Equilibrium Models of Personal Bankruptcy: A Survey

Kartik Athreya

Academic research aimed at understanding the consumer default decision has grown rapidly over the past decade. The genesis of this research is the product of three broad sets of forces. First, advances in pure theory gave economists a better understanding of the implications of allowing default for consumer welfare and credit market outcomes. Second, the relatively striking growth of unsecured consumer debt and default in the 1990s spurred the interests of applied researchers in explaining the default decision.[1] Third, and most recently, advances in computational technology have allowed economists to map the insights from pure theory into models capable of confronting observed data and yielding quantitative implications.

This article documents the evolution of recent work on personal bankruptcy. The questions addressed by this research range from the role of income uncertainty in driving financial distress to the roles of statutes allowing households to shelter wealth and of decreased moral commitment to repay debts.

The extant literature consists of a variety of economic approaches to policy analysis. For example, several recent papers employ detailed analyses of observed data. Some of these analyses—notably Gropp, Scholz, and White (1997); Elul and Subramanian (2002); and Grant (2003)—cleverly exploit the near-natural experiments provided by interstate variation in bankruptcy law. At the other end of the spectrum are the so-called "equilibrium" approaches, typified in the work of Athreya (2002, 2004, forthcoming), Chatterjee, Corbae, Nakajima, and Rios-Rull (2002); Li and Sarte (forthcoming); Livshits,

I thank Huberto Ennis, Pierre Sarte, and John Weinberg for very helpful comments. The views expressed are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Richmond or the Federal Reserve System. All errors are mine.

[1] Throughout this article, the words "bankruptcy" and "default" are used interchangeably.

MacGee, and Tertilt (2003); and others. These approaches have in common settings in which households optimize, and in which equilibrium conditions such as those implied by competition, market clearing, and resource feasibility are imposed. Such researchers explicitly solve household optimization problems parameterized through "calibration" or estimation, and then employ simulation to understand bankruptcy policy. What follows documents most directly the contributions of these "equilibrium" models.

The article is organized as follows. Section 1 presents the intuition captured in purely theoretical models of default. Section 2 documents some of the empirics pertaining to personal bankruptcy. Section 3 discusses a set of quantitative equilibrium models based on the theoretical foundations discussed in the first section. Each of these models analyzes the role of default at both the individual and aggregate levels in a manner that respects salient features of U.S. data. The final section concludes.

1. BASIC THEORY

The Role of Limited Insurance

Proponents of current bankruptcy law have long argued for the role of debt forgiveness in helping those hit by unexpected hard times. The "honest but unfortunate debtor" is now folkloric. Central to this view is a belief that life is characterized by a nontrivial amount of *uninsurable* risk. In particular, instead of pooling risks through explicit insurance contracts, households may do some risk-management by saving in good times and borrowing in bad times. The spectre of persistent misfortune, as it may leave households in financial straits, raises the possibility that occasionally allowing default may be useful. In other words, the default option may act like an insurance policy against very bad luck.

Formalizing the insurance role of the option to default is the subject of several recent works. The seminal references are the papers of Dubey, Geanakoplos, and Shubik (2005) and Zame (1993), who both study outcomes in stylized two-period models in which all uncertainty resolves at the terminal date. The critical assumption made in their work is that insurance markets against second-period risk are incomplete. Of course, without this imperfection, allowing default simply corrodes the ability of borrowers to commit to repaying debts, and can therefore only make matters worse. The key result of their work is that in such a world, allowing default subject to a finite penalty can improve allocations, even if it means that default occurs in equilibrium. The intuition is that default, by allowing repayment to partially suit the immediate needs of a household, can provide a "state-contingency" to debt that aids smoothing. The following elegant example from Zame (1993) makes the point more clearly.

Example 1: Default Can Be Especially Useful

The analysis in Zame (1993) contains two central insights. The first insight is that allowing default can very nearly create a fully insured income stream in situations where *completely* insuring against risk may require promising to occasionally pay very large amounts to a counterparty. The second insight is that if the underlying structure of assets is such that their payoffs are "similar" to each other, then households will necessarily be forced to promise very large payments to counterparties in some states in order to obtain insurance against other states. In such a setting, the addition of more assets with similar structure will *not* help. That is, default is uniquely useful in overcoming insurance-market incompleteness.

Consider a world with two agents with identical, risk-averse, standard, expected utility functions, and two dates, 1 and 2. On the first date, both agents are endowed with a unit of the good. On the second date, each agent draws stochastic income in the form of a single consumption good that depends on an aggregate state that takes values from a countably infinite set, $\Omega = \{1, 2, ...\}$. The endowment of Agent 1, (where the first entry is the date-1 endowment), is denoted e^1. The endowment of each agent depends on the aggregate state as follows: $e^1 = \{1; 1, 7, 1, 1, 1, ...\}$. Similarly, Agent 2 faces endowment $e^2 = \{1; 7, 1, 1, 1, 1, ...\}$. Next, we impose a probability distribution, $\mu(\omega)$, over income levels, $\omega = \{1, 2, ...\}$. Specifically, let $\mu(1) = \mu(2) = 1/4$, and $\mu(\omega) = 3^{-\omega+2}$ for $\omega > 2$. This probability distribution has the property that the probability of receiving income from a state with a relatively high index, ω, is far lower than the probability of receiving income from a state with a relatively low index.

A complete-markets Walrasian equilibrium for this economy is Pareto efficient, whereby there is perfect risk sharing between the two agents. Therefore, given the symmetry in preferences and endowments, the equilibrium allocation is one in which consumption for each agent, c^i, $i = 1, 2$, is an equal share of aggregate income. Under the assumed aggregate endowment process, equilibrium consumption is therefore given by $c^i = \{1; 4, 4, 1, 1, 1, ...\}$, for $i = 1, 2$. Notice that this means that individual-level consumption varies only with the aggregate endowment of the economy, and not with any agent-specific changes in endowments.

To achieve this allocation, Agent 1 gives up three units of date-2, state-2 consumption, in return for three units of date-2, state-1 consumption, with Agent 2 taking the other side of the transaction. Lastly, consider the following finite set of assets, A_1,A_2,...A_K, that will allow some, but not complete, insurance. Let each row in the matrix below describe the state-contingent date-2 payoffs for a given asset. The following asset structure does not permit perfect smoothing.

$$
\begin{array}{lcccccccccc}
A_1 = & 1 & 2 & 0 & 0 & 0 & . & . & . & . \\
A_2 = & 0 & 1 & 2 & 0 & 0 & 0 & . & . & . \\
A_3 = & 0 & 0 & 1 & 2 & 0 & 0 & . & . & . \\
. & 0 & 0 & 0 & 1 & 2 & 0 & 0 & . & . \\
. & . & . & . & . & . & . & . & . & . \\
. & . & . & . & . & . & . & . & . & . \\
A_K = & . & . & . & . & . & . & .1(k^{th}\text{ position}) & 2 & .
\end{array} \tag{1}
$$

In this case, notice that to replicate the payoffs of the Pareto efficient outcome in the first four states, for example, Agent 1 needs to hold the following portfolio:

$$F_4 = 3A_1 - 9A_2 + 18A_3 - 36A_4.$$

While this allocation replicates the Pareto optimum, by construction, for states $1-4$, the liability that Agent 1 incurs in state 5 is -72 units of consumption. This liability far exceeds Agent 1's endowment in state 4 (which is 1). If default were disallowed, then the given asset structure could not replicate the optimal allocation for even the first four states. Moreover, the essence of Zame's example is that the addition of more assets with the payoff structure here beyond $A_1 ... A_K$ will not help, so long as we continue to require that all liabilities be satisfied with certainty. As seen, with securities $A_1 ... A_{2N}$, to achieve the optimal risk sharing for states $1 - 2N$, Agent 1 would owe $-9(2^{2N-1})$ in state $2N+1$. Given the endowment structure, this is again infeasible. However, if debt forgiveness is allowed, with a finite penalty, λ—that is, assumed proportional to the liability—matters change. For default in state $2N$, the expected penalty is $9\lambda(\frac{2}{3})^{2N-1}$. Therefore, the possibility exists for λ small enough that Agent 1 will be willing to hold portfolio F_{2N}, while Agent 2 finds this acceptable. The key, to repeat, is that default can void the need for households to hold very large liabilities, and can also facilitate trades that additional assets, however many, may not be able to achieve.

