Borrowing Premia in Unsecured Credit Markets*

Kyle Dempsey†
The Ohio State University

Felicia Ionescu‡
Federal Reserve Board

Abstract

Administrative data reveal large spreads in the U.S. unsecured credit market that are in excess of and vary sharply with default risk. Incorporating the empirically correct incidence of these borrowing premia allows standard unsecured credit models to match key granular credit moments, including the distribution of credit balances by default risk. The incidence and response of premia to aggregate shocks influence cross-sectional patterns in and aggregate dynamics of borrowing and default. An extended model infers lending standards from observed shifts in borrowing premia; we find that standards tightened only for risky borrowers at the start of the recession in 2020.

JEL Codes: E21, E32, E44, E51, G12, G21, G22
Keywords: Bankruptcy, borrowing premia, consumer credit, business cycles

*The views expressed here are those of the authors and do not necessarily represent the views of the Federal Reserve Board of Governors or the Federal Reserve System. We thank: Andrew Castro, Mike Carter, and Quinn Danielson for outstanding research assistance. We thank conference participants at the Barcelona GSE Summer Forum, IMF Macro-Finance conference, Midwest Macro, SAET, SED, System Macro Conference, Virtual Macro Conference, and seminar participants at Bank of Canada, Bank of Lithuania, Bureau of Economic Analysis, Federal Reserve Board, the Federal Reserve Bank of Richmond, Ohio State, and the Office of Financial Research at the Treasury for helpful comments. We especially thank Dan Cao, Dean Corbae, John Driscoll, Lukasz Drozd, Kyle Herkenhoff, Aubhik Khan, Kurt Mitman, Makoto Nakajima, Kjetil Storesletten, Julia Thomas, and Christian Wolf.

†The Ohio State University, Department of Economics. 1945 North High Street, Columbus, OH 43210. Email: dempsey.164@osu.edu.

‡Federal Reserve Board of Governors. 1801 K St. NW Washington, D.C. 20006. Email: felicia.ionescu@frb.gov.

The U.S. unsecured credit market has experienced extraordinary transformation in recent decades, with credit balances and bankruptcies increasing substantially.[1] A growing literature has found that households, particularly those with low income and wealth, rely heavily on unsecured credit, and a rich body of quantitative research has made significant advancements in understanding this markets in environments with household heterogeneity.[2] A particular focus of this existing work has been on how loan size choices and the evolution of idiosyncratic and aggregate states map into borrower risk, i.e. probabilities of default. Such theories deliver a linear relationship between loan prices and default probabilities (yellow line, Figure 1) in which default risk premia account fully for the spread between the benchmark interest rate and the loan rate.

In stark contrast, using administrative data, we document that the slope of interest rate spreads on credit card loans with respect to default probability is much smaller than such models predict (black dashed line, Figure 1). This implies the presence of large borrowing premia – interest rate spreads *in excess of* those required to compensate for default risk – in the unsecured credit market. For example, we estimate that these premia averaged 11.3 percentage points (pp) in 2019 and vary significantly across borrowers. These premia point to key pricing frictions which are absent from existing quantitative models of unsecured credit. How do these borrowing premia shape borrowing, saving, and default behavior in the cross-section and through time?

To address this question, we embed empirically consistent borrowing premia into an otherwise standard heterogeneous agent model with defaultable debt. By incorporating these premia, the relationship between interest rates and default probabilities in our model matches its empirical counterpart exactly (blue line, Figure 1). This involves striking a balance in terms of the slope of the interest rate schedule: a model in which default risk alone drives spreads yields too high a slope (yellow line), but a model with no default risk premia and a pure fixed cost of borrowing delivers too low a slope (red line). We use our quantitative model to analyze the effects of the cross-sectional incidence and dynamic evolution of borrowing premia on credit outcomes.

Our analysis delivers four main takeaways. First, borrowing premia in the data: (i) are large on average; (ii) decline with borrower risk; (iii) cannot be explained by other borrower-, account-, or bank-specific factors; and (iv) increase in downturns, especially for high risk borrowers. Second, we demonstrate that incorporating these premia into our model allows it to match not only standard aggregate credit moments, but also an important set of granular moments, such as the distribution of loan balances across levels of borrower risk and income. An alternative

[1] See, for example, Livshits et al. (2010).

[2] For example, Gross and Souleles (2002) and Sullivan (2008) find that unsecured credit is used to smooth consumption: households whose income declines accumulate more debt and declare bankruptcy more often. Furthermore, Herkenhoff et al. (2016) find that self-employment increases with credit scores and limits. After bankruptcy, individuals are more likely to start a new business and borrow. For quantitative work in stationary settings, see, for example, Athreya (2002), Athreya et al. (2009), Livshits et al. (2007, 2010), Chatterjee et al. (2007), Chatterjee et al. (2020). For work in non-stationary settings, see Nakajima and Ríos-Rull (2019).

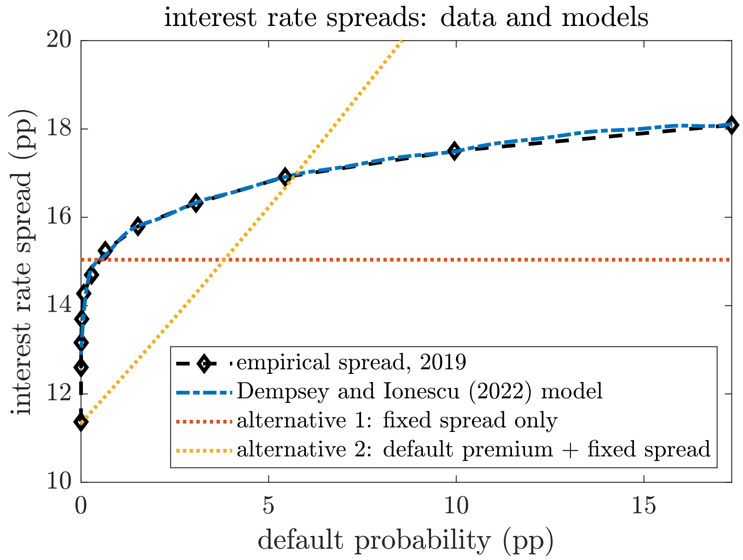

Figure 1: **Interest rate spreads in the data and alternative models**

Notes: The black line uses the Y-14M / Equifax sample described in Section 1.2. The blue line is our model's schedule of spreads. The yellow line uses the workhorse framework described in Section 1.1, and the red line imposes a fixed interest rate spread that does not depend on default risk.

model which captures the average *level* but not the cross-sectional *incidence* of these borrowing premia performs worse with respect to these disaggregated features, despite matching aggregate moments just as well by construction. Third, we show that both the cross-sectional incidence of borrowing premia and their response to aggregate shocks play a crucial role in shaping aggregate credit market dynamics. For example, in response to a negative aggregate shock, we show that the tightening of borrowing premia induces a more pronounced drop in total borrowing, while the declining relationship between borrowing premia and default risk promotes borrowing among riskier agents. Fourth, we show that an extension of our model allows us to use observed shifts in borrowing premia to infer a measure of lending standards. This provides a more concrete, quantitative interpretation of survey evidence that is widely used to proxy credit supply. Using this approach, we find that standards were essentially unchanged for all but the highest risk borrowers at the outset of the Covid-19 recession in 2020:Q2. We proceed as follows.

In Section 1 we combine two large administrative data sets – Y-14M and the FRBNY Consumer Credit Panel (Equifax) – to document the relatively flat relationship between interest rate spreads and default probability in Figure 1. Using a "wedge" measurement approach adapted to workhorse unsecured credit pricing theories,[3] we use these spreads to quantify borrowing premia which are large on average – far exceeding pure default risk premia – and decline with borrower risk. We estimate a balance-weighted average borrowing premium of 11.3 pp, with a range from 14.5 pp for FICO scores near 800 to 0.5 pp for FICO scores near 620. This declining pattern is puzzling

[3]This wedge approach was developed in the context of business cycle accounting in Chari et al. (2007).

because it implies a relatively small difference between credit prices faced by low and high risk borrowers. In this paper, we do not solve this puzzle in the sense of determining why these premia arise; rather, we document it empirically and measure its impact on borrower behavior.[4] Notably, though, we establish that systematic differences by credit score in loan recovery rates – a common free parameter in the literature – cannot explain our results. We show further that this pattern is robust to controlling for income and other borrower or account characteristics, allowing for bank-specific effects, and comparing "normal" and "crisis" periods.

