Evaluating default policy: The business cycle matters

Grey Gordon
Department of Economics, Indiana University

More debt forgiveness directly benefits households but indirectly makes credit more expensive. How does aggregate risk affect this trade-off? In a calibrated general equilibrium life-cycle model, aggregate risk reduces the welfare benefit of making default very costly when the costs are borne by all households at all times. The result does not necessarily extend to state-contingent policies. The Bankruptcy Abuse Prevention and Consumer Protection Act of 2005 in particular generates a small welfare loss with or without aggregate risk.

1. Introduction

Consumer default law faces a trade-off. Specifically, while more debt relief directly benefits indebted households, it also makes credit more expensive: To cover default-related losses, creditors charge a premium. Recognizing this trade-off, the most recent revision of bankruptcy law, the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005 (BAPCPA), prevents households with above-median income from filing for Chapter 7 bankruptcy. This law has had a drastic short-term effect on Chapter 7 filing rates per household: Filing rates, which averaged 0.9% from 2000 to 2004, rose to 1.4% in 2005, averaged only 0.4% from 2006 to 2007, and have since recovered to 0.7% over 2008–2013. Passed just before the Great Recession, it is important to understand whether this decrease in debt relief improved the balance of debt relief and credit or worsened it.

More generally, this trade-off should be clearly understood because it has wide-reaching ramifications. Bankruptcy laws differ greatly both across countries and across time, and the United States in particular has had major revisions to bankruptcy law on average every 40 years.[1] Not only is there large variation in laws, but, as evidenced by

[1]Robe, Steiger, and Michel (2006) document the variation in bankruptcy law across time and location. The present form of U.S. bankruptcy law was codified in 1898 and has since had major revisions in 1938, 1978, and 2005 (BAPCPA). Some key provisions of the 1978 law were not enacted until 1984.

filing rates before and after BAPCPA, households respond to these laws using debt forgiveness when it is an option. In fact, the option value of debt forgiveness appears to be significantly higher in recessions. While it is difficult to tell from the data precisely how cyclical default rates are, from 1960 to 1984, when bankruptcy laws were mostly constant, (log) default rates had a correlation with (log) output of -0.31 and a standard deviation of 0.075.[2]

To properly assess the trade-off inherent to bankruptcy law, one must have an accurate understanding of the risks households face. While the literature has considered earnings risk and "expenditure risk" (i.e., risk coming from shocks such as uninsured medical bills), one important source that the literature has ignored is aggregate risk. As evidenced by countercyclical default rates, aggregate risk plays a substantial role in consumer default. Not only this, but, as mentioned above, BAPCPA made debt forgiveness more difficult just before the Great Recession. Was BAPCPA welfare improving in the context of aggregate risk? More generally, how does aggregate risk affect the consequences of restricting or eliminating default?

I find that aggregate risk substantially reduces the welfare and debt associated with high default-cost regimes (when that cost is uniform across households) but has little effect on low default-cost regimes. When aggregate risk is introduced, households have two ways to insure themselves against it: They can either default more in recessions or reduce their debt. When default is very costly, households choose the second option. This result naturally extends to infinitely costly default (a case frequently investigated in the literature) where most households can *only* choose the second option. While aggregate risk can have a large effect on credit usage, it tends to have a small effect on credit prices despite a volatile and countercyclical default rate. This is because business cycles typically have short durations, and so the potential loss creditors face if a recession is realized is mostly offset by the potential gain from an expansion.

In the case of BAPCPA, where high default costs are applied only to above-median households, I find aggregate risk has very little effect on the main outcomes: With or without it, there is a small welfare loss for newborn households, a large expansion in credit, a small decrease in total filings, and a sharp reduction in filings by above-median households. BAPCPA's means test, together with a hump-shaped earnings profile, effectively makes bankruptcy costs age-dependent. Since young households are overwhelmingly below median, their default costs, and, consequently, available credit, are mostly unchanged. Rather, it is middle-aged households who face higher default costs, whose credit opportunities increase, and who borrow more. Welfare is slightly lower, despite expanded credit, because BAPCPA does a poor job of insuring households. BAPCPA does

---

[2]If one takes the sample as 1960–2004, the filing rate is surprisingly acyclical with a correlation of just -0.01. However, the filing rate data appear to have a structural break around 1984: The average filing rate from 1960 to 1984 was 0.25%, the rate steadily increased until around 1997, and from 1997 to 2004, the average rate was 0.96% (Livshits, MacGee, and Tertilt (2010), investigate the rise, attributing it to lower default and lending costs). For the small sample of 1997–2004 (with only one recession), filing rates are very countercyclical at -0.94 with a standard deviation of 0.082. Both the filing rates and output series are annual and are logged and Hodrick–Prescott (HP) filtered (with smoothing parameter 100). For a description of the data, see Appendix A.

particularly poorly at insuring against expenditure shocks since the means test does not account for debt: When hit by expenditure shocks, most above-median households avoid a costly bankruptcy process, but they have to drastically cut consumption to do so.

I also investigate who bears aggregate risk and to what extent bankruptcy law can change this. I find that while bankruptcy policy can help insure households against aggregate risk, it cannot change who bears aggregate risk, namely, the young. However, if all households have access to a complete set of aggregate-state contingent, defaultable Arrow securities, then the young are the most insured. In fact, in this case, their consumption growth ends up larger in recessions.

Aware that bankruptcy laws vary substantially and affect many households, the literature has examined the trade-off of debt forgiveness and credit. Without expenditure shocks, virtually all the papers have found the extreme of making default *infinitely* costly vastly improves welfare. Athreya (2002, 2008), Athreya, Tam, and Young (2009a, 2009b), and Chatterjee and Gordon (2012) all find this result, and they do so in a wide variety of environments.[3] Athreya, Tam, and Young (2009b) in particular find the result holds for numerous specifications of earnings risk and preferences. In Appendix C, this paper presents a similar finding: Without expenditure shocks, infinite default costs produce a welfare gain of 4.22% that aggregate risk reduces to 3.05%. With expenditure shocks, Livshits, MacGee, and Tertilt (2007) find the U.S. system is preferable to a European-like system that vastly restricts default. The present paper, which considers an infinite-cost regime for those not hit by expenditure shocks, finds a similar result after accounting for aggregate risk: Aggregate risk reduces the welfare gain of the infinite-cost regime from 0.21% to a welfare *loss* of 0.13%.

The literature has also looked at BAPCPA and has found mixed results. Athreya (2002), Li and Sarte (2006), and Nakajima (2008) have found modest changes in welfare and allocations from it, while Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007) and Mitman (2011) have found sizable welfare gains. Athreya, Sánchez, Tam, and Young (2014) find BAPCPA likely restrained bankruptcy rates in the last recession, which is similar to this paper's findings.

A small but growing literature has looked at default and business cycles. Nakajima and Ríos-Rull (2005, 2010) examine how default amplifies or smooths aggregate shocks. Fieldhouse, Livshits, and MacGee (2014) investigate how well the Livshits, MacGee, and Tertilt (2007) framework captures business cycle regularities. Herkenhoff (2013) examines how historical changes in credit have affected the business cycle durations. Since these papers focus on default's effect on aggregate dynamics rather than aggregate dynamics' effect on default, they are complementary to the present study. Additionally, Athreya et al. (2014) examine BAPCPA's effects over a steady-state to steady-state transition path with aggregate shocks during the transition, but they do not consider welfare. A technical contribution of this paper is to model the economy in a way that ensures creditors make zero profits loan-by-loan and that loans are priced by no arbitrage.

The quantitative framework I use is from Chatterjee et al. (2007) and Livshits, MacGee, and Tertilt (2007) extended to incorporate aggregate risk. In addition to the

---

[3]To my knowledge, only Li and Sarte (2006) find the opposite result, but the welfare criterion they use is very sensitive to transitional dynamics, and these are not computed.

idiosyncratic earnings and expenditure shocks of those papers, I allow for aggregate risk of three types: Changes in total factor productivity, changes in earnings variance à la Storesletten, Telmer, and Yaron (2004), and changes in exogenous labor supply. A description of the data, additional calibration and computation details, and extensive robustness exercises are available in the Appendices, available in a supplementary file on the journal website, http://qeconomics.org/supp/372/supplement.pdf and http://qeconomics.org/supp/372/code_and_data.zip.

2. Model

The model is set up recursively using $\mathcal{S} = (z, \mu, K)$ as the aggregate state, where z is total factor productivity (TFP), μ is a distribution of households, and $K > 0$ is aggregate capital holdings.[4] Productivity $z \in Z$ evolves according to a finite-state Markov chain $F(z'|z)$. The aggregate state evolves according to a law of motion Γ with $\mathcal{S}'_{z'} = \Gamma(z', \mathcal{S})$ denoting the next period's aggregate state conditional on a z' realization.

2.1 *Basic environment and preferences*

The economy is populated by a unit mass of households that die with certainty after T years. Households differ in the productive efficiency e of their unit time endowment. Efficiency is independently and identically distributed (i.i.d.) conditional on TFP z and "characteristics" s. The density function of e is denoted $f(e|s, z)$, and it has support in $\mathbb{R}^{++}$ for all s. Characteristics, which include persistent components of earnings, expenditure shocks, and *the household's age*, lie in a finite set S and evolve according to a conditional distribution $F(s'|s, z')$. Age, of course, evolves deterministically. Households face an age-dependent conditional probability of survival ρ_s. Households that die are replaced by "newborn" households that have zero assets, efficiency distributed according to $\hat{f}(e|s, z)$, and characteristics distributed according to $\hat{F}(s|z)$.

Preferences over consumption c are time separable with discount factor $\beta > 0$ and a period utility function $U(c, s)$ given by

$$U(c, s) = (c/\theta_s)^{1-\sigma}/(1-\sigma), \quad \sigma > 0, \sigma \neq 1, \tag{1}$$

with $U(c, s) = \log(c/\theta_s)$ in the case of $\sigma = 1$. Changes in household composition are captured by θ_s, an age-dependent effective number of household members. Initially endowed with zero assets, $a = 0$, households accumulate debt $a < 0$ or savings $a \geq 0$ over time. Households are subject to i.i.d. expenditure shocks $x \geq 0$ that directly affect their net worth, $a - x$. The expenditure shock support is positive, finite, and includes zero. The probability of state x being realized is denoted π_x.

A neoclassical production firm operates the production technology $zK^\alpha N^{1-\alpha}$ with $\alpha \in (0, 1)$ that uses capital K rented at rate $r(\mathcal{S})$ and labor N hired at wage $w(\mathcal{S})$ as inputs. Capital depreciates at a rate $\delta \in (0, 1]$.

---

[4]The inclusion of K in $\mathcal{S}$ is for convenience. An earlier version of this paper (available by request) had $\mathcal{S} = (z, \mu)$.

