on the characteristics of every loan in the pool, including the monthly mortgage payment, the interest rate, the remaining principal balance, and the delinquency status. Investors had access to important data, and access to computational tools that allowed them to accurately price these securities by coding all of the rules from a prospectus concerning the allocation of cash flows to different tranches of a deal.

Fact 8: Investors understood the risks.
Following from fact 7, lenders and issuers supplied investors with sufficient data to enable them to predict how MBSs and related securities would fare under a variety of macroeconomic scenarios. A prime example is an August 2005 analyst report issued by Lehman Brothers (one of the firms famously to go under during the financial crisis) showing the predicted losses on a pool of subprime mortgages issued that year, given a variety of different assumptions about the future path of U.S. house prices. The three most likely scenarios, ranging from "base" to "aggressive," predicted annual losses between 1 and 6 percent. Two adverse scenarios, labeled "pessimistic" and "meltdown," assumed annual near-term house price growth of 0 and –5 percent, respectively, with corresponding losses of 11.1 percent and 17.1 percent. The report notes that the meltdown scenario would lead to massive losses in all but the highest-rated tranches. Analysts at other banks reached similar conclusions, and these documents clearly delineate that investors recognized the potential risk inherent in subprime deals if house prices declined.

Investors understood the risks inherent in subprime deals and understood just how much they would lose if house prices fell but assigned a very small probability to such falls.

Fact 9: Investors were optimistic about house prices.
Investors understood the risks inherent in subprime deals if adverse house price scenarios were to come to pass, but assigned a very small probability to a severe decline occurring. The Lehman report's meltdown scenario, the only one generating losses that would threaten repayment of the AAA-tranches, received a 5 percent probability, while the more benign pessimistic outcome had a 15 percent probability. The top two price scenarios, assuming at

least an 8 percent annual house price appreciation, received probabilities that taken together amounted to 30 percent. This optimism was characteristic of many analyst reports, and offers real-time evidence that investors continued to purchase subprime securities based on expectations that U.S. house prices would continue to appreciate

Fact 10: Mortgage market insiders were the biggest losers.
The insider/outsider interpretation of the foreclosure crisis contends that insiders, those most closely associated with mortgage origination and securitization, had informational advantages that allowed them to profit at the expense of those more removed from the process. Yet the most compelling evidence against an "inside job" is that investors closely tied to the mortgage industry suffered massive, if not catastrophic, losses. Bear Stearns, the investment bank most closely associated with the subprime market, was heavily involved in every aspect of the mortgage market, from origination to securitization to loan servicing. The firm's executives were major investors in two hedge funds managed by Bear Stearns, which in June 2007 began to report enormous losses associated with subprime securities.

Fact 11: Mortgage market outsiders were the biggest winners.
Unlike the Bear Stearns hedge fund managers who bet on a continuing upward trend in house prices, John Paulson, a hedge fund manager with no ties to the mortgage industry, gambled on bearish bets that the U.S. house price boom was not sustainable, and he profited from credit protection on subprime MBSs when these investments suffered huge losses. The insider/outsider story is not supported. Rather, the more useful narrative is the division between those analysts who thought house prices would continue to rise and those who were willing to bet that house prices would fall.

Fact 12: Top-rated bonds backed by mortgages did not turn out to be "toxic." Top-rated bonds in collateralized debt obligations did.
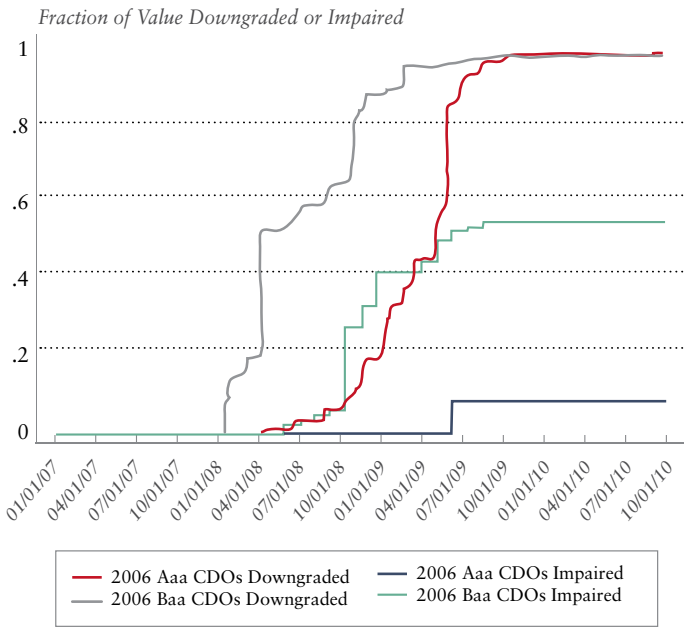Private-label AAA-rated subprime securities did not suffer major losses, as credit protection largely spared investors in these securities. Rather, the securities that were created from lower BBB-rated tranches of subprime MBSs, such as the ABS CDOs discussed above, proved to be the "toxic" mortgage-related securities that helped cause the financial crisis. Whereas the AAA-rated tranches of the original MBSs suffered losses under 10 percent, losses occurred on 90 percent of the ABS CDOs. Part of this disparate performance can be traced to the two very different methods used to rate ABSs and CDOs. The loss probabilities on subprime ABSs were modeled by mortgage industry analysts using individual-level data on borrowers, and this structural risk analysis proved to be quite accurate, as it examined how correlation in individual default probabilities might arise (such as if U.S. housing prices fell). CDOs were originally constructed from various corporate bonds, and CDO analysts rated the performance of this type of security using historical correlations, which proved quite accurate for corporate bonds. As noted in fact 6, the CDO evolved to include mostly subprime mortgages, and in the case of BBB-rated tranches of subprime MBSs, CDO analysts had no way to model the effect of a national decline in house prices because the past data did not encompass such a decline.

Taken together, these 12 facts consistently point to high house price expectations as the fundamental explanation for why credit expanded during the housing boom, which exhibits the hallmarks of a classic asset bubble. Viewed in this way, the decisions of both borrowers and lenders are understandable, as the asset bubble story explains why investors thought subprime mortgages were a good investment, and why credit-constrained households might

> *The most compelling evidence against an "inside job" is that investors closely tied to the mortgage industry suffered massive, if not catastrophic, losses.*

Downgrades and Impairments among Mortgage-Backed Securities (MBSs) and Collateralized Debt Obligations (CDOs)

2006 Vintage MBSs

Fraction of Value Downgraded or Impaired

2006 Vintage CDOss

Fraction of Value Downgraded or Impaired

— 2006 Aaa CDOs Downgraded — 2006 Aaa CDOs Impaired
— 2006 Baa CDOs Downgraded — 2006 Baa CDOs Impaired

— 2006 Aaa CDOs Downgraded — 2006 Aaa CDOs Impaired
— 2006 Baa CDOs Downgraded — 2006 Baa CDOs Impaired

2007 Vintage MBSs

Fraction of Value Downgraded or Impaired

2007 Vintage CDOs

Fraction of Value Downgraded or Impaired

— 2006 Aaa CDOs Downgraded — 2006 Aaa CDOs Impaired
— 2006 Baa CDOs Downgraded — 2006 Baa CDOs Impaired

— 2006 Aaa CDOs Downgraded — 2006 Aaa CDOs Impaired
— 2006 Baa CDOs Downgraded — 2006 Baa CDOs Impaired

Source: Financial Crisis Inquiry Commission (2010). Tables 12, 13, 17, and 18 and Moody's Structured Finance Default Risk Services.
Note: The two panels on the left show that among private-label MBSs, lower-rated tranches suffered massive losses. However, while a large fraction of AAA-rated tranches were downgraded, the vast majority of these tranches paid off, as few of them suffered actual impairments. The two panels on the right show that the same is not true for CDOs. Because the these bonds tended to be backed by lower-rated tranches of private-label MBSs, both the AAA-rated and the lower-rated tranches of CDOs suffered significant impairments.

have taken on more housing debt than proved wise in hindsight. When the value of the underlying collateral is expected to rise rapidly, such decisions are rational—at the time.

IMPLICATIONS

Economic thought does not have a robust explanation for why asset bubbles form, but the experience of the recent boom and bust in the U.S. housing market suggests that economists should make a serious attempt to better understand how beliefs are formed about the prices of long-lived assets. It is clear that asset prices move in ways not yet understood. Bearing this in mind, institutions can be designed to better withstand extreme shocks. The authors suggest two central questions for evaluating future policies bearing on the U.S. housing market. One, can financial institutions withstand a serious house price shock, such as a 20 percent drop in value, and not suffer liquidity problems? Two, can individual borrowers withstand a substantial fall in house prices? It may not be possible to recognize or deal with asset price bubbles in real time, but it should be possible to make market structures more resilient to the adverse effects a bubble may inflict.

p-12-3

The Supplemental Security Income Program and Welfare Reform

by Lucie Schmidt

MOTIVATION FOR THE RESEARCH

Over the past 20 years, the Supplemental Security Income (SSI) program, which provides federally funded income support for disabled individuals, has become one of the most important means-tested cash aid programs in the United States. The number of disabled adult SSI recipients increased by 89 percent between 1990 and 2010, and the number of child SSI cases quadrupled over this same time period. However, existing research tells us little about the determinants of SSI caseloads, which vary dramatically both across states and over time.

During this same period, the United States enacted major welfare reform. The passage of the Personal Responsibility and Work Opportunity Reconciliation Act (PRWORA) in 1996 replaced the Aid to Families with Dependent Children (AFDC) program with the Temporary Assistance for Needy Families (TANF) program, a change that then-President Bill Clinton said would "end welfare as we know it." Welfare reform coincided with substantial increases in labor supply and earnings for a number of former welfare recipients, and with unprecedented decreases in the number of AFDC/TANF recipients.

Understanding variation in SSI caseloads is particularly important in the post-welfare reform era for a number of reasons. While SSI is targeted at the disabled and AFDC/TANF is targeted at single-parent families, there is some degree of substitutability between the two programs. Previous research provides evidence that some portion of the increase in SSI caseloads can be attributed to efforts to reform the AFDC/TANF program over the same period (Schmidt and Sevak 2004). Evidence also suggests that some localities actively attempted to move TANF recipients who faced time limits to SSI (Pavetti and Kauff 2006). If these reasons account for the increased SSI caseload, then the SSI program might represent an alternative safety net for

former welfare recipients who are disabled. Work by Duggan and Kearney (2007) suggests that SSI benefits have become an important source of income for economically disadvantaged families and can reduce the incidence of family poverty. Despite the growth of the SSI program, little research has been done on the factors that determine SSI caseloads. From a policy perspective, understanding what causes these caseloads to rise has become increasingly important.

Research Approach

In this paper, the author uses regression analysis and state panel data, exploiting variation both across states and over time, to determine what factors affect SSI caseloads involving disabled individuals. Adult and child disabled cases are analyzed separately. She examines the relative contribution of a number of factors, including economic conditions, demographic variables, health conditions, and relative program generosity. She then examines the effect of the 1996 federal welfare reform, as well as the effect of variation in welfare policies across states, such as time limits and sanctions for noncompliance. Given previous research that provides evidence of interactions between the SSI program and other welfare programs that provide income support to single-parent families, the author also examines how the effect of the factors listed above has changed since major welfare reform was enacted in 1996. These findings could be particularly important in the context of the Great Recession, as evidence suggests that cash aid through the AFDC/TANF program has become less cyclical after the Clinton era welfare reform (for example, Bitler and Hoynes 2010).

SSI-Disabled Adults, Selected States, 1980–2010
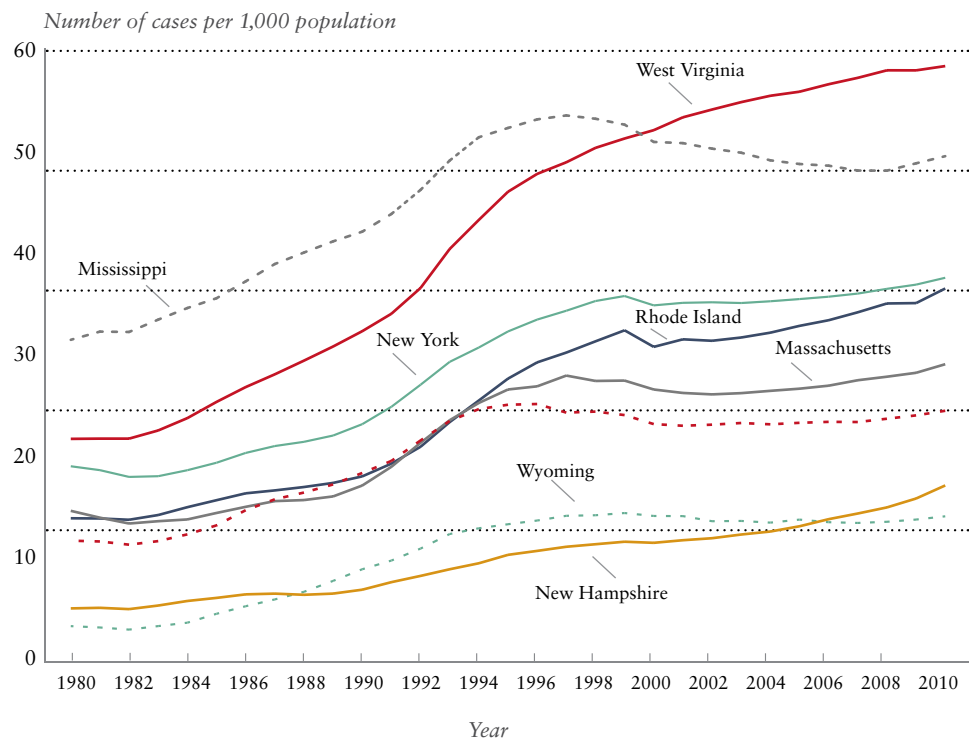
Number of cases per 1,000 population

Source: Annual Statistical Supplement to the Social Security Bulletin, various years, and Census Bureau population estimates, various years.