The Role of Limited Commitment

The opponents of the above arguments argue that any benefits that potentially arise from default are confronted by even larger costs. Perhaps most commonly, opponents of default and bankruptcy have argued that debt forgiveness and other forms of limited liability may simply encourage profligacy, sloth, and impose costs on other, more judicious borrowers.[2] Moreover, lenders and insurers themselves may be wary of entering into contracts with households

[2] Arguments against bankruptcy tend to center around the spill-overs that may accompany default. For example, the argument that easy default reduces thrift is only germane if it changes the opportunities that others may have. Similarly, if easy bankruptcy encourages shirking at the

endowed with the right to "walk away." In other words, an important possibility is that it is misleading to take incompleteness as a given when studying bankruptcy, precisely because *incompleteness may be caused by the option to default* in the first place. The crux of this argument is twofold. First, the willingness of households to commit to repayment is driven by, among other things, the extent to which they can self-insure. If bankruptcy offers good self-insurance, then participation and the demand for formal insurance may be diminished. That is, private credit with default may be crowding out formal insurance, leading observers to wrongly conclude that markets for risk remain highly incomplete. This strand of the literature, typified by the theoretical work of Kocherlakota (1996), begins by allowing a full set of insurance contracts to be traded, and then studies the extent to which limited commitment for repayment "shrinks" the feasible set of contracts. In this manner, these models produce incomplete insurance as an *outcome*, as opposed to merely asserting it to begin with. In these settings, because insurance markets are rich enough to allow full risk sharing under unlimited liability, the welfare consequences of introducing consumer bankruptcy are unambiguously negative. The following simplified example illustrates the manner in which the possibility of default limits insurance provision.

Example 2: Default Can Cause Incomplete Insurance

Consider a risk-averse consumer who faces income risks. Assume that the consumer can contract with a perfectly diversified insurance company who *is* committed to honoring all contracts. Denote a finite set of income realizations by $y \in \{y_1, y_2, ...y_N\}$, where $y_i < y_j$ if $i < j$. Denote the relative likelihoods of various income realizations by a probability distribution π_1, π_2,...π_N, whereby $\sum_{i=1}^{N} \pi_i = 1$. Let household utility be given by $u(C(y_i))$, where $C(y_i)$ denotes consumption in state-i. Let mean income be normalized to one, and let $P(y_i)$ denote the transfers made by the consumer to the insurance company. Negative values are interpreted as transfers to the consumer, while positive values represent payments by the consumer.

In this setting, the problem of the household is the following:

$$\max \sum \pi_i u(C(y_i)),$$

s.t.

$$C(y_i) = y_i - P(y_i).$$

---

workplace, or reduces search effort when unemployed, this too is relevant for social well-being only to the extent that such behavior places a burden on an unwitting employer or unemployment insurance scheme. The ability of bankruptcy to impose such costs depends fundamentally on what aspects of individual behavior may be observed—and, hence, priced—in competitive markets. For example, if shirking on the job and search effort among the unemployed were costly to observe, the distortion to behavior created by a default option might be socially quite harmful. I will return later to the role of observability in altering the desirability of allowing bankruptcy protection.

The solution to this problem is perfect insurance. Namely, the consumer will be left, net of insurance premiums, with a perfectly smooth profile of consumption equal to the mean income of unity, regardless of the income realization. To achieve this profile, the consumer will pay a premium equal to the excess of current income over mean income: $y_i - 1$, whenever income is larger than the mean, while receiving the difference between realized income and mean income whenever the former is smaller.

The limited commitment is now introduced as follows. Assume that the consumer cannot credibly commit to repay any more than a fraction, $\psi < 1$. Intuitively, it may help to think of a consumer who can leave town if asked to pay more, and, further, that the consumer has no way of *credibly* promising to the insurance company that he will remain in the contract. To facilitate comparison with the full-commitment problem, we see that the household still solves:

$$\max \sum \pi_i u(C(y_i)),$$

s.t.

$$C(y_i) = y_i - P(y_i).$$

However, an insurance contract must now satisfy:

$$P(y_i) \leq \psi y_i,$$

which represents limited commitment or liability. While the solution is slightly cumbersome and, therefore, not presented here, the intuition for the solution of this problem is straightforward.[3] If the insurance company wishes to break even, it must not contract with the consumer in a way that leaves it vulnerable to default. In other words, it must not rely on large payments from the consumer in good states to offset payments it makes to consumer in bad states. Instead, it must *limit* payments in bad states in a way that allows it to break even, given that it will not be able to collect in good states. In other words, limited commitment can lead to incomplete insurance.

Despite the possibility that limited commitment may create market incompleteness, a caveat is in order. One must be careful to distinguish the repudiation of a true "insurance" contract, whereby payments are negatively correlated with one's income, from repudiation of a "debt" contract, whereby payments are uncorrelated (up to bankruptcy) with one's income. Notably, default in the former occurs when income is *high,* while default in the latter occurs when income is low. Moreover, even if it is limited commitment that is responsible for incomplete insurance, it may still be useful to have default as an option. In particular, if there was incompleteness of insurance contracts,

[3] The interested reader is referred to Obstfeld and Rogoff (1996, Chapter 6) for the detailed solution to this problem and to Ljungqvist and Sargent (2000, Chapter 15) for the more general dynamic limited commitment problem modeled on Kocherlakota (1996).

regardless of whether incompleteness arose from limited commitment, the incentive for households to trade contracts to share risk (e.g., through assets such as those specified in Zame's example) would remain.

Default and Moral Hazard

In the previous example, information was assumed to be perfect, and the incompleteness of insurance emerged solely from the inability of the insured to commit to the full-insurance contract. If enough information is available to limit opportunistic behavior by insured households, bankruptcy will be essentially unnecessary. Moreover, in such a setting, any distortions to allocations caused by bankruptcy will be limited to the level of borrowing that may take place, since the interest rate on loans will need to reflect default risk. If creditors have few ways of punishing defaulters, as appears to be the empirically relevant case, one might expect default risk to be high in such a world. In particular, exclusion from borrowing, the primary long-run consequence of bankruptcy, may not be an effective deterrent to default when insurance markets are complete. However, purely intertemporal smoothing, such as that undertaken by the young to finance education, will still be greatly hindered if borrowing is very expensive. Such a potentially important distortion may persist even with relatively rich insurance markets.