In Section 2 we augment a standard quantitative model of unsecured credit to include borrowing premia in a flexible way. Our model economy closely resembles Chatterjee et al. (2007). In this workhorse model, borrowers' idiosyncratic states and loan choices determine repayment probabilities, which in turn pin down loan prices via simple discounting. In our model, states and loan choices determine repayment probabilities in the exact same way, but the mapping from repayment probabilities into loan prices is changed to reflect a given schedule of borrowing premia.

We consider two versions of this model. The first captures the heterogeneous incidence of borrowing premia estimated in Section 1 exactly; we call this our baseline or "heterogeneous premia" (HP) economy. The resulting pricing system implies higher costs for low risk loans and lower costs for high risk loans than in Chatterjee et al. (2007) because borrowing premia decline as default risk premia increase, leading to a relatively flat profile of interest rates. The second version is a "fixed premium" (FP) economy which matches the average level of borrowing premia but not their incidence; all loans bear the same premium in excess of default risk. This economy features one-for-one drops in loan discount prices (increases in loan rates) with default probabilities, yielding the familiar steep profile of interest rates from the literature.

In Section 3 we calibrate both versions of the model to a standard set of aggregate unsecured credit moments. Comparing these two models under the same parameterization is illustrative, but we calibrate each version to emphasize how the HP model adds additional explanatory power even conditional on matching key aggregates. Both calibrated economies have the exact same average interest rate spread and borrowing premium by construction. This implies that the differences between these two economies stem from cross-sectional variation in credit prices, not differences in their overall level. To match U.S. credit data, the FP model requires an implausibly high degree of impatience, as well as lower default costs than the HP model does.

We compare key dis-aggregated, cross-sectional properties of these two model economies in Section 4. The primary insight from this analysis is that the HP model closely matches the empirical joint distribution of credit balances and default risk, whereas the FP model overstates the riskiness of the safest 75% of loans in the economy. Moreover, the HP model more closely captures the profile of average loan spreads and leverage by income quartile, especially outperforming

[4]We address this important question empirically in an ongoing companion paper, Castro et al. (2020).

the FP model for the highest and lowest earners. The mechanism underlying these findings is straightforward: by accurately capturing the marginal costs of borrowers' loan choices on loan prices, the HP model more accurately explains their actual loan choices.

We study aggregate dynamics in Section 5. To fix ideas, in Section 5.1 we simulate an impulse response to a negative aggregate shock. The relatively high costs of increasing leverage in the FP economy dampen agents' ability to smooth consumption, leading to a 0.7 pp larger drop in total debt, 4.8 pp larger rise in bankruptcies, and a 9.1% larger drop in aggregate consumption relative to the HP economy. We then study the effect of the upward shift in borrowing premia we observe on impact of a recession. Since borrowing premia increase most for high risk loans, and high risk loans comprise a higher share of total loans in a recession, credit usage declines 0.2 pp more and bankruptcies increase 2.6 pp more than in the case in which premia do not shift in response to the shock. Section 5.2 expands these analyses to study business cycles. We show that the HP model more closely matches key cyclical properties of the unsecured credit market than the FP model, and that accounting for cyclical shifts in borrowing premia explains the high volatility of interest rate spreads over the cycle.

We conclude our analysis in Section 6 by introducing a simple extension of our baseline model which endogenizes borrowing premia as the product of two components. The first component varies endogenously with overall credit market tightness, and the second exogenously proxies banks' pricing decisions via a weighting function. We consider this function a measure of lending standards, allowing us to relate difficult-to-quantify changes in "standards" to observed changes in prices, augmenting our understanding of this widely used proxy for credit supply. Specifically, we impose that the model matches observed shifts in borrowing premia in response to a (negative) aggregate shock, and use the model to control for the associated shift in aggregate credit demand, allowing us to infer changes in standards. Our analysis delivers two key results, both consistent with survey evidence. First, the increase in the level of borrowing premia is driven by aggregate credit market tightening, i.e. the increase in demand for credit relative to the supply of savings. Second, controlling for this first effect, lending standards only become tighter for the riskiest borrowers, those with credit scores below approximately 680.[5]

Related literature We contribute to two broad areas in the literature: quantitative studies of unsecured credit and empirical studies of loan supply frictions. Our contribution to the former is both empirical and methodological. To our knowledge we are the first to incorporate pricing data from Y-14M into a quantitative unsecured consumer credit model. Methodologically, it is useful to think of loan pricing in two stages: how do borrower states and choices map into default

[5]We apply our method to the Covid-19 shock in the second quarter of 2020 due to the short history of Y-14M data, but our approach is generalizable to any episode given data on borrowing premia.

probabilities, and how do default probabilities map into loan prices? Our paper focuses on the latter question and nests workhorse models with regards to the former.[6] Correctly capturing the prices consumers face is critical, and we help reconcile existing theory with data on this front.[7]

Several recent papers have examined the role of specific pricing and market features to explain key trends in credit markets. For example, Raveendranathan (2020), Herkenhoff and Raveendranathan (2020), and Greenwald et al. (2020) all depart from the standard pricing paradigm and model credit contracts explicitly as long-term lines of credit. The former two papers yield pricing relationships that are relatively flat with respect to default risk, consistent with our findings.[8] Another example is Drozd and Kowalik (2019), which uses account-level supervisory data to study promotional pricing, linking credit card deleveraging to the collapse in promotional activity.

Our work also provides further insights into recurring questions in the consumer credit literature. Since the pricing relationships in our model induce persistently larger, riskier debts for borrowers, we offer an additional channel explaining the persistence of financial distress (Athreya et al. (2018)). In the spirit of Nakajima and Ríos-Rull (2019), we find that cyclical properties of interest rate spreads shape credit outcomes over the cycle, but we introduce the notion that not all spread variation is attributable to default risk. Symmetrically, we complement recent studies on how credit access shapes individual outcomes over the cycle (e.g. Herkenhoff (2019)) by demonstrating how borrowing premia influence such access.

Lastly, our paper contributes to the empirical literature on loan supply frictions linked to lending standards (e.g. Bassett et al. (2014) at the bank level and Lown and Morgan (2006), Schreft and Owens (1991) in the aggregate). For example, in line with our model, Bassett et al. (2014) find that tightening of credit supply is associated with a substantial decline in the ability of households to borrow from banks.[9] We contribute to these studies by using our model to link lending standards and borrowing premia and to assess the impacts of observed changes in lending standards on credit outcomes. We remain silent on the drivers of these changes beyond aggregate credit market tightness.[10]

[6] Examples of workhorse models in this literature include Athreya (2002), Livshits et al. (2007, 2010), Chatterjee et al. (2007), and Nakajima and Ríos-Rull (2019).

[7] Several studies have found similar pricing phenomena. For example, Agarwal et al. (2015) and Galenianos and Gavazza (2020) document related pricing patterns on credit cards in the US, with a focus on the effects of regulatory interventions, such as the 2009 CARD Act. Allen et al. (2014) and Allen et al. (2019) document and explore price dispersion in the Canadian mortgage market.

[8] Herkenhoff and Raveendranathan (2020) also documents extensive markups in the U.S. credit card industry. Our borrowing premia are consistent with such markups, and we measure how they vary across borrowers. Greenwald et al. (2020) uses Y-14Q data (the counterpart of Y-14M for business loans) to show that credit lines reproduce the observed flow of credit toward less constrained firms after adverse shocks.

[9] In related research, Chen et al. (2021) find that uncertainty in macroeconomic outlook rather than banks' balance sheet positions led banks to tighten standards during the financial crisis. They build a DSGE model consistent with this fact, arguing that credit supply shocks drive cyclical movements of lending standards.

[10] Our empirical companion paper Castro et al. (2020) fills this gap by combining SLOOS data with bank- and

1 Measuring Borrowing Premia

In this section, we document the relationship between loan rates and borrower risk in the U.S. credit card market. We describe our borrowing premium measure in Section 1.1 and our data in Section 1.2. Section 1.3 reports our main empirical results.

1.1 Measurement approach

Standard unsecured credit pricing To fix ideas, we highlight the key elements of workhorse unsecured loan pricing models.[11] Consider an economy with competitive lenders that offer a variety of contracts to households. Lenders can borrow at the equilibrium interest rate $i(s)$, where s is the aggregate state. A loan contract specifies a size $\ell < 0$ and a discount price q; the household pays the lender $q\ell$ today in order to receive ℓ tomorrow. Lenders choose how many contracts of size ℓ to issue to households with individual state x in aggregate state s. Next period, households may choose to repay the loan or default. In default, the lender recovers a fraction $\xi \in [0,1]$ of the principal ℓ. This canonical framework delivers the loan pricing equation

$$q(\ell;x,s) = \frac{p(\ell;x,s) + \xi(1-p(\ell;x,s))}{1+i(s)}, \text{ where } p(\ell;x,s) = \mathbb{P}\left(\text{repay } \ell \text{ tomorrow}|x,s\right) \quad (1)$$

is the expected probability of repayment. Expected repayment probabilities fully determine loan prices, which are linear in p with slope governed by ξ and i, and low risk (high p) loans have lower interest rates (higher prices). In light of this, much of the literature focuses on how loan size choices and the evolution of idiosyncratic and aggregate states drive default decisions, assuming that prices are determined according to (1).