Key Findings

- For adults, economic variables have significant effects on the SSI caseload share, but not always in the direction that would be expected if SSI is a substitute for earned income. Higher per capita personal income is associated with a lower SSI caseload share, significant at the 1-percent level. The coefficient estimate suggests that a 10 percent increase in per capita personal income would be associated with a 6 percent decrease in the SSI caseload share. However, higher unemployment rates are also associated with a significantly lower SSI caseload share, such that a 1 percentage-point increase in the unemployment rate would lead to a 2.4 percent decrease in caseload share. This unemployment rate effect is the opposite of findings by Stapleton et al. (1998, 1999) but is consistent with results from Garrett and Glied (2000) and Schmidt and Sevak (2004).

- The share of nonmarital births is positively and significantly associated with the adult disabled SSI caseload share. A one-standard-deviation increase in the share of nonmarital births would lead to an increase in SSI caseload share of approximately 8 percent. For adults, higher AFDC/TANF benefits for a family of three are negatively associated with the SSI caseload share. The obesity rate is not significantly associated with adult SSI participation. Consistent with work by Kubik (2003), unexpected state-level deficit shocks significantly affect the adult SSI caseload share, with the effect for a negative shock roughly twice the magnitude of the effect of a positive shock.

- While the point estimate on TANF implementation is positive, it is not statistically different from zero. However, the major welfare waivers implemented pre-PRWORA are positively and significantly associated with SSI caseload share, consistent with work by Schmidt and Sevak (2004). State sanction policies for TANF recipients are positively and significantly associated with a higher SSI caseload share among adults, consistent with evidence that the disabled were more likely to be removed from the TANF rolls. The presence of a TANF sanction policy is associated with a 4.8 percent increase in SSI caseload share.

- Welfare reform variables (both indicators for implementation as well as specific TANF time limit and sanction policies) have a stronger effect for women than for the overall adult disabled SSI caseload.

After the 1996 welfare reform, SSI participation appears to have become more cyclical for women and children.

- Economic conditions have similar effects on the child caseload share as on the adult caseload share, with both log per capita personal income and the unemployment rate negatively and significantly associated with child caseload share. The share of nonmarital births is positively and significantly associated with the child SSI caseload share, as is the share of the caseload population that is black.

- The relative generosity of AFDC/TANF and SSI benefits affects the child SSI caseload share, with higher AFDC/TANF benefits and lower SSI supplements both reducing the child SSI caseload share. As in the adult regressions, TANF sanction policies also significantly increase the SSI child caseload share.

- Regressions were estimated that allow the effects of the model variables to differ in the post-PRWORA period by interacting variables with the indicator for TANF implementation. There is some evidence that the relative magnitudes of AFDC/TANF benefits and SSI supplements matter less for children after welfare reform, which would be consistent with the weakening of the safety net provided through TANF. Results also suggest that the

SSI-Disabled Recipients: 1980–2010
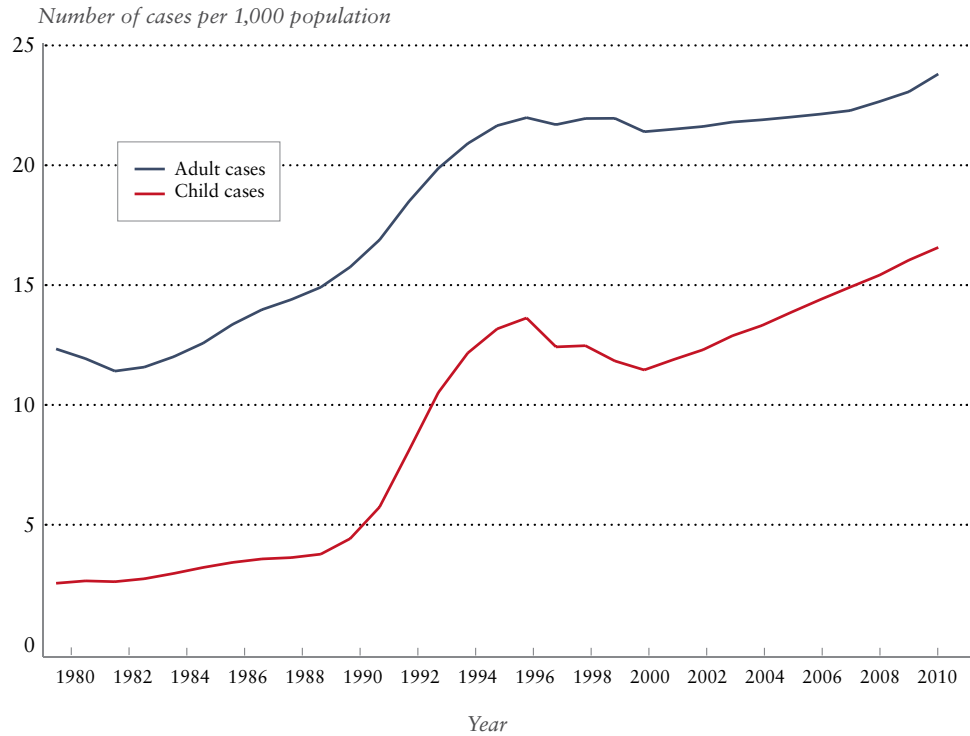
Number of cases per 1,000 population

Legend:
— Adult cases
— Child cases

Source: Annual Statistical Supplement to the Social Security Bulletin, various years, and Census Bureau population estimates, various years.

Total Federal SSI-Blind and SSI-Disabled Payments
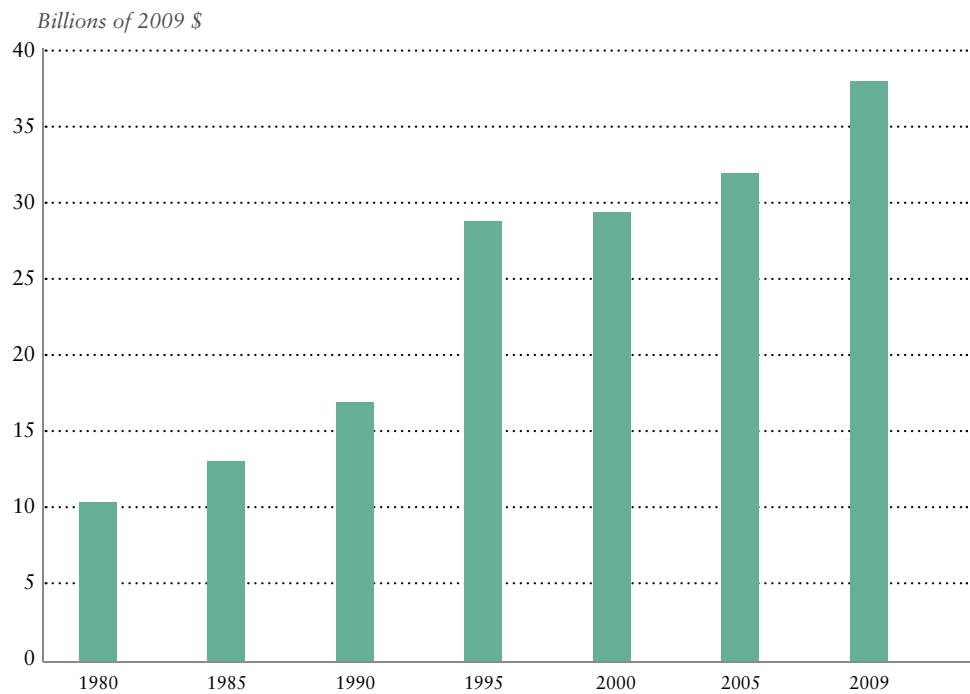
Billions of 2009 $

Source: Annual Statistical Supplement to the Social Security Bulletin, various years, and Census Bureau population estimates, various years.

effect of a state's having a Democratic governor on SSI caseload share has become more positive since welfare reform, and that the positive association between the black population share and the SSI caseload share has become weaker after welfare reform.

• Interestingly, despite the somewhat counterintuitive relationship between the unemployment rate and the SSI caseload share over the full time period, the interaction between unemployment rates and TANF implementation is positive, suggesting that for all groups except adult males, the SSI caseload share has become more cyclical post-PRWORA.

Implications

This paper provides preliminary evidence about the role of economic conditions and policy variables on disabled adult and child SSI caseloads, and about how those effects may have changed in the post-welfare reform era. Preliminary results suggest that higher levels of per capita income reduce the SSI caseload share for both adults and children, and that higher percentages of nonmarital births are associated with greater SSI participation. The welfare waivers implemented in the early 1990s have had a significant effect on SSI participation among adult women, and TANF sanction policies significantly increase the SSI caseload share for both adults and children. Results from a specification that allows effects to vary post TANF-implementation suggest that after the 1996 welfare reform, the SSI participation has become more cyclical for adult women as well as for children and the presence of a Democratic governor is more positively associated with the SSI caseload share.

The robust negative relationship between unemployment rates and the SSI caseload share for all groups is puzzling, although the results suggest that this relationship has become significantly less negative since welfare reform. Further work is necessary to fully understand this relationship. One possibility is that it is related to the distinction between stocks and flows. The dependent variable used in this analysis represents the stock of individuals on the SSI program, but economic conditions should affect transfer program rolls primarily through the flow of individuals onto and off of the program (for example, Grogger et al. 2003; Klerman and Haider 2004). The fact that individuals on SSI are likely to remain on SSI for long periods of time suggests looking directly at application rates. The next version of this paper will incorporate these rates into the analysis.

The evidence presented here suggests a direct relationship between elements of welfare reform and SSI participation rates among women and children. Furthermore, the increased cyclicality of the SSI program is consistent with existing evidence suggesting that after the passage of welfare reform, cash benefits provided through AFDC/TANF offer less recessionary protection than before and evidence that other programs such as food stamps have become more cyclical since welfare reform was enacted (Bitler and Hoynes 2010). These findings suggest that SSI is, to some extent, playing the role of an alternative safety net in the post-welfare reform era. As a result, the program could have important implications for the wellbeing of low-income families, particularly given the sustained high unemployment rates during and following the Great Recession.

w-12-1

Are American Homeowners Locked into Their Houses? The Impact of Housing Market Conditions on State-to-State Migration

by Alicia Sasser Modestino and Julia Dennett

MOTIVATION FOR THE RESEARCH

Although the Great Recession that began in December 2007 technically ended in June 2009, the U.S. unemployment rate remains high. One factor contributing to the recession was the unprecedentedly steep decline in house prices that occurred in many parts of the country. This sharp fall in house prices and the resulting widespread and persistent weakness in the housing market raises the question of whether some people are unable to relocate to better job markets because they have negative equity in their homes, meaning that they owe more on their mortgages than the houses are currently worth, and thus are "locked in" to their existing residences.

While negative equity positions were not a significant factor in previous postwar recoveries, it is possible that the depth and duration of the Great Recession and the severe and prolonged weakness in house prices have combined to depress the U.S. labor market more than might otherwise be expected. Historically, state-to-state migration has been associated with procyclical, employment-related moves, although over the past 25 years there has been a steady downward trend in interstate migration. Mobility in the United States picked up somewhat during the economic boom that preceded the Great Recession, but began to fall in 2006, a date that precedes the recession but corresponds closely to the rise in the share of households with nonprime mortgages experiencing a rise in negative equity. CoreLogic, a mortgage servicer, estimates that at the end of 2011:Q3, 22.1 percent of all U.S. residential mortgages had a negative equity position and an additional 5 percent of U.S. homeowners had less than 5 percent positive equity in their homes. Together, these amount to 13.1 million properties, or 27.1 percent of all residential mortgages as of 2011:Q3.