2.2 *Legal environment*

Households have a credit history $h \in \{0, 1\}$. Households in good standing, $h = 0$, have the right to file for bankruptcy, $d = 1$. The bankruptcy option is designed to resemble a Chapter 7 bankruptcy, which is often referred to as a "fresh start." If a household files, their debts are discharged in exchange for all their assets, they may not save or borrow, they face a pecuniary cost from default equal to a fraction $\chi(e) \in [0, 1)$ of their income, and they are subsequently in bad standing, $h' = 1$. As is common in the *sovereign* default literature (see, for instance, Arellano (2008)), the cost $\chi(e)$ has both a flat and a progressive component, and I use the functional form $\max(0, \chi_0 - \chi_1 e^{-1})$ with $\chi_0 \in [0, 1]$ and $\chi_1 \geq 0$. The cost is deadweight loss. Households with a bad credit history $h = 1$ are not allowed to borrow but may save, and they face the pecuniary cost $\chi(e)$. This history is removed, that is, h' becomes zero, with probability $1 - \lambda$. Households begin life with $h = 0$.

The inclusion of expenditure shocks means some households with a bad credit record, despite having positive assets a, will need to default. To handle this case, I allow a household to obtain a discharge only if they have negative net worth, $a - x < 0$.[5] Because U.S. law forbids households from filing for a Chapter 7 bankruptcy until 6 years have passed from a previous Chapter 7 filing, I interpret the case of $h = 1$ and $d = 1$ as a default by other means (either a Chapter 13 filing or an informal default). Throughout the paper, reported filing rates refer only to the case of $h = 0$ and $d = 1$, because filing rates in the data are measured using Chapter 7 filings.

2.3 *Asset markets*

Households borrow or save using z-contingent contracts that resemble Arrow securities. A face value $a'_{z'}$, which is constrained to lie in a finite set A, is to be delivered (if positive) or repaid (if negative) if and only if the next period's productivity shock is z'. The price $q_{z'}(a'_{z'}, s, h; \mathcal{S})$ depends on all factors that can influence next period's default decision (e does not appear because it is i.i.d. conditional on s and z). The prices $q_{z'}(\cdot, s, h; \mathcal{S})$ define a "price schedule."

Giving households access to z-contingent contracts offers a number of theoretical advantages, which are discussed in Section 2.7. However, Section 3.4 shows that if access to these contracts is completely unrestricted, debt and default end up counterfactually procyclical and volatile: Households smooth consumption by pledging to repay in expansions; when an expansion occurs, there are many households in debt and consequently many file for bankruptcy. To improve the model's business cycle predictions, I restrict household portfolios $\{a'_{z'}\}$ to lie in a set $P(s)$. The exact portfolio restrictions are described in Section 3.2.

---

[5]Chatterjee et al. (2007) make a similar assumption, but they force a household with negative net worth to default (which they typically want to do anyway).

2.4 *The household problems*

Taking the law of motion and prices as given, households solve the following problems. Let $V(a, e, s, h; \mathcal{S})$ denote the value function of a household. A household in good standing $h = 0$ that can repay its debt solves

$$V(a, e, s, h = 0; \mathcal{S}) = \max_{d \in \{0,1\}} (1 - d) \cdot V^R(a, e, s, 0; \mathcal{S}) + d \cdot V^D(e, s; \mathcal{S}), \tag{2}$$

where the value of repaying is

$$V^R(a, e, s, h = 0; \mathcal{S}) = \max_{c \geq 0, \{a'_{z'}\} \in P(s)} U(c, s) + \beta \rho_s \mathbb{E} V(a'_{z'}, e', s', 0; \mathcal{S}'_{z'}),$$

$$c + \sum_{z'} q_{z'}(a'_{z'}, s, h = 0; \mathcal{S}) a'_{z'} = w(\mathcal{S})e + a - x \tag{3}$$

and the value of defaulting is

$$V^D(e, s; \mathcal{S}) = \max_{c \geq 0, \{a'_{z'} | a'_{z'} = 0\}} U(c, s) + \beta \rho_s \mathbb{E} V(0, e', s', 1; \mathcal{S}'_{z'}),$$

$$c = w(\mathcal{S})e(1 - \chi(e)). \tag{4}$$

A household in bad standing $h = 1$ solves

$$V(a, e, s, h = 1; \mathcal{S}) = \max_{d \in \{0, \mathbf{1}[a-x<0]\}} (1 - d) \cdot V^R(a, e, s, 1; \mathcal{S}) + d \cdot V^D(e, s; \mathcal{S}), \tag{5}$$

where the value of repaying is

$$V^R(a, e, s, h = 1; \mathcal{S}) = \max_{c \geq 0, \{a'_{z'} | a'_{z'} \geq 0\} \in P(s)} U(c, s) + \beta \rho_s \mathbb{E} V(a'_{z'}, e', s', h'; \mathcal{S}'_{z'}),$$

$$c + \sum_{z'} q_{z'}(a'_{z'}, s, h = 1; \mathcal{S}) a'_{z'} = w(\mathcal{S})e(1 - \chi(e)) + a - x, \tag{6}$$

with h' stochastic in this case. Any household that cannot repay its debt must default. All the expectations are conditioned on s, $\mathcal{S}$, and surviving to the next period. The associated policy functions are denoted $d(a, e, s, h; \mathcal{S})$, $a'_{z'}(a, e, s, h; \mathcal{S})$, and $c(a, e, s, h; \mathcal{S})$.

2.5 *The intermediary's problem*

The counterparty of the debt–savings contracts is a financial intermediary. The intermediary maximizes the net present value of dividends using contracts, capital K', and Arrow securities $A'_{z'}$. The Arrow securities, although in zero net supply, make the intermediary's problem well defined by providing Arrow security prices $\bar{q}_{z'}(\mathcal{S})$ that can be used to discount future dividends.[6] As discussed in detail in Section 2.7, the prices

---

[6]In a related environment, Carceles-Poveda and Coen-Pirani (2010) show this discounting causes an equivalence between households choosing capital investment and infinitely lived firms choosing it. Carceles-Poveda (2009) explores other types of discounting and finds macroeconomic aggregates are sensitive to the discounting used.

$\bar{q}_{z'}(\mathcal{S})$ are endogenously determined by market clearing conditions. A contract—which households view as a defaultable Arrow security $a'_{z'}$—is to the intermediary an asset costing $q_{z'}(a'_{z'}, s, h; \mathcal{S})a'_{z'}$ units of the consumption good in state $\mathcal{S}$ and paying out $p(a'_{z'}, s, h; \mathcal{S}'_{z'})a'_{z'}$ units of the consumption good in state $\mathcal{S}'_{z'}$ (and zero in every other state). In equilibrium, p is a repayment rate that is consistent with household default decisions and stochastic transitions.

In stating the intermediary's problem, it is useful to think of the intermediary as choosing portfolios of contract holdings $l_{z'} : A \times S \times \{0, 1\} \to \mathbb{R}$, one for each z'. Then the intermediary's problem can be written as

$$P(W; \mathcal{S}) = \max_{K', \{l'_{z'}\}, \{A'_{z'}\}} D + \sum_{z'} \bar{q}_{z'}(\mathcal{S})P(W'_{z'}; \mathcal{S}'_{z'}), \tag{7}$$

$$D + \sum_{z', a', s, h} q_{z'}(a', s, h; \mathcal{S})a'l'_{z'}(a', s, h) + K' + \sum_{z'} \bar{q}_{z'}(\mathcal{S})A'_{z'} = W, \tag{8}$$

$$W'_{z'} = \sum_{a', s, h} p(a', s, h; \mathcal{S}'_{z'})a'l'_{z'}(a', s, h) + (1 + r(\mathcal{S}'_{z'}) - \delta)K' + A'_{z'}, \tag{9}$$

where W is the intermediary's wealth and D is a dividend. Let the associated policies functions be denoted $D(W; \mathcal{S})$, $A'_{z'}(W; \mathcal{S})$, $K'(W; \mathcal{S})$, $l'_{z'}(a, s, h, W; \mathcal{S})$, and $W'_{z'}(W; \mathcal{S})$. Since the household problem was specified without dividends, equilibrium requires that the intermediary chooses $D = 0$ and $W'_{z'} = 0$ (for each z') when $W = 0$. This generates zero dividend payouts for any history of shocks as long as $W = 0$ initially, which I assume.

2.6 *Equilibrium*

A recursive competitive equilibrium is a collection of price functions r, w, $\bar{q}_{z'}$, and $q_{z'}$, repayment rates p, policy functions c, d, $a'_{z'}$, K', $A'_{z'}$, $l'_{z'}$, D, and $W'_{z'}$, value functions V and P, and a law of motion Γ such that the following conditions hold:

1. The policies and value functions solve the household problems.

2. The policies and value function solve the intermediary's problem.

3. Factor prices are given by their marginal products (ensuring the production firm optimizes).

4. The labor, Arrow security, and contract markets clear: for each z', a, s, and h,

$$N = \int e \, d\mu, \tag{10}$$

$$A'_{z'}(W = 0; \mathcal{S}) = 0, \tag{11}$$

$$l'_{z'}(a', s, h, W = 0; \mathcal{S}) + \int \mathbf{1}[a' = a'_{z'}(a, e, s, h; \mathcal{S})]\mu(da, de, s, h) = 0. \tag{12}$$

5. The goods market clears (as ensured by Walras' law).

6. Repayment rates are consistent: for each z', a', s, and h,

$$p(a', s, h; \mathcal{S}'_{z'}) = \rho_s \mathbb{E}(1 - d(a', e', s', h'; \mathcal{S}'_{z'})), \tag{13}$$

where the expectation is conditioned on s and h.

7. Dividends are zero: $D(W = 0; \mathcal{S}) = 0$.

8. The intermediary saves using capital: $K'(W = 0; \mathcal{S}) > 0$.

9. The intermediary has zero wealth next period: for each z', $W'_{z'}(W = 0; \mathcal{S}) = 0$.

10. The law of motion Γ is consistent with stochastic transitions, household policies, and the intermediary's policies.

2.7 Why this asset structure?

The asset structure, both at the household and intermediary level, carries a number of advantages and is similar in many respects to the one in Krusell, Mukoyama, and Şahin (2010).[7] As mentioned in Section 2.5, the presence of Arrow securities makes the intermediary's problem well defined by providing Arrow security prices that the intermediary can use to discount future dividends. This section demonstrates the other benefits of the asset structure and also provides some equilibrium characterizations.