By contrast, if information about the actions taken by insured households is not easily available, then bankruptcy may impose costs that are rather large. In the context of unemployment insurance schemes for example, bankruptcy may greatly exacerbate the incentives to shirk and even take on additional risks. That is, by enhancing the ability to self-insure, the provision of explicit insurance may become undesirable. This is an instance of the more general problem of "moral hazard," whereby an insured party's incentives to take risks in a manner unobservable to the insurer are enhanced. In recent work, Athreya and Simpson (forthcoming) argue that as a quantitative matter, the ability of the U.S. public unemployment insurance system to improve welfare is seriously limited by the availability of bankruptcy.

In sum, the principal insights of the theoretical work presented above are as follows:

- *Default may enhance the functioning of asset markets in helping consumers hedge risks.*
- *Limited insurance, though it may justify allowing the option to default, may itself arise from limited commitment.*
- *The availability of default can be expected to increase moral hazard.*

2. SOME FACTS

Beginning in the late 1980s and early 1990s, events in unsecured credit markets attracted the attention of economists. It was a period characterized by three features: increased credit availability, increased indebtedness, and rapidly increasing personal bankruptcy rates. Striking facts such as these begged explanation, and prompted several analyses. With respect to the trends mentioned above, the interested reader is referred to Sullivan, Warren, and Westbrook (1989, 2000). This article, however, will focus less on accounting for these facts and more on documenting aspects of bankruptcy that have remained more stable over time. In particular, the scrutiny triggered by the growth of credit and default-related variables led to the discovery of a variety of other facts that seem to be relatively time-invariant. It is these more "long-run" phenomena that are the subject of several recent papers discussed below.

The Facts: Bankrupts "R" Us, but Not Quite

Are bankruptcy filers educated or uneducated? Are they rich or poor, young or old, sick or healthy, employed or unemployed, entrepreneurs or workers? The extant body of work has taught us much about "who" bankruptcy filers are. The starting point in research on the demographics of bankruptcy filers is the seminal work of Sullivan et al. (1989, 2000). The 1989 study surveyed bankruptcy filers in the 1980s, while the second captured data through 1997.

With respect to educational attainment, the distributions of attainment are strikingly similar, (Sullivan et al. 2000, 53) with bankruptcy filers even more likely than the general population to report having attained "some college." However, this is potentially misleading in that bankruptcy filers are somewhat less likely than the rest of the population to hold either a college or advanced degree.

More interesting differences appear when comparing the earnings of filers and non-filers for a given level of educational attainment. In Sullivan et al. (2000), the earnings of college-educated non-filers to filers is 1.4, and is most extreme for those with advanced degrees, where the ratio is roughly 2. This fact suggests that bankruptcy filers are disproportionately the recipients of degrees which offer lower returns. Foreseeable aspects of human capital acquisition may be the cause of this situation as well as the possibility that bankruptcy filers are often those for whom the uncertain nature of post-college earnings have been resolved in favor of poor earnings. Lastly, in the year that a filing occurs, incomes among filers are substantially lower than for non-filers. For example, a sample of five U.S. states shows that the median income for filers is often only half that of non-filers, while the mean is even more skewed in favor of non-filers. Moreover, recent work of Sullivan et al. (2000) and Bermant and Flynn (2002) suggests that the unemployment rate among filers is between three and four times as high as is for the rest of the population,

Figure 1 Influence of Total Consumer Debt on Bankruptcy Filings Trends by Year, 1980–2003

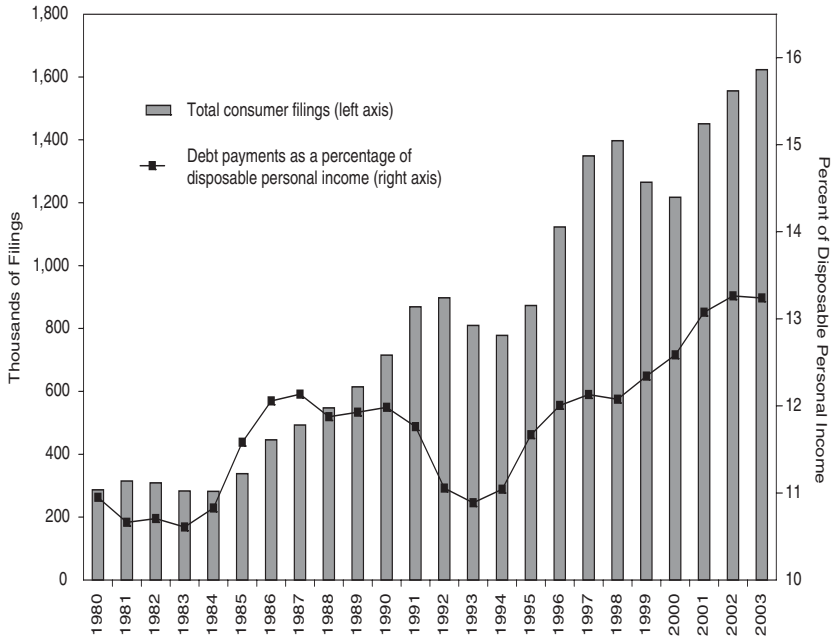

giving fillip to the narrative that places shocks to earnings front and center in explaining personal bankruptcy.

Net indebtedness is much higher for the population of filers than for the rest. The only reason that indebtedness is not an obvious feature of the data is that the median asset exemption as of 1998, measured by Grant (2003), was $44,000. In other words, households could hold large stocks of wealth while holding unsecured debts that were dischargeable in bankruptcy. Nonetheless, few households in bankruptcy hold large stocks of wealth. Median net worth among filers was -$10, 500 in 1991, while it was $36,000 for the U.S. population as a whole. However, the income levels reported above suggest that it may simply be the mismatch between income and debts, rather than the size of debts, that matters for bankruptcy. Figure 1 is suggestive of the role of debt relative to income in accounting for recent trends in bankruptcy.

With respect to age, bankruptcy filers are also similar in age to the general population. Sullivan et al. (2000) find in their sample that, as of 1997, the median age of filers was only slightly lower than the national median of 41.7 years. However, the median hides the feature that roughly 85 percent of filers are younger than 55. A final dimension along which bankruptcy filers may be

distinguished is by their prior employment experience. Specifically, entrepreneurship is featured prominently in the histories of bankruptcy filers. Sullivan et al. (1989) document that fully 20 percent of filers in their data from the 1980s report themselves to be self-employed in a business venture, even though entrepreneurs account for only slightly more than 10 percent of the U.S. population. In sum:

Bankruptcy filers are similar to the overall U.S. population along the dimensions of education and median age, but differ from the overall U.S. population in that are they more likely to be young, low net-wealth, sick, unemployed, and self-employed.

3. QUANTITATIVE THEORY: RUNNING A HORSE RACE

More recently, the theoretical contributions of Dubey et al. (2005), Zame (1993), Kocherlakota (1996), and others have been critical in motivating dynamic general equilibrium models aimed at evaluating the consequences of bankruptcy. The distinguishing features of these models are much greater attention to institutional richness pertaining to personal bankruptcy statutes and a reliance on numerical solutions. Notably, the households in these settings can be distinguished clearly by age, debt portfolios, insurance policies, and other ways observable in the data. The payoff to the introduction of these complications is to allow for *quantitative* evaluations of the relative strengths of the forces of incomplete insurance and limited commitment.