Theoretical predictions about how negative equity affects geographic mobility are ambiguous. Some households with negative equity may be liquidity constrained and unable to move because they lack the funds for a downpayment on a new home. Given this situation, in order to relocate they might have to default on the mortgage—an unpalatable choice that might result in losing other assets or inflicting permanent damage to their credit rating. Other homeowners may be loath to sell their houses for less than they paid, and prefer to wait until prices recover. A competing theoretical prediction holds that at a particular threshold of negative equity, households might engage in a strategic, or deliberate, default, especially if lenders are unable to recover losses by seizing the borrowers' other assets.

Evidence from prior empirical research on whether negative equity induces "house lock," and thus reduces geographic mobility and state-to-state migration, has been inconclusive. Work conducted prior to the recent housing market episode has concluded that house lock

restricts geographic mobility, but the majority of these studies focus on a particular region or demographic group. Newer studies examining the house price decline that began in 2006 have been hampered by an inability to match recent trends in negative equity with accurate data on state-to-state migration patterns. Longitudinal datasets tracking individual migration, such as the Survey of Income and Program Participation and the Panel Study of Income Dynamics, contain no information on mortgage debt and home values. While the American Housing Survey is a longitudinal study that tracks home values, it is based on houses, not households, and so lacks information on where the former homeowners may have moved. Other recent studies have used cross-sectional data to examine variation over time, but lack concurrent data to match recent trends in negative equity with state-to-state migration. Since the data on negative equity are limited, some recent work uses proxy measures to match trends in house price changes to the migration patterns of homeowners versus renters, but this method suffers from insufficient time-series data on migration patterns based on homeownership status. Other studies control for nationwide economic conditions, but no study has accounted for the differences in relative economic conditions and local amenities that exist between origin and destination states, despite findings in prior work showing that these variables are significant determinants of migration.

This paper seeks to provide a more definitive response to the issue of whether house lock has played a role in the persistently high U.S. unemployment rate that has hampered the economic recovery to date. The authors investigate three main questions: 1) Does negative housing equity reduce state-to-state migration? 2) What is the relative importance of negative equity versus other economic factors in reducing state-to-state mobility? 3) How much impact does the housing market have on the current recovery in the U.S. labor market?

Research Approach
The authors examine the period from 2006 through 2009, which encompasses the Great Recession and is the most recent period for which state-to-state migration data are available. To overcome some of the shortcomings in previous studies, the authors use two proxies to capture the variance in negative equity across states and over time and the changes in house prices between 2006:Q1 and 2009:Q4. The first proxy is calculated by the Government Accountability Office and is based on their analysis of CoreLogic's active nonprime mortgage data and state-level house price index, available quarterly from 2006:Q1 to 2009:Q4. The second proxy is the change in house prices during this period, captured by the Federal Housing Finance Agency's quarterly data on the nominal house price index. Both measures reflect changes in state-level housing market conditions present during the residential real estate crash that are to some degree correlated with the prevalence of negative equity in a given state.

While state-to-state migration generally moves in tandem with the business cycle, during the Great Recession some states fared better than others as a result of strong local economies and an absence of house price volatility. To untangle the factors that might impact migration from one state to another, the authors construct a regression model that allows them to control for movements in relative economic conditions between origin and destination states in order to accurately estimate the separate effect of negative housing equity on state-to-state migration. Following the literature, the basic model is a logistic specification, where individuals are assumed to choose from among a finite number of destinations the location that yields the highest expected net discounted return on migration. The logistic specification assumes that individuals compare each potential destination state with the origin state in a pair-wise fashion, so the economic conditions in states beyond the pair have no effect on the

choice to migrate. The authors' model has several distinct features. First, it examines gross migration patterns by analyzing both inflows and outflows, rather than trying only to explain net flows. Calculated from Internal Revenue Service data, the state-to-state migration rates reflect the number of individuals moving from the origin state to another state in a given year as a percentage of the total number of people initially residing in the origin state in that same year. Second, the model controls for relative economic conditions in the origin state versus the destination state as measured by labor market conditions, per capita incomes, and housing affordability. Third, the model controls for different propensities to migrate among the origin populations as well as for unobservable amenities unique to individual states that do not change over time, such as climate, culture, and recreational features. These controls ensure that the authors do not overestimate the impact that negative housing equity has on declining out-migration trends.
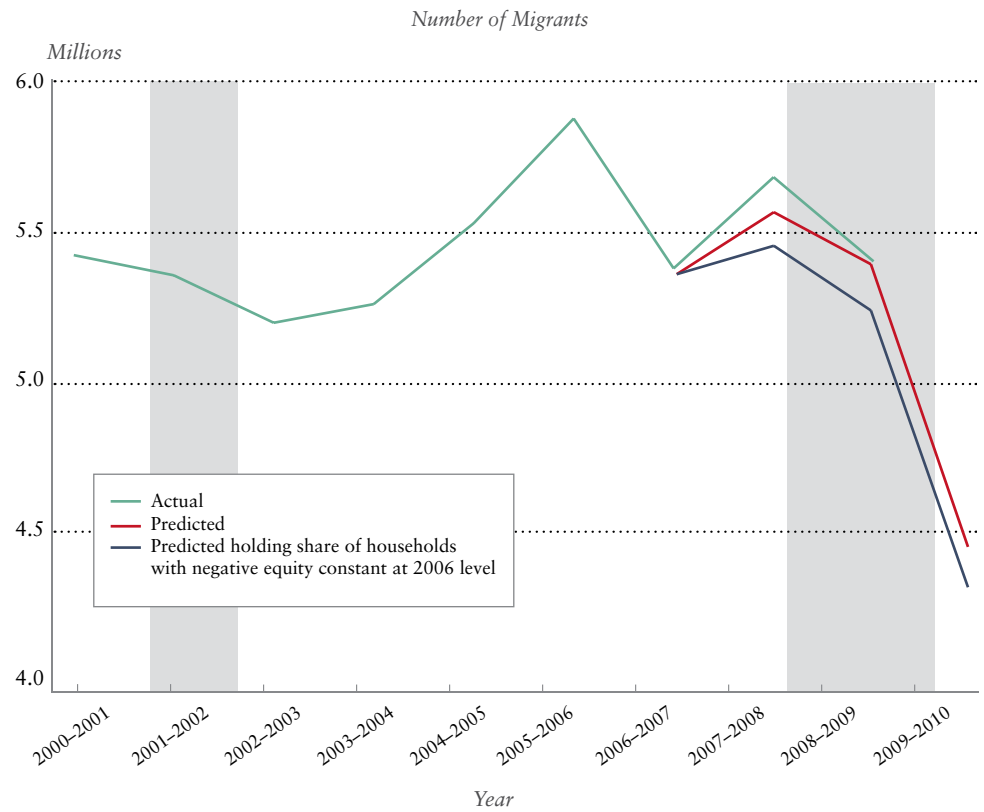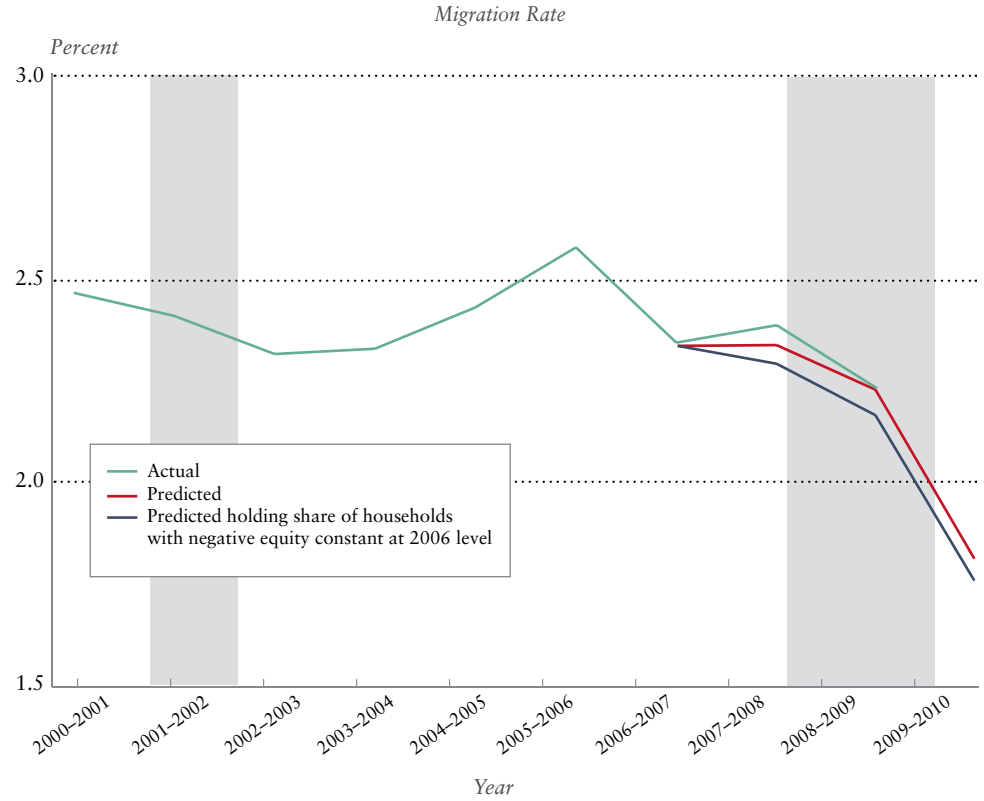
To measure the impact of negative equity on the national labor market, the authors simulate mobility under two alternative scenarios. The first scenario predicts migration for the 2006–2009 period using observed data on negative equity, relative economic conditions, and demographics. The alternative scenario predicts migration over the same period, but holds constant the share of nonprime households with negative equity at the levels observed in 2006, in order to generate a counterfactual path of migration that would have occurred if housing prices had not deteriorated across the country.

Finally, the authors perform a robustness check to determine the channel by which negative equity affects migration. Using the American Community Survey, they generate separate migration rates for homeowners versus renters and estimate the same logistic model. If the housing bust primarily affected homeowners, the expectation would be that the share of nonprime households with negative equity would have had an impact on the interstate migration rates of homeowners, but that the migration rates of renter households would not have been affected.

Key Findings

- Using the share of nonprime households, the authors' results indicate that negative housing equity had a small but significant impact on state-to-state migration, even when controlling for relative economic and demographic conditions. A one-standard-deviation increase in the share of underwater households in the origin state reduced the outflow of migrants from the origin to the destination state by 2.93 percent. For the average origin-destination pair of states, this effect decreased the mean rate of out-migration for every 1,000 initial residents living in the origin state from 0.595 to 0.578. This result translates into a reduction of 85 migrants per year. Summed over all possible destination states, this would reduce the annual outflow from the average origin state by around 4,000 residents.

- Higher rates of foreclosure in the origin state are associated with an increase in the out-migration rate, all else being equal. This result suggests that the impact of negative equity on out-migration operates primarily at low-to-moderate levels of negative equity rather than at extremely high levels, where individuals are more likely to strategically default and thus be free to move across state lines.

- Negative equity has a significant impact on the state-to-state migration of homeowners but no detectable impact on renters. A one-standard-deviation increase in the origin state's nonprime negative equity share decreases the out-migration rate among homeowners by

Actual Versus Predicted Migration Under Alternative Scenarios, United States

Migration Rate

Percent

Legend:
- Actual
- Predicted
- Predicted holding share of households with negative equity constant at 2006 level

X-axis (Year): 2000–2001, 2001–2002, 2002–2003, 2003–2004, 2004–2005, 2005–2006, 2006–2007, 2007–2008, 2008–2009, 2009–2010

Year

Number of Migrants

Millions

Legend:
- Actual
- Predicted
- Predicted holding share of households with negative equity constant at 2006 level

X-axis (Year): 2000–2001, 2001–2002, 2002–2003, 2003–2004, 2004–2005, 2005–2006, 2006–2007, 2007–2008, 2008–2009, 2009–2010

Year

Source: Author's calculations from the Internal Revenue Service state migration data.
Note: Shading represents National Bureau of Economic Research U.S. recession periods.

16.4 percent, but negative equity has no significant impact on the state-to-state migration of renters.

- The reduced mobility across states attributable to negative housing equity is small relative to the annual number of U.S. migrants moving across state lines. Homeowners account for roughly 20 percent of all state-to-state migrants in a given year, so the impact of their reduced mobility on the national labor market is negligible at best. When aggregated across all possible origin and destination pairs, the reduction in the national state-to-state migration rate caused by negative housing equity over the 2006–2009 period was only 0.05 percentage points; this represents about 110,000 to 150,000 fewer individuals migrating across state lines in any given year. Compared with the annual number of migrants typically observed—roughly 5.6 million individuals in 2008–2009—this reduction amounts to the proverbial drop in the bucket.

- The relatively small effect of negative housing equity on state-to-state migration translates into a negligible impact on the national unemployment rate. If all the would-be interstate migrants who were constrained from relocating as a result of negative housing equity had been able to move, this would have reduced the nation's unemployment rate by at most 0.10 percentage points annually between 2006 and 2009. The cumulative effect over this period would have yielded an unemployment rate of 9.0 percent versus 9.3 percent in 2009. Since not all interstate migrants relocate for job-related reasons or were previously unemployed, this effectively means that the national unemployment rate has not been measurably impacted by negative housing equity.