One principal benefit of giving the intermediary access to Arrow securities is that they allow contracts to be priced according to no arbitrage. In particular, no arbitrage requires

$$q_{z'}(a'_{z'}, s, h; \mathcal{S}) = \bar{q}_{z'}(\mathcal{S}) p(a'_{z'}, s, h; \mathcal{S}'_{z'}) \tag{14}$$

because each contract can be replicated by $p(a'_{z'}, s, h; \mathcal{S}'_{z'})a'_{z'}$ units of an $A'_{z'}$ Arrow security and the contract's price is $q_{z'}(a'_{z'}, s, h; \mathcal{S})a'_{z'}$.[8] Without Arrow securities, or some other set of securities spanning aggregate risk, the risk neutrality of the intermediary would play a role in contract pricing.

Another benefit is that the intermediary is indifferent over all feasible allocations as long as prices satisfy no arbitrage conditions. To see this, first note that the return on a unit of capital can be replicated by a choice of $A'_{z'} = 1 + r(\mathcal{S}'_{z'}) - \delta$ for each z'. Consequently, no arbitrage requires

$$1 = \sum_{z'} \bar{q}_{z'}(\mathcal{S})(1 + r(\mathcal{S}'_{z'}) - \delta). \tag{15}$$

Because the first order conditions of the intermediary's problem are precisely (14), (15), and $\bar{q}_{z'}(\mathcal{S}) = \bar{q}_{z'}(\mathcal{S})$, the intermediary is indifferent over feasible allocations.

---

[7]Both here and in Krusell, Mukoyama, and Şahin (2010), the asset structure makes the firm problems well defined. In Krusell, Mukoyama, and Şahin (2010), dividends are nonzero because firms earn profits. Here, the structure makes the intermediary's problem well defined and also results in zero profits. The only portfolio restriction in Krusell, Mukoyama, and Şahin (2010) is that asset holdings be nonnegative.

[8]The value of $q_{z'}(0, s, h; \mathcal{S})$ is a normalization (the contract price is zero regardless).

Since the benefits listed so far come from the inclusion of zero net supply Arrow securities, why not have a more standard portfolio choice problem (such as just choosing a bond or just choosing capital) at the household level? There are two reasons. First, without household access to flexible portfolios, in general, it is impossible for the intermediary to make "zero profits" in the sense of always having zero wealth and never distributing a dividend. To see this, consider the equilibrium requirement $W'_{z'} = 0$ for each z'. Using contract and Arrow security market clearing with the definition of $W'_{z'}$ in (9), this requires

$$K'(W = 0; \mathcal{S}) = \frac{\int p(a'_{z'}(a, e, s, h; \mathcal{S}), s, h; \mathcal{S}'_{z'}) a'_{z'}(a, e, s, h; \mathcal{S}) \, d\mu}{1 + r(\mathcal{S}'_{z'}) - \delta} \qquad (16)$$

for each z'. Consequently, the right-hand side cannot vary with z'. If there is no default, this is satisfied if every household portfolio replicates capital, $a'_{z'} = k'(1 + r(\mathcal{S}'_{z'}) - \delta)$ for some k' and each z'. With default and more than one TFP value, this is hopeless: Default rates, and consequently p, naturally vary with the aggregate state $\mathcal{S}'_{z'}$ in a nontrivial way.

The second reason for giving at least some households flexible portfolios is that, in this case, (16) determines Arrow security prices. This is most easily seen in the case of only two TFP states, g and b. Then (16) requires

$$\frac{\int p(a'_g(a, e, s, h; \mathcal{S}), s, h; \mathcal{S}'_g) a'_g(a, e, s, h; \mathcal{S}) \, d\mu}{\int p(a'_b(a, e, s, h; \mathcal{S}), s, h; \mathcal{S}'_b) a'_b(a, e, s, h; \mathcal{S}) \, d\mu} = \frac{1 + r(\mathcal{S}'_g) - \delta}{1 + r(\mathcal{S}'_b) - \delta}. \qquad (17)$$

For a given K', the relative price $\bar{q}_g / \bar{q}_b$ uniquely determines $\bar{q}_g$ and $\bar{q}_b$ from (15). As $\bar{q}_g / \bar{q}_b \downarrow 0$, saving using a'_g becomes arbitrarily cheap while saving with a'_b does not. If some households have flexible portfolios, then this causes the left-hand side of (17) to rise. As $\bar{q}_g / \bar{q}_b \uparrow \infty$, the reverse is true. In the general case, (16) imposes $\#Z$ conditions that would be satisfied with K' and the $\#Z - 1$ relative prices $\bar{q}_{z_1}/\bar{q}_{z_{\#Z}}, \ldots, \bar{q}_{z_{\#Z-1}}/\bar{q}_{z_{\#Z}}$ with the prices in levels determined by (15).[9]

3. Calibration and baseline properties

This section discusses the calibration and the baseline model's properties. The data, along with some less important aspects of the calibration, are described in Appendix A.

3.1 *Parameters chosen a priori*

The model period is a year. Households begin life at age 20, retire at 65, and live to at most 85. The coefficient of relative risk aversion σ is 2, the capital share α is 0.36, the depreciation rate δ is 0.10, and the probability of a bad credit record remaining is $\lambda = 0.9$. All of these are from Chatterjee et al. (2007). The mortality profile ρ_s is from Hubbard,

---

[9]I do not provide a proof of existence. Chatterjee et al. (2007) prove existence in the case of $\#Z = 1$ in a very similar model.

Skinner, and Zeldes (1994). The household-size profile θ_s is calibrated using Fernández-Villaverde and Krueger (2007) and is similar to the one in Livshits, MacGee, and Tertilt (2007).

The calibration with aggregate risk has a two-state TFP process with $Z = \{g, b\}$. The process is symmetric with $F(g|g) = F(b|b) = 2/3$ implying an average business cycle duration of 3 years. The values $g = 1.0224$ and $b = 0.9776$ generate the 2.24% unconditional standard deviation in Cooley (1995). The calibration without aggregate risk has $Z = \{1\}$.

The expenditure shock values and probabilities, which represent the costs and likelihood of uninsured health expenditures, having children, or experiencing a divorce are taken from Livshits, MacGee, and Tertilt (2007) (whose model period is 3 years) and converted to annual values.[10] This results in a "small" shock $x = 0.762$ hitting with probability 2.42% and a larger shock $x = 2.37$ hitting with probability 0.15% (relative to average earnings, the magnitudes are roughly 0.92 and 2.86). I assume they continue to hit households in retirement.

Because earnings, credit, and the value of a default option are closely connected in the model, I try to capture three potentially important findings in the literature. First, earnings shocks are much less persistent early in life (Karahan and Ozkan (2011)). Second, the variance of persistent earnings shocks increases in recessions and decreases in expansions (Storesletten, Telmer, and Yaron (2004)). Last, the earnings distribution has a thick right tail (Castañeda, Díaz-Giménez, and Ríos-Rull (2003)).

To do this, I use two efficiency processes for working households. The efficiency process for the majority of working households, which I refer to as the log process, is governed by

$$
\begin{aligned}
e_{h,z} &= \nu \psi_z \phi_h \exp(u_h + \varepsilon), \\
u_h &= \gamma_{h-1} u_{h-1} + \eta_{h,z}, \qquad u_0 = 0, \\
\eta_{h,z} &\sim N\big(0, \sigma^2_{\eta,h,z}\big), \qquad \varepsilon \sim N\big(0, \sigma^2_\varepsilon\big),
\end{aligned}
\tag{18}
$$

where h denotes age. This process has a deterministic earnings profile ϕ_h, a persistent component u_h, a transitory shock ε, and an "aggregate labor supply shifter" ψ_z. As labor is supplied inelastically, the supply shifter is used to match the cyclical volatility of hours worked. The persistence of the shock $\eta_{h,z}$ is determined by γ_{h-1}, which is age-dependent, as is the variance $\sigma^2_{\eta,h,z}$. The economy-wide average e is normalized to 1 using ν. The log-process parameters $(\gamma_h, \sigma^2_{\eta,h,1}, \sigma^2_\varepsilon)$ are from Karahan and Ozkan (2011). For countercyclical earnings variance, the ratio $\sigma_{\eta,h,g}/\sigma_{\eta,h,b}$ is assumed to be age-independent and equal to 0.59, the value in Storesletten, Telmer, and Yaron (2004), with $0.5\sigma_{\eta,h,b} + 0.5\sigma_{\eta,h,g} = \sigma_{\eta,h,1}$. The labor supply shifter ψ_z is calibrated to match the 1.74% standard deviation of log hours worked from Castañeda, Díaz-Giménez, and Ríos-Rull (1998), resulting in $(\psi_g, \psi_b) = (1.023, 0.977)$ (with $\psi_1 = 1$). The earnings profile ϕ_h is from Hubbard, Skinner, and Zeldes (1994).

---

[10]To do this conversion, I assume the annual shocks are also i.i.d. and that the magnitudes are the same as in Livshits, MacGee, and Tertilt (2007). The probabilities are set so that the probability of being hit with an expenditure shock over a 3 year period is the same as in Livshits, MacGee, and Tertilt (2007).

While most households have this process, some have a "right-tail" process

$$e_{h,z} = \nu \psi_z \phi_h v,$$

$$v \sim \left(\frac{v - \underline{v}}{\overline{v} - \underline{v}} \right)^\xi \quad \text{with support } [\underline{v}, \overline{v}].$$

(19)

A similar process is used in Chatterjee et al. (2007) to successfully generate the right-tail of both the earnings and the wealth distribution. Households are born into the right-tail process with probability $\hat{\pi}_r$, move to the log process with probability π_{rl}, and transit back with probability π_{lr}. When transiting to the log process, they draw u_h from $N(0, \sigma^2_{\eta,1,z})$. To limit the degrees of freedom, I set $\hat{\pi}_r$ to 0.20 and choose π_{lr} such that, given π_{rl}, the measure of working right-tail households is constant at $\hat{\pi}_r$. Consequently, 20% of working age households have the right-tail process and 80% have the log process. Loosely speaking, this also means households with the right-tail process are in the top 20% of earners and households with the log-process are in the bottom 80%.

Households with the log process at age R retire in the following period, with efficiency

$$e_z = \kappa_F \nu \psi_z \phi_R \exp(u_R) + \kappa_G \psi_z$$

(20)

from then on. This process is very similar to the one in Livshits, MacGee, and Tertilt (2007) and Athreya, Tam, and Young (2009a). Households that reach retirement with the right-tail process have efficiency

$$e_z = \kappa_F \nu \psi_z \phi_R \mathbb{E} u + \kappa_G \psi_z.$$

(21)

The retirement parameters (κ_F, κ_G) are set to $(0.35, 0.15)$, giving an average replacement rate of roughly 50%.[11] Robustness checks for these parameters are conducted in Appendix C.

