Consumer Finance and the "Richmond" View

Kartik Athreya
Federal Reserve Bank of Richmond
March 5, 2007

In the 1930s, economist Irving Fisher posed what would become the canonical consumer finance problem on which nearly all later work is based. A household faces an earnings stream over its lifetime that does not line up with its desired spending pattern: The fundamental "problem" of consumer finance is, therefore, one of using assets and liabilities to match earnings with desired spending patterns. Without such an obstacle, families would neither borrow nor save, and consumer finance would be irrelevant. However, given the empirical relevance of both predictable and unpredictable variations in their income, households must indeed address the question of how to use financial instruments to maintain a stable lifestyle.

Fisher put it this way in his seminal *Theory of Interest* (1930): "There is an eternal conflict between the impulse to spend and the impulse to invest. The impulse of a man to spend is caused by his impatience to get enjoyments without delay, and his impulse to invest is caused by the opportunities to obtain by delay relatively more enjoyment either for himself or others." This conflict is resolved not by looking at "statistical records of the past" but by turning to the future. "What is to be one's future income stream … becomes of supreme importance." The emphasis Fisher placed on expectations about the future as a determinant of current actions was prescient and has greatly influenced the course of modern research on consumer finance.

In this essay, we discuss the evolution of consumer finance theory, from Fisher to the present. Economists' understanding of consumer finance has advanced considerably over the years. Particularly, in recent years, the ability of a class of models to account for a variety of observed dimensions of household spending and savings behavior has improved significantly. The story of this advance is as much a story of technology, both mathematical and computational, as anything else.

Interpreting historical behavior and predicting future patterns first requires a theory about how consumers make financial decisions. What are people considering when they choose how much to spend, how much to borrow, and how much to save? By themselves, the data tell us little. For example, if we see the pace of consumer defaults rising, what might we conclude? One possibility is that improvements in financial intermediation have made credit-granting decisions easier and have led to greater borrowing by risky groups previously denied credit. Another is that nothing has changed in the lending industry but that households anticipated rapid future income

growth, which led them to borrow. However, for those whose income failed to grow as expected, default proved useful, leading overall bankruptcy rates to rise. Still another explanation is that neither lender behavior nor income expectations have changed, but instead that there is no longer any "shame" in defaulting on debts.

Each of these explanations may partially account for the facts, and some may fail altogether. The purpose of a theory of consumer finance is to have a systematic way of organizing the data to allow us to evaluate competing narratives, such as the three above. Modern economics develops theories in the form of mathematical models of household and firm decisionmaking in which their collective behavior is required to be consistent with the feasibility requirements imposed by the model. This is known as an "equilibrium" approach. The Federal Reserve Bank of Richmond's Consumer Finance Web site showcases analyses that satisfy this approach, rather than attempting to provide exhaustive review in a vast field.

Equilibrium analysis may be clearly contrasted with an alternative still prevalent in consumer finance, one that places far less emphasis on modeling explicit decisionmaking. The latter approach instead relies on summarizing observed features of the data, usually using regression analysis, and treating the correlations as being informative for the effects of policy.

Recent research, including some by Richmond Fed economists, has used theory explicitly to understand aspects of household borrowing. With respect to rising consumer default rates, in particular, this body of work suggests that improvements in lending technologies are a promising candidate for explaining both borrowing and default behavior in the past two decades, while mere reductions in "stigma" are not able to match the data.

In what follows, we first describe the modern approach to thinking about consumer finance, especially the interaction between households and financial institutions, and the forces affecting them both. We then discuss the key methodological improvements that allow economists to say more than ever before on a topic in which discussions are heavily influenced by larger societal concerns of equity and distributional justice.

THE PLAYERS: 1) HOUSEHOLDS

Today's households must manage income to meet the demands of mortgage payments, credit card bills, equity loans, car loans, out-of-pocket health expenses, college education, and a variety of other obligations. Also, the income households receive is determined in complex ways.