Comparing model-implied and empirical loan rate spreads How, though, do default probabilities $(1-p)$ map into loan prices *in the data*? The literature has been silent on this front. We attempt to fill this gap using a simple approach motivated by the workhorse theory. In a standard unsecured credit model, the spread that a borrower pays over the equilibrium interest rate (i.e. the one implied by default risk only) is computed using equation (1):

$$\tilde{R}(\ell;x,s) = \frac{q^{-1}(\ell;x,s)}{1+i(s)} = \frac{1}{\xi + (1-\xi)p(\ell;x,s)} \implies \tilde{R}(p) = \frac{1}{\xi + (1-\xi)p}, \quad (2)$$

where $q^{-1}(\ell;x,s)$ is the gross interest rate on the loan. This representation makes clear that the spread in this class of models depends on individual characteristics only through the repayment

loan-level data from Y-14M and Call Reports to investigate why and how banks change standards.

[11] See, for example, Athreya (2002); Chatterjee et al. (2007); Livshits et al. (2007).

probability; that is, p is a "sufficient statistic" for borrower characteristics x, the choice of ℓ, and aggregate conditions s. For transparency in our measurement ahead of our data description, assume we can directly observe interest rate spreads and repayment probabilities (R_{it}, p_{it}) for each borrower i and period t.

We define the "borrowing premium" b_{it} as the percentage difference between the observed spread R_{it} and the theory-implied spread $\tilde{R}(p_{it})$ from equation (2):[12]

$$b_{it} = \frac{R_{it}}{\tilde{R}(p_{it})} - 1. \tag{3}$$

This measure captures all additional costs of borrowing *in excess of* those implied by measurable default risk.[13] While these costs likely arise at least in part from supply frictions, such as lender market power or constraints on loanable funds, we do not attempt to determine the origin of these additional costs in the present paper, but leave this topic for future work. Rather, we measure these costs and replicate them in a model to assess their impacts on individuals' behavior in the unsecured credit market.

1.2 Data and implementation

To provide evidence on how borrowing premia change with borrower risk, we combine two data sources: (i) Y-14M, a detailed account-level panel data set built from the portfolios of large bank holding companies in the United States, collected by the Federal Reserve Board as part of the Comprehensive Capital Analysis and Review; and (ii) the FRBNY Consumer Credit Panel (Equifax), a nationally representative five percent sample of all credit files for U.S. borrowers. Combining these data sources provides the most complete picture possible of how terms of credit, borrower characteristics, and borrowing and default behavior vary across the universe of borrowers in the U.S.[14] Specifically, this approach allows us to uncover the relationship between borrowing premia on credit card loans and borrowers' default probabilities. We now briefly describe the key

[12]We measure spreads in ratio terms (as opposed to simple differences) for ease of comparison with the extended model presented in Section 6. We describe our measurement of p in the next section. Our empirical approach could in principle be generalized to allow borrower choices, borrower states, and the aggregate state to matter separately from their impact on p.

[13]This "wedge" measurement is similar in approach and interpretation to that presented in Chari et al. (2007), and the notion of excess premium is consistent with Gilchrist and Zakrajsek (2012). The latter constructs a credit spread index based on two components, one that captures firm-specific information on expected defaults and a residual component (the excess bond premium), and uses it to predict economic activity and credit supply. Although different in scope, our study builds on the same principle of recognizing the importance of both risk and borrowing premia in movements in spreads on consumer loans.

[14]Focusing on terms of credit set by U.S. banks does not meaningfully limit our analysis. According to the FRB G.19 release, the bank share of total revolving unsecured consumer credit is about 90%. Within the banking sector, both the Y-14M sample) include the largest issuers of credit card loans and are representative of the U.S. unsecured credit market.

features of each data set and how we combine them to construct borrowing premia per equation (3), with additional details on sample and variable construction provided in Appendix A.1.

Y-14M Y-14M is a monthly, loan-level data set built from reporting by banks as part of their annual stress tests. This data is available starting in 2012 and covers the 35 largest U.S. banks. It includes detailed borrower characteristics and loan terms not available in other data sets used in the literature; in particular, it contains terms of credit (e.g. interest rates, credit limits) and measures of both the extensive and intensive margins of credit usage that can be dis-aggregated by borrower characteristics such as credit score and income. Y-14M's main shortcoming for our purposes is that it contains insufficient information on default behavior at the account level.

Equifax Consumer Credit Panel We supplement Y-14M with the Equifax Consumer Credit Panel (henceforth Equifax) because Y-14M does not contain reliable measures of default at the borrower level.[15] The data set is a nationally representative anonymous random sample from Equifax credit files, collected quarterly dating back to 1999. It contains rich descriptions of consumer credit behavior, including various measures of default and outstanding balances by loan type. Its main shortcoming is that it lacks information on income and terms of credit.

Combining the data sets To compute average borrowing premia by probability of default according to (3), we combine Y-14M and Equifax using measures of borrower credit scores in the two data sets as the common identifier. Our procedure does not and cannot match individual borrowers across the two data sets. Instead, we aggregate borrowers into bins and appeal to the huge samples in both data sets to match information at the bin level. We proceed in four steps.

First, we group borrowers in each data set by vigintiles (5% bins) based on their credit scores in 2019.[16] The credit score measure is the Equifax Risk Score (ranging from 280 to 850) in Equifax and the FICO score (ranging from 300 to 850) in Y-14M.[17]

Second, we compute the average likelihood of default for each of these 20 groups using Equifax. Our measure of default includes bankruptcy and severe derogatory, both of which affect access to credit and interest rates on credit card loans.[18] Importantly, both imply that the lender has removed the debt from its books and thus represents the best counterpart to the model

[15] While information about banks' assessment of each account's probability of default is collected in Y-14M, this is an optional reporting item for most banks (except those subject to the "advanced approaches" rule). As a result, the variable is quite sparse, with less than 20% of banks reporting it.

[16] Our results are robust to finer partitions, for example into 2% bins.

[17] FICO scores come from the Y-14M data at the account level.

[18] "Severe derogatory" refers to any delinquency paired with a repossession, foreclosure, or charge-off. As described in Appendix A.1, our results are robust to both broader and narrower definitions of default, but this definition most closely matches the default institution in our model in Section 2.

definition of default, as opposed to including delinquent loans, even with longer term past due.

Third, we compute average interest rates conditional on the median range of debt levels for each borrower risk group using Y-14M. We control for debt level in order to control for the effects of loan sizes on loan prices.[19] We use the statutory rate, as opposed to computing an effective rate from Y-14M, given that the retail APR represents a better measure for issuer pricing strategies. While levels between statutory and effective rates are slightly different, the patterns across borrower risk are the same.[20] We transform raw interest rates into spreads using an average prime rate (data analog of $i(s)$) of 5.28% for 2019.[21]

Fourth, we map borrowing premia from Y-14M to likelihood of default from Equifax. This mapping requires two underlying assumptions. First, the population of borrowers in the two data sets is the same. Given the dominance of the largest banks in the credit card market, Y-14M captures almost the entire universe of credit card loans and borrowers (as represented by Equifax), and so we view this assumption as reasonable. Second, Equifax's Risk Score and Y-14M's FICO score assess borrowers' likelihood of default in equivalent ways. There is ample evidence that this latter assumption holds as well.[22]

A note on recovery rates Our borrowing premia specification (3) assumes that the recovery rate ξ does not vary systematically with borrower characteristics. Our data does not allow us to observe recovery rates conditional on default at the borrower level. Using Call Report data, however, we document that there is no correlation between recovery and charge-off rates at the bank level. This suggests that lenders do not experience systematically different recovery rates from borrowers of different risk levels.[23] In addition, we use our measurement system (1) - (3) to run an experiment in which we assume a fixed borrowing premium ($b_{it} = \bar{b}_t$ for all i) to try to uncover a *schedule* of recovery rates that would rationalize observed interest rate spreads by borrower risk. As Appendix Figure A.5 shows, there is no such viable schedule of recovery rates. Details of these empirical and model-based tests of recovery rates are provided in Appendix A.3.