3.2 *Portfolio availability*

Recall that the portfolio $P(s)$ available to households is allowed to vary with their characteristics s. I now let $P(s)$ equal $\{(a'_g, a'_b) \in A \times A | a'_g = a'_b \text{ if } a'_g < 0 \text{ or } a'_b < 0\}$ for households with the right-tail process and $P(s)$ equal $\{(a'_g, a'_b) \in A \times A | a'_g = a'_b\}$ for households with the log process.[12] Roughly speaking, this means the bottom 80% of earners only have access to a bond $a'_g = a'_b$, while the top 20% have access to a bond for borrowing but can save using any (a'_g, a'_b) combination.[13] Section 3.4 shows that these portfolio restrictions bring the model's cyclical properties closer to the data's.

---

[11]Athreya, Tam, and Young (2009a) use $(\kappa_F, \kappa_G) = (0.35, 0.20)$ but do not have the right-tail of the earnings distribution.

[12]Following Livshits, MacGee, and Tertilt (2007), I also assume a usury law prevents households from choosing any portfolio with an interest rate greater than 100%. Without this assumption, households sometimes find it optimal to borrow huge amounts at very high interest rates, skewing the interest rate and debt statistics. Because of this assumption, P should technically be a function of $\mathcal{S}$. To simplify notation, this is omitted.

[13]Kennickell (2009) demonstrates that the top 20% of the income distribution hold a disproportionate share of their portfolio in businesses and other nonhousing wealth, while the bottom 80% hold primarily

3.3 *Estimated parameters and baseline properties*

The seven remaining parameters $(\beta, \chi_0, \chi_1, \underline{v}, \overline{v}, \xi, \pi_{rl})$ are used to minimize the distance between seven steady-state model and target statistics. The targets used are standard: The debt–output ratio, the percentage of indebted households, the filing rate, the capital–output ratio, and select wealth and earnings statistics. The targeted values are given in Table 1 and, with the exception of the filing rate, are from Chatterjee et al. (2007). The targeted filing rate, 0.93%, is the average Chapter 7 filing rate for households from 1999 to 2003. In Appendix C, I conduct a robustness check when larger debt statistics, similar to those in Livshits, MacGee, and Tertilt (2007), are targeted. Information on the computation is provided in Appendix B.

The results from the calibration are listed in Table 1. The model does well at delivering the targeted values. This is in contrast to most of the bankruptcy literature, and the discrepancy is mostly due to the flexible nature of the default cost. The model also does fairly well in terms of untargeted statistics. For example, the percentage of filers with below-median income is predicted to be 66%, while in the data it is 69%.[14] The calibration does fail in some respects. For instance, the average interest rate on debt (7.4%) is low relative to the data (12.7%).[15] The calibration also underpredicts the average debt–income ratios of filers, although the numbers are closer for medians.[16] The model also predicts that, on average, filers have a positive asset position a (households only file if $a - x < 0$, but the expenditure shocks are quite large).

The model's untargeted predictions for cyclical properties are reported in Table 2. The model's output, consumption, and investment series inherit the usual real business cycle properties, although the leads and lags are off, in part due to the two-state TFP process. In terms of the bankruptcy and debt statistics, the model comes close to matching the volatilities of debt and interest rates while underpredicting the volatility of filing rates and discharged debt. The model correctly predicts that debt is procyclical and that interest rates (on debt) are nearly acyclical. While filings are much too countercyclical for this sample, the filing rate in the data is more countercyclical over 1960–1984 (−0.31) and over 1997–2004 (−0.96). The discrepancies with respect to filing rates and

---

housing wealth. Additionally, Campbell (2006) shows the portfolios of households with the least assets contain virtually only safe ones. The chosen portfolio restrictions are similar to those in Chien, Cole, and Lustig (2011).

[14]The discharged debt to output ratio is not as good as it appears. The ideal measure of discharged debt in the data would be negative net worth discharged in bankruptcy. However, I do not have data on this, and so for discharged debt, I use the charge-off rate on credit card debt times the stock of revolving consumer credit. The benchmark calibration predicts a large charge-off rate for $-a + x$ and a low (in fact, negative) charge-off rate for $-a$ debt. Appendix C presents a high-debt calibration where the charge-off rates are more reasonable.

[15]One way to fix this would be to include transaction costs as in Livshits, MacGee, and Tertilt (2007). Transaction costs are not included because they complicate the model (especially in general equilibrium) and because they likely are not invariant to bankruptcy policy. For example, if recovery rates on debt were close to 100%, this would presumably reduce transaction costs.

[16]Bermant and Flynn (1999) report that the mean–median ratio of unsecured debt for filers is 1.86.

TABLE 1. Model targets, statistics, and select parameters.

Statistic	Data	Model	Parameter	Value
Targeted Statistics				
Capital–output ratio	3.08	3.08	β	0.956
Debt–output ratio $\times$ 100	0.67	0.61	χ_0	0.918
Population filing (%)	0.93	0.97	χ_1	8.601
Population in debt (%)	6.7	7.2	π_{rl}	0.730
Earnings share of top 20%	60.2	55.5	$\underline{u}$	1.664
Earnings mean–median	1.57	1.56	$\overline{u}$	66.08
Wealth mean–median	4.03	4.01	ξ	0.054
Untargeted Statistics				
Wealth share of top 20%	81.7	74.6		
Wealth share of 4th quintile	12.2	18.5		
Wealth share of 3rd quintile	5.0	5.6		
Wealth share of 2nd quintile	1.3	1.3		
Wealth Gini	0.80	0.72		
Earnings share of 4th quintile	22.9	17.4		
Earnings share of 3rd quintile	13.0	12.8		
Earnings share of 2nd quintile	4.0	8.6		
Earnings Gini	0.61	0.50		
Average interest on debt (%)	12.7	7.4		
Discharged debt–output ratio $\times$ 100*	0.32	0.31		
Discharged $-a$–output ratio $\times$ 100*	0.32	−0.08		
Debt–income of filers	1.62	0.90		
Debt–income of below-median filers	1.66	1.17		
Debt–income of above-median filers	1.58	0.68		
Percentage of filers below median	68.8	66.2		
Population with $d = 1$, any h (%)		1.11		
Right-tail population filing		0.01		
Right-tail debt–output ratio $\times$ 100		0.12		

Note: Model debt is measured as $-a + x$, a filing is measured as $h = 0$ and $d = 1$, and discharged debt is $-a + x$ when $h = 0$ and $d = 1$. Statistics marked with an asterisk (*) have debt in the data measured with revolving consumer credit.

debt discharge may also be due to expenditure shocks playing too large a role in causing default.[17]

3.4 *Cyclical properties without portfolio restrictions*

Without portfolio restrictions, the model's cyclical properties (which are also presented in Table 2) are far from the data's. In particular, the volatilities of debt and interest rates are extremely high, debt and default rates are strongly procyclical, and consumption's excess smoothness is worsened. Additionally, while I have no time-series data on the

---

[17]As already mentioned, the average asset position for filers, measured in terms of a, is actually positive. In the data, only about 1% of households have any assets that are eventually distributed to creditors (Bermant and Flynn (1999)). By overstating the importance of expenditure shocks (which are not time-varying), the model underpredicts the volatilities of filing rates and discharged debt.

TABLE 2. Business cycle properties: Data and model.

Variable (x)	St. Dev. x (in %)	Correlation of Lagged x With y				
		$x(-2)$	$x(-1)$	x	$x(+1)$	$x(+2)$
U.S. (1985–2004)						
Output (y)	1.53	0.12	0.64	1.00	0.65	0.14
Consumption	1.11	−0.30	0.24	0.85	0.78	0.45
Investment	5.78	0.48	0.78	0.86	0.37	−0.24
Hours worked	1.63	0.34	0.69	0.84	0.37	−0.25
Population filing	11.23	0.31	0.18	−0.06	0.04	0.30
Debt*	4.24	0.55	0.47	0.39	0.31	0.05
Discharged debt*	15.35	0.44	0.18	−0.12	0.09	0.45
Interest rates	6.50	0.04	−0.17	−0.01	0.42	0.30
Model						
Output (y)	2.94	−0.16	0.10	1.00	0.10	−0.16
Consumption	0.79	−0.32	−0.10	0.84	0.47	0.17
Investment	8.08	−0.12	0.14	0.99	0.01	−0.23
Labor supply	1.75	−0.09	0.16	0.98	−0.04	−0.26
Population filing	6.84	0.19	−0.06	−0.95	−0.26	0.09
Debt	3.64	0.05	0.23	0.65	−0.64	−0.43
Discharged debt	1.99	0.04	0.24	0.75	−0.49	−0.40
Interest rates	6.54	0.13	0.03	−0.32	−0.14	−0.05
Population in debt	5.72	0.20	0.20	−0.02	−0.96	−0.37
Model without portfolio restrictions						
Output (y)	2.93	−0.16	0.10	1.00	0.10	−0.16
Consumption	0.60	−0.37	−0.18	0.73	0.52	0.26
Investment	8.56	−0.12	0.14	0.99	0.03	−0.22
Labor supply	1.75	−0.09	0.17	0.98	−0.04	−0.27
Population filing	5.16	−0.16	0.08	0.92	0.19	−0.14
Debt	18.83	−0.16	0.10	1.00	0.10	−0.14
Discharged debt	7.40	−0.16	0.10	1.00	0.13	−0.13
Interest rates	19.04	0.19	0.87	−0.33	−0.26	−0.14
Population in debt	24.49	−0.15	0.11	0.99	0.11	−0.16

Note: Statistics marked with an asterisk (*) are measured with revolving consumer credit.

population in debt, the model's prediction of an extremely volatile and procyclical series is hard to believe. These features are all consistent with households smoothing consumption by choosing $a'_g \ll a'_b$: When an expansion occurs, there are more households in debt and so more households default. Overall, portfolio restrictions do a better job at capturing the data's cyclical properties, and so they are adopted in the baseline.

4. Aggregate risk and bankruptcy reform

This section examines how aggregate risk affects the consequences of restricting access to bankruptcy.

4.1 *NFS*

In the baseline, all households in good standing have access to a default option resembling a Chapter 7 bankruptcy. As in the model, Chapter 7 bankruptcy provides a "fresh start," forgiving debt at low cost. Relative to the baseline, which I refer to as FS, I begin by looking at a "no fresh start" (NFS) environment.[18] In NFS, bankruptcy is only an option if a household has been hit by an expenditure shock, that is, $x > 0$, in which case they face the same filing cost $\chi(e)$ as in the benchmark. Households that do default are "never forgiven" in that they forever have a bad credit record $h = 1$ and hence can never borrow again. The bankruptcy literature has found that high default-cost regimes, like NFS, lead to large increases in debt and welfare. However, all existing studies have been done absent aggregate risk, and so it is worthwhile to see whether this result still stands. Additionally, NFS is closest to a standard incomplete markets model (such as Aiyagari (1994)), where default is not allowed.[19] Consequently, by considering how aggregate risk's impact on FS and NFS differ, future research can be informed as to when including a default option may be important. Last, if there is a policy interest in restraining default, NFS provides a bound on what is possible.