Households earn labor income – their paychecks – but many also hold interest-bearing assets, which have become much more widely held. In 1950, the median U.S. household did not directly hold any stocks. Today, many hold stocks directly, as well as other more complicated instruments, such as options. Still more hold such assets indirectly via mutual funds, insurance policies, and so on.

In order to account for the borrowing, saving, and spending decisions of households who face such time-varying and uncertain income processes, mainstream economics employs so-called "utility-maximizing" households who are intelligent and purposeful in the pursuit of their objectives. Many analyses assume further that households are fairly self-interested. We use the word "fairly" because, most often, the entity being analyzed is the household as a whole. In particular, the behavior of individuals within a household is difficult to observe; an implicit assumption, therefore, is that familial conflicts are resolved in a way to allow us, as economists, to think of a single, purposeful, and self-interested decisionmaking unit.

To be more concrete, we need to specify what we mean by the "needs" of households. Economists have typically stuck to four premises on the nature of the household goals and preferred outcomes. These are, respectively: 1) self-interest 2) impatience 3) dislike of risk, and 4) a dislike of a lifestyle that involves large swings in the quality of life. With respect to the last property of preferences, households are assumed to dislike changes in lifestyle, regardless of whether they arise from poorly insured risks to income or from well-anticipated changes in earnings capacity (such as retirement). These premises are natural, and it is somewhat remarkable how powerful they can be in explaining observed behavior.

THE PLAYERS: 2) FINANCIAL INSTITUTIONS

The second key ingredient in the study of consumer finance is the set of financial intermediaries. These are the firms who channel the savings and investments of one set of households into loan products for another set of households. The list of such organizations is lengthy. It includes banks, savings and loans, thrifts, credit card companies, and check cashing outlets. To model the behavior of financial services providers, economists start by assuming that they are both savvy and single-minded in their pursuit of profits. This too should seem a natural premise to many.

THE TOOLS

The list of financial products offered by the entities mentioned above is similarly lengthy. For borrowing, households may simultaneously have access to credit cards, home equity lines of credit, reverse mortgages, overdraft allowances, and so forth. Households may also simply lower their stocks of accumulated wealth. Through the use of each of these products, households can enjoy consumption at a level that is temporarily in excess of income.

For savings, an equally large array of instruments exists, including IRAs, direct stock market participation, and the accumulation of home equity. Each of these products performs the task of allowing households to defer purchasing power to a later time or circumstance.

THE FORUM: COMPETITIVE MARKETS

The predominant arena in which borrowers and lenders interact is a large, impersonal market in which the terms of loans are sharply governed by competitive forces. In particular, there is strong evidence that the "suppliers" of credit compete intensely with each other. We believe that analysis of consumer finance, to be useful to policymakers and interested observers, must respect the observed level of competition is most likely. One implication of competition is that it substantially mitigates the potentially adverse effects of profit-seeking by lenders. Indeed, the threat of competing for profits is often what prevents "exploitative" arrangements from persisting. Conversely, therefore, when competition fails, there should be no presumption that profit maximization by lenders will serve households best.

THE CONSUMER FINANCE "PROBLEM"

Economists will have understood the consumer finance problem when they are able to account simultaneously for observed consumption, as well as for the size and composition of balance sheets over the entire lifetimes of households.

To proceed, we therefore require a model of household decisionmaking over time and under uncertainty. As Fisher pointed out long ago, consumers must predict future income in order to decide how to spend and save in the present. However, households may face considerable uncertainty about what their future holds. The central challenge for consumer finance research is

to explain how market participants use assets and liabilities to provide for consumption over time and across various contingencies. It is precisely in modeling these two facets of decisionmaking where substantial progress has occurred in the past two decades.