[19] This is an economically important step for consistency between our model and the data. Credit card contracts in the data take the form of credit lines, as opposed to a full loan-specific rate schedules. For excellent examples of quantitative models using the credit line paradigm, see Raveendranathan (2020) and Herkenhoff and Raveendranathan (2020).

[20] See, for example, the 2019 CFPB Report on The Consumer Credit Card Market. In fact, our finding that rates flatten quite a lot for the high probability of default which holds across effective and statutory rate measures, suggests that non-price terms seem to drive that portion of the market, justifying our censoring of the data below the 4th credit score vigintile. Of course, the interaction between price and non-price terms is important, but is it outside of the scope of this paper.

[21] This number is the evenly-weighted average of end-of-month prime rates for 2019.

[22] Given that multiple scoring models are used in determining applicant's creditworthiness, one might argue that the algorithm used to compute different credit scores might not produce identical measures for the same consumer. However, findings in a 2012 CFPB report reveal that the scores produced by different models provide similar information about consumers' relative creditworthiness.

[23] See Appendix Figure A.4. Section 3 provides further details on how we compute recovery rates in our data.

1.3 Empirical results

Borrowing premia by default risk Figure 1 documents the stark divergence between interest rate spreads observed in our combined Y-14M / Equifax data set and those implied by workhorse unsecured credit models. As default probability increases, loan rates increase in both settings, but this rise is far steeper in standard models than in the data and is fully driven by risk premia. Applying our measurement scheme to data from 2019 and the onset of the Covid-19 pandemic in 2020, panel (a) of Figure 2 confirms the presence of large borrowing premia in excess of these risk premia in the unsecured credit market. The solid lines show our bin-specific borrowing premia, and the dashed lines correspond to the balance-weighted average, $\bar{b}$.

Figure 2 contains four primary findings. First, there is a large wedge between the risk-free rate and the loan rate even for risk-free loans. Borrowing premia for loans with no default risk face an average borrowing premium of 11.4%. This figure reaches as high as 14.6% for the highest-premium bin, whose default probability is a mere 0.65%. This distorts credit usage along the *extensive* margin, since there is a discrete jump in interest rates from savings to even the safest loan. Second, borrowing premia are (statistically significantly) different across credit score bins and are non-monotone and non-linear in default probability. For all but the lowest risk bins, borrowing premia dissipate as risk premia increase. This distorts credit usage along the *intensive* margin since the pricing penalty for a higher risk loan is much smaller in the data than in standard models. As default risk premia rise, borrowing premia tend to decline to offset them. Third, this pattern is preserved at the onset of the recession in 2020, despite an upward shift in premia across the board. Premia increased most in proportional terms for high risk borrowers with this shift. Fourth, panel (b) highlights that these patterns are preserved across income quartiles. Small shifts in default probability correspond to much larger differences in borrowing premia than even large differences in income. Conditional on default risk bin, the differences in premia in the top three income quartiles relative to the bottom are below 1 percentage point (pp). In contrast, the difference between low risk and high risk borrowers is just under 14 pps.

Determinants of borrowing premia A crucial assumption in our measurement approach is that the benchmark spread (2) is summarized completely by default risk, which we map from credit scores. Therefore, variation in borrowing premia (3) must also be explainable by variation in credit score. Is this assumption valid? Our income-specific measures from panel (b) of Figure 2 address this question to some extent, suggesting that indeed default probability is the central determinant of borrowing premia for credit card loans. Here we conduct a more formal regression analysis to document the relative importance of default risk, income, and other observable factors in explaining variation in interest rate spreads and borrowing premia.

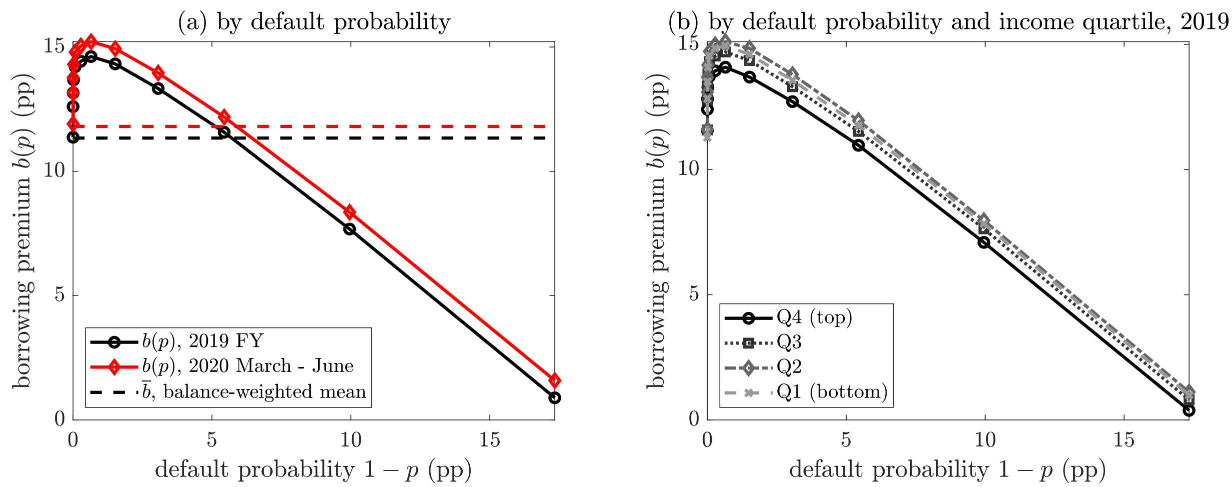

Figure 2: **Borrowing premia by default probability and income quartile**

Notes: This figure is constructed using the combined Equifax and Y-14M sample described in Section 1.2, with $b(p)$ defined in (3). The balance-weighted average premium is $\bar{b} = 11.3\%$. for 2019 and 11.8% for 2020.

Specifically, we estimate regressions of the following form via OLS:[24]

$$b_{it} = \alpha + \beta_1 \text{FICO}_{it} + \beta_2 \text{income}_{it} + \gamma_1 X_{it} + \gamma_2 Y_t + \gamma_3 Z_{j(i,t)} + \varepsilon_{it} \tag{4}$$

where: b_{it}, FICO_{it}, and income_{it} are the borrowing premium, credit score, and reported income, (all in logs) for account i in quarter t, respectively; X_{it} is a vector of other borrower or account characteristics; Y_t is a quarter fixed effect; $Z_{j(i,t)}$ is a bank fixed effect for the bank j at which account i is held at time t; and ε_{it} is an error term. Examples of characteristics included in X_{it} are: (i) whether the account is a revolver; (ii) whether the borrower holding the account has multiple relationships with the same lender; and (iii) whether the account is newly originated. Our primary objects of interest from these regression analyses are the coefficients β_1 and β_2, as well as the increase in explanatory power from including income, other account-specific characteristics X_{it}, and bank and time fixed effects.

Table 1 presents our estimates for a range of sub-specifications from equation (4), providing three main takeaways. First, FICO scores explain most of the variation in borrowing premia. Our simplest specification with FICO only and no controls (column [1]) explains 60.7% of the variation of borrowing premia. When adding income (column [2]), the explanatory power increases by only 0.1 pp. This finding is robust across alternative sets of controls, such as when additional account characteristics like revolver status and multiple relationships are included (columns [3] and [4]).[25] Additionally, the coefficients on FICO and income are all statistically significant at the

[24] Details on sample and variable construction are in Appendix A.

[25] Alternatively, one could imagine starting with income and adding in FICO; in that exercise, the explanatory power of the income-only analog of column [1] from Table 1 is only 0.7%.

	dependent variable: borrowing premium					
	[1]	[2]	[3]	[4]	[5]	[6]
FICO score	0.909	0.913	0.913	0.918	0.918	0.913
	(2e-4)	(2e-4)	(2e-4)	(2e-4)	(2e-4)	(2e-4)
income		-0.004		-0.005	-0.005	-0.006
		(2e-5)		(3e-5)	(3e-5)	(3e-5)
revolver			X	X	X	X
multiple relationships			X	X	X	X
new account			X	X	X	X
quarter FE					X	X
bank FE						X
R^2	0.607	0.608	0.609	0.610	0.610	0.621

Table 1: **Determinants of borrowing premia: regression analysis**

Notes: Number of observations is 14,426,760. All estimates are significant at 1 percent level. All specifications are variations on estimating equation (4). Column [1] imposes $\beta_2 = \gamma_1 = \gamma_2 = \gamma_3 = 0$; [2] imposes $\gamma_1 = \gamma_2 = \gamma_3 = 0$; [3] imposes $\beta_2 = \gamma_2 = \gamma_3 = 0$; [4] imposes $\gamma_2 = \gamma_3 = 0$; [5] imposes $\gamma_3 = 0$; and [6] estimates all coefficients. All continuous variables are in logs. Details of variable and sample construction, as well as alternative specifications, may be found in Appendix A.1.