With or without aggregate risk, restricting bankruptcy through NFS results in a large increase in debt, a large increase in the population in debt, a lower filing rate (recall NFS's filing rate is nonzero because households can file if an expenditure shock hits), and a lower capital–output ratio. This is borne out in Table 3, which lists statistics for

TABLE 3. Effects on allocations of restricting default.

Economy	K/Y	Debt$/Y \times 100$	% in Debt	% Filing	% Filing $e \leq \psi_z \tilde{e}$
Steady State					
FS	3.08	0.61	7.22	0.97	0.64
NFS	2.99	2.11	15.15	0.64	0.39
BAPCPA	3.07	2.05	12.63	0.87	0.82
Business Cycle					
FS	3.09	0.56	6.83	0.98	0.66
NFS	3.00	1.47	12.78	0.63	0.39
BAPCPA	3.08	1.88	11.91	0.90	0.85
FS flex	3.09	0.61	7.65	0.97	0.63
NFS flex	3.00	1.67	13.47	0.63	0.39

[18]This terminology is taken from Livshits, MacGee, and Tertilt (2007), but their NFS is very different. In particular, in their model, NFS means households can never obtain a discharge (although they are allowed to die in debt). Additionally, unpaid debt rolls over at interest and results in garnishment.

[19]In Appendix C, I examine a model without expenditure shocks, which does reduce to a standard incomplete markets (SIM) model. While expenditure shocks (of sufficient magnitude) make default a necessity, NFS is still very close to a SIM model for households that are never hit with an expenditure shock. The primary difference is that interest rates bear a small risk premium. Virtually all the results of how aggregate risk affects the allocations and welfare of NFS carry over to the environment without expenditure shocks.

FS and NFS. However, another feature in Table 3 is that NFS has a significantly muted impact on debt when one accounts for aggregate risk. For instance, the debt–output ratio, which before increased 350% (from 0.0061 to 0.0211), now increases "only" 260% (from 0.0056 to 0.0147).

This raises two questions. First, why does NFS result in so much more debt? Second, why is this effect reduced by aggregate risk? The basic answer to these questions is that because of earnings uncertainty, impatience, and a hump-shaped earnings profile, households in both economies have *incentive* to borrow. However, only the NFS economy gives households the *opportunity* to borrow large amounts, and this opportunity is diminished by the inclusion of aggregate risk.

The best way to see this is to consider borrowing limits. In the FS economy, the maximum possible loan size is restricted by the price of debt. In particular, the largest loan size is $\max_{a'} \sum_{z'} q_{z'}(a', s, 0; \mathcal{S})(-a')$, which depends on a household's type and the aggregate state. In the NFS economy, the price of debt does not constrain borrowing as the recovery rate p is bounded below (implying debt prices are as well) since households cannot default if $x = 0$. However, as in Aiyagari (1994), household borrowing is constrained by a natural borrowing limit. In particular, because a household can file for bankruptcy only if $x > 0$, the household must not borrow more than they can repay conditional on $x = 0$.

These borrowing limits, averaged across types and aggregate states, are presented in Figure 1 with one unit on the vertical axis being roughly \$50,000.[20] As is clear, the maximum loan size in the NFS economy is uniformly and typically much higher than in the FS economy. Because of this and because households have incentive to use debt, the NFS economy is much more indebted. However, aggregate risk noticeably reduces the natural borrowing limit while leaving the FS limit virtually unchanged. This reduces debt usage in the NFS economy more than in FS.

While Figure 1 answers two questions, it raises two others: Why do the borrowing limits have such different shapes and why does aggregate risk have a differential effect on them? Both answers lie in that the limits are determined by completely different factors: The FS limit is determined by household *willingness* to repay on *average* and the NFS limit is determined by *ability* to repay in the *worst* circumstances. As it turns out, the minimum ability of households to repay is larger than their average willingness, but it is also reduced more by aggregate risk.

To see this, first consider the FS economy. For a small equity premium, the price of an Arrow security is roughly $F(z'|z)$ times the price of risk-free bond, $\bar{q}_B(\mathcal{S})$. Using this relationship, the price of a bond (i.e., $a'_{z'} = a'$ from some a' and all z') for an $h = 0$ household is approximately[21]

$$\bar{q}_B(\mathcal{S})\rho_s \mathbb{E}\big[1 - d\big(a', e', s', 0; \mathcal{S}'_{z'}\big)|s, z\big]. \tag{22}$$

---

[20]Because the average e is normalized to 1 and the wage in the FS economy is roughly 1.2, average earnings are around 1.2 in the FS economy. Equating this with average U.S. household earnings of \$60,000 gives a value of \$50,000.

[21]Combining (13) and (14), the exact price is $\sum_{z'} \bar{q}_{z'}(\mathcal{S})\rho_s \mathbb{E}[1 - d(a', e', s', 0; \mathcal{S}'_{z'})|s, z, z']$.
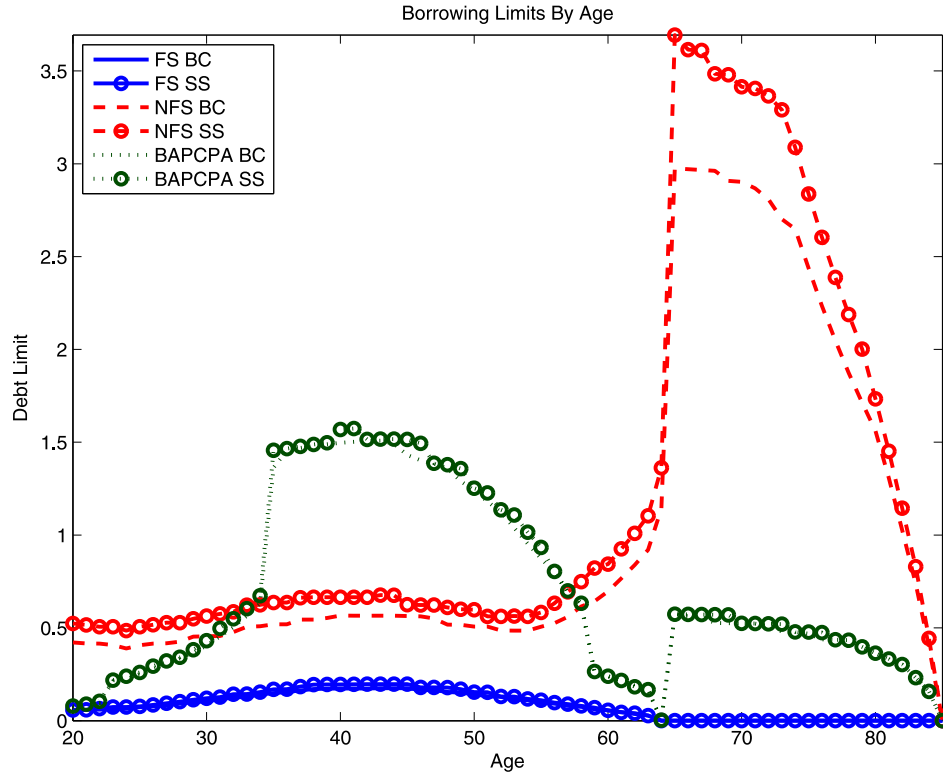
FIGURE 1. Average borrowing limits in FS, NFS, and BAPCPA.

The household's willingness to pay is seen in the default decision and their average willingness in the expectation over it. Credit in the FS economy is not affected much by aggregate risk in large part because recessions are short-lived. Because default rates are very countercyclical (cf. Table 2), if it were known that next period would be a recession, credit would be expensive: The expectation would effectively be conditioned on $z' = b$, making the bond price close to zero. However, since business cycles last for a relatively short time, credit prices reflect both a decreased willingness to pay in recessions and an increased willingness to pay in expansions. In fact, since $F(z'|z)$ is not far from $1/2$, the bond price is not far from $\bar{q}_B \rho_s \mathbb{E}[1 - d|s]$, which is analogous to the steady-state pricing equation. Because of this, the FS borrowing limit is changed little by aggregate risk.

Now consider the NFS economy. There, the borrowing limit can be mechanically calculated as the net present value of future earnings. For a newborn household, the formula is

$$\min_{\{s_t, z_t\}} \sum_{t=2}^{T} \left(\prod_{j=2}^{t-1} \bar{q}_B(\mathcal{S}_j) \rho_{s_j} \pi_0 \right) w(\mathcal{S}_t) \underline{e}_{s_t, z_t}$$

$$\text{s.t.} \quad F(s_{t+1}|s_t, z_{t+1}) > 0 \quad \text{and} \quad \mathcal{S}_{t+1} = \Gamma(z_{t+1}, \mathcal{S}_t) \quad \text{with } s_1, \mathcal{S}_1 \text{ given,}$$

(23)

where $\underline{e}_\xi$ denotes the lowest efficiency conditional on ξ and $x = 0$.[22] Relative to the steady-state limit where z and $\mathcal{S}$ are time-invariant, three factors can make this limit decrease: Changes in the support of e, changes in the support of s, and changes in prices. The support of e changes in the business cycle because of the labor supply shifter ψ_z. Specifically, $\psi_b = 0.977$ causes the limit to decrease by 2.3%. The support of s also changes slightly due to numerical precision, specifically rounding in the normal cumulative distribution function (c.d.f.) computation. However, it is in fact changes in factor prices that account for most of the change. While fluctuations in w and $\bar{q}_B$ are small, compounding—via the $\prod \bar{q}_B \rho_s \pi_0$ term—plays a large role because $\underline{e}_{s_t, z_t}$ is close to zero until retirement. This is most obvious when looking at the net present value of guaranteed retirement earnings for newborn households: In steady state, this is 0.192; in a lifelong recession, it is only 0.105.

It is worth briefly interpreting this result. When default is eliminated, creditors extend any amount of debt at a risk-free rate. While creditors offer any amount, households avoid taking on debt beyond what they can repay in the worst case scenario. Absent aggregate risk, the worst case scenario is bad efficiency shocks forever. With aggregate risk, the possibility of a protracted recession makes the worst case scenario worse.

Aggregate risk's differential impact on credit and debt in FS and NFS changes how well insured households are against shocks. Insurance coefficients, introduced by Blundell, Pistaferri, and Preston (2008), are one useful way to measure household insurance. An insurance coefficient against shock ξ is defined by

$$\phi^\xi = 1 - \frac{\text{Cov}(\log(c_{it}/c_{i,t-1}), \xi_{it})}{\text{Var}(\xi_{it})}, \tag{24}$$

where $\{c_{it}\}$ and $\{\xi_{it}\}$ are simulated panel data, with i denoting a household and t denoting time. It measures how much consumption responds to a ξ shock, with $\phi^\xi = 1$ meaning household consumption does not respond and $\phi^\xi = 0$ meaning consumption responds one-for-one.[23] More precisely, the term $1 - \phi^\xi$ would be the coefficient on ξ in a regression of ξ and a constant on log consumption growth. To make ϕ^x similar to the earnings shock coefficients, I report it as the coefficient for a normalized shock $-x_{it}/(w_t e_{it})$.