BETTER MODELS OF DECISIONMAKING OVER TIME

By the 1950s, the so-called "permanent income" and "life-cycle" hypotheses were born. The former was formalized by Milton Friedman in his1957 book, "A Theory of the Consumption Function," building on Fisher's observations about the consumer finance problem. Friedman found that current income matters less in consumption than "permanent" income, by which he meant a long-run average of anticipated income. More to the point, Friedman saw that people's consumption tended to be smooth throughout their lifetimes, often based on how much they expected to earn over long periods.

Also in 1957, Japanese economist Albert Ando and Italian economist Franco Modigliani published "Tests of the Life Cycle Hypothesis of Saving," which provided a natural refinement of Fisher's work.

Friedman's work, and that of Ando and Modigliani, provided economists with the first set of rules regarding how time, uncertainty, and the path of predictable changes (such as retirement) each altered the savings and spending behavior of households throughout life. Ando and Modigliani tested the prediction that the young consume in excess of their current income, the middle-aged consume less than their income and thereby save, and the old deplete these savings. (This theory was further refined over the years by Modigliani and his frequent collaborator, Richard Brumberg.) However, less recognized was that the versions of the theories that were formally explicated and tested constituted cases of what was potentially a far richer framework in which to think about saving and spending choices. They were not representative of the real environment in which households make spending and saving decisions.

Unfair Rejections of Theory?

The attention given to specific versions of the permanent income/life-cycle hypotheses was deliberate. These versions were tractable enough for economists to use the model to make predictions. Unfortunately, the results from this approach generally didn't match real-world data.

After 20 years of repeated rejections of what were really special cases, economists seemed to have little more ability in describing household spending and savings decisions.

In one study, for example, the life-cycle model explained only one-fifth of total U.S. wealth. Through the 1970s, empirical work seemed to show that many households actually would spend all of their current income rather than save or spend based on their expected lifetime earnings. A famous paper by Robert Hall and Frederic Mishkin (1982) found that as much as 20 percent of consumers defied the permanent income hypothesis because their consumption was much more sensitive to current income than Friedman had predicted it should be. In turn, these authors argued that inability to borrow was pervasive, something entirely absent from Friedman's analysis.

In all of the preceding work, economists had to compromise on descriptive realism in return for computational tractability. In addition to the absence of potentially binding borrowing constraints, this compromise included some strong assumptions about how well families could predict their future labor incomes and how they reacted to increases in risk. In particular, two common assumptions made were that households 1) had perfect foresight about the future and 2) cared most about averages and little about extremes.

The End of Certainty Equivalence

All this began to change in the late 1980s, when computer power grew enough to both allow, and merit, the exploration of a more general class of models that held the promise of being able to account for more of what was observed. In particular, this new class of models could accommodate the intuitively appealing idea that increases in the uncertainty facing households might well convince them to put off spending, and moreover, save in anticipation of more risky times ahead.

This new class of models included, perhaps most importantly, the work of Stephen Zeldes (1989) and Christopher Carroll (1991), who were among the first to analyze the consumer finance problem in an environment without certainty equivalence and without liquidity constraints. Their work showed that, even in the absence of binding constraints on borrowing, the presence of precautionary motives deterred borrowing and made households partially defer consumption in the face of uncertainty and misfortune.

Liquidity Constraints

A second improvement was the accommodation of yet more common sense into the basic model. Now, models included limits to borrowing that were "tight" enough that households might sensibly choose to borrow all the way up to them. Earlier models did not allow for constraints on borrowing that might "bind," and especially not ones that might bind only occasionally. Basically, these more primitive models assumed that people's credit lines were large enough to capitalize all future labor income, an assumption that doesn't appear to fit with how credit markets operate in the real world.

The seminal paper of Angus Deaton (1991) was among the first to develop a model that helped explain why empirical studies kept rejecting the permanent income hypothesis. Among other things, Deaton's work showed that the very presence of limited borrowing capacity could lead a household to use debt sparingly. In turn, this meant absorbing some, though not all, blows to income by lowering consumption temporarily.