To provide intuition as to how the particular quantitative aspects of incomplete insurance operate, consider the following: As a quantitative matter, the relevance of Dubey et al.'s (2005) justification depends on the precise difficulties imposed by incomplete insurance markets that, in turn, depend crucially on the nature of income risk faced by households. Labor income is important as it accounts for most household income and is typically not directly insurable (perhaps for the obvious complications created by moral hazard). For example, consider an environment in which households that face risk to income may borrow and save, but have no option to default. That is, all debts have to be honored, regardless of the circumstances facing the borrower. Suppose also that the household was susceptible to both short-term and very long-lasting shocks. Relative to a short-term shock, receiving a long-term shock can much more seriously alter the present value of future income. If such a shock were negative, the present value of future income could fall substantially, making the acquisition of debts to smooth consumption in the face of short-term shocks potentially very expensive.

A key implication of standard models of consumption and savings is that households borrow (or "dissave") progressively less as shocks grow more persistent. In other words, the willingness of households to effectively self-insure

temporary income disturbances might reasonably depend on their likelihood of receiving more persistent shocks. This is one sense in which the possibility of default can "grease the wheels" of credit markets. This line of reasoning is addressed more fully in Athreya and Simpson (forthcoming).

The relative strength of the effects of the forces listed above are ultimately key to deciding on the possibility of a beneficial role for bankruptcy. While there is a good deal of empirical work describing the salient facts, questions of welfare are not addressed in that literature. The interested reader is referred to the thorough review of empirical work contained in the Congressional Budget Office (2000). I focus instead on a handful of representative quantitative models, all of which contain forward-looking households that face uninsurable risks. An advantage of this approach is that welfare can be easily evaluated for equilibriums allocations. The first group of environments contain a single asset, attached to which is a bankruptcy option. These models are essentially direct extensions of the work of Aiyagari (1994) and Huggett (1993). However, the payoffs from the single asset do not vary for any other reason, which precludes analysis of aggregate risk and also constrains insurance possibilities in a particularly strong manner. The second class of models allows for more than one asset, as well as for richer insurance possibilities beyond noncontingent unsecured consumer debt. For convenience, Table 1 presents a stylized taxonomy of the models discussed below.

Single-Asset Models

Chronologically, the first quantitative evaluation of bankruptcy appears to be Zha (2001). The key features of this environment are (1) a large number of risk-averse households and (2) the presence of uninsurable, but purely idiosyncratic, risk. That is, aggregate activity in this economy does not fluctuate through time, making it incapable of addressing questions related to business cycles. Households are free in this setting to save by accumulating capital, which in turn is used to produce output, or by borrowing to invest in (idiosyncratically) risky capital equipment to produce output. Borrowers operate investment projects that allow this uncertain output to be seen only at a cost.

An important assumption made in Zha (2001) is that contracts are static, a feature used later in several other papers (e.g., Athreya [2002, 2004, forthcoming]; Livshits et al. [2004]; Chatterjee et al. [2002]; and Li and Sarte [2002]). It is known from Townsend (1979) that in a setting where a lender faces a cost of verifying the "state" (income, in this case), the optimal static contract is one resembling "debt"—that is, a contract where the borrower announces output (truthfully) and repays a constant amount unless output is relatively low. In such cases, the household announces (truthfully) that output is low, and the lender seizes a fraction of the output. The amount of wealth that may be given

Table 1 Comparing Some Equilibrium Models

	Athreya (2002)	Athreya (2004, forthcoming)	Livshits et al. (2004)	Li/Sarte (2004)	Chatterjee et al. (2002)	Mateos-Plannas/Seccia (2004)
Credit Limits/ Observability	Exogen./Only Unconditional Moments	Endogen./ Total Debt	Endogen./ Total Debt, Current Prod., Age	Exogen./Only Uncondtional Moments	Endogen./ Total Debt, Current Prod.	Endogen./Only Unconditional Moments
Expense/ Preference Shocks?	No	No	Yes	No	Yes	No
Flexible Labor Supply/ Production?	No	No	Yes	Yes	No	No
Can the "No-Bankruptcy" Case Be Studied?	Yes	No/Yes	No	Yes	No	Yes
Finite Horizon or Infinite Horizon?	Infinite	Infinite	Finite	Infinite	Infinite	Infinite
Allows Chapter Choice?	No	No	No	Yes	No	No
Single or Multiple Asset?	Single	Single/ Multiple	Single	Multiple	Single	Single
Partial Equilibrium (PE) or General Equilibrium (GE)?	GE	PE	PE	GE	PE	GE

up to satisfy unsecured creditors is restricted in practice by "exemptions" on bankruptcy. Zha (2001) finds that increasing the value of exemptions does have a positive effect on welfare. Zha (2001) does not, however, compare the welfare consequences of repealing all default to a setting with default.

Bankruptcy reform efforts as of the late 1990s centered around two aspects of the law. First, existing bankruptcy law did not require "means-testing" for would-be filers. Therefore, potentially high-income filers would be able to discharge substantial debt without giving up income. Note carefully the distinction between income and wealth. Exemptions are rules that explicitly cover the levels of wealth that debtors may retain in bankruptcy. Means-tests did not directly apply to these provisions. Athreya (2002) addressed the issue of means-testing in a simpler environment than that studied by Zha (2001). The key features of the environment, as in Zha (2001), were a large number of infinitely lived households facing income risk. Unlike Zha (2001), creditors were assumed to issue credit cards with fixed "lines" of credit, where the interest rate on loans did not move with the total debt level. The advantage of this assumption is that contracts were simple fixed-rate instruments that bore a closer resemblance to the predominant form of unsecured debt held by U.S. households. The disadvantage is that credit conditions are restricted to appear

only through prices, whereby bankruptcy law would have no effects on credit limits.

Athreya (2002) finds that the means-testing provisions embodied in the Bankruptcy Reform Act of 1999 would not greatly alter allocations. The intuition is that, given the costs of bankruptcy implied by the data, most households would not file unless already poor. On the other hand, in Athreya (2002), the elimination of all bankruptcy was found to be quite beneficial. This is important because, as mentioned earlier, the main costs of bankruptcy in Athreya (2002) come from (1) interest rate "externalities" arising from the pricing of loans to cover average, as opposed to personalized, repayment rates and from (2) the assumption that all costs of bankruptcy were deadweight in nature. In other words, easy bankruptcy law generated a shift of households toward the use of bankruptcy that necessitated the more frequent use of deadweight penalties. On balance, the frequency of filing under lax penalties for bankruptcy did not help households smooth consumption in a manner that offset society's cost of imposing court costs and possibly "stigma" on filers.

The estimate in Athreya (2002) is, however, prone to producing a downwardly biased estimate of the benefits of eliminating default, at least for the income process employed.[4] This bias arises because with only marginally more observability, whereby creditors could see total household debt, one would have expected a large reduction in the cost of credit. Therefore, if the initial credit limits assumed in Athreya (2002) were lax enough that the expansion in credit did not exceed the limit initially assumed, then the estimate may be reasonable. However, if the initially assumed debt level was narrower than the endogenous debt level that would emerge under partial observability of debt, then the estimate of the benefits of strict bankruptcy are too low.

Recognizing that the willingness of creditors to lend will be related to the characteristics that they may observe about a borrower, Chatterjee et al. (2002), Livshits et al. (2004), and Mateos-Planas and Seccia (2004) each allow for more observability than Athreya (2002). Chatterjee et al. (2002) use an infinite-horizon setting where current period income and current wealth are both observable, in principle. In addition to allowing for endogenous limits on debt, Chatterjee et al. (2002) also contribute by establishing theoretical properties associated with recursive representations of household optimization problems, specifically for the case of i.i.d. income shocks. Livshits et al. (2004), on the other hand, use a life-cycle model and augment the set of debtor observables to include age. In contrast to both preceding papers, Mateos-Planas and Seccia (2004) assume a lack of observability beyond population averages. They appeal to an institutional structure whereby banks finance lending by issuing securities backed by repayments on the unsecured

[4] We revisit this issue when discussing the alternative earning/expense process used by Livshits et al. (2004).

loans they make. The "pools" backing these securities are large and aggregate a variety of types of households. This route was originally taken in Dubey et al. (2005) in order to conform to the predominant form of financing used by credit card and mortgage lenders. Mateos-Planas and Seccia (2004) then go beyond the exogenous limits to borrowing used in Athreya (2002) and derive borrowing limits consistent with zero profits under this restriction on observability.