1% level and have the expected signs.[26] Second, other borrower and account characteristics have little independent contribution to borrowing premia, as evidenced by the small gain in explanatory power in columns [3] and [4] relative to [1] and [2]. These first two findings suggest that borrower risk indeed encapsulates all the relevant observable borrower-level (i.e. idiosyncratic) information that determines borrowing premia in excess of risk premia. Third, macroeconomic and bank-specific (i.e. aggregate) characteristics – captured via quarter and bank fixed effects in columns [5] and [6] – add little explanatory power.

Robustness One cause for concern for our estimates might be the linearity of equation (4) in FICO score, given the shape of the borrowing premium schedule in Figure 2. An alternative quadratic formulation (estimated in Panel A of Appendix Table A.1), however, highlights that our findings are robust to such an alternative specification.

Another concern might be that the results in Table 1 are mechanical, since our binning procedure already uses FICO to construct our borrowing premia estimates. To assuage this concern, we repeat the analysis of equation (4) and Table 1 with APR (in spread terms, R_{it}) as

[26] The significance is unsurprising given the large sample size. A positive β_1 indicates that higher FICO scores (i.e. lower default probabilities) correspond to higher borrowing premia on average. A negative β_2 implies that higher income corresponds to a lower borrowing premium.

the dependent variable, rather than borrowing premium (b_{it}). Results for this analysis are in Panel B of Appendix Table A.1. Consistent with our baseline analysis, we find that borrower risk explains far more of the variation in APRs than any other borrower- or account-level characteristics, though the total amount of variation explained declines in this case. Notably, we find that bank fixed effects add appreciable explanatory power to these regressions.

Summary We find empirical evidence for borrowing premia in excess of pure default risk premia that are large on average and vary with borrower risk, tending to decline as borrower risk increases. Borrower risk is a sufficient statistic for these premia; other characteristics like income and the relationship between the borrower and the lender, as well as macroeconomic and bank-specific factors, are secondary. The model we build in Section 2 is designed to incorporate this evidence into existing theory, delivering empirically consistent terms for consumer credit.

2 A Model of Unsecured Credit with Borrowing Premia

The model adapts Chatterjee et al. (2007) to match empirically observed interest rates by probability of default. Time is discrete, and there are no aggregate shocks. There is a single good used for consumption and investment and a single asset available for saving and borrowing. All quantities are measured in real terms. There are three types of agents. First, there is a unit mass of infinitely lived households who decide each period how much to consume, how much to borrow or save, and whether or not to default on existing debt. Second, there is a representative firm which produces according to a constant returns to scale production function and hires capital and labor in competitive factor markets. Third, there is a representative, competitive lender who offers menus of borrowing and saving contracts to households.

2.1 Legal environment and market arrangements

Bankruptcy Following Chatterjee et al. (2007), the bankruptcy procedure in the model captures Chapter 7 bankruptcy in the U.S. A household's credit history is summarized by a "bankruptcy flag" $f \in \{0, 1\}$. A status of $f = 1$ ($f = 0$) indicates the presence (lack of) default in the household's recent credit history. A household without a bankruptcy flag has access to the credit market and may borrow or save, choosing wealth tomorrow $a' \lesseqgtr 0$; in contrast, a household with a bankruptcy flag cannot borrow and must choose $a' \geq 0$.

A household who enters the current period with debt must choose whether to file for bankruptcy (BK) or repay (R). In the event of a bankruptcy, the household: (i) cannot borrow or save in the current period ($a' = 0$); (ii) begins the next period with a bankruptcy flag ($f' = 1$); (iii) incurs a fixed pecuniary cost $\kappa > 0$; (iv) incurs a fixed non-pecuniary utility cost ("stigma") $\chi > 0$; and (v) pays back only a fraction $\xi \in [0,1]$ of its debt. A household with a bankruptcy flag loses the flag and regains access to the credit market with i.i.d. probability $\theta \in [0,1]$, yielding an average duration of bad credit standing of $1/\theta$. The institution of bankruptcy implies that competitive, profit-maximizing lenders must price loans today to reflect the probability of default tomorrow. In the case of default, lenders receive the fraction ξ of the loan that is repaid.[27]

Production and factor markets The single good is produced by the representative firm according to the constant returns production function $Y = F(K,N) = K^\alpha N^{1-\alpha}$, where: Y is aggregate output; K is aggregate capital; N is aggregate labor supply (in effective units); and $\alpha \in (0,1)$ is a parameter measuring the capital share in production. Since there are no government expenditures, all production must be channeled toward: aggregate consumption, C; aggregate investment, $I = K' - (1-\delta)K$, where $\delta \in [0,1]$ is the depreciation rate on capital; and aggregate default costs. The rental rate on capital r and the wage rate w are determined competitively, delivering standard factor prices $r = \alpha \left(\frac{K}{N}\right)^{\alpha-1}$ and $w = (1-\alpha)\left(\frac{K}{N}\right)^{\alpha}$. The equilibrium interest rate is $i = r - \delta$.

Lending and saving markets The competitive, representative lender offers a variety of contracts to households to maximize expected profits. A contract specifies a size ℓ and a discount price q, whereby the household pays the lender $q\ell$ today in order to receive ℓ tomorrow. The appropriate sign convention thus defines contracts with $\ell < 0$ as loans and those with $\ell > 0$ as savings. In the current period, the lender must choose a mass of contracts, $m'(\ell; x)$, of size ℓ to issue to households with state x. Lenders can borrow at the equilibrium risk-free interest rate i and are fully diversified against individual risk given the assumption of a continuum of households.

2.2 Households

Preferences Households are risk averse, value consumption flows according to the CRRA utility function $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, and supply their full labor productivity inelastically. They discount the future at rate $\beta \in [0,1)$, which evolves stochastically according to Γ^β. Furthermore, households receive a pair of additively separable i.i.d. shocks, $\nu = \{\nu^{BK}, \nu^R\}$, which are attached to the options to declare bankruptcy or repay and are drawn from a type one extreme value distribution with scale parameter ζ and location parameter $\overline{\nu}$. These shocks serve three purposes. First, they capture the fact that many defaults are associated not with income shocks, but with events such as marital disruptions and medical expenses which we do not model explicitly. Second,

[27] Including this partial recovery is important for consistency with the data, since empirically a bankrupt or severely derogatory account does not imply a zero return for the lender.

these shocks imply a positive probability of default on even the safest loans, consistent with data. Third, they smooth out individuals' repayment probability functions – and therefore loan price schedules –, which eases computation of the model.

Endowments and states Households begin the current period with net worth $a \in \mathbb{R}$. Each household has labor productivity $\epsilon_1\epsilon_2\epsilon_3$, where $\epsilon_1 \in \mathbb{R}_+$ is a permanent component or type, $\epsilon_2 \in \mathbb{R}_+$ is a persistent AR(1) component with persistence ρ_{ϵ_2} and innovation variance σ_{ϵ_2}, and $\epsilon_3 \in \mathbb{R}_+$ is a transitory component. We denote the transition process over labor productivity by Γ^{ϵ}.[28] The full idiosyncratic state of the household is $x = (a, \beta, f, \epsilon = (\epsilon_1, \epsilon_2, \epsilon_3), \nu)$, and the endogenous distribution of households over x is $\mu(x)$.