Very recently, in the context of credit card borrowing, David Gross and Nicholas Souleles (2002) found some fairly direct evidence of the behavior suggested by Zeldes (1989). These authors showed that households are indeed most interested in maintaining a level of "spare" borrowing capacity. Therefore, when their credit limits are increased, households increase their indebtedness by an amount that leaves their spare borrowing capacity the same as before.

Combining Precautionary Savings and Liquidity Constraints

A key aspect of Deaton's work is that it incorporated both precautionary motives and liquidity constraints. Strikingly, Deaton's results allowed economists to account "…for important aspects of reality that are not explained by traditional life-cycle models." Deaton's work was soon extended by S. Rao Aiyagari (1994), who embedded Deaton-style households into a general equilibrium model. This step allowed economists to measure, for the first time, the extent to which purely individual-level economic uncertainty increased the size of the stock of aggregate U.S. savings as households prepared for a rainy day.

Still, challenges remained. Prominent among them was the observation that households with relatively low expected lifetime incomes saved almost nothing. This fact was addressed in the landmark paper of Glenn Hubbard, Jonathan Skinner, and Steven Zeldes (1995), who found that the incentives created by so-called "means-tested" transfer programs were powerful enough

to discourage saving by those who, if only they remained wealth-poor, would remain eligible for such transfers.

In addition to the low levels of savings seen among some households, a second puzzle was that, at any given point in time, households with relatively high incomes saved much larger *proportions* of their incomes than their low-income counterparts. Mark Huggett and Gustavo Ventura (2000) studied a model of household decisionmaking in which aggregate outcomes were required to be consistent with individual decisions. These authors found that the particular combination of income uncertainty and the U.S. social security system generated incentives that naturally accounted for savings rate differentials.

Even more recently, Gourinchas and Parker (2002) analyzed a rich environment to account for the life-cycle consumption decisions of households partitioned by education levels. A distinguishing feature of their work was the careful and formal statistical estimation of the parameters governing household behavior. This approach provided a contrast to the more informal approach of assigning values to parameters know as "calibration." Also, in "Accounting for the U.S. Earnings and Wealth Inequality" (2004), several economists modeled two additional reasons for households to save: providing for retirement and endowing their estates. Their results successfully accounted for the previously unexplained distribution of wealth in the United States. In sum, the inclusion of liquidity constraints and precautionary motives allowed for the first time a reconciliation of theory and data.

MODELING UNSECURED CONSUMER CREDIT MARKETS

Once it became apparent that liquidity constraints could be critical in reconciling theories with facts, it became incumbent on economists to produce a quantitatively disciplined account of where such constraints came from. Recent research has placed emphasis on the idea that limits to borrowing arise because borrowers may be able in some cases to avoid repayment. Such concerns are particularly relevant for discussions about the inability of those without collateral to access credit.

The literature on consumer default and its consequences is nascent (see Athreya [2005] for a detailed review). However, this work has already been useful in advancing our understanding of credit constraints, as well as the interaction of seemingly disparate policies, such

as those providing debt-relief and those providing old-age security, in shaping household balance sheets.

With these improved models, we have begun to see a dramatic increase in the congruence between theory and data. Many of the phenomena observed in the world that seemed anomalous from the perspective of the earlier generation of models now seemed to be exactly as theory predicted.

Policy Analysis and the Importance of Rational Expectations

A key feature of the preceding evolution of theory is the interplay between data and models. It was the lack of fit of the earlier models that, along with technological progress, directly motivated economists to make new extensions. The evolution of this work is a central example of "quantitative" equilibrium theory; that is, reasoning about economic forces using specifically parameterized models that are restricted to reproduce real-world outcomes under prevailing policies and constraints.

An example relevant to consumer finance is the case of understanding how the deductibility of mortgage interest affected household borrowing. Quantitative equilibrium theory of such a change would proceed by first locating (1) behavioral parameters, such as aversion to risk and impatience, and (2) technological parameters, such as the resource cost of lending, whereby the resulting model of household and firm behavior matched salient features of the data *prior* to the imposition of mortgage interest deductibility. These parameters are likely to reflect "deep" attributes of household and firm behavior. That is, they are unlikely to change if policies do, allowing us to more confidently study the implications of changes in policy. Having used the model to infer deep parameters of preferences and technology, a quantitative equilibrium approach to policy evaluation would impose the proposed deductibility policy, recompute the household's optimization problem, and then evaluate the outcome.