Table 4 provides the insurance coefficients for both idiosyncratic and aggregate shocks under FS and NFS.[24] For the efficiency shocks, η, ε, and v, NFS insures better than FS, which is consistent with what the bankruptcy literature has found (especially Athreya, Tam, and Young (2009b)). However, aggregate risk decreases the insurance coefficients for NFS and weakly increases them for FS. The NFS coefficients

---

[22]The term π_0, the probability of $x = 0$ occurring, appears as long as $\min\{x | x > 0\}$ is large relative to the lowest efficiency levels (as it is in the calibration).

[23]When ξ is an innovation of a log-efficiency process, it measures how growth in current earnings translates into growth in current consumption. So, it makes sense to compare the magnitudes across the η, ε, and v shocks. For expenditure shocks and the level (rather than the innovation) of TFP shocks, the coefficients are not directly comparable as they do not change a household's cash-at-hand in the same way. However, for any shock, insurance coefficients are comparable across bankruptcy regimes.

[24]The sample is restricted to working households. For the persistent shock η, the transitory shock ε, and the right-tail process shock v, the sample is also restricted to households that had the same efficiency process in the preceding period.

TABLE 4. Insurance coefficients and welfare by bankruptcy regime.

| | Insurance Coefficient | | | | | Welfare Gain (%) | |
| | | | | | | Relative to | |
	$\phi^{\log(z)}$ TFP	ϕ^{η} Pers.	ϕ^{ε} Trans.	ϕ^{υ} Top 20	ϕ^{x} Exp.	FS SS	FS BC
Steady State (SS)							
FS	–	0.548	0.780	0.761	0.933	0.00	0.20
NFS	–	0.581	0.797	0.773	0.931	0.21	0.41
BAPCPA	–	0.542	0.813	0.757	0.904	−0.11	0.09
Business Cycle (BC)							
FS	0.814	0.548	0.782	0.761	0.933	−0.20	0.00
NFS	0.821	0.577	0.794	0.773	0.931	−0.33	−0.13
BAPCPA	0.796	0.542	0.814	0.757	0.905	−0.35	−0.15
FS flex	1.053	0.553	0.781	0.761	0.933	−0.00	0.20
NFS flex	1.064	0.580	0.794	0.773	0.931	−0.13	0.08

worsen because credit worsens (as seen in Figure 1). Why the FS coefficients improve is not very clear, but is perhaps explained by aggregate risk inducing a slightly higher capital–output ratio and hence higher wages. Additionally, NFS insures households better against aggregate shocks. While this is investigated more in Section 5.1, the principal reason is that aggregate risk, as modeled, is primarily earnings risk: Countercyclical earnings variance, the labor supply shifter ψ_z, and TFP all cause fluctuations in earnings.

While NFS insures better against earnings risk, FS insures slightly better against expenditure shocks, and this is not affected by aggregate risk. FS does better for two reasons. First, when an expenditure shock hits, bankruptcy is often the only option. This can be seen in Figure 1: The expenditure shock magnitudes 0.762 and 2.37 are both larger than the average FS and NFS borrowing limits for young households. Second, conditional on defaulting, it is better to do so in FS where exclusion from credit markets is temporary rather than in NFS where it is permanent. Aggregate risk does not change this because its effect on credit markets is unimportant to bankrupts (who are excluded from credit markets in the short run).

Given that NFS does not always provide better insurance than FS and that the amount of insurance is affected by aggregate risk, the welfare effects of implementing NFS are far from clear. The picture becomes even more muddled once one considers that the insurance coefficients only measure changes in consumption and not levels. So, to assess welfare, I use the percent increase in consumption that a newborn household would need in every state of the FS economy to be indifferent between living in the FS economy and moving to the NFS economy (consumption equivalent variation). For $\sigma = 2$, this is

$$\frac{\int_{\mathcal{S}} \left(\sum_s \hat{F}(s|z) \int_e V_{\text{FS}}(0, e, s, 0; \mathcal{S}) \hat{f}(e|s, z) \, de \right) dG_{\text{FS}}(\mathcal{S})}{\int_{\mathcal{S}} \left(\sum_s \hat{F}(s|z) \int_e V_{\text{NFS}}(0, e, s, 0; \mathcal{S}) \hat{f}(e|s, z) \, de \right) dG_{\text{NFS}}(\mathcal{S})} - 1, \tag{25}$$

where G_X is the implied ergodic distribution over $\mathcal{S}$ induced by household policies and stochastic transitions (including z, which is part of $\mathcal{S}$) in economy X. While G could be a complicated object, its computational equivalent is straightforward to calculate (see Appendix B for details).

Table 4 reports the welfare numbers. Without aggregate risk, the welfare gain of implementing NFS is 0.21%. Once aggregate risk is added, this turns into a small welfare *loss* of 0.13%. Built into these welfare numbers are general equilibrium effects. In partial equilibrium, that is, assuming factor and Arrow security prices in NFS are the same as in FS, the welfare gains are 1.80% without aggregate risk and 1.31% with it. General equilibrium effects make NFS look worse because it is associated with lower wages. It should be noted that while the levels of the welfare gains are not robust to different modeling assumptions, aggregate risk's effect on them is.[25]

4.2 *BAPCPA*

BAPCPA, the Bankruptcy Abuse and Consumer Protection Act of 2005, is perhaps the most important recent example of restricting default. Among many changes (see White (2007), for a comprehensive list), the reform has made it typically impossible for households with above-median income in their state to file for Chapter 7 bankruptcy. At the same time, BAPCPA did not eliminate debt forgiveness for these households because they may still file for Chapter 13. In contrast to the fresh start offered by Chapter 7, Chapter 13 requires households to forfeit future income for 3–5 years.[26] BAPCPA is mapped into the model through two assumptions: First, households with efficiency less than $\psi_z \tilde{e}$, where $\tilde{e}$ is the median efficiency in steady state, may default as before; second, households with efficiency above $\psi_z \tilde{e}$ may forfeit a fraction χ_{13} of their earnings to creditors in the period of default to obtain a discharge. I assume that when χ_{13} is paid, this replaces the default cost $\chi(e)$ in the period of filing.[27] I set $\chi_{13} = 0.75$ to loosely represent a 15% annual contribution of earnings over the course of 5 years.[28] While there are much better ways to model Chapter 13 bankruptcy (a good example being Li and Sarte (2006)), this formulation captures the quintessential feature of Chapter 13, that debt is forgiven in exchange for income, without overcomplicating the model.

Table 3 reveals that BAPCPA induces a substantial long-run increase in debt with or without aggregate risk. In fact, the debt–output ratio increases more under BAPCPA

---

[25]The levels can be either large and positive or large and negative. For instance, Appendix C shows that eliminating expenditure shocks leads to a 4.22% welfare gain of implementing NFS without aggregate risk and a 3.05% gain with it. Appendix C also shows that recalibrating the model to match higher debt targets results in a welfare gain of -4.76% without aggregate risk and -5.12% with it.

[26]Households may keep their assets in a Chapter 13 discharge. However, Bermant and Flynn (1999) report that out of 975,370 Chapter 7 filings ending in 1998, only around 10,000 had assets that were eventually distributed to creditors. Hence the number of households with assets that would be pushed into a Chapter 13 filing is likely very small.

[27]These assumptions change the repayment rates in (13). Now $p(a', s, h; \mathcal{S}'_{z'})$ equals $\rho_s \mathbb{E}(1 - d(a', e', s', h'; \mathcal{S}'_{z'}) + \mathbf{1}_{e' > \psi_{z'} \tilde{e}} d(a', e', s', h'; \mathcal{S}'_{z'}) \chi_{13} w(\mathcal{S}'_{z'}) e'/(-a'))$.

[28]In the law, households in this situation would have to contribute 100% of "disposable income'— income exceeding allowances defined by the IRS—for 5 years. For a two person household, these allowances can exceed $17,500.

than NFS once aggregate risk is taken into account. At the same time, the filing rate decreases and BAPCPA almost completely eliminates bankruptcy filings by above-median households.[29]

A natural question is whether these findings are consistent with the post-BAPCPA experience in the United States. Unfortunately, this is a difficult question to answer because these are long-run predictions and the 8 years of post-BAPCPA data include a severe recession. The first issue can be addressed to some extent by comparing the steady-state to steady-state transition path induced by BAPCPA (absent aggregate risk) with the data. This is done in Figure 2. As BAPCPA was partially anticipated, I assume BAPCPA is announced 6 years in advance, but that the announcement itself is a surprise.[30] The debt series in the model and the data are initially far apart because negative net worth data are not available annually (and so the data's debt series is revolving consumer credit rather than net worth). The data and model both show an increase in debt from 2005 to
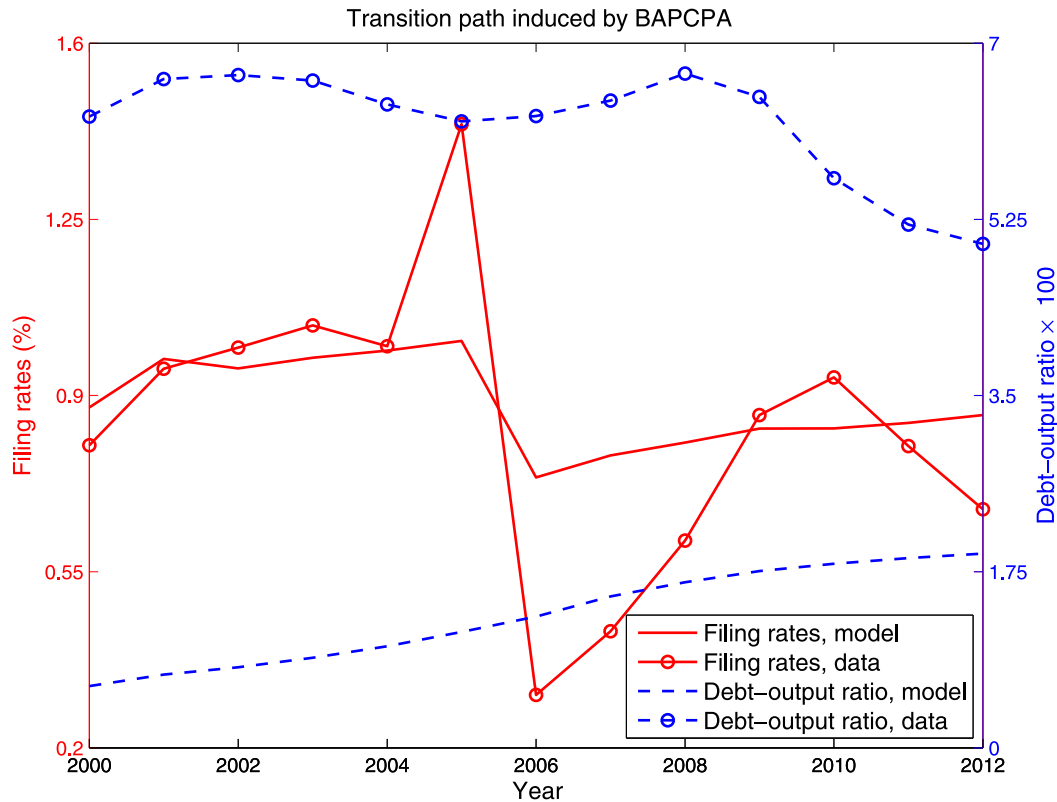

Figure 2. The transition path induced by BAPCPA.