The policy experiments pursued in Chatterjee et al. (2002) study the desirability of (1) means-testing and (2) a reduction in the length of time for which credit reporting agencies may retain the record of a bankruptcy filing. They find that nearly all households would prefer to restrict the option of bankruptcy to those households with above-median earnings, and few at all are in support of a reduction in the legal limit on the length of time that a past bankruptcy filing may be recorded. Similarly, with respect to whether bankruptcy should be allowed or not, Athreya (2002) finds that prohibiting bankruptcy is preferable to permitting it as currently practiced. The intuition for these findings is that the increased costs of credit apply to all borrowers (although not with equal force) and is an externality that offsets the better consumption insurance allowed by bankruptcy in isolation. Athreya (2002) also finds that means-testing is not very important, while the availability of bankruptcy is. Notably, at present there are few restrictions on the current earnings of a household at the time of a Chapter 7 bankruptcy filing. Nonetheless, the income of filers is rarely higher than the median.

Athreya (2002) abstracts from certain types of catastrophic risks, such as severe health shocks, legal costs of divorce, and the inability to collect mandated child support. The shock processes used in the preceding models affect only income. However, to the extent that such severe shocks are important, as Sullivan et al. (2000) argue is true for approximately 20 percent of all filings, income shocks underrepresent the risk faced by households. Livshits et al. (2004) therefore propose a more extreme feature—that of shocks to "expenses," or net worth. Households are assumed to be subject to occasional, but large, reductions in their asset positions. These shocks are meant to capture large unanticipated expenditures on inelastically demanded types of goods, such as hospital care or child care. The advantage created by this is to allow for real risks that appear important in at least a portion of personal bankruptcy filings (for example, 19 percent of filings involved large medical expenses, according to Sullivan et al. [2000]). However, the presence of expense shocks that may make current resources deeply negative under the assumption of standard constant-relative-risk-aversion preferences implies that without the possibility of default, no one would borrow. This extreme outcome arises because the full enforcement of debts would mean forcing the household to zero consumption, i.e., literally forcing it to starve in the face of large negative shocks. Knowing this, the household would never *voluntarily* assume debts.

The disadvantage created by this condition is that the comparison of outcomes with and without a bankruptcy option is not easy to evaluate. On the other hand, the results of Livshits et al. (2003) are a useful, if stark, benchmark for how the availability of default facilitates risk-sharing through asset-markets. Namely, since prohibiting default in their environment would mean not borrowing at all, a portion of the welfare gains from bankruptcy in their model comes from the complementary role that default plays in enhancing the value of asset markets.

Multiple-Asset Models

As mentioned earlier, all of the preceding models feature a single asset that when held in negative quantities, implied borrowing. The absence of any distinction between assets means that borrowing and negative net worth are the same thing. This restriction does not match well with the complicated portfolios households have in the data. For example, Sullivan et al. (2000) document that while many households in bankruptcy do have substantial unsecured debt, many also do not have negative net worth. In large part, this restriction of trade to a single asset was driven by tractability. Some recent work has relaxed this restriction. The payoffs to accommodating this richness is that it (1) allows for a closer mapping from model outcomes to observable data, and (2) allows us to correctly measure the ability to smooth consumption and (3) allows us to analyze policies aimed at allowing households to retain wealth while discharging debts.

Tractability

With respect to tractability, the chief complication of solving a multiple asset problem is that of storing all the information the household needs to solve its optimization problem. Specifically, dynamic programming techniques are typically the tools to solve household optimization problems in environments complicated by uncertainty and "discrete" choices (e.g., to file or not to file). Dynamic programming becomes radically more cumbersome as the amount of information required of the household increases. In the context of bankruptcy exemptions, for example, a household would like to know its feasible options as it enters a period. However, these options may depend on the precise composition of household debts and assets. Moreover, even if one is creative and can define a single summary statistic that allows households entering a period to solve their optimization problem, it may be at the expense of complicating the "within-period" problem of the household. For example, in Athreya (forthcoming), households may hold secured and unsecured debt. Nonetheless, "total wealth" serves as the single variable that households must be aware of within any period.

Analyzing Exemptions: The Payoff to Allowing Multiple Assets

Despite the computational burdens involved, multiple-asset models are valuable, as they allow for the study of households that may be saving and borrowing simultaneously. In turn, such models allow researchers to meaningfully study important policy questions pertaining to bankruptcy, most notably, the rules defining exemptions. The latter govern the extent to which wealth may be held while debts are discharged. Exemptions have attracted a great deal of attention recently, both academically and in the public discussion on bankruptcy. Exemptions are, most generally, allowances for certain types of wealth that bankruptcy filers may retain after bankruptcy. Any wealth in excess of exemptions must be surrendered to satisfy unsecured lenders.

The idea that substantial wealth may be protected from seizure by creditors has long been controversial (see, e.g., Moss [forthcoming]). In support of exemptions is the following intuition: Because they partially govern how much "state-contingency" is truly embedded in an unsecured loan, it may be useful for a household to be able to hold at least one asset with payoffs that it can manipulate to serve its needs. For example, even if a loan is not explicitly collateralized, it may be implicitly so, simply because a household in bankruptcy would be required to transfer any nonexempt wealth to unsecured creditors. Thus, an unsecured loan may be effectively collateralized, even if wealth is not explicitly pledged by the debtor. By contrast, a large exemption will allow most unsecured borrowing to be cleanly discharged in bankruptcy if the household finds itself in difficult circumstances. Proponents also argue that, after bankruptcy, it is important to allow a household to run its affairs without becoming destitute and, potentially, a recipient of publicly funded transfers (see, e.g., Baird [2001]). Opponents, on the other hand, argue that those who are wealthy along some dimensions (though not along net-worth perhaps) should not be excused from debt obligations. A final concern with more ambiguous welfare implications is that to the extent that they are aware of a household's assets when making unsecured loans, creditors will price this risk. If observability is low, all unsecured borrowers will face higher borrowing costs, and if not, wealthy households alone will be able to obtain unsecured credit. We will return to the role of observability later in the article.

In a model that allows for simultaneous holdings of both debt and equity, Athreya (forthcoming) finds that, conditional on allowing bankruptcy, high exemptions are actually useful, even if they make unsecured debt more expensive. A ramification of high exemptions is that the cost of unsecured debt will be higher under high exemptions. This is precisely what data analyzed in Gropp et al. (1997) suggest. Their work is an important study on exemptions that utilizes the natural experiment provided by interstate variations within the United States. They document strong evidence supporting the view that exemptions make unsecured credit more expensive. More recently, Grant (2003) also exploits interstate variation in exemptions and finds significant support

for the risk-sharing role of exemptions. His estimates indicate that exemptions noticeably reduce growth in the cross-sectional variation of income. The latter can be interpreted as an improvement in "market completeness," or insurance possibilities (see Deaton and Paxson [1994] for details).