Why is it useful to find parameter values of a model to match the data? After all, could we not simply stare at data, perform a purely statistical analysis, and maybe learn from the results? Since the seminal work of Robert E. Lucas, Jr., in the 1970s, economists have become sensitive to the pitfalls of using history to learn about the effects of future policies, especially for those that are novel and far-reaching. The so-called "Lucas critique" (1976) pointed out that

presumably structural relationships in the economy were in part products of past policy. Lucas' work forced economists to push expectations to the forefront of consumption research.

The argument is simple and powerful. As discussed earlier, what we see in the data is usefully interpreted as the outcome of purposeful decisionmaking by the principal actors in the economy: households and firms. However, this decisionmaking depends first on the incentives created by public policy in place at the time decisions are made, and second on expectations about future policy.

Consider a football game. If painstaking data analysis reveals that the instances in which teams gained the most yardage were on passing plays, would it make sense for teams to drastically increase the amount they passed? A little reflection suggests that it probably won't. Any opponent would alter its behavior in order to accommodate this change in strategy.

While seemingly unrelated to economic policy analysis, this analogy teaches us that 1) the data are an outcome of optimization under a given policy regime, and 2) when policies change, so might behavior. This is a potentially serious problem for empirical work in macroeconomics. After all, we do not have the luxury, in most cases, of running highly controlled experiments on citizens to learn how they'd respond. Instead, we must be clever and insist that our models match observed behavior under current policy. Consequently, to predict how policies would alter outcomes, we must explicitly re-analyze household decisionmaking under a proposed policy, and then compare the results. The outcome of this process thereby overcomes the thorny problem of using data to learn about the effects of proposed, but historically novel, policy changes. The Richmond view is that work on consumer finance, especially work dealing with predicting the effects of policy, must respect Lucas' insight.

`What We See Is Influenced by What We Don't See

The modern approach to consumer finance, in addition to overcoming the difficulties facing empirical policy analysis, also proves useful in overcoming a related problem of "selection-bias." For example, consider a setting in which individuals can choose to work outside the home or give up labor income to work in the home. Next consider how changes in tax policy alter behavior. Ideally, sufficiently rich variation in a sample of households could help us solve this problem. However, we generally observe in the data only the characteristics of those households for whom it was better to work outside the home than otherwise. If there are systematic differences between

these households and the ones who stayed home, then it will be difficult to learn how people as a whole would respond to any change in working conditions.

More concretely, if working outside the home involved a large fixed cost, for example, it may well be true that households currently working might not change their behavior, as they are disproportionately those for whom such costs are either low or worth bearing. However, over time, many other households may simply choose never to enter the labor force. Any statistical analysis restricted only to those who were working at the time of a policy change could be seriously misleading.

In contrast to a statistical analysis, an equilibrium model is advantageous in that it will deliver the full range of decisionmaking for all conceivable situations that may face households and firms. In turn, we can learn exactly what drives people to (in this case) stay at home, and then have a clearer view of how those who currently do not, might change their behavior if we changed policy.

WHAT ABOUT CONSUMER *FINANCE*?

Our tour of the literature to this point still leaves a central question: How *do* households structure their borrowing and their savings in order to attain the consumption path they most prefer? How do they finance their consumption?

In other words, the careful reader will have noticed that we have been silent thus far on the specific instruments through which households borrow and save. Instead, we have studied a literature that, with few exceptions, simplified the "finance" part of the consumption/savings problem to the point where households chose net holdings of a single asset. Negative holdings of such assets were interpreted as borrowing, and positive holdings as saving. It is precisely this long oversimplified area that is now the subject of a rapidly growing area of high-quality research.