[29]A previous version of the paper, which did not have expenditure shocks, predicted that filing rates would increase in the long run. Expenditure shocks change this conclusion because, as will be discussed, BAPCPA does a very poor job of insuring above-median households against these shocks. This poor insurance induces precautionary savings that lead to a higher average asset position and reduces filing rates.

[30]BAPCPA was, of course, not perfectly anticipated. However, BAPCPA grew out of several pieces of legislation introduced in 1998 and 1999. In fact, these pieces of legislation were analyzed in Athreya (2002).

2008, but then diverge with debt in the data falling sharply. While the model misses most of the "rush to file" in 2005, it captures some of the subsequent decline. The model's predicted filing rates from 2008 to 2012 look plausible, but may be a bit high.

The discrepancies between the data and the model's short-run predictions can be partly explained by aggregate risk and the characteristics of the last recession. For instance, debt is procyclical, and the recession would have depressed debt accumulation.[31] Since the model's standard deviation of debt is only 3.6%, this effect is likely not strong enough to bring the series in line. However, if there was a substantial "credit crunch" in the last recession, in the sense of a wedge between default rates and debt prices, this could both decrease debt and decrease filing rates. There is anecdotal evidence for such a credit crunch: Ben Bernanke revealed that he was recently unsuccessful in an attempt to refinance his mortgage (Campbell and Woellert (2014)). BAPCPA's long-run prediction of an increase in debt seems reasonable: Higher default costs should lead to more credit availability. That this has not occurred yet provides some evidence that indeed there has been a credit crunch.

At any rate, taking the model's predictions at face value, it is important to understand why BAPCPA increases debt and why, in contrast to NFS, this is not affected much by aggregate risk. The proximal reason is BAPCPA's borrowing limits, which are plotted in Figure 1. The average borrowing limits are uniformly larger in BAPCPA than in FS and are much larger for middle-aged households. Borrowing limits are larger for middle-aged households because BAPCPA effectively makes default costs a function of age: Since average earnings are hump-shaped over the life cycle, middle-aged households are far more likely to be above-median; consequently, their cost of default increases, making them less likely to default (conditional on debt) and expanding borrowing limits. Additionally, because BAPCPA's borrowing limits are based on expected default rates rather than natural borrowing limits, aggregate risk has a muted effect on BAPCPA as it did on FS.

In terms of welfare, the model suggests BAPCPA is worse than both FS and NFS. In particular, it produces a welfare loss relative to FS of 0.11% without aggregate risk and a loss of 0.15% with it. The reason BAPCPA lowers welfare is revealed by the insurance coefficients in Table 4: BAPCPA has lower insurance than FS and NFS against every shock except transitory shocks.

BAPCPA insures well against transitory earnings shocks and poorly against persistent ones because of its defining feature, the means test. Consider an above-median household that receives a negative earnings shock that pushes it below median. If the shock is transitory, then the household will typically be above median next period and face a high default cost, which translates into good credit prices in the current period. If the shock is persistent, then the household will typically be *below* median next period and face a low default cost, which translates into poor credit prices. Since a negative earnings shock is precisely when a household would like to borrow, this endogenous tightening of credit comes at precisely the wrong time.

---

[31]BAPCA has cyclical properties similar to FS except that discharged debt and interest rates both become strongly procyclical (in the long run).

TABLE 5. Welfare gains of making default more costly.

Statistic	$\chi_0 =$					
	0.00	0.25	0.50	0.75	0.90	0.99
Welfare gain relative to FS, SS	0.58	2.33	3.32	3.03	1.72	−3.65
Welfare gain relative to FS, BC	0.62	1.86	2.80	2.60	0.89	−7.64
Capital–output ratio, SS	3.15	3.07	3.04	3.05	3.08	3.21
Capital–output ratio, BC	3.15	3.09	3.07	3.07	3.10	3.24
Debt–output ratio $\times$ 100, SS	0.3	6.7	11.9	12.5	9.9	2.5
Debt–output ratio $\times$ 100, BC	0.3	5.4	9.5	10.6	8.4	1.8
Population in debt (%), SS	1.2	25.8	30.9	30.9	27.7	12.8
Population in debt (%), BC	1.3	23.7	28.7	29.1	25.7	9.4
Population filing (%), SS	0.92	0.31	0.12	0.05	0.03	0.004
Population filing (%), BC	0.95	0.32	0.12	0.06	0.03	0.005

Note: SS means steady state and BC means business cycle.

BAPCPA also does a poor job of insuring against aggregate shocks. This result is not driven by TFP per se, but by the countercyclical earnings variance associated with it. The standard deviation of persistent shock innovations is 69% larger in recessions than in expansions. Because BAPCPA insures poorly against these particular shocks, consumption responds strongly when a recession occurs.

BAPCPA does particularly poorly at insuring against expenditure shocks, and this is because the means test does not consider a household's debt. When an above-median household, say with zero assets and earnings of 1, is hit with an expenditure shock, their debt–earnings ratio goes from zero to either 76% or 237%. Under FS, they could discharge this at little cost. Under BAPCPA, this cost is much larger. While the household may not find it optimal to file (given that Table 3 shows filing rates of above-median households are very low), rolling over this debt and eventually paying for it is expensive.[32]

4.3 *Making bankruptcy uniformly more costly*

The previous sections have considered two alternative bankruptcy regimes, NFS and BAPCPA. Aggregate risk induced a large change in NFS, an infinite default-cost regime (for households with $x = 0$), but almost no change in BAPCPA, and a high default-cost regime (for above-median households). This section explores aggregate risk's effects on many regimes differentiated by the bankruptcy costs $\chi(e) = \max\{0, \chi_0 - \chi_1 e^{-1}\}$. In particular, I consider six regimes corresponding to $\chi_0 \in \{0, 0.25, 0.5, 0.75, 0.9, 0.99\}$ and $\chi_1 = 0$.[33] Select statistics for these regimes, including welfare gains relative to FS, are presented in Table 5.

---

[32]Gordon (2014), who characterizes the optimal bankruptcy rule in a similar environment, finds that a cutoff rule allowing bankruptcy for high debt–endowment ratios nearly implements the optimum. BAPCPA strays far from this rule by not accounting for debt levels.

[33]As in the benchmark, the costs are deadweight loss.

With the exception of the filing rate, all the statistics are nonmonotonic in χ_0. Welfare is one striking example, attaining a maximum for $\chi_0 = 0.50$ with or without aggregate risk, but the capital–output ratio is another, first declining in χ_0 but eventually reaching a maximum for $\chi_0 = 0.99$. These features can be explained by bankruptcy costs' effect on both borrowing limits and precautionary savings. For instance, as χ_0 moves from 0.00 to 0.75, the debt–output ratio steadily rises because credit markets expand and households willingly using credit. They willingly use credit knowing that if they are very unlucky, bankruptcy can be used at low cost. As χ_0 moves from 0.75 to 0.99, credit markets continue to expand. However, bankruptcy's increasing cost induces households to use it only as a last resort, resulting in households deleveraging. In the extreme of $\chi_0 = 0.99$, household savings drastically increase, resulting in a large capital–output ratio and little debt.

Aggregate risk reduces the welfare gain, debt–output ratio, and population in debt for each $\chi_0 > 0$, with the $\chi_0 \geq 0.9$ regimes experiencing the largest declines in welfare. It does this by effectively inducing precautionary savings for high-cost regimes. To insure themselves against aggregate risk, households have two options: They can default and they can reduce their debt. When bankruptcy comes at high cost, households naturally favor the second option. NFS is just an extreme version of this where reducing debt is the *only* option. In this sense, NFS and high default-cost regimes are not very different.

5. Aggregate risk and insurance

The previous section examined how aggregate risk affects the consequences of bankruptcy reform. This section attempts to understand who bears aggregate risk and to what extent bankruptcy policy can insure against it.

5.1 *Who bears aggregate risk?*

Figure 3 presents the insurance coefficients $\phi^{\log z}$ by age, household characteristics, bankruptcy regime, and portfolio restrictions. Only FS and NFS are considered as these regimes are very different and so, in some sense, bound what bankruptcy policy might accomplish. With the benchmark portfolio restrictions, overall it is the young who bear aggregate risk: Insurance coefficients, conditioned only on age (the top left panel), steadily rise over the life cycle as households accumulate assets. Whether the bankruptcy regime is FS or NFS seems to have little effect. Conditioning on households with the log process (19) reveals a similar pattern (seen in the bottom left panel). These households, which are restricted to hold a bond, are exposed to aggregate risk when young.