Decision problem A household in good credit standing ($f = 0$) who is in debt ($a < 0$) first decides whether or not to default, solving:

$$V_0(a, \beta, \epsilon, \nu) = \max\left\{V_0^{BK}(a, \beta, \epsilon) + \nu^{BK}, V_0^{R}(a, \beta, \epsilon) + \nu^{R}\right\}, \tag{5}$$

where V_0^{BK} and V_0^{R} are the values associated with bankruptcy and repaying, respectively:

$$V_0^{BK}(a, \beta, \epsilon) = u\left(\xi a + w\epsilon_1\epsilon_2\epsilon_3 - \kappa\right) - \chi + \beta\mathbb{E}\left[V_1(0, \beta', \epsilon')\right] \tag{6}$$

$$V_0^{R}(a, \beta, \epsilon) = \max_{a' \in \mathcal{F}_0(\beta, a, \epsilon)} u\left(c(a'; \beta, \epsilon)\right) + \beta\mathbb{E}\left[V_0(a', \beta', \epsilon', \nu')\right] \tag{7}$$

The bankruptcy value (6) reflects the fact that a defaulting household can neither borrow nor save, incurs pecuniary and non-pecuniary default costs, loses good credit standing, and must still pay back a fraction ξ of the loan. The value of not defaulting (7) reflects the fact that the household can either borrow or save, choosing a' from the set of feasible choices $\mathcal{F}_0(a, \beta, \epsilon) = \{a' \in \mathcal{A} : c(a'; a, \beta, \epsilon) \geq 0\}$, which accounts for equilibrium prices via the budget constraint:

$$c(a'; a, \beta, \epsilon) = \begin{cases} a + w\epsilon_1\epsilon_2\epsilon_3 - q(a'; \beta, \epsilon)a' & \text{if } a' < 0 \\ a + w\epsilon_1\epsilon_2\epsilon_3 - \overline{q}a' & \text{if } a' \geq 0 \end{cases} \tag{8}$$

The equilibrium price of a loan, $q(a'; a, \beta, \epsilon)$, depends on the household state and loan size, whereas the equilibrium price of a savings contract $\overline{q}$ is a scalar.

[28] This earnings process is tailored to accommodate the implementation of Storesletten et al. (2004), but of course it may be extended to other processes.

Since ν follows a type one extreme value distribution, the probability of a bankruptcy is

$$g_{BK}(a,\beta,\epsilon) \equiv \mathbb{P}\left(d = BK | a,\beta,\epsilon\right) = \begin{cases} 0 & \text{if } a \geq 0 \\ \left[1 + \exp\left\{\frac{V_0^R(a,\beta,\epsilon) - V_0^{BK}(\beta,\epsilon)}{\zeta}\right\}\right]^{-1} & \text{if } a < 0 \end{cases} \tag{9}$$

The same analysis shows that the ex-ante value of being in good credit standing is

$$\overline{V}_0(a,\beta,\epsilon) = \begin{cases} V_0^R(a,\beta,\epsilon) & \text{if } a \geq 0 \\ \zeta \log\left[\exp\left\{\frac{V_0^{BK}(\beta,\epsilon)}{\zeta}\right\} + \exp\left\{\frac{V_0^R(a,\beta,\epsilon)}{\zeta}\right\}\right] - \zeta\log(2) & \text{if } a < 0 \end{cases} \tag{10}$$

where the branching reflects that agents not in debt do not receive the extreme value shocks associated with the default decision.[29]

A household in bad credit standing ($f = 1$) can neither default nor borrow. Therefore, such a household decides only how much to save:

$$V_1(a,\beta,\epsilon) = \max_{a' \in \mathcal{F}_1(\beta,a,\epsilon)} u(c(a';a,\beta,\epsilon)) + \beta\mathbb{E}\left[(1-\theta)V_1(a',\beta',\epsilon') + \theta\overline{V}_0(a',\beta',\epsilon')\right] \tag{11}$$

The feasible set $\mathcal{F}_1(a,\beta,\epsilon) = \{a' \in \mathcal{A}_+ | c(a';a,\beta,\epsilon) \geq 0\}$, with $c(a';a,\beta,\epsilon)$ given by the bottom expression in (8), reflects that borrowing is forbidden. The optimal savings policy is $g_a(x)$.

2.3 Loan pricing

To determine the price of a loan, the lender must determine the probability that a loan of size ℓ to a borrower in state x will be repaid tomorrow. Taking into account decision rules and transitions over exogenous idiosyncratic states, this probability is

$$p(\ell;\beta,\epsilon) = \int_{\mathcal{B}\times\mathcal{E}} \left[1 - g_{BK}(\ell,\beta',\epsilon')\right]\Gamma^{\beta}(d\beta'|\beta)\Gamma^{\epsilon}(d\epsilon'|\epsilon). \tag{12}$$

In order to determine loan price schedules given the repayment probability (12), the lender must: (i) account for the recovery rate, ξ; (ii) discount at the equilibrium rate of return, i (the discount price of a riskless savings contract is $\overline{q} = (1+i)^{-1}$); and (iii) incorporate any borrowing premia.

[29]These results are standard in discrete choice. For recent applications to macroeconomic models with default, see Chatterjee et al. (2020); Dvorkin et al. (2019). It can be shown that for the case when $a < 0$

$$\overline{V}_0(a,\beta,\epsilon) = \int_{\mathcal{V}} V_0(a,\beta,\epsilon,\nu)dF(\nu) = \overline{\nu} + \zeta\gamma_E + \zeta\log\left[\exp\left\{\frac{V_0^D(\beta,\epsilon)}{\zeta}\right\} + \exp\left\{\frac{V_0^{ND}(a,\beta,\epsilon)}{\zeta}\right\}\right]$$

where γ_E is the Euler-Mascheroni constant and $\overline{\nu} = -\zeta(\gamma_E + \log 2)$ is the location parameter for $F(\nu)$. This normalization of $\overline{\nu}$ insures that, *a priori*, there is no additional value to being in debt associated with the option value of receiving the extreme value shocks.

In our baseline version of the model, these borrowing premia are modeled as exogenous wedges. Combining these three facets, the loan price schedule is

$$q(\ell;\beta,\epsilon)=\frac{p(\ell;\beta,\epsilon)+\xi(1-p(\ell;\beta,\epsilon))}{(1+i)\,(1+b(p(\ell;\beta,\epsilon)))} \tag{13}$$

The numerator is the expected repayment per unit of principal (adjusting for default risk and the recovery rate in the case of default), and the denominator reflects that borrowers pay an additional premium b above and beyond the base cost of funds adjusted for expected repayment. Inverting (13) and filtering out the interest rate confirms that this wedge maps exactly into our measurements from Section 1. Consistent with our empirical analysis, we assume that the borrowing premium depends on the size of the loan and the idiosyncratic state of the borrower only insofar as these objects affect the repayment probability p.

2.4 Equilibrium

A stationary recursive competitive equilibrium in this model is a list of: (i) value functions $V(x)$ for households, with associated default and savings policies $g_{BK}(x)$ and $g_a(x)$; (ii) factor prices r, w, and i; (iii) contract prices $\overline{q}$ and $q(\ell;x)$; (v) aggregate quantities N, K, and Y; and (vi) a distribution of households over idiosyncratic states $\mu(x)$ such that

1. **Households optimize:** The value function and associated optimal policies are consistent with the household problem represented by equations (5) through (11);
2. **Firms optimize:** r, w, and i are consistent with the firm's optimization problem;
3. **Lenders optimize:** $q(\ell;x)$ satisfies (13), and $\overline{q}=(1+i)^{-1}$;
4. **Consistency:** the distribution $\mu(x)$ is stationary and consistent with household decisions;
5. **Markets clear** for: (i) labor: $N=\int \epsilon_1\epsilon_2\epsilon_3\mu(dx)$; (ii) capital: $K=\int a\mu(dx)$; and (iii) contracts: $m'(\ell;x)=\int \mathbf{1}[\ell=g_a(x)]\mu(dx)$. (Goods clear via Walras' Law.)

2.5 An illustration of pricing and borrowing premia in the model

Panel (a) of Figure 3 plots the empirical borrowing premium schedule against default probability (estimated in Section 1), as well as the fit to that schedule which we use in our baseline model.[30] We call this baseline a heterogeneous premia model (HP) because the level of the borrowing premium varies with default risk. We compare our baseline HP model to two common versions in

[30] See Section 3.1 for details.

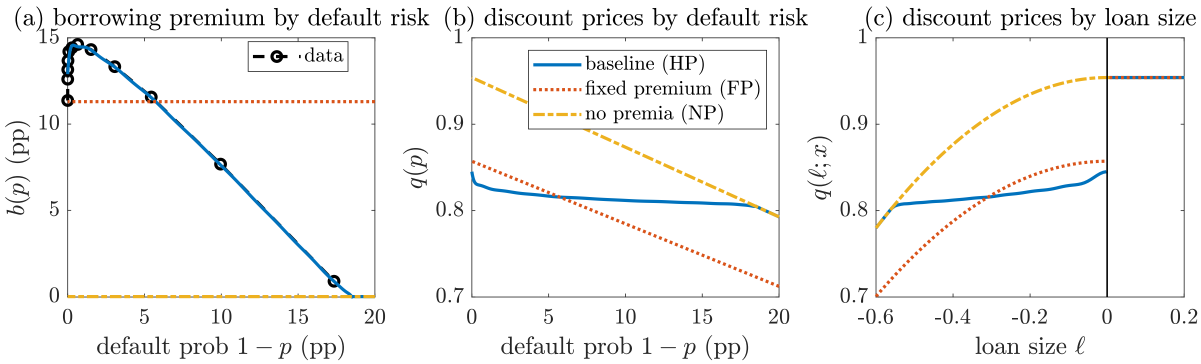

Figure 3: **An illustration of price schedules in the model**

the literature: a fixed borrowing premium model (FP), and no borrowing premium model (NP). Both the FP and NP models are special cases with $b(p) = \bar{b}$ for all p: the fixed premium case has $\bar{b} > 0$, while the no premia case imposes $\bar{b} = 0$. Both the HP and FP models feature more expensive credit for all borrowers than the NP model. Low risk borrowers face higher borrowing premia in the HP economy than in the FP economy, but high risk borrowers face lower premia.