PORTFOLIOS

One aspect of the attention to unsecured credit discussed earlier was that it forced economists to confront the observation that many households held debt in unsecured forms at the same time that

they held debt in secured forms, such as mortgages. Understanding the *portfolio* decisions of households with access to multiple sources of funding became critical. The fresh attention given to the composition of household liabilities has also uncovered seemingly anomalous behavior. A nice example of this is the widespread use of high interest-rate credit card debt, including that by households with positive levels of wealth in the form of bank accounts or home equity. Why don't these households simply draw down their equity and repay their credit cards?

On the asset side, the simplicity of the models cited earlier left them unable to account for a variety of unintuitive observations on the savings behavior of households. A key portfolio choice problem comes in the form of the so-called "equity premium puzzle," in which theory would seem to suggest that consumers should keep most, if not all, of their money in the stock market. Another example is the observation of higher savings rates among those with good insurance coverage relative to those without. After all, if savings is partly a "rainy day" fund, one might expect the reverse. At a minimum, one requires a model in which both insurance contracts and savings may be chosen.

At present, consumer finance is finally getting its due. Households are modeled in ongoing work as having precautionary motives, facing rich arrays of assets and liabilities, as well as complex constraints on the portfolios they may hold. The profile of household finance has grown dramatically in recent years, so much so that John Y. Campbell, a leading finance scholar, dedicated his 2006 Presidential Address to the American Finance Association to "Household Finance."

In this address, Campbell draws on a growing body of research that attempts to account for the balance sheets of households over the life cycle. In "Consumption and Portfolio Choice over the Life Cycle," (2005), a trio of economists – Cocco, Gomes, and Maenhout – appear to be the first to integrate three key elements of the household finance problem: portfolio choice, uninsurable and nontradable labor income risk, and borrowing constraints. Relatedly, Steven Davis, Felix Kubler, and Paul Willen (2006) also demonstrate how borrowing costs inhibit the demand for equity over the life cycle, and thereby help explain the (still) limited participation of many U.S. households in the stock market.

Borrowing to buy a home is typically the largest contributor to household debt, while home equity is typically the largest household asset. Recent work has grappled with the type of mortgage products that households choose. Campbell and Cocco (2003) are among the first to systematically analyze the tradeoffs to mortgage choice and link these choices to the precise nature of constraints and risks faced by households. A useful guide to how economic theory on

the structure of household portfolios compares with the data can be found in the edited 2003 volume, *Household Portfolios.*

As this body of work progresses, economists will be able to claim, without reservation, to have understood consumer finance and will be able to perform sophisticated policy analysis that avoids the pitfalls that affected previous work.

CONCLUSION: THE FUTURE OF CONSUMER FINANCE RESEARCH

In this essay, we have described both a brief history of methodological improvements in the models available for addressing the data on household consumption and savings, and also a general approach to evaluating policy questions in consumer finance. We have argued that it is critical to study settings in which the true decision problem faced by purposeful households and profit-seeking, competitive financial intermediaries is both modeled explicitly and forced to meet consistency requirements. Peril lurks in relying wholly on statistical summaries of data for policy analysis.

In our view, there is no alternative to the "equilibrium" approach for meaningful policy analysis. However, our insistence on modeling purposeful actors whose competing needs are reconciled within the model does not inherently bias our conclusions in any obvious way. Rather, we believe that the adherence to some rules for the game simply allows us to learn precisely what is at work in the process that turns real-world policies into real-world outcomes.

Right now, what we have are quantitatively reasonable stories that allow our models to reproduce the path of consumption, net worth, and, increasingly, portfolio allocations for many households. It is perhaps fair to say that the current period is one in which economists are dealing very directly, arguably for the first time, with consumer finance in a theoretically and empirically satisfactory way. Exciting times lie ahead.

For references, please see Essential Readings and Living Bibliography on the Richmond Fed's Consumer Finance Research site.