The picture looks substantially different once conditioned on households with the right-tail process (19) (as seen in the middle left panel). In particular, now young households are the *most* insured against aggregate risk. Because these households have access to flexible portfolios, the Arrow security prices endogenously determine which of them
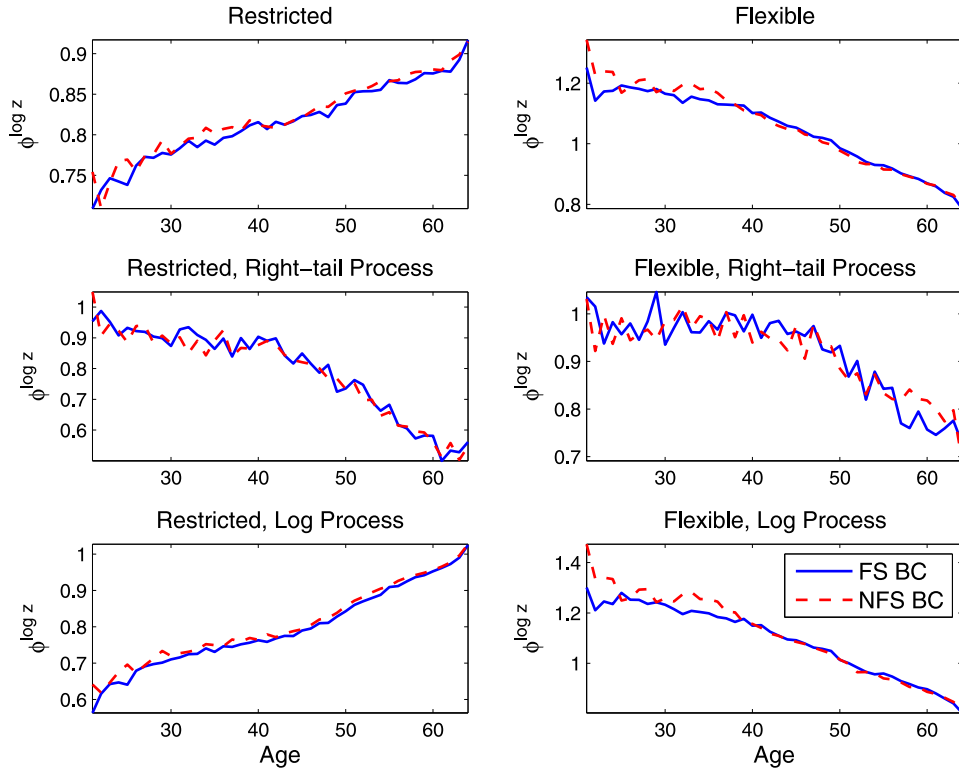
FIGURE 3. Aggregate shock insurance coefficients: Flexible versus benchmark.

are most exposed to aggregate risk. Absent default and mortality risk, the equilibrium condition in (16) (which is implied by market clearing) requires

$$\int \left(a'_g - a'_b\right) d\mu^{\text{flex}} = \frac{(r(\mathcal{S}'_g) - r(\mathcal{S}'_b))K'}{\pi^{\text{flex}}}, \tag{26}$$

where μ^{flex} and π^{flex} are the conditional distribution and the measure of households with no portfolio restrictions, respectively. Hence, right-tail households in the aggregate must hold a leveraged position in capital. However, some of the right-tail households can insure themselves by having $a'_b > a'_g$. The young, who are poorer on average, are the ones who ultimately do this. Since almost all right-tail households are savers, this result does not depend on the bankruptcy regime in place.

5.2 *Insuring by removing portfolio restrictions*

With the benchmark's portfolio restrictions, overall it is the young who bear aggregate risk. While NFS insures aggregate risk marginally better than FS does (at least conditioned on log-process households), bankruptcy policy does not seem capable of overturning this general pattern. However, what can overturn this pattern is the lifting of portfolio restrictions.

Flexible portfolios do an excellent job of insuring households against aggregate shocks. As can be seen in the top right panel of Figure 3, now it is the youngest households that are most insured. In fact, they have insurance coefficients greater than 1, which implies their consumption growth is larger in *recessions*. Even when not conditioned on age, the flexible portfolio versions of FS and NFS produce aggregate insurance coefficients of 1.053 and 1.064, respectively.[34]

Given that NFS, with its expanded credit markets, should let households take greater advantage of flexible portfolios than FS, it may be surprising that FS with flexible portfolios generates a 0.20% welfare gain (relative to the benchmark), while NFS with flexible portfolios only generates a 0.08% gain. However, flexible portfolios can overcome some of the deficiencies in FS credit markets: In a recession, households have less earnings and benefit from borrowing (which would tend to favor NFS), but, anticipating that, households can bring less debt into a recession. Interestingly, the welfare cost of business cycles is 0.00%, that is, zero, for FS when portfolios are flexible (as can be seen in Table 4). Intuitively, this is because the business cycle opens up an additional channel by which rich and poor households can share risk.

6. Conclusion

How does aggregate risk affect the trade-off between debt forgiveness and credit? Aggregate risk reduces the debt and welfare associated with high default-cost regimes when that cost is uniform across households. This is not necessarily the case when default costs are state-contingent: BAPCPA induces a small welfare loss and large increase in debt with or without aggregate risk.

While aggregate risk reduces the welfare of high default-cost regimes, bankruptcy policy, by itself, has a limited impact on who bears aggregate risk. With portfolio restrictions, both FS and NFS have aggregate risk borne by young households. When portfolio restrictions are lifted, aggregate risk is instead borne by the old.

References

Aiyagari, S. R. (1994), "Uninsured idiosyncratic risk and aggregate savings." *Quarterly Journal of Economics*, 109 (3), 659–684. [809, 810]

Arellano, C. (2008), "Default risk and income fluctuations in emerging economies." *American Economic Review*, 98 (3), 690–712. [799]

Athreya, K., J. M. Sánchez, X. S. Tam, and E. R. Young (2014), "Labor market upheaval, default regulations, and consumer debt." Working Paper 2014-002C, Federal Reserve Bank of St. Louis. [797]

Athreya, K. B. (2002), "Welfare implications of the bankruptcy reform act of 1999." *Journal of Monetary Economics*, 49 (8), 1567–1595. [797, 815]

Athreya, K. B. (2008), "Default, insurance and debt over the life-cycle." *Journal of Monetary Economics*, 55 (4), 752–774. [797]

---

[34]While not conditioned on age, the sample is restricted to working-age households. This is done because ϕ^η, ϕ^ε, and ϕ^v are only well defined for working-age households.

Athreya, K. B., X. S. Tam, and E. R. Young (2009a), "Unsecured credit markets are not insurance markets." *Journal of Monetary Economics*, 56 (1), 83–103. [797, 805]

Athreya, K. B., X. S. Tam, and E. R. Young (2009b), "Are harsh penalties for default really better?" Working Paper 09-11, Federal Reserve Bank of Richmond. [797, 812]

Bermant, G. and E. Flynn (1999), "Incomes, debts, and repayment capacities of recently discharged Chapter 7 debtors." Available at http://www.justice.gov/ust/eo/public_affairs/articles/docs/ch7trends-01.htm. Accessed: October 27, 2014. [806, 807, 814]

Blundell, R., L. Pistaferri, and I. Preston (2008), "Consumption inequality and partial insurance." *American Economic Review*, 98 (5), 1887–1921. [812]

Campbell, E. and L. Woellert (2014), "You know it's a tough market when Ben Bernanke can't refinance." Available at http://www.bloomberg.com/news/2014-10-02/you-know-it-s-a-tough-market-when-ben-bernanke-can-t-refinance.html. Accessed: October 10, 2014. [816]

Campbell, J. Y. (2006), "Household finance." *Journal of Finance*, 61 (4), 1553–1604. [806]

Carceles-Poveda, E. (2009), "Asset prices and business cycles under market incompleteness." *Review of Economic Dynamics*, 12, 405–422. [800]

Carceles-Poveda, E. and D. Coen-Pirani (2010), "Owning capital or being shareholders: An equivalence result with incomplete markets." *Review of Economic Dynamics*, 13, 537–558. [800]

Castañeda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull (1998), "Exploring the income distribution business cycle dynamics." *Journal of Monetary Economics*, 42 (1), 93–130. [804]

Castañeda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull (2003), "Accounting for the U.S. earnings and wealth inequality." *Journal of Political Economy*, 111 (4), 818–857. [804]

Chatterjee, S. and G. Gordon (2012), "Dealing with consumer default: Bankruptcy vs garnishment." *Journal of Monetary Economics*, 59, S1–S16. [797]

Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Ríos-Rull (2007), "A quantitative theory of unsecured consumer credit with risk of default." *Econometrica*, 75 (6), 1525–1589. [797, 799, 803, 805, 806]

Chien, Y., H. Cole, and H. Lustig (2011), "A multiplier approach to understanding the macro implications of household finance." *Review of Economic Studies*, 78 (1), 199–234. [806]

Cooley, T. F., ed. (1995), *Frontiers of Business Cycle Research*. Princeton University Press, Princeton, NJ. [804]

Fernández-Villaverde, J. and D. Krueger (2007), "Consumption over the life cycle: Facts from the consumer expenditure survey data." *Review of Economics and Statistics*, 89 (3), 552–565. [804]

Fieldhouse, D., I. Livshits, and J. MacGee (2014), "Aggregate fluctuations, consumer credit, and bankruptcy." Report. [797]

Gordon, G. (2014), "Optimal bankruptcy code: A fresh start for some." CAEPR Working Paper 2014-002, Indiana University. [817]

Herkenhoff, K. (2013), "The impact of consumer credit access on unemployment." Report. [797]

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1994), "The importance of precautionary motives in explaining individual and aggregate saving." *Carnegie-Rochester Conference Series on Public Policy*, 40 (1), 59–125. [803, 804]

Karahan, F. and S. Ozkan (2011), "On the persistence of income shocks over the life cycle: Evidence and implications." PIER Working Paper 11-030, University of Pennsylvania. [804]

Kennickell, A. B. (2009), "Ponds and streams: Wealth and income in the U.S., 1989 to 2007." Working Paper 2009-13, Finance and Economics Discussion Series. [805]

Krusell, P., T. Mukoyama, and A. Şahin (2010), "Labour-market matching with precautionary savings and aggregate fluctuations." *Review of Economic Studies*, 77 (4), 1477–1507. [802]

Li, W. and P. Sarte (2006), "U.S. consumer bankruptcy choice: The importance of general equilibrium effects." *Journal of Monetary Economics*, 53 (3), 613–631. [797, 814]

Livshits, I., J. MacGee, and M. Tertilt (2007), "Consumer bankruptcy: A fresh start." *American Economic Review*, 97 (1), 402–418. [797, 804, 805, 806, 809]

Livshits, I., J. MacGee, and M. Tertilt (2010), "Accounting for the rise in consumer bankruptcies." *American Economic Journal: Macroeconomics*, 2 (2), 165–193. [796]

Mitman, K. (2011), "Macroeconomic effects of bankruptcy and foreclosure policies." PIER Working Paper 11-015, University of Pennsylvania. [797]

Nakajima, M. (2008), "Equilibrium default and temptation." Report. [797]

Nakajima, M. and J.-V. Ríos-Rull (2005), "Default and aggregate fluctuations in storage economies." In *Frontiers in Applied General Equilibrium Modeling*, Chapter 7, Cambridge University Press, Cambridge. [797]

Nakajima, M. and J.-V. Ríos-Rull (2010), "Bankruptcy filings and aggregate fluctuations in growth economies." Report. [797]

Robe, M. A., E.-M. Steiger, and P.-A. Michel (2006), "Penalties and optimality in financial contracts: Taking stock." Working Paper 2006-013, Humboldt-Universität zu Berlin. [795]

Storesletten, K., C. Telmer, and A. Yaron (2004), "Cyclical dynamics in idiosyncratic labor market risk." *Journal of Political Economy*, 112 (3), 695–717. [798, 804]

White, M. J. (2007), "Bankruptcy reform and credit cards." *Journal of Economic Perspectives*, 21 (4), 179–199. [814]