Over the 2006–2009 period, negative housing equity caused a reduction in the national state-to-state migration rate of 0.05 percentage points, which would exert a very small impact on the U.S. unemployment rate.

IMPLICATIONS

Since conditions in the U.S. housing and labor markets in 2011 and 2012 have remained largely unaltered from the conditions that prevailed in 2009, this paper's results are relevant for ongoing policy discussions aimed at reducing the nation's high unemployment rate. It seems reasonable to conclude that policymakers should continue to focus on measures designed to stimulate aggregate demand in order to reduce the nation's unemployment rate. Policies designed to reduce the impact of negative equity on underwater homeowners may help some individual households but are unlikely to exert a measurable impact on the employment rate. Increased efforts to alleviate the housing sector's drag on the economy, such as helping more homeowners refinance into more affordable mortgages and/or stemming the tide of foreclosures, may be more effective at stimulating the economy and reducing the high unemployment rate.

w-12-2

How Consumers Pay: Adoption and Use of Payments

by Scott Schuh and Joanna Stavins

MOTIVATION FOR THE RESEARCH

In a previous paper (Schuh and Stavins 2010), the authors addressed the question of what determines consumer payment behavior by using the results of a consumer survey conducted in 2006. They found that payment method characteristics affect payment use more than the

demographic attributes of the consumers who conduct the transactions. In particular, the cost and convenience of payments were found to contribute substantially to the decline in the use of paper checks. Clearly, the perceptions of payment characteristics vary across individuals: one person may consider online banking convenient, while another may find it cumbersome. Nevertheless, measuring these attributes is important for estimating the demand for payment methods and for predicting future changes in the use of paper, card, and electronic payment methods. This paper uses newer and better survey data to extend the Schuh and Stavins (2010) two-step Heckman model approach to estimate the adoption (first stage, extensive margin) and use (second stage, intensive margin) of seven payment methods.

Research Approach
The authors employ the 2008 Survey of Consumer Payment Choice, a nationally representative survey of U.S. consumers designed by the Federal Reserve Bank of Boston and administered by the RAND Corporation, which improves significantly on the 2006 version. They test the robustness of their methodology by using a variety of specifications.

While the 2008 survey is similar in content to the 2006 survey used in Schuh and Stavins (2010), several important improvements were made in the 2008 survey that allow for better estimation. First, the 2008 survey collected data on nine different payment instruments rather than seven. The 2008 survey asked about four paper instruments: cash, check, money orders, and traveler's checks; three payment cards: credit cards, debit cards, and prepaid cards; and two types of online payments: online banking bill payment (OBBP) and bank account number (BAN) payments. Second, the survey included consumers' ratings of the characteristics of payment instruments along several dimensions, by both adopters and non-adopters of each payment method. For each payment instrument respondents assessed the characteristics on an absolute scale of 1 to 5, where 1 was the least desirable (for example, slowest or most expensive) and 5 was the most desirable (fastest or cheapest). Third, a much more extensive set of questions gathered more information on the survey respondents.

The authors expand on the previous consumer payment behavior literature in several ways. They are the first to model the number of payment instruments adopted by a consumer, conditional on bank account adoption. The number of payment options available to unbanked consumers is obviously very limited, compared with the number available to those with bank accounts. Therefore, the authors estimate a two-step model: bank account adoption, and the number of payment instruments adopted conditional on bank account adoption. They then estimate a set of regressions for adoption and for use, conditional on the adoption, for each payment instrument separately. In this paper, using more comprehensive data than were available for the research underlying Schuh and Stavins (2010), they are able to include payment characteristics in the adoption stage and test various estimation techniques and model specifications.

Key Findings
- Although demographic variables explain some of the variation in consumer payment behavior, the perceived characteristics of payments are significant for both the adoption and the use of payment instruments: setup and recordkeeping are especially important in payment adoption, while convenience, cost, and security affect payment use.

- Following the Federal Reserve's 2011 announcement of the new interchange fee policy, some large banks announced new debit card fees but later retracted these plans after

widespread customer outrage. However, it is clear that changes to debit card fees can lead to an increase in the cost of debit cards to consumers. The authors find that both the adoption of debit cards and the use of debit cards—conditional on adoption—are sensitive to debit card cost. This finding indicates that consumers may reduce their reliance on debit if banks raise the cost of setting up or using debit cards.

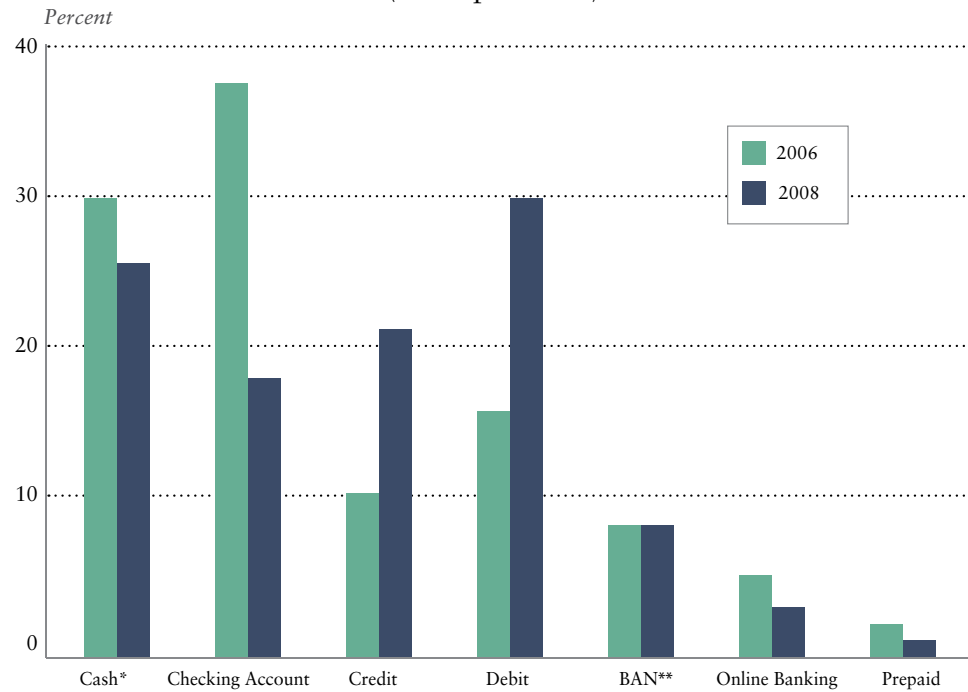- The authors analyze how bank account adoption affects payment behavior in order to show how unbanked consumers' payment choices differ from the choices of those who have bank accounts. Approximately 6 percent of respondents did not have any bank accounts. Because most payment instruments require bank account adoption, the unbanked held—on average—approximately one payment method, compared with over five payment instruments per banked consumer. Not surprisingly, unbanked consumers rely on cash much more heavily than bank account holders do: 76 percent of their transactions were conducted in cash, compared with 25 percent for consumers with a bank account. Low-income and black respondents were less likely than other consumers to have a checking account.

- The average consumer held 5.1 of the nine payment instruments and used 4.2 of these in a typical month. However, consumers were quite heterogeneous in the combination of payment instruments held.

- Online banking, BAN payments, and debit cards experienced the highest increases in adoption over the two years between the 2006 and the 2008 surveys.

- Cash adoption was almost universal: 98 percent of respondents were cash adopters.

- The rate of check adoption was almost as high as that for cash. Over 90 percent of the sample had adopted checks. Check adoption was higher for older, higher-income, or more educated respondents than for those who were younger, had lower incomes, or were less educated. It was lower for single or separated respondents than for those who were married or widowed, and it was lower for blacks than for white or Asian respondents.

- The overall rate of credit card adoption was 78 percent, slightly above the 2006 rate of 74 percent. Similar to the adoption of checks, the credit card adoption rate was higher for older, more educated, higher-income, and wealthier respondents; was much lower for blacks than for whites or Asians; and was lower for single or separated people than for those who were married or widowed. Men had a higher credit card adoption rate than women.

- For the first time since the inception of the SCPC, credit card adoption fell slightly below debit card adoption, which was 80 percent. However, the distribution within the sample differed substantially between the two payment methods. In contrast to credit cards, the adoption of debit cards was greater for the young than for the old and was not greater for highly educated consumers (although it was lowest for those with the lowest level of education). Married respondents were more likely to have a debit card than respondents in any other category, especially those who were single, and blacks were less likely to adopt debit cards than were respondents of any other race. Even though debit adoption was lowest for those earning an annual income below $25,000, there was no discernible difference among the remaining income groups. Adoption of prepaid cards was lower in the 2008 SCPC than in the 2006 survey, possibly because the survey questions differed.

Payment Method Adoption Rates, 2006 and 2008

Percent

Legend:
- 2006
- 2008

Categories: Cash*, Checking Account, Credit, Debit, BAN**, Online Banking, Prepaid
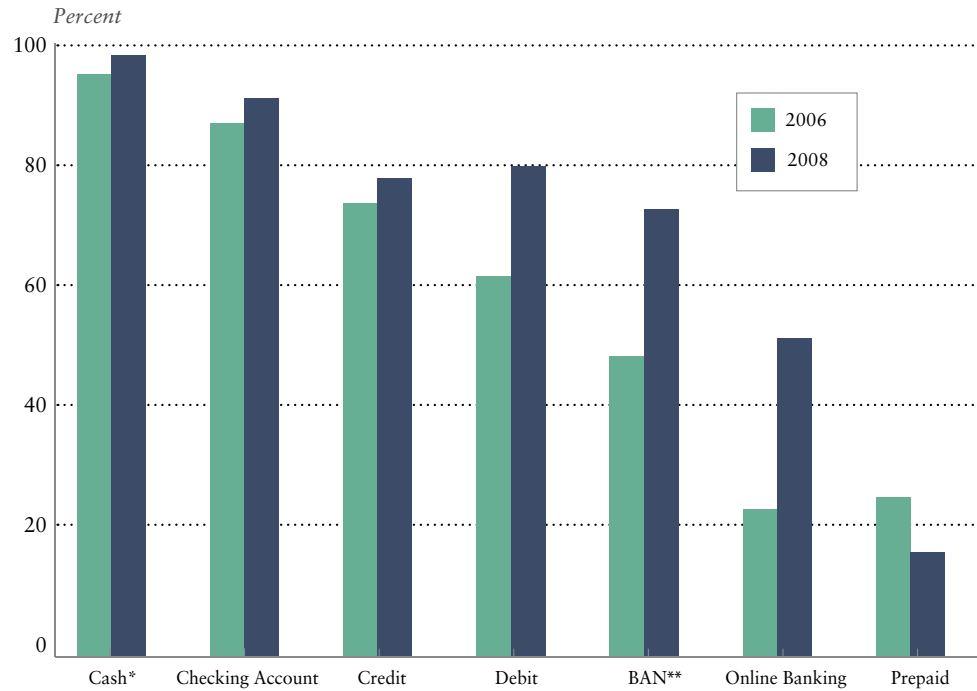
Share of Monthly Payments, 2006 and 2008
(all respondents)

Percent

Legend:
- 2006
- 2008

Categories: Cash*, Checking Account, Credit, Debit, BAN**, Online Banking, Prepaid

Sources: Survey of Consumer Payment Choice (2006 and 2008).
*A respondent adopted cash if he/she had cash on property, or if he/she gets or uses cash at least once in a typical year.
**BAN are bank account number payments. The 2006 SCPC did not include a BAN adoption question, so Automatic Bill Payment was used in its place for 2006. They are not directly comparable, however.

- By far the largest change between the 2006 and the 2008 survey results was in the adoption and use of electronic payments. The rate of adoption of BAN payments was 73 percent in this survey, compared with 49 percent in the 2006 version. The adoption of BAN payments did not exhibit strong demographic patterns, other than being lowest for the youngest, lowest-income, black, and least educated respondents. Because BAN payments are often used for housing-related payments, such as mortgage and utility payments, some of these differences are probably due to the lower rate of homeownership among these respondent groups. The adoption of online banking bill payment increased from 24 percent in 2006 to 52 percent in 2008—the fastest growth of any payment method included in the survey. Similar to the adoption rate of debit cards, the OBBP adoption rate was lower for older and less educated respondents, highest for married people, and lowest for blacks and those with annual income below $25,000.

- For the whole sample, debit cards were the most intensively used payment method, accounting for 35 percent of all transactions. Credit cards and cash were used almost equally, while checks—at 16 percent of all transactions—ranked fourth. These numbers contrast with the 2006 results, when checks constituted 38 percent of all transactions and were the most popular payment method, while cash was second with 30 percent of transactions.