Because bankruptcy exemptions vary across states, Elul and Subramanian (2002) document the extent to which households move across state lines to avail of more generous bankruptcy exemptions. Given that bankruptcy filers typically have low incomes at the time of filing, interstate moves may be a useful option for a household, given the fall in the market value of its time. The authors find that "forum shopping" (the explicit decision to search for a friendly set of laws) accounts for roughly 1 percent of all interstate moves to a state with a higher exemption.

Two additional recent studies deserve discussion. First, Li and Sarte (forthcoming) employ a setting in which two simplifications used in all prior work are relaxed. Their first innovation is to model the choice of bankruptcy "chapter," and the second is to accommodate production decisions. With respect to the former, in the discussion so far, I have implicitly combined all forms of personal bankruptcy. However, there are typically two forms of bankruptcy available to a household debtor. These are Chapters 7 and 13 of the U.S. Bankruptcy Code. The former is most familiar, and simply removes all unsecured debt in exchange for the surrender of all wealth above the exemption. Chapter 13, by contrast, is (at least in principle) less extreme. In particular, Chapter 13 is a form of debt rescheduling whereby a debtor agrees to repay a portion of his unsecured debts over time. This form of bankruptcy is particularly useful for households that hold wealth substantially in excess of the exemption allowed to them. However, the repayment plan poses two challenges. First, repayment over time may act like a tax on effort, and lead households to change effort. To capture the consequences of such reductions in effort, it is useful to study settings which allow for production of output. Second, current law allows households to convert a Chapter 13 filing into a Chapter 7 at any time. The latter clearly limits the extent to which repayment can be extracted from debtors. Li and Sarte (forthcoming) find notably that allowing production and capital accumulation overturns the stark results of Athreya (2002) in that eliminating bankruptcy lowers welfare. To elaborate, easy bankruptcy lowers precautionary savings and thereby lowers output, while eliminating Chapter 7 bankruptcy altogether hinders risk sharing to an inefficient degree, as Chapter 13 does not provide the same allowance for contingent repayment.

As Li and Sarte (forthcoming) allow for production, the accumulation of capital augments production and also allows for an additional asset with which households may smooth consumption. At the household level, wealth is still described by financial claims, as the capital in their model is useful only for producing the consumption good, and does not itself generate

a flow of services. However, a key aspect of many exemptions applicable to bankruptcy is that they often apply only to special classes of nonfinancial assets, which typically provide a flow of services to households. The largest and most famous exemptions are those applying to home equity (the Homestead Exemption), to equity in cars, and to "tools of trade." Pavan (2003) moves research forward by explicitly modeling the services that durable goods provide their owners. This allows for more precise welfare analysis. Unlike Athreya (2002, 2004), Pavan uses a life cycle model, and estimates its parameters. Pavan's work is noteworthy for its emphasis on the use of formal statistical inference in guiding the selection of parameters. By contrast, while Athreya (2002, 2004, forthcoming), Li and Sarte (forthcoming), and others do assign values to parameters, the procedures used are less informed by formal statistical practice, and rely more on the informal matching of key features of the data. Unlike Li and Sarte (forthcoming) and Athreya (2002), however, Pavan uses a partial equilibrium model whereby the costs of funds are held fixed. Given that Li and Sarte (forthcoming) find that ignoring general equilibrium is not innocuous for studying exemptions, this distinction is worth keeping in mind.

The Role of Observability

The "insurance" value of bankruptcy depends critically on the extent to which debt prices do *not* vary with default risk. To see why, consider a world in which competing creditors were able to view a common set of factors associated with a debtor's default probability. Competition requires that cross-subsidization of some borrowers by other borrowers must not occur along any dimension that is commonly observed by creditors. This is simply because if any two borrowers have observably different characteristics, the relatively less risky borrower would be offered a cheaper rate.

As discussed earlier, assumptions on observability matter for predicting the response of lenders in the face of changes in bankruptcy policy. Therefore, these assumptions must also matter from a welfare perspective. The intuition is as follows. An insurance contract works by allowing the buyer to diversify risk and thereby transfer purchasing power from contingencies where additional consumption is valued less, to those contingencies in which it is valued more. Accurate risk-based pricing and competition will make the default option more expensive, as pricing will more closely reflect marginal default risk as opposed to the average risk that would be accounted for in the absence of such observability. However, the benefit of strong observability is that explicit insurance contracts can play an important role, voiding the need for households to rely on the relatively clumsy implicit insurance of credit with a default option.

The work of Edelberg (2003) is the first to document the changes in the pricing of credit to reflect risk at the household level. Edelberg (2003) uses

data primarily from the Survey of Consumer Finances (SCF), which contains detailed information on household balance sheets, especially the level and terms of borrowing. She finds that the period beginning in the mid-1990s witnessed a sharp increase in the cross-sectional variance of interest rates charged to consumers. This work casts some doubt on the results arising in models such as Athreya (2002), where a single interest rate on unsecured credit applied to all households. According to Moss and Johnson (1999), the latter approach may have been a reasonable assumption for the 1980s and early 1990s, but advances in recordkeeping and other intermediation technologies have dropped the costs of differentiation of borrowers. The issue of intermediation costs will play an important role, as we will see in the next section. Before proceeding, however, it is useful to summarize the following provisional conclusions reached by the works cited above:

- *Under income processes that allow for large shocks to net worth, bankruptcy can play a role in improving welfare, but not without them.*
- *As a quantitative matter, moral hazard needs to be taken seriously in evaluating bankruptcy provisions.*
- *Exemptions tend to distribute credit away from the asset-poor to the asset-rich, but can improve welfare.*
- *Observability is an important determinant of how the supply side of credit markets and household welfare respond to bankruptcy policy.*

A Quantitative Equilibrium Approach for Explaining Recent Trends

Despite the now relatively large quantitative equilibrium literature on bankruptcy, the project of accounting for the time-series of the 1990s has thus far been almost exclusively tackled via purely empirical approaches. A main reason is that only recently has quantitative theorizing on bankruptcy become tractable, but even within these models, long-run "stationary" states have proved far easier to characterize than ongoing aggregate dynamics such as those observed in the 1990s. In particular, the technical problems facing quantitative equilibrium approaches in dealing with the 1990s are twofold. First, when aggregates move over time, so will prices such as the interest rate on savings or loans. Unfortunately, incompleteness of insurance is a precondition for evaluating the trade-offs associated with default. These models, in turn, contain households whose fortunes diverge with time, and as a consequence, feature wealth holdings that diverge through time. Precisely because households become heterogenous in their wealth holdings, their response to individual level or aggregate uncertainty also becomes heterogenous. In turn, future prices, which depend on aggregate wealth accumulation, are determined

by the entire distribution of wealth and make problems very hard to analyze numerically. Nonetheless, Athreya (2004), described below, makes a first, and admittedly simple, attempt to account for the 1990s.