Of course, no borrower cares solely about borrowing premia; rather, she cares about the all-in cost of a loan. Therefore, panel (b) of Figure 3 plots the loan price schedule against default risk, $q(p)$, which includes both risk premia and borrowing premia. The conclusion is straightforward: borrowers face lower marginal penalties for taking on high risk loans in our baseline HP model than either alternative. The HP price schedule is nearly flat in default risk, while the cost of a loan increases sharply as risk increases in both the FP and NP economies.

Panel (c) presents panel (b) using the typical convention from the literature, charting a specific individual's price schedule $q(\ell; x)$. This panel illustrates the effect of leverage choices alone on prices by fixing the individual's state x as well as the repayment probability function $p(\ell; x)$ across all three model variants. We highlight two main features. First, both the HP and FP economies feature large wedges between the return on savings and the rate on a riskless loan. Second, our baseline HP model features a price schedule that closely – but not exactly – approximates linear pricing, with interest rates rising only slightly as leverage and default risk increase from zero. By contrast, both alternative models feature sharp increases in interest rates as leverage increases. This property is critical to our quantitative analysis.

3 Mapping Models to Data

This section describes our calibration approach. We proceed in three phases. First, we directly estimate the schedule of borrowing premia over default risk from our Y-14M / Equifax data. Second, we assign a set of standard parameters reflecting with technology, the legal environment, and idiosyncratic labor productivity shocks. Third, we calibrate the remaining parameters via simulated method of moments (SMM) to match a set of credit market moments widely targeted in the literature. Notably, in this third phase, we calibrate *both* our baseline model with heterogeneous borrowing premia (HP) *and* an alternative with a fixed borrowing premium (FP). These two calibrated models allow us to demonstrate how the empirical incidence of borrowing premia shapes credit market outcomes. We re-calibrate the FP model – rather than simply apply the parameters from our baseline – so that the differences we highlight cannot be explained by aggregate differences between the credit markets in the two economies.

3.1 Borrowing premia estimation

The goal of our borrowing premia estimation is to match the empirical schedule as closely as possible since borrowing premia are exogenous in our baseline model. To construct the borrowing premium schedule for our baseline HP model, we seek a smooth function defined for all p (interpolated between the measured bins). To achieve this, we fit a 20th-order polynomial to the schedule for FY2019 in Figure 2. A high order polynomial is required to match the sharp curvature in the schedule at very low levels of default risk. The resulting premium schedule matches the empirical one almost exactly. For details on fit and parameter estimates, see Figure B.1 and Table B.1 in Appendix B.1.

3.2 Parameters assigned outside the model

Technology Panel A of Table 2 describes the key parameters which we specify outside the model. Risk aversion, capital share, and depreciation parameters are standard from the literature. The direct cost of default κ is set to 1.52% of median earnings as in Chatterjee et al. (2020). The probability of regaining access to the credit market, $\theta = 1/7$, is consistent with an average duration of limited credit access upon default of seven years. The recovery rate, $\xi = 16\%$, matches estimates from Call Report data between 1990 and 2020.[31]

Labor productivity The individual labor productivity process is taken from Storesletten et al. (2004). We choose this process since it captures the rich heterogeneity in earnings across individuals and allows for countercyclical earnings variance, which is useful in later experiments with aggregate dynamics. The permanent and transitory components of this process are distributed log-normally around zero, with variances $\sigma_{\epsilon_1} = 0.448$ and $\sigma_{\epsilon_3} = 0.351$, respectively. The persistent component of the process follows an AR(1) process in logs, where the standard deviation of

[31] Recall from Section 1 that we find no evidence that borrowers of different risk levels have systematically different recovery rates. See Appendix A.3 for details on how we compute recovery rates.

shocks is $\sigma_{\epsilon_2} = 0.129$ and the persistence coefficient is $\rho_{\epsilon_2} = 0.957$.[32] Appendix Figure B.2 shows that our earnings process yields an economy-wide earnings distribution which closely matches the one observed in Y-14M data.

3.3 Parameters estimated within the model

Conditional on these externally assigned parameters, we estimate the remaining ones by iteratively searching the parameter space using the Nelder-Mead simplex approach.[33] We target a set of statistics widely used in the unsecured credit literature: bankruptcy and charge-off rates, debt to income ratio, fraction of households in debt, and the average loan rate. Given the general equilibrium nature of the model, we target the capital-output ratio as well. Lastly, following Herkenhoff and Raveendranathan (2020), we use the share of default coming from events other than job loss or income shocks from the Panel Survey of Income Dynamics (PSID) to capture defaults not driven purely by income shocks in our model.

Our approach yields an overidentified system with six parameters and seven targets. When reporting out estimates, we provide intuition for how these parameters shape model moments and, in turn, how our empirical targets determine our parameter estimates. We also describe the overall fit of the model and precision of our estimates to the extent possible. After performing this analysis for the baseline HP model, we compare our estimates to the FP case in the next subsection. See Appendix B.2 for an extended discussion of these issues.

Discount factors Following Athreya et al. (2018) and Krusell and Smith (1998), we assume there are two β levels governed by transition matrix Γ^{β}.[34] The discount factor process has four parameters: (i) the average β level, $\overline{\beta} = 0.876$; (ii) the difference between high and low β levels, $\Delta\beta \equiv \beta_H - \beta_L = 0.379$; (iii) the ergodic share of the population with high β, $\mu_H = 0.704$; and (iv) the transition probability from β_L to β_H, $\Gamma^{\beta}_{LH'} = 0.077$. This process implies high and low β levels of $\beta_H = 0.988$ and $\beta_L = 0.609$, and a more persistent high β, $\Gamma^{\beta}_{HL'} = 0.032$.

[32] The number for σ_{ϵ_2} comes from averaging the estimates for recessions and expansions in Storesletten et al. (2004). The permanent and transitory components are discretized with 5 points. The persistent component is discretized into a 10-state Markov chain using the Tauchen method, which allows the grid points for ϵ_2 to remain the same as the transition matrix changes over the cycle by manipulating the "step size."

[33] Given the many nonlinearities and complexity of solving our model, a derivative-based approach is infeasible. The Nelder-Mead method (attempted for many sets of initial conditions) yields consistent, stable results.

[34] Athreya et al. (2018) find that discount factor heterogeneity is crucial to capture the persistence of financial distress at the borrower level. Krusell and Smith (1998) (and others since) find that this heterogeneity also helps replicate the empirical distribution of wealth in the U.S. economy. We find this as well; see Appendix Figure B.2.