- Based on the Heckman two-step regression results, cash and debit card use (conditional on adoption) was higher for younger, lower-income, less educated, and poorer respondents, and was highest for single people. In contrast, credit card use was higher for older, higher-income, more educated, and wealthier consumers. Check use was higher for older people, but did not show any other strong patterns. The use of BAN payments was fairly similar across the demographic cohorts, while the use of OBBP among adopters was moderately higher for older and higher-income respondents.

- Looking at average shares for all respondents (not just adopters) based on the 2008 and 2006 survey data, the largest increase occurred in the use of debit cards, while the largest decline was in the use of checks and cash. Most of the transactions took place at the point of sale, and the composition of payment methods used varied depending on the type: most point-of-sale transactions were conducted with cash or debit, while paper checks dominated bill payments.

Setup and recordkeeping are especially important in payment adoption, while convenience, cost, and security affect use.

- There is little variation across consumers in the way they assess payment characteristics: the mean ratings ranged from 3.3 for prepaid cards to 3.8 for cash and debit cards, on a 1-to-5 scale. On the other hand, there is more variation across the characteristics, ranging from a 2.9 mean rating for security of payments to a 4.0 mean rating for acceptance. One characteristic that does vary across the payment instruments is cost: cash stands out as the least costly instrument, while credit cards are considered the most expensive. Cash is also rated as the fastest and the easiest to set up, but also as least secure and the worst for recordkeeping. Adopters rated each payment method higher than did nonadopters, especially in terms of the cost and setup of payments.

- Ratings by both adopters and nonadopters allow us to infer the major barriers preventing consumers from adoption. The greatest discrepancies in ratings between adopters and nonadopters were in cost, setup, and ease of use, suggesting that these were the main reasons consumers had not adopted certain payment instruments. Because the perceived

payment characteristics varied even within each sociodemographic cohort, including the payment characteristics in the regressions of payment behavior helps to explain consumer decisions, as Schuh and Stavins (2010) demonstrated.

Implications

Payment characteristics are found to be even more important in 2008 than in 2006. In particular, security is especially significant in the payment use regressions, while setup and recordkeeping are significant in payment adoption regressions. Cost was significant in both the adoption and use of debit cards. Following issuance of the rule on debit card interchange fees in 2011, several large banks announced new fees for debit card use in order to recover their lost revenues from debit card transactions. It is not clear whether debit card fees will be instituted, but the authors' results indicate that consumers are likely to reduce their reliance on debit if these fees are implemented.

w-12-3

Valuable Cheap Talk and Equilibrium Selection

by Julian C. Jamison

Motivation for the Research

Coordination games and other situations with multiple possible outcomes have received increasing attention in the game theory literature, but the description of equilibrium selection in such games has remained relatively informal, relying on concepts like focal points or initial conditions. While rarely formally justified, the standard Nash equilibrium idea holds that if two or more players communicate before a game, they should converge on a stable outcome—meaning that each player makes the best decisions he or she can, based on the other player's choices, and no one's outcome can be unilaterally improved.

The concept of "cheap talk" presents an intuitive method for formally describing the equilibrium selection problem and the concept of Nash equilibrium. Cheap talk describes a type of pregame communication that is defined as nonbinding, nonpayoff-relevant preplay interaction. In practice cheap talk has mainly been used in the study of signaling games, in repeated learning and game environments, and in certain applied settings. In this paper the author takes up the challenge of constructing a more comprehensive model of cheap talk that can potentially offer a more formal justification of equilibria, equilibrium selection, and Nash equilibrium.

Research Approach

The author develops a formal model of cheap talk that centers on an unlimited communication session, called a conversation, which takes place before play begins in a standard game. The model assumes that the players have full information, in order to abstract from any signaling incentives during the conversation. Each player begins with a common forecast about what actions he or she will take in the upcoming game. These expectations can be interpreted as vague initial ideas about how the game might be played, based perhaps on societal conventions or focal points, meaning any behavior that stands out as salient along a certain dimension.

This common forecast is slowly updated during the conversation phase, when players make advance announcements of what actions they plan to take in the upcoming game. An announcement is deemed credible only if it self-committing, meaning that if the other players believed it and best responded to it, the announcer would still carry through with the action. This requirement is equivalent to being part of some Nash equilibrium for the action game. If there is no external justification for believing an announcement's accuracy, it is judged to be untrustworthy and is disregarded. In this manner the conversation proceeds indefinitely and recursively, possibly but not inexorably toward some limit, and the common forecast is updated by each credible announcement. As beliefs are updated, the initial forecast may be discarded and only the actual credible announcements taken into account to form an average forecast, which constitutes a player's appearance. It is important to stress that in the author's model, players have a choice over what to say, as this is the hallmark of a conversation. The players may ignore what they themselves are "expected" to do, although they may take into account the influence this expectation has on how the other players will perceive them. This freedom of choice, along with the lack of payoffs until the game concludes, are what differentiate this paper's model from an evolutionary learning model.

The author's model is closest to Rabin (1994), in that both seek a notion of optimality rather than equilibrium in the analysis of the extended game, and adhere to the full rationality paradigm of classical game theory and previous work on cheap talk. The specific form of cheap talk used in the author's model differs from that of Rabin with respect to the element of choice between strategies against which to best respond credibly. Moreover, while Rabin's model only allows for finite communication, the author's model allows for infinite communication.

Key Findings

- The first main result is that if the conversation converges toward a limit, this limit must be a Nash equilibrium of the underlying action. Since arriving at any Nash equilibrium forms a possible limit to the conversation, this result can be interpreted as stating that any meaningful preplay communication can lead only to Nash equilibrium outcomes.

- The paper's second main result states that the optimal pregame play in the conversation stage leads to an efficient outcome, and that any efficient outcome is a possible result of such strategic conversation. In other words, rational or thoughtful speech by the players leads to an efficient outcome. Stated somewhat differently, why would players agree in advance to an inefficient equilibrium outcome for a game if another potential equilibrium outcome was available that conferred better payoffs to everyone?

- The implications of the second main result contrast with the "babbling" results presented in the previous literature. Contrary to the author's paper, this alternative view contends that it is impossible to select among the set of Nash equilibria because players ignore all pregame communication. The key to the difference is that previous studies looked for equilibria of the extended communication game as a whole, for example by assuming that the full strategies of all players are known. This full-knowledge assumption allows for equilibrium strategies in which no value is placed on seemingly mutually informative communication, whereas the author's model of cheap talk assumes that the beneficial pregame exchange of information among players will not be ignored, and that such beliefs will inform the action game.

- Taken together, the two main results form a complete theoretical connection among the concepts of cheap talk, as modeled in this paper, Nash equilibria, and Pareto optimality.

Implications

The author's model presents one possible resolution to the question of equilibrium selection, as well as to the older question of justifying the Nash equilibrium concept. The model provides a decisive solution to these two issues within the context of a single model and also applies to games with more than two players or to games where players do not necessarily exhibit common interests. Yet the model as currently specified has several drawbacks. One, the results do not prove that convergence must take place, only that if it does, it then takes a certain form. Two, since not all applications allow for preplay communication, this model cannot serve as a general justification for the Nash equilibrium concept. Three, the model imposes restrictions on the belief formation process, in the sense that over the long run it requires that some small amount of trust be attributed to credible announcements.

The author's model could be extended to include coordinated equilibrium and to introduce a stochastic element into the conversation. It would also be worthwhile to pursue experimental studies of extended cheap talk, as there is little work on this concept to date. Such a pursuit could address both general applications and examine the author's concept of stable efficiency for n-player games.

Investment in Customer Recognition and Information Exchange

by Oz Shy and Rune Stenbacka

....

Motivation for the Research

In some industries, having established customer relationships makes it possible for a firm to learn the characteristics of its individual customers, thereby enabling the firm to institute customer-specific pricing. This may be particularly true in service industries, with banking and insurance as prominent examples. In industries like these, firms often design institutions to facilitate the exchange of such customer-specific information.

Consumers may benefit from information exchange among firms because this practice facilitates the design of aggressive poaching offers (price cuts intended to lure a consumer to switch brands). For example, a consumer who initially exhibits a higher preference for a rival brand might realize a welfare benefit by accepting a poaching offer that is sufficiently competitive to compensate for the switching costs. On the other hand, within the framework of an established customer relationship, an incumbent firm facing competition from a rival firm seeking to lure its customers by targeting them with differential pricing offers based on their type-specific preferences (meaning preferences that are common within specific groups of customers) can also adjust its own type-contingent prices offered to existing customers so as to maximize the extraction of consumer surplus. General economic intuition is insufficient to evaluate the effects that the exchange of customer-specific information has on industry profits and consumer surplus, so an analytical study is needed. The authors analyze how the exchange

of customer-specific information affects industry profits and consumer welfare, as well as how such an exchange influences firms' incentives to invest in learning their customers' preferences.

Research Approach

Although a firm might be able to distinguish its own inherited customers from those of its rival at a relatively low or even negligible cost, in practice the cost of acquiring information about its customers' preferences is significantly higher. Therefore, the firm faces the optimization problem of whether to acquire customer-specific information to facilitate setting individualized prices, or whether to set prices contingent only on whether it already has an established customer relationship with an individual consumer. In addition, the incentives to acquire customer-specific information depend crucially on whether the firms have committed themselves to a system of information exchange. In order to investigate the effect of information exchange on firms' incentive to invest in customer recognition, the authors design a duopoly model in which consumers are differentiated by their switching costs. In the model, each consumer holds an individual valuation (high or low) for the two competing brands.

The authors begin by characterizing the firms' incentives to invest in learning their customers' idiosyncratic valuations regarding their (each firm's) own brand and the competing brand, and they investigate how these incentives are affected by the costs of acquiring information. The authors frame the set of decisions facing the firms when they are considering whether to engage in customer-information acquisition as a three-stage game with the following sequence of decisions: 1) each firm decides whether or not to share customer-related information, 2) each firm decides whether or not to invest in information acquisition, and 3) firms engage in price competition. The authors then conduct an equilibrium analysis of this three-stage game.

Key Findings

- Both firms invest in information acquisition when the costs of information gathering are sufficiently low, whereas neither firm invests when the information gathering costs are sufficiently high. When investment costs fall in an intermediate range, both firms invest in learning their customers' preferences provided that this information is not exchanged between the firms. The exchange of acquired, proprietary, customer-specific information harms industry profits.

- A firm's acquisition and use of information regarding customer-specific preferences as a basis for type-contingent pricing always exacts a welfare loss to the consumer. Information sharing between firms further magnifies the loss to consumer welfare.

- The case of no information sharing between the two firms supports a subgame perfect equilibrium for such a three-stage game. Furthermore, the equilibrium with no information sharing is efficient from the perspective of total welfare. Finally, the market equilibrium supports excessive investment in information acquisition for a low investment cost.

Consumers are worse off when firms acquire information about their preferences.

Implications

Firms value the informational advantage associated with learning their customers' preferences. The authors' analysis implies that a firm has no incentive to relinquish this advantage through information exchange even though such an exchange would broaden the knowledge of consumer-specific preferences to include the rival firm's customers. A central reason for this conclusion is that such an information exchange would intensify price competition and lower firm profits.

The analysis implies that information exchange among firms is not likely to be observed in industries satisfying the general features covered by the authors' oligopolistic model in which a few firms dominate an industry. Still, information exchange is observed in some industries, such as banking and insurance, where many firms compete. Indeed, as studies focusing on these types of industries emphasize, information exchange can enhance profits and serve as an efficient mechanism to overcome significant problems associated with moral hazard and adverse selection under circumstances where some types of customers may cause firms to suffer severe losses. However, as the authors' general analysis suggests, information exchange is typically an inefficient practice unless such industry-specific conditions prevail. But most firms have no incentive to engage in a voluntary exchange of customer information with their competitors.

These arguments hold true under the assumption that firms engage in noncooperative price competition. Information exchange may very well serve as a device to facilitate tacit or explicit collusion. As a policy conclusion, these findings suggest that there is merit in monitoring the firm-to-firm exchange of customer-specific information and having antitrust authorities challenge this practice when warranted. However, firms that do engage in exchanging customer-specific information should be given the opportunity to present arguments for the enhanced efficiency associated with this information exchange.

This study identifies only one potential market failure when firms make low-cost investments in acquiring customer-specific information and adjust prices accordingly. Under such circumstances the welfare loss to consumers outweighs the gains in industry profits. However, this result by itself seems insufficient to warrant restrictions on information acquisition, as firms may have to reward consumers for revealing information about themselves. Rewards can be granted in the form of points or discounts on future purchases. Indeed, one can interpret the cost parameter, c, as including the costs associated with such customer loyalty programs, but the authors' model does not capture the process by which consumers respond to such programs. Future research could enrich the model by incorporating mechanisms for how savvy customers reveal information about their types.