The Roles of Stigma and Technology

Detecting stigma matters for the welfare analysis of bankruptcy statutes. Let stigma be defined to mean all costs of social disapproval associated with filing for bankruptcy. Stigma matters because it is a penalty suffered by filers *after* a filing has taken place, and most importantly, one that hurts the filer, but does not directly help anyone else. If bankruptcy is an activity that society seeks to limit, then society must be aware of the possibility that if a penalty is too weak at the individual level to deter bankruptcy, imposing it very often may make it undesirable. Conversely, severe censure of filers by society may be bad for at least two reasons. First, if the incidence of bankruptcy does not fall substantially, society will find itself imposing a large "after-the-fact" punishment far too frequently. Second, if bankruptcy has a potential role in risk sharing, severe social sanctions may stunt consumption smoothing. The optimal social stigma strikes a balance between these two concerns. The preceding argument is also normative in the sense that it may tell us how strictly we wish to deal with bankruptcy filers along dimensions that merely punish them without helping others.[5]

With respect to the time path of stigma, the stunning rise in per-capita filing rates in bankruptcy, from 0.3 percent in 1980 to 1.6 percent at present, has captured the attention of many researchers. Gross and Souleles (2002) is perhaps the best known of these studies. The essence of their exercise is to ask if debtors who look similar along a multitude of financial dimensions have differentially large likelihoods of filing for bankruptcy recently relative to the past. They find the answer to be "yes," and conclude that falling stigma is a plausible story for recent data. Similarly, Fay, Hurst, and White (2002) argue that, even after controlling for state-level fixed effects, a rise in the bankruptcy rate in a given state predicts a further increase in the following year. Interpretation of empirical regularity is tricky, however. In particular, to the extent that households learn the bankruptcy process from each other and decide that it is easier than they believed, currently high filing rates may lead to even higher filing rates. From the viewpoint of predicting bankruptcy rates, whether one interprets the data as evidence of either falling stigma or learning matters little.[6] However, it may matter much more from a welfare perspective, for the reasons discussed above.

---

[5] The obvious analogy here is to crime and punishment, with its focus on rehabilitation against deterrence.

[6] Indeed, Gross and Souleles (2002) do not restrict their interpretations to falling stigma, but allow for the possibility that information flow has improved.

Athreya (2004) addresses the issue of stigma via an alternative route, by studying a quantitative, dynamic equilibrium model of borrowing. The article conducts two experiments. In the first case, stigma is initially "calibrated," i.e., set to a level that allows the model to match debt and bankruptcy filing rate data as of 1991. The value of this stigma is then lowered to zero. In the second experiment, stigma is fixed at the level consistent with data for 1991. The transactions costs associated with unsecured consumer lending is then sharply lowered. The latter assumption is motivated chiefly by the work of Edelberg (2003) discussed above.

The exercises in Athreya (2004) suggest that the elimination of stigma will produce increases in bankruptcy on the order observed in the data but will also result in sharply lower debt loads carried by households, a situation distinctly at odds with the data. However, the experiment of lowering transactions costs is able to match not only the observed increase in filing rates, but also the increases in debt/income ratios observed in U.S. data. The key intuition driving these results is first that in a low-stigma environment, borrowers will have few incentives to repay loans. To the extent that lenders recognize this, the amount lenders will be willing to lend will decrease. In the extreme case where there is neither stigma nor the possibility of a bad credit rating following a bankruptcy, unsecured lending cannot occur, as households will happily take any loan offered to them and then promptly default. By contrast, to the extent that Edelberg (2003) is accurate, credit "supply" (i.e., the willingness to lend for a given interest rate) should have expanded for purely technological reasons, allowing for the simultaneous increase in debt and filing rates observed in the data.

Even with lowered stigma, the expansionary pressure on credit supply arising from cheaper intermediation is present, though to a more limited degree. That is, stigma may indeed have fallen, only to be overwhelmed by the opposing force of technological innovation. Therefore, it may still be possible to partially reconcile reductions in stigma with observed data. In sum:

Falling stigma is not a convincing explanation for the recent rise in bankruptcy rates.

Bankruptcy and its Relation to Other Forms of Insurance

As documented by Fisher (2003), and Shepard (1984), a large portion of filers report receiving publicly funded transfers in the year in which they filed for bankruptcy. Maybe most tellingly, Sullivan et al. (2000) find that more than two-thirds of all bankruptcy filers report experiencing an income disruption near the time of filing.

The empirical work of Sullivan et al. (2002) argues that not only are many of those in bankruptcy receiving unemployment insurance, but also that a dis-

proportionate share of those receiving unemployment insurance also file for bankruptcy. Specifically, Sullivan et al. (2002) find that the unemployment rate among bankruptcy filers is between three to four times the national average.[7] The bankruptcy rate among the unemployed is also much higher at roughly 3.5 percent, or four times higher than the overall population rate (approximately 1 percent annually in recent years). Data from the Panel Study of Income Dynamics show that 12 percent of bankruptcy filers in 1995 lost their jobs between 1994–1995, as opposed to just 2.15 percent of non-filers. Bermant and Flynn (2002) argue that bankruptcy filers also have shorter job tenure than non-filers whereby job tenure at the median for the bankrupt population is only two years, less than half of the 4.7 years of tenure for the non-bankrupt population. Sullivan et al. (2000) find that more than two-thirds of all households that file for bankruptcy report job-related income disruptions.

The quantitative evaluation of the effects of this form of insurance on behavior relating to other forms is new. The work of Livshits et al. (2003) compares bankruptcy in the United States and Germany, taking as given the structure of public insurance policies in each country. The main finding is that "fresh-start" bankruptcy is far less desirable in the presence of the comprehensive public insurance present in Germany. Athreya and Simpson (forthcoming) study an environment in which the public insurance system, including programs like the U.S. unemployment insurance system and welfare, coexist with the more implicit insurance that unsecured debt with bankruptcy may provide. Unlike prior work, Athreya and Simpson (forthcoming) not only allow more generous insurance to affect credit markets, but also to allow credit markets to affect behavior in public insurance programs. They find that the U.S. bankruptcy code obstructs public insurance provision in the United States. More surprisingly, they find that a more comprehensive social safety net provided, for example, through publicly funded insurance, will actually encourage risk-taking and reduce job search among the unemployed to such an extent that bankruptcy filings rise. More generally, the preceding makes the following clear:

Bankruptcy should be analyzed jointly with available insurance programs, because these systems can interact in perverse ways.

4. CONCLUDING REMARKS AND FUTURE WORK

Recent work has produced substantial progress in revealing the conditions prevailing at, before, and after bankruptcy. Such work has also revealed that interstate variation in rules, particularly exemptions, matter for household

[7] Bermant, Flynn, and Bakewell (2002) estimate a rate of 19 percent in 2002 from self-reported unemployment data. Athreya and Simpson (2004) find a rate of 16 percent from PSID data.

decisions and for the availability and distribution of unsecured credit. The findings established by careful empirical work has proved critical for quantitative equilibrium work that aims to account for the observations. The latter has reached some provisional conclusions.

- *Under income processes that allow for large shocks to net worth, bankruptcy can play a role in improving welfare. Without large shocks, however, bankruptcy is difficult to justify.*
- *As a quantitative matter, moral hazard needs to be taken seriously in evaluating bankruptcy provisions.*
- *Exemptions tend to distribute credit away from the asset-poor to the asset-rich but do not appear to lower welfare.*
- *Falling stigma is not a convincing explanation for the recent rise in bankruptcy rates.*
- *Bankruptcy should be analyzed jointly with available insurance programs, because these systems can interact in perverse ways.*

Despite the progress evident from the work done so far, many interesting questions remain, the resolution of which is critical for forming a definitive view of the role for personal default. From a theoretical perspective, perhaps the most useful work will be to endogenize fully the exclusion that seems to affect those who have filed for bankruptcy. Specifically, in a competitive setting, exclusion can only take place if, after a bankruptcy filing, it is optimizing to restrict lending to such a borrower. Several mutually nonexclusive possibilities arise. First, and perhaps most naturally, bankruptcy may reveal something more "permanent" about a filer. The most obvious is that a bankruptcy was triggered by a persistent shock to household earnings-generating capacity. In such an event, a household may have a very low ability or willingness to service debt. However, the desire of a household to smooth intertemporally may also be very limited. In this case, observing that a household does not borrow following bankruptcy does not pin down why it does not borrow. The problem of rationalizing exclusion may in part stem from the inability of credit market data and income data to provide high quality targets that an equilibrium model must match. Recent progress on this dimension is the work of Yue (2004), who allows for a limited form of "renegotiation" in debts between creditors and borrowers.