In order to be able to highlight the central economic mechanisms in a transparent way, the authors' model makes a number of simplifying assumptions. The robustness of the results can be questioned in light of the generality of these assumptions. In this respect, the model suggests the following topics for future research: To what extent are the results robust to an alteration of the cost structure of information acquisition, such as increasing marginal cost with respect to the number of the firm's own customers? Following the literature on information exchange, the authors have assumed that firms reveal their information in a truthful way. But can the model be extended to capture strategic information exchange, in which firms may not reveal truthfully the full extent of the information they gather? Finally, it should be emphasized that the authors' analysis has not incorporated behavioral aspects, according to which the collection and exchange of customer-specific information could potentially induce consumers to modify their behavior in attempts to, for example, defend individual privacy. To incorporate such features the model could be extended to include elements from behavioral economics.

Selecting Public Goods Institutions: Who Likes to Punish and Reward?

by Michalis Drouvelis and Julian C. Jamison

....

Motivation for the Research

Institutions reflect and reinforce social norms, so understanding how these structures are chosen and established, and what factors help to predict this process, is of great interest to economists and other social scientists. A number of real-life public goods settings, such as tax compliance, charitable donations, tipping in restaurants, and participation in group actions, have incentive structures where people's individual and collective goals are at odds, causing tensions that are exacerbated by the incentive to free ride on others. Identifying what forces determine the acceptable standards of behavior embodied in these institutions can illuminate the proximate sources of human cooperation and improve our inadequate understanding of how social norms arise and are enforced—insofar as these norms arise from self-selection into groups that prefer certain behavioral rules and norms. A burgeoning experimental literature investigates individuals' voting preferences concerning institutions and the specific rules that govern these institutions, but these studies offer mixed evidence regarding what institutions people favor. (It is generally observed, however, that democratically selected institutions perform better than institutions that are exogenously imposed, both in terms of average contribution levels and in terms of efficiencies as measured by net earnings.) The existing literature lacks studies investigating the relationship between risk tolerance and social preferences, or exploring the possibility that preferences other than standard risk preferences, such as loss and ambiguity aversion, may predict social preferences. Essentially, the literature fails to address two important issues pertaining to public goods settings: which institutions do people actually prefer and what individual characteristics have predictive power over the choice of institutions and behaviors in these institutional frameworks?

Research Approach

In this paper the authors design a controlled economic experiment that distinguishes between conflicting personal and collective gains in order to provide a complete analysis of the processes underlying the way that people choose among public goods institutions and make decisions based on the rules of these institutions. This experimental method allows the authors to elicit a number of variables that may influence subjects' choice of institutions, but the central concern rests on the individual subjects' preferences regarding risk, loss, and ambiguity aversion. The laboratory setting allows the authors to collect data on the variables of interest, a task that often cannot be performed in naturally occurring environments. Using an incentive-compatible design, the experiment elicits choices in order to construct four primary preference measures—risk aversion, loss aversion, ambiguity aversion, and ambiguity aversion over losses—and offers the first comprehensive analysis of how these measures can predict subjects' reciprocal behavior when punishing or rewarding their peers. The experimental design includes all four possible public goods environments: a standard public goods game, a public goods game with punishments, a public goods game with rewards, and a public goods game with both punishments and rewards. The subjects are randomly assigned to the various environments, so the authors can separate the effect of selection from the effect of the institutional rules per se. The comparison of behavior among the three institutions that have

punishment and/or reward mechanisms enables the authors to disentangle which particular institutional aspect is important for sustaining norms of high cooperation and maximizing individuals' overall welfare.

The experiment has a two-part design. In the first part, the authors elicited the individual subjects' preference levels by showing them a table with seven rows and asking them to choose between a safe option and a lottery option in each row. The safe option amount of £6 was the same in all seven rows, while the lottery option amount, with a 50 percent chance of either receiving nothing or winning, increased in each succeeding row, with respective amounts of £11, £12, £13, £14, £16, £18, and £20. After a subject made a decision for each row, it was randomly determined which row became relevant for payoff. Subjects learned their lottery payment at the end of the experiment, a procedure that guaranteed that each decision was incentive-compatible. The number of times a subject chose the safe option indicated his or her attitude toward risk, so the more times the sure payoff of £6 was chosen, the greater the subject's risk aversion. Since public goods environments in the experiment (and in real life) involve payoffs that are ambiguous and may entail losses, the experiment elicited individual attitudes towards loss aversion, ambiguity with loss aversion, and ambiguity aversion with losses. The authors adapted the risk preference procedure to elucidate these attitudes. Individual attitudes toward losses were elicited by using the same table for the risk preference choice, except that the payoff amounts shifted downwards by £3 in each row, so the lottery payoffs included a 50 percent chance of losing £3 and a 50 percent chance of winning a positive amount. Loss aversion measured the frequency with which a subject chose the safe option. To gauge ambiguity aversion with and without losses, the probability of each outcome, made explicit in the lottery option, was replaced with a question mark to indicate uncertainty. To control for order effects, subjects were shown the four tables in a random order.
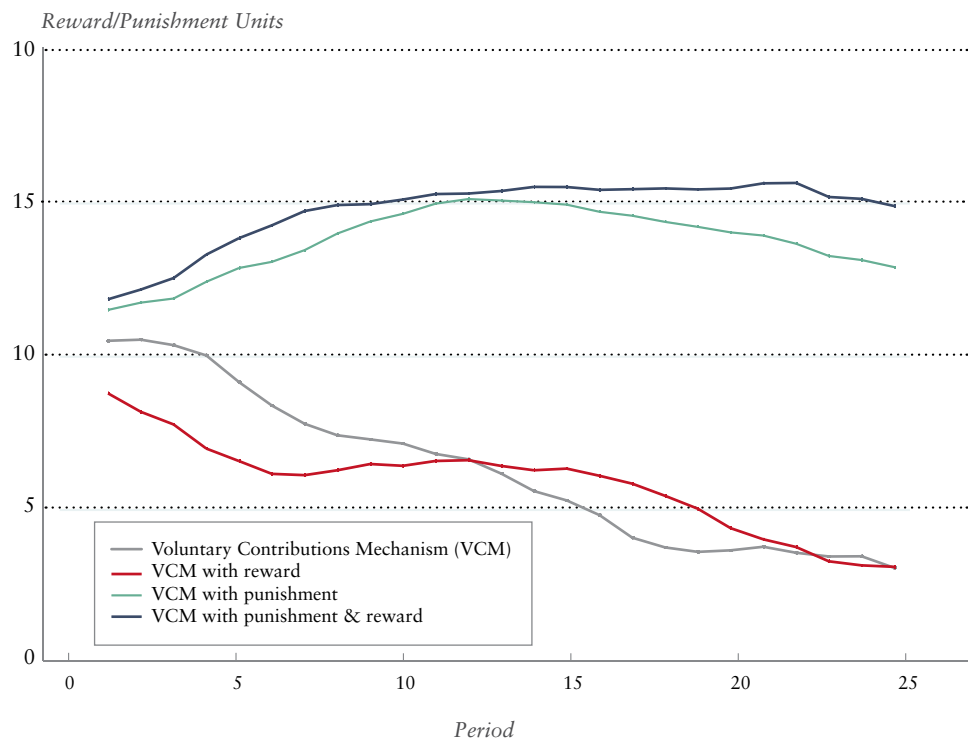
The objective in the second part of the experiment was to elicit the subjects' preferences over the four public goods environments. Subjects were randomly assigned to the same treatment in three groups of four to play a repeated voluntary contribution game. Each experimental session involved only one treatment, but all four treatments were used in the end, since there were multiple sessions. The literature on public goods games suggests that there are different payoff implications among these four public goods institutions, depending on the rules governing the particular game. Using a voluntary contribution mechanism (VCM), the four treatments correspond to the four public goods environments, respectively: the VCM, the VCM with punishment, the VCM with rewards, and the VCM with punishment and rewards. The baseline VCM treatment captures the conflict between private and social interests, and has linear payoffs. To conform to the standard approach used in the literature, subjects received tokens, and had to decide how to use an initial endowment of 20 tokens, meaning how many to keep and how many to contribute to a public goods project. Each token retained earned a subject one money unit, while for each token contributed to the project, each subject earned a return of 0.4 money units, resulting in a gain of 1.6 tokens for the group as a whole. This baseline treatment measures the extent of the subject's self-interested behavior, as a selfish member has an economic incentive to contribute nothing to the project and free ride on others, while the socially efficient choice requires all group members to contribute their entire endowment to the project, which would mean that each group member would ultimately receive a payoff equal to 32 money units, an amount that exceeds the initial endowment.

The VCM with punishment treatment was identical to the baseline treatment, except that it had one additional stage. After subjects made their contribution decisions, the other

members' contribution profiles were revealed at the start of a second stage, but their identities remained anonymous. Each subject could then assign between 0 and −5 penalty points to the other group members. Each negative point cost the punisher one money unit and the punished member three units. At the end of this second stage, the subjects learned the cost they may have incurred for assigning penalty points, the total number of penalty points they received, and their earnings from each period. No information was given about the number of adjustment points assigned to the other group members, so they learned nothing about possible social norms regarding punishment. The VCM with rewards treatment had a two-stage structure similar to the VCM with punishment, but in the second stage the entire group learned the individual contributions made by other group members during the first stage. Then, each subject had the opportunity to assign positive points to the other group members. Assigning positive points cost the donor but benefited the recipient. The VCM with punishment and rewards combined the two separate punishment and reward treatments. After subjects made their contribution decisions in the first stage and the group's contribution profile was revealed, each subject had the opportunity to assign up to five negative or five positive points to each of the other members. Each negative point cost the punisher one money unit and the punished member three money units, while each positive point reduced the donor's earnings by one unit and increased the recipient's earning by one unit.

Before each treatment was administered, the subjects read the instructions and then answered a number of computerized control questions to ensure they understood the decision stakes and the payoff calculations. The experiment did not start until everyone was clear on the procedure. Subjects were then asked to indicate on a percentage scale how much they expected to earn relative to the maximum potential payoff, considering only earnings from the 25

Average Contributions in Each Period by Treatment

Reward/Punishment Units

Legend:
- Voluntary Contributions Mechanism (VCM)
- VCM with reward
- VCM with punishment
- VCM with punishment & reward

Period

Source: Authors' calculations.

rounds of the baseline public goods game. This assessment was used to gauge how overconfidence on the part of a subject might influence self-selection into a particular public goods game. After answering the overconfidence question, the subjects were then asked to indicate which public goods institution treatment they preferred to participate in by quantifying how much each institution was worth to them, expressed in a monetary sum. They were instructed that if assigned to their preferred institution, the monetary amount they indicated would be subtracted from their final payment. If they were randomly assigned to an institutional setting that they indicated they would need to be paid to participate in, this amount was added to their final payment. The allowable payment range was between –£5 and £5, and the amounts given for all four institutions were required to sum to zero. Order effects were controlled for by randomizing the onscreen presentation of each institutional setting. This incentive mechanism allowed the participants to truthfully express the ordinal ranking of their institutional preferences as well as the strength of their preferences. The subjects were then told which of the four environments they were assigned to. They played the game for 35 rounds, then were told of their payoffs from the lottery task, the overconfidence measure, and their earnings from the public good game. Finally, the authors collected data on the individual subjects' demographic characteristics (gender, age, nationality, marital status, father's level of education, and political and religious affiliations) and on a self-control task correlated with cognitive outcomes.

The experiment consisted of 16 sessions, four each for each treatment. A total of 192 subjects participated, 48 in each treatment. All of the subjects were recruited at the University of York in the United Kingdom, and most were undergraduates with various majors.

Key Findings

- The four institutional preference measures are significantly correlated. The results provide experimental evidence that a subject's individual characteristics, such as age and political and religious affiliation, help to explain the subject's economic preferences regarding risk, loss, and ambiguity. In turn, these preferences affect behavior in public goods settings, but surprisingly do not affect the initial institutional preferences. Age is a statistically significant determinant and is related to lower aversion to risk, loss, and ambiguity. Political affiliations also affect preferences over ambiguity; relative to those with no party affiliation; those self-identified as affiliated with the Conservative party are more ambiguity averse, whereas those affiliated with a party other than the four major ones in the United Kingdom (Conservative, Labour, Liberal Democrats, and Green) are less ambiguity averse. Subjects who report a religious affiliation (other than Catholic or Protestant, for which there is no effect) are less risk averse than those with no religious affiliation.

- Surprisingly, the institutions that individuals prefer are not influenced by preference measures, although other individual traits do have some explanatory power. Risk aversion is positively (negatively) correlated with points assigned in the VCM with punishment treatment (VCM with punishment and rewards treatment), whereas loss aversion is positively related with the point assignment in the VCM treatment with punishment and rewards.

- Institutions with punishment opportunities are best able to able to maintain cooperative norms. Relative to institutions without punishment options, institutions that permit sanctions incur enforcement costs that lower overall welfare in the short run but increase overall efficiency in the long run. This result confirms previous experimental findings.

A subject's individual characteristics, such as age and political/religious affiliation, help to explain the subject's economic preferences regarding risk, loss, and ambiguity

- Positive and negative reciprocity are significantly correlated with the preference measures. Negative and positive deviations also determine the expected punishment and reward responses. Other individual characteristics, such as nationality, gender, age, and whether the subject was an economics/business major, play a major role in how individuals actually use the sanctions and rewards.

Implications

To the best of the authors' knowledge, this paper is the first experimental study to use an incentive-compatible method to elicit preferences regarding public goods institutions and the intensity of these preferences. Better understanding the self-selection process can help to improve the design of public institutions that promote social welfare and cooperation, as the paper demonstrates the significance of individual traits for preference measures. A natural avenue for further research is to investigate what happens when subjects have amassed some degree of experience with the various institutional environments. This paper provides evidence that risk tolerance and social preferences should be incorporated into economic analysis, as these are related notions that help in understanding certain aspects of economic behavior. However, additional work remains to be done to adequately explain how individual traits interact with economic preferences.

Public Policy Briefs

b-12-1

Long-Term Inequality and Mobility

by Katharine Bradbury

Motivation for the Research

While the Occupy Wall Street movement and speeches by President Obama and Alan Krueger (chair of the Council of Economic Advisers) have raised considerable controversy regarding income mobility and inequality, most observers agree that the United States has a problem of low upward mobility for those at the bottom of the economic ladder. Although most of the recent discussion has concerned intergenerational mobility (how well adults fare economically relative to the families in which they grew up), in this brief the author provides some new data documenting the problem of low upward mobility within a 10-year period. The up and down movements in family income experienced by working-age family heads and spouses over this shorter interval help to indicate the degree to which families can improve their economic prospects and/or retain their advantages. This research is related to intergenerational mobility because if the family circumstances in which children grow up help to determine their economic prospects as adults, then the prospects of children from families with persistently low incomes will be more limited than if their parents' incomes were higher during at least some portion of their childhood.

Economic inequality across families has increased substantially in the United States since the 1970s. Some analysts argue that because overall U.S. mobility is so high we need not worry

about rising inequality. This viewpoint contends that even though the wealthiest income group is now much richer than it was 30 years ago when compared with middle- or low-income groups, no one remains poor or rich for long, so it all evens out over time. Short spells in poverty are deemed problematic if they result in serious deprivation, but become an even greater concern if those who are poor fail to escape poverty in two or five or 10 years. In other words, the real concern is not with single-year inequality but with long-term inequality across individuals' average income over a longer period, such as a decade. This longer-term inequality reflects year-to-year changes in income (mobility) as well as short-term inequality.

If there is any mobility at all, the long-term distribution of income is more equal than the short-term income distribution because when people's incomes change from one year to the next, these up and down moves mitigate the inequality of any one year's income. And the long-term distribution is more equal: for example, according to the Gini coefficient—a widely used measure of inequality—the 10-year average post-government family income for family heads and spouses between the ages of 16 and 62 during the 1996-to-2006 period was 0.28, while the Gini coefficient averaged across these same individuals' single-year family incomes during the same years was 0.33.

But even though long-term inequality is lower than single-year inequality, if single-year inequality is rising (as it has been in the United States in recent decades), then mobility (year-to-year income changes) would have to increase in order to prevent a rise in long-term inequality. This brief investigates whether mobility has increased or fallen, analyzing the mobility and income situation of family heads and spouses who have low long-term incomes, where long-term refers to average family income over a 10-year period.

Research Approach

The author analyzes family income data from the Panel Study of Income Dynamics (PSID) and Cornell University Cross National Equivalent File (CNEF) covering the years 1976 to 2006, to understand the underlying longer-term inequality and mobility patterns these data reflect. She categorizes all U.S. households by their long-term income quintile (fifth of the U.S. family income distribution gauged over a 10-year period) for three periods: 1976–1986, 1986–1996, and 1996–2006. She does this by averaging post-tax, post-transfer incomes of U.S. families, observable in the data for each year within the 10-year period, and dividing the resulting distribution of 10-year average incomes into fifths from poorest to richest. She then looks at whether most families' 10-year average income reflects many year-to-year moves smoothed out by averaging, or very few moves, with single-year income changing little from year to year. Because data are available only for every other year, each 10-year period includes six observations of single-year income.

Key Findings

- For the poorest and richest U.S. families—those in the poorest or richest quintile of the long-term distribution—mobility is quite low. During the 1996–2006 period, 51 percent of the poorest-quintile individuals remained in the poorest quintile in all or all-but-one of the six single-year income observations; the corresponding figure for the long-term richest-quintile individuals was 54 percent.

- Furthermore, those households in the long-term poorest or richest quintile who moved outside the corresponding single-year quintile in one or more years did not typically move far. Over three-quarters (78 percent) of the members of the long-term poorest quintile

Mobility of the Long-Term Poorest and Richest

Percentage of people in the poorest or richest long-term quintile who spend five years or six years in the single quintile

Percentage of people in the poorest or richest long-term quintile who spend five years or six years in that or an adjacent single quintile
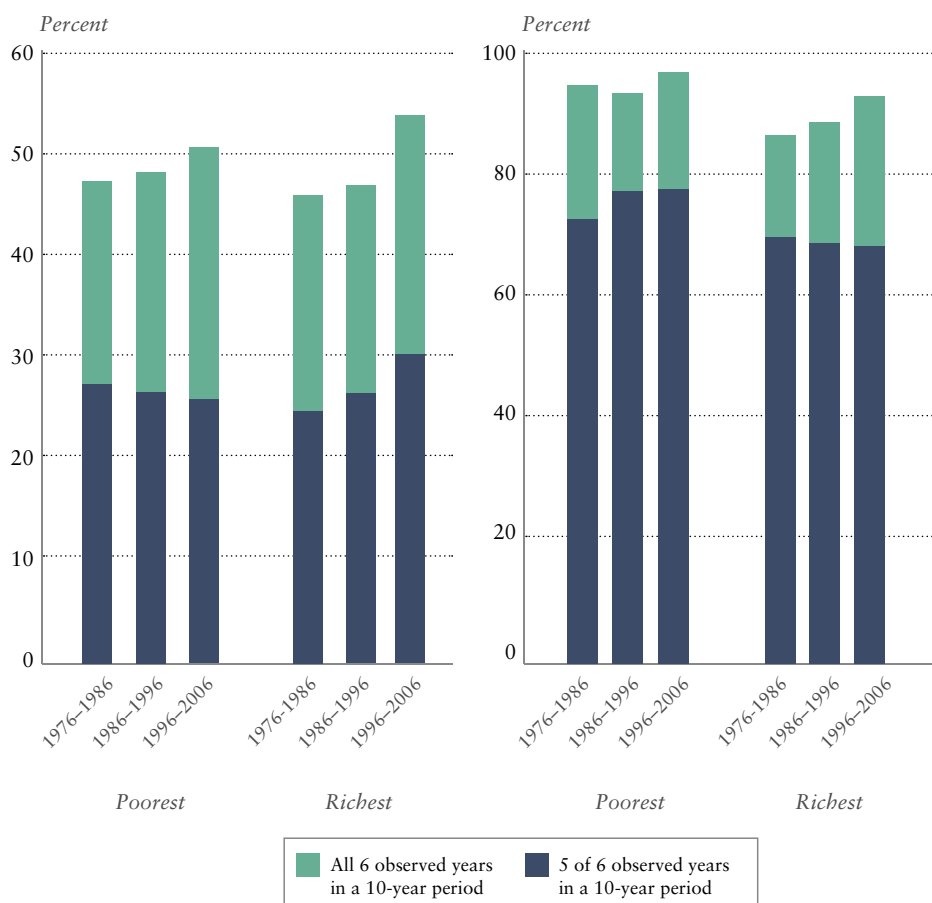
Source: Author's calculations based on data from Panel Study of Income Dynamics (PSID) and Cornell University Cross-National Equivalent File (CNEF).

Most of the long-term poor are stuck at the bottom, and most of the long-term rich have a strong grip on the top.

were inside the single-year poorest or adjacent second-poorest quintiles in all six observations during the 1996–2006 period, and 97 percent were in the poorest two quintiles for at least five of the six observations. For those in the richest long-term quintile, the corresponding figures were somewhat lower: 69 percent and 93 percent. Thus, a miniscule 3 percent of the poorest long-term quintile saw incomes rise beyond the second-poorest quintile (above the 40th percentile) in more than one of the six years, and only 7 percent of the richest long-term quintile experienced below-60th-percentile income in more than one of the six years. Not surprisingly, this lack of mobility in relative (quintile) terms shows up in average income changes as well: averaging the real income changes experienced by all heads and spouses in the poorest long-term quintile of the 1996–2006 period shows an average rise of 0.053 in log income (approximately 5.3 percent increase in real family income) between 1996 and 2006, while those in the richest 1996–2006 long-term quintile saw their single-year log incomes rise by an average of 0.276 (27.6 percent).

• Most of the immobility figures for the richest and poorest long-term quintiles over the three 10-year periods—1976–1986, 1986–1996, and 1996–2006—were higher in the last

decade than they were 10 or 20 years earlier (see figures). The bottom line: most of the long-term poor are stuck at the bottom, most of the long-term rich have a strong grip on the top, and each of these two groups is somewhat more entrenched than the same groups were 20 years earlier. Even families in the three middle quintiles have become less likely to see a range of income positions during a 10-year period.

- A similar time pattern emerges for income changes within a 10-year period, meaning between the first and last year of the period. Members of the poorest long-term quintile in the 1976–1986 period saw slower real growth in single-year income between 1976 and 1986 than members of the richest long-term quintile, but that gap was not as large as it was 20 years later for corresponding quintile members in the 1996–2006 period. It is also the case that the fraction of the poorest long-term quintile enjoying faster income growth than the overall average was lower in 1996–2006 than in 1976–1986. The gap between that fraction and the percentage of richest-quintile members enjoying above-average growth widened from 1976–1986 to 1986–1996 (from 7 percentage points to 10 percentage points) and widened again (to 16 percentage points) from 1986–1996 to 1996–2006. In all three 10-year periods, more than half (53 to 59 percent) of the members of the richest long-term quintile saw their incomes grow faster than average during the period, and less than half (43 to 47 percent) of poorest-quintile members experienced above-average income growth.

- Not only are the rich and poor somewhat less likely to move far afield from year to year in the relative (quintile) sense, but they have moved markedly further apart in income levels, as rising inequality implies. As discussed above, this fact is well known in the cross-section of single-year incomes and it is true even in terms of 10-year-average incomes. The ratio of the median of the richest-quintile of the long-term income distribution to the median of the poorest quintile of the long-term income distribution climbed from 3.0 in 1976–1986 to 4.3 in 1996–2006.

- Of course, there will always be a poorest long-term quintile—that is the nature of relative rankings. However, the situation of the poorest long-term quintile in the United States has worsened compared with 20 years earlier: in terms of real income and income growth, members of this quintile have become more isolated (have seen less relative movement from year to year) as well as poorer. These changes have left families with the lowest long-term (10-year average) income in 1996–2006 in a troubling situation: half of them spent either no years or only one-sixth of the years in this period in a quintile of the single-year income distribution that was higher than the poorest. Those families who escaped the poorest quintile typically did not go far up the income ladder, as three-quarters of these households spent every year in the poorest or second-poorest quintile, and fully 97 percent spent at most one-sixth of this time outside the poorest or second poorest quintile. The long-term poorest in the 1996–2006 decade had less than one-quarter of the median long-term income of the richest quintile in this same period.

Implications

Although these data are purely descriptive, the reduced income mobility of the poorest U.S. family heads and spouses and their relatively stagnant real income levels strongly suggest limited individual economic opportunity. Furthermore, these findings suggest that low-income parents are less able to raise their children's prospects, even over a 10-year time span. While we would like to gain a better understanding of what factors have caused this deterioration in economic opportunity, policymakers concerned about equal opportunity will want

to investigate strategies to loosen the tight connection between single-year and long-term income at the bottom of the U.S. income distribution.

r-12-1

When the Tide Goes Out: Unemployment Insurance Trust Funds and the Great Recession, Lessons for and from New England

by Jennifer Weiner

MOTIVATION FOR THE RESEARCH

The unemployment insurance (UI) program in the United States is a federal-state program that was established by the Social Security Act of 1935. Its primary objectives are: (1) to provide temporary, partial compensation for the earnings individuals lose when they become unemployed through no fault of their own and (2) to serve as a stabilizer by injecting additional resources into the economy in the form of benefit payments that are likely to be spent immediately during an economic downturn.

Each state, as well as the District of Columbia, Puerto Rico, and the Virgin Islands (collectively referred to as "the states"), operates its own UI program within federal guidelines. In these programs employers pay state taxes, which in turn are deposited in trust fund accounts maintained by the federal government. Monies in the accounts are then used to pay benefits to the unemployed. Employers also pay a separate federal UI tax which is used to support program administration, to pay for extended benefits in times of high unemployment, and to provide loans to states that have exhausted their trust funds.