A second open issue is the role of stigma and an assessment of its importance in determining outcomes. While it seems unconvincing to argue that stigma has fallen, it seems quite plausible that some stigma or societal disapproval exists for bankruptcy. However, there is not a single, universally accepted measure of this. Nonetheless, knowing the extent to which shame

and stigma matter is necessary for making accurate statements about the quantitative usefulness of bankruptcy to improve welfare.

A third issue is the role of bankruptcy in exacerbating moral hazard in insurance programs such as public unemployment assistance, as well as the role of the latter in encouraging bankruptcy. More work along the lines of Athreya and Simpson (forthcoming) will be useful in determining the extent to which bankruptcy and other insurance schemes (both private and public), confound each other.

A fourth issue based on the arguments above, bodes the question, why is default, regardless of whether it is legally mandated or privately contracted, useful? In the United States, however, the Constitution preserves the right of the federal government to promulgate uniform nationwide bankruptcy law. It is primarily the legal provision for bankruptcy that raises the issue of whether it improves welfare or not.

A final issue is to reconcile the two possible rationales for market incompleteness. Namely, is insurance incomplete precisely because of limited commitment arising from bankruptcy, or because of standard complications arising from asymmetric information? Future work that is able to deal effectively with these issues will greatly advance our understanding on the desirability of personal bankruptcy.

REFERENCES

Aiyagari, S. R. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109 (3): 659–84.

Athreya, Kartik B. 2002. "Welfare Implications of the Bankruptcy Reform Act of 1999." *Journal of Monetary Economics* 49 (8): 1567–95.

_____. 2004. "Shame As It Ever Was: Stigma and Personal Bankruptcy." *Federal Reserve Bank of Richmond Economic Quarterly* 90 (Spring): 1–19.

_____. Forthcoming. "Fresh Start or Head Start? Uniform Bankruptcy Exemptions and Welfare." *Journal of Economic Dynamics and Control.*

_____, and Nicole B. Simpson. Forthcoming. "Unsecured Debt with Public Insurance: From Bad to Worse." *Journal of Monetary Economics*.

Baird, Douglas G. 2001. *Elements of Bankruptcy*, 3rd Edition. New York: Foundation Press.

Bermant, Gordon, and Ed Flynn.. 2002. "Just Recently Hired: Job Tenure Among No-asset Chapter 7 Debtors." *American Bankruptcy Institute Journal*. May.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and Jose-Victor Rios-Rull. 2002. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." Federal Reserve Bank of Philadelphia: Working Paper 02-6.

Congressional Budget Office. 2000. "Personal Bankruptcy: A Literature Review." Available at http://www.cbo.gov/ftpdocs/24xx/doc2421/Bankruptcy.pdf. (Accessed 25 March 2005).

Deaton, Angus M., and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy* 102 (3): 437–67.

Dubey, Pradeep, John Geanakoplos, and Martin Shubik. 2005. "Default and Punishment in General Equilibrium." *Econometrica* 73 (1): 1–38.

Fay, Scott, Erik Hurst, and Michelle White. 2002. "The Household Bankruptcy Decision." *American Economic Review* 92 (3) 706–18.

Edelberg, W. 2003. "Risk-based Pricing of Interest Rates in Household Loan Markets." Federal Reserve Board Working Paper 2003-62. Available at http://www.federalreserve.gov/pubs/feds/2003/200362/200362pap.pdf (accessed 17 March 2005).

Elul, R., and N. Subramanian. 2002. "Forum-shopping and Personal Bankruptcy." *Journal of Financial Services Research* 21 (3): 233–55.

Fisher, Jonathan D. 2003. "The Effect of Government Unemployment Benefits, Welfare Benefits, and Other Income on Personal Bankruptcy." Manuscript, Bureau of Labor Statistics .

Grant, C., 2003. "Evidence on the Effect of U.S. Consumer Bankruptcy Exemption." Mimeo, European University Institute.

Gropp, R., J. K. Scholz, M. J. White. 1997. "Personal Bankruptcy and Credit Supply and Demand." *Quarterly Journal of Economics* 112 (1): 217–51.

Gross, David B., and Nicholas Souleles. 2002. "Explaining the Increase in Bankruptcy and Delinquency: Stigma vs. Risk Composition." University of Pennsylvania, Wharton Financial Institutions Center Working Paper 98-28-B.

Huggett, Mark. 1993. "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." *Journal of Economic Dynamics and Control* 17 (5–6): 953–69.

Kocherlakota, Narayana. 1996. "Implications of Efficient Risk Sharing without Commitment." *Review of Economic Studies* 63 (4): 595–609.

Li, Wenli, and Pierre-Daniel Sarte. Forthcoming. "The Macroeconomics of U.S. Consumer Bankruptcy Choice: Chapter 7 or Chapter 13?" *Journal of Monetary Economics*.

Livshits, Igor, James MacGee, and Michele Tertilt. 2003. "Consumer Bankruptcy: A Fresh Start." Federal Reserve Bank of Minneapolis, Working Paper 617.

Ljungqvist, Lars, and Thomas J. Sargent. 2000. *Recursive Macroeconomic Theory*. Cambridge, Mass.: MIT Press.

Mateos-Planas, Xavier, and Giulio Seccia. 2004. "Welfare Implications of Endogenous Credit Limits with Bankruptcy." Mimeo, Department of Economics, University of Southampton, UK.

Moss, David A., and Gibbs Johnson. 1999. "The Rise of Consumer Bankruptcy: Evolution, Revolution, or Both?" *American Bankruptcy Law Journal* 73 (2): 311–51.

Obstfeld, Maurice, and Kenneth Rogoff. 1996. *Foundations of International Macroeconomics*. Cambridge, Mass.: MIT Press.

Pavan, M. 2003. "Consumer Durables and Risky Borrowing: The Effects of Bankruptcy Protection." Mimeo, Boston College.

Shepard, Lawrence. 1984. "Personal Failures and the Bankruptcy Reform Act of 1978." *Journal of Law and Economics* 27 (2): 419–37.

Sullivan, Teresa A., Elizabeth Warren, and Jay Lawrence Westbrook. 1989. *As We Forgive Our Debtors : Bankruptcy and Consumer Credit in America*. New York: Oxford University Press.

_____. 2000. *The Fragile Middle Class: Americans in Debt.* New Haven: Yale University Press.

Townsend, R. M. 1979. "Optimal Contracts and Competitive Markets with Costly State Verification." *Journal of Economic Theory* 21 (October): 265–93.

Yue, Vivian Z. 2004. "Sovereign Default and Debt Renegotiation." Mimeo, University of Pennsylvania.

Zame, William R. 1993. "Efficiency and the Role of Default When Security Markets Are Incomplete." *American Economic Review* 83 (5) 1142–64.

Zha, Tao. 2001. "Bankruptcy Law, Capital Allocation, and Aggregate Effects: A Dynamic Heterogeneous Agent Model with Incomplete Markets." *Annals of Economics and Finance 2* (November): 379–400.