Between the onset of the Great Recession in 2007:Q4, and 2011:Q2, at least 35 states borrowed from the federal government in order to continue paying UI benefits after depleting their trust funds. Among the New England states, only Maine's trust fund remained solvent throughout this period. By mid-2011, 30 states, including Connecticut, Rhode Island, and Vermont, continued to carry outstanding loan balances totaling a combined $42 billion. With principal and interest payments on these loans now coming due, many states are raising taxes on employers, potentially slowing the economic recovery.

This paper examines why, during the Great Recession and its aftermath, some state UI programs became insolvent and needed to borrow funds from the federal government while others did not. It places special emphasis on the New England states, examining the solvency of their trust funds over time and the reforms they have proposed or enacted. The paper draws on lessons from the states to identify options for policymakers that may help to strengthen the solvency of UI trust funds during future downturns.

Trust Fund Flows Relative to Total Wages, Average of all State UI Programs, 1970–2009

Percent of total wages

Legend:
— Inflows (Contributions)
— Outflows (Benefits)
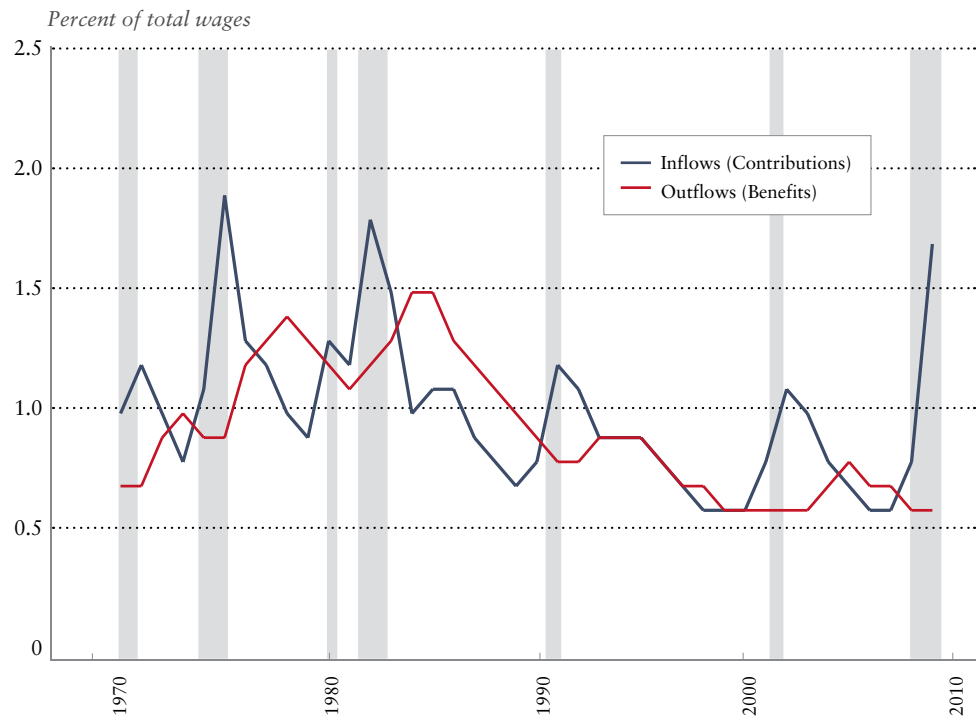
Source: Department of Labor Employment and Training Administration.
Note: Shaded areas approximate official recessions, as set by the National Bureau of Economic Research.

Research Approach

The author provides background on both the federal and the state components in the UI program, discusses how the Great Recession impacted the program and whether it should be a concern if state UI trust funds are depleted during an economic downturn, analyzes factors associated with insolvency, and draws key lessons from the New England states. She then focuses more sharply on New England, examining the region's experience in the years leading up to, during, and following the Great Recession. She then draws conclusions from the paper's analysis and discusses policy options. In an appendix, the author provides individual state narratives for each of the New England states.

Key Findings

- There is a strong relationship between a state's borrowing activity during or after the Great Recession and the financial status of its trust fund at the start of the downturn. As economic growth receded over the course of the Great Recession, it became clear that some states were ill prepared for even a milder downturn. The states that borrowed most heavily also faced higher unemployment, on average, than other states. All the borrower states had, on average, lower ratios of taxable to total wages than states that did not take out loans, but the borrower states did not necessarily have more generous UI benefits.

- An erosion of a state's taxable wage base—that is, the portion of an employee's wages that is subject to UI taxes—appears to have been an important contributing factor to the solvency issues faced by many states, including those in New England, during the Great Recession. When the taxable wage base does not grow with average wages it can lead to

UI Trust Fund Borrowing in or after the Great Recession

by Duration of Borrowing

Legend:
- No loans
- Fewer than 8 quarters
- 8 or more quarters
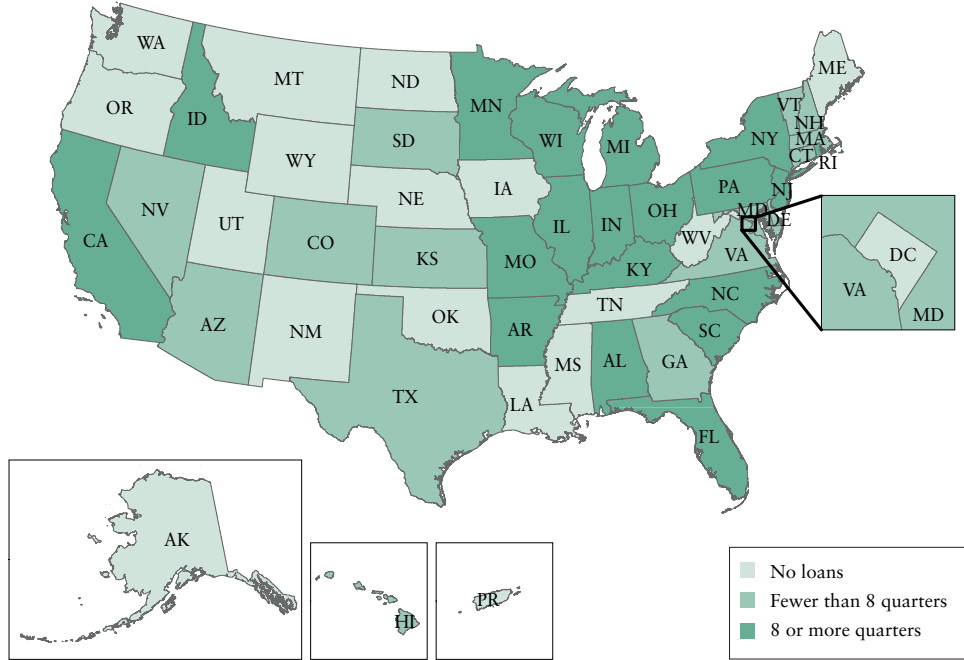
Note: Duration measured by the number of quarters between 2007:Q4 and 2011:Q2 in which state had an outstanding loan balance in peak quarter. The Virgin Islands, not pictured, fall in the heaviest borrowing group.

by Magnitude of Borrowing

Legend:
- No loans
- Less than 3.3 percent
- 3.3 percent or greater
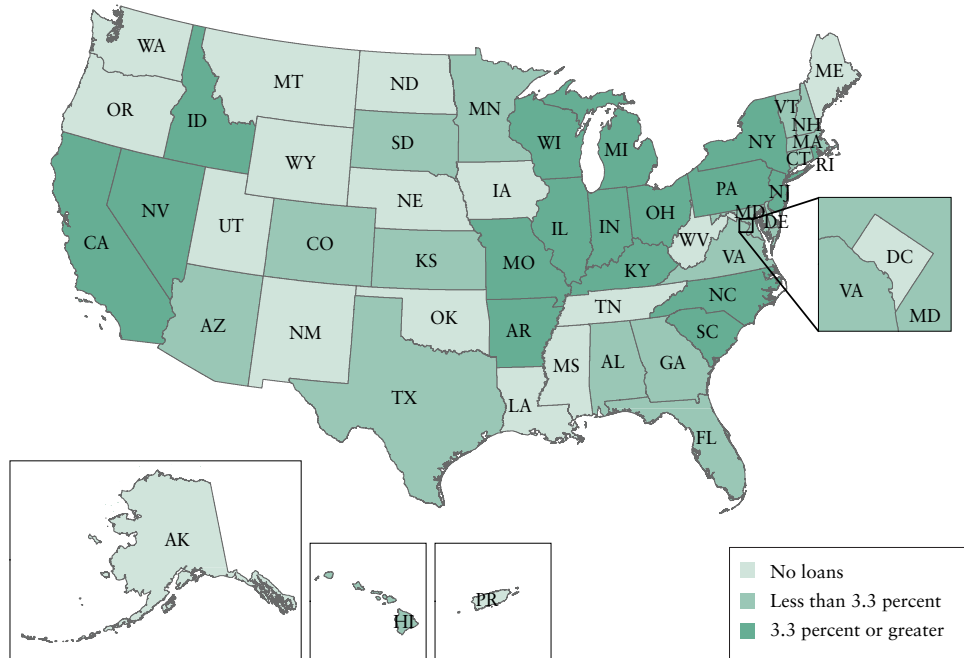
Note: Magnitude measured as a state's peak loan balance between 2007:Q4 and 2011:Q2, measured as a percent of total state wages in the peak quarter. The Virgin Islands, not pictured, fall in the heaviest borrower group.

Source: Department of Labor Employment and Training Administration.

a structural imbalance between taxes flowing into the UI trust fund and benefits flowing out, as the latter are based on unemployed workers' previous earnings.

- Examples from New England also illustrate how unbalanced reforms—that is, those that cut taxes without reducing benefits or those that increase benefits without also raising tax revenues—as well as low trust fund targets can lead to solvency problems.

- Maine's ability to weather the Great Recession may be credited to reforms undertaken in the late 1990s, when the economy was performing well. The state raised its taxable wage base and introduced a new method of assigning employer tax rates that spreads contributions more evenly across employers and gives the state more control over the amount of revenue flowing into its trust fund in a given year.

IMPLICATIONS

While a state's aim should not necessarily be to build a UI trust fund sufficient to withstand even the most severe economic downturn, taking measures to strengthen long-term UI solvency can help to promote the stabilizing impacts of the program and may limit the need to borrow from the federal government during future downturns. The author's findings suggest that to strengthen UI trust fund solvency and reduce the possibility of having to make up shortfalls in future downturns by borrowing from the federal government, states should consider: 1) increasing and indexing the taxable wage base, 2) avoiding unbalanced reforms, and 3) re-examining employer tax rates.

Contributing Authors

Katharine Bradbury is a senior economist and policy advisor in the research department of the Federal Reserve Bank of Boston.
http://www.bostonfed.org/economic/econbios/bradbury.htm

Julia Dennett is a research associate with the New England Public Policy Center, part of the research department of the Federal Reserve Bank of Boston.

Michalis Drouvelis is a lecturer in the department of economics at the University of Birmingham.

Christopher L. Foote is a senior economist and policy advisor in the research department at the Federal Reserve Bank of Boston. http://www.bostonfed.org/economic/econbios/foote.htm

Kristopher S. Gerardi is an economist and an assistant policy advisor at the Federal Reserve Bank of Atlanta.

Fumiko Hayashi is a senior economist in the payments system function of the economic research department of the Federal Reserve Bank of Kansas City.

Julian C. Jamison is a senior economist in the research department of the Federal Reserve Bank of Boston. http://www.bostonfed.org/economic/econbios/jamison.htm

Alicia Sasser Modestino is a senior economist in the New England Public Policy Center, part of the research department of the Federal Reserve Bank of Boston.
http://www.bostonfed.org/economic/econbios/sasser.htm

Lucie Schmidt is an associate professor of economics at Williams College and a visiting scholar in the New England Public Policy Center, part of the research department of the Federal Reserve Bank of Boston.

Scott Schuh is director of the Consumer Payments Research Center and a senior economist and policy advisor in the research department of the Federal Reserve Bank of Boston.
http://www.bostonfed.org/economic/econbios/schuh.htm

Oz Shy is a senior economist with the Consumer Payments Research Center, part of the research department of the Federal Reserve Bank of Boston.
http://www.bostonfed.org/economic/econbios/shy.htm

Joanna Stavins is a senior economist and policy advisor with the Consumer Payments Research Center, part of the research department of the Federal Reserve Bank of Boston.
http://www.bostonfed.org/economic/econbios/stavins.htm

Rune Stenbacka is a professor of economics at the Hanken School of Economics in Helsinki.

Jennifer Weiner is a senior policy analyst in the New England Public Policy Center, part of the research department of the Federal Reserve Bank of Boston.

Paul S. Willen is a senior economist and policy advisor in the research department of the Federal Reserve Bank of Boston. http://www.bostonfed.org/economic/econbios/willen.htm