Panel A: parameters assigned directly and common across models

Parameter		Value	Notes
(i) technology and legal			
capital share	α	0.360	standard
depreciation rate	δ	0.072	standard, annual model
risk aversion	γ	3.0	CRRA preferences
bankruptcy filing cost	κ	0.0152	author's calculations
prob. regain credit access	θ	0.143	7-yr avg. exclusion
recovery rate	ξ	0.160	midpoint of estimates [0.157 - 0.163]
(ii) labor productivity			
standard deviation, ϵ_1	σ_{ϵ_1}	0.448	permanent component, log-normal
persistence, ϵ_2	ρ_{ϵ_2}	0.957	persistent component, AR(1) in logs
standard deviation, ϵ_2	σ_{ϵ_2}	0.129	std. dev. of innovations
standard deviation, ϵ_3	σ_{ϵ_3}	0.351	transitory component, log-normal

Panel B: parameters estimated internally for each model

Parameter		HP	FP	Notes
average discount factor	$\overline{\beta}$	0.876	0.871	$\beta_L = \overline{\beta} - \mu_H \Delta\beta = \beta_H - \Delta\beta$
difference, high to low β	$\Delta\beta$	0.379	0.531	$(\beta_L^{VP}, \beta_H^{VP}) = (0.609, 0.988)$
high β share	μ_{β_H}	0.704	0.831	$(\beta_L^{FP}, \beta_H^{FP}) = (0.430, 0.961)$
low to high β trans. prob.	$\Gamma^{\beta}_{LH'}$	0.077	0.052	$\Gamma^{\beta}_{HL'} = \Gamma^{\beta}_{LH'}(\mu_H^{-1} - 1)$
utility cost of default	χ	0.727	0.548	
extreme value scale	ζ	0.123	0.104	

Panel C: target moments for model calibration

Moment (pp)	Data	HP	FP	Source
bankruptcy rate	0.404	0.374	0.390	US Bankruptcy Courts
fraction in debt	11.7	12.3	12.2	Equifax
debt to income	4.30	4.34	4.32	Y-14M
average loan rate	19.6	21.1	21.1	Y-14M
capital-output ratio	3.00	3.09	3.09	BEA
charge-off rate	3.70	3.79	3.69	Call Reports
suboptimal bankrupt share	44.8	45.5	45.7	PSID
SSE	-	**3.07**	**3.15**	

Table 2: **Baseline and fixed premium model: parameterization and targets**

Notes: The model period is annual. The labor productivity process is taken from Storesletten et al. (2004), Table 2 Row B. All moments are evenly weighted, and the sum of squared errors (SSE) is computed using absolute deviations. Section 3.1 descibes the calibration of the borrowing premium schedule for the HP model; the fixed borrowing premium for the FP model is the balance-weighted average from the HP model, 12.2%.

The average level of discounting describes overall patience in the economy and therefore has first order effects on the capital-output ratio and bankruptcy statistics. Conditional on this average level, the other three parameters describe higher order moments of the discount factor process. While all model parameters are jointly determined, they are closely tied to the fraction of households in debt and the debt to income ratio. Intuitively, credit quantity moments (along both the extensive and intensive margins) help inform the distance between and population shares of the two β levels because our model is a closed economy and we target the average level of wealth via the capital-output ratio.

Other preference parameters We estimate a non-pecuniary cost of default, or "stigma," of $\chi = 0.727$. The extreme value scale parameter is $\zeta = 0.123$. Both parameters are sensitive to the size of the debt market and determine the bankruptcy rate. In addition, both parameters are pinned down by the charge-off rate and the suboptimal bankruptcy share.[35] Our model matches empirical bankruptcy and charge-off rates, and is also consistent with the empirical fact that around 45% of bankruptcies have a primary cause other than income or job loss.

Model fit and quality of estimates We obtain a close fit of the model to our empirical targets, with a sum of squared errors of 3.07. While our parameter estimates are consistent with related estimates in the literature, one would ideally evaluate the precision of our estimates by obtaining standard errors. We face two difficulties in calculating these. First, since our empirical moments come from several different data sources, obtaining an empirical variance-covariance matrix requires strong assumptions on the correlations between moments computed from different data sets. Second, some of our moments (e.g. bankruptcy and average interest rates) are computed from almost the full "universe" of observations, suggesting that there is essentially no variation in the calculated "sample" moment.

Still, to provide insight into both the precision of our estimates and into what features of the data drive our estimates, we conduct two analyses in Appendix B.2 which we summarize here. First, we compute the Jacobian matrix around our estimated parameters in order to measure the elasticity of each model moment with respect to a deviation in each estimated parameter. This reveals which parameters drive the largest changes in particular moments as discussed above.

Second, we compute the sensitivity matrix Λ from Andrews et al. (2017) around our estimates to measure the elasticity of each parameter estimate with respect to each empirical moment. This analysis reveals which moments drive the estimates of each parameter. We highlight three main findings. First, the average discount factor is the most precisely estimated of our parameters.

[35] The suboptimal bankruptcy share is the share of bankruptcies accounted for by agents for whom bankruptcy is not the action delivering the highest "fundamental" value. See Appendix A.5 for a formal description of this and other model moments.

Second, the other three parameters of the discount factor process are sensitive to measures of the fraction of households in debt and the debt to income ratio. Since we have no reason to believe that there is substantial mismeasurement in these moments, we are confident in the precision of our estimates of these parameters. Third, the stigma and extreme value scale parameters are the parameters most strongly affected by the share of suboptimal default and the charge-off rate. While this speaks to the merit of targeting both these moments, it is reasonable to suspect that these are less precisely measured and therefore that the default preference parameters are less precisely estimated than the others.[36]

3.4 Parameter differences between HP and FP economies

We re-calibrate the FP model so that any differences we document with respect to our baseline HP model are not attributable to gaps in aggregate targets. We keep all externally assigned parameters (besides borrowing premia) the same in the baseline and FP economies and set the fixed borrowing premium, $\bar{b}$, equal to 12.2% so that the two economies have the exact same balance-weighted average borrowing premium.[37] This feature insures that all differences we document stem from differences in the incidence of borrowing premia, not in their average level. We focus on two main differences between the estimated parameters for these model economies.

First, while we find very similar *average* discount factors across the two economies, the discount factor *processes* differ meaningfully. Relative to the baseline, the FP economy features: (i) lower β's for both high and low levels; (ii) a 40% larger difference between the high and low levels; (iii) an 18% higher ergodic share of high β; and (iv) more persistence in β overall (68% lower transition probability for β_H, 33% lower for β_L). Second, the utility cost of default and extreme value scale parameters are 25% and 16% lower in the FP economy, respectively.

These findings have a common economic driver which is central to our analysis. As highlighted in Figure 3, a key difference is that loans with high default probabilities are much more expensive in the FP economy than the baseline, while loans with low default probabilities are modestly cheaper. This has two important corollaries. First, *across households*, those whose idiosyncratic states make them fundamentally less risky – i.e. those with higher discount factors or higher levels of labor productivity – face lower borrowing costs in the FP economy. By contrast, households in riskier states face higher borrowing costs in the FP economy. Second, *for a given household*, the FP economy features a higher elasticity of the loan rate to leverage: that is, marginal increases

[36]The share of suboptimal bankruptcies is a necessarily difficult object to measure since there can be many reasons operative simultaneously in a household's decision to file. Moreover, we measure this moment using survey, rather than administrative, data. Charge-off rates have the potential to be noisy since recovery is a lengthy and idiosyncratic process.

[37]While exogenous in the FP model, this average is endogenous in our baseline HP economy because the borrowing premium on a given loan depends on the default risk of that loan, which is endogenous.

in the size of the debt have a larger effect on credit prices than in the baseline.[38]

How do these forces manifest in our estimates? In order to replicate the share of agents in debt and their leverage, the FP economy features implausibly more impatient borrowers who are willing to bear these higher costs, explaining the downward shift in both β_H and β_L. By construction, though, the FP economy must deliver the same average wealth, and so the gap between β levels and the share of high β households increase to preserve the average level of discounting.

While these discount factor adjustments bring the FP model in line with debt quantity moments, they also increase the riskiness of borrowing in the economy *ceteris paribus*. As a result, there must be offsetting adjustments in the utility cost of default and extreme value scale parameter to match loan rate spreads and our bankruptcy and charge-off targets. The discount factor shifts, however, must be considered in tandem with the fundamental decrease in riskiness implied by households' responses to the additional steepness of interest rate schedules in the FP economy. This makes the required changes in χ and ζ ambiguous *ex ante*. Our findings suggest that striking the right balance requires that both these parameters be reduced relative to the baseline. Reducing stigma lowers the costs of default, increasing the riskiness of borrowing all else equal. Reducing the variance of the taste shocks on default lowers bankruptcies and charge-offs on average, but also makes the default decision more predictable.

4 Credit Composition and Borrowing Premia

To this point, we have shown that both the heterogeneous premia and fixed premium economies can match key credit aggregates equally well, though each requires meaningfully different parameterizations. We now analyze how the incidence of borrowing premia shapes key cross-sectional features of the unsecured credit market by comparing a set of untargeted moments across the HP and FP economies. We consider two ways of disaggregating the market: by default risk (Section 4.1) and by income (Section 4.2). These two sets of model-to-data comparisons allow us to further illustrate the differences between the HP and FP models, and also provide us a set of untargeted empirical tests which are useful for weighing the relative merits of each model.

4.1 The distribution of debt with respect to default risk

As discussed in Section 3.4, the key difference between the HP and FP economies is how default risk is reflected in credit prices. An immediate question arises: how do these prices translate into

[38] Of course, both effects are relevant for all agents. We systematically find, though, that agents with low discount factors and low labor productivities face significant default risk on any loans they might take, and so they never benefit in practice from the cheapness of low risk debt in the FP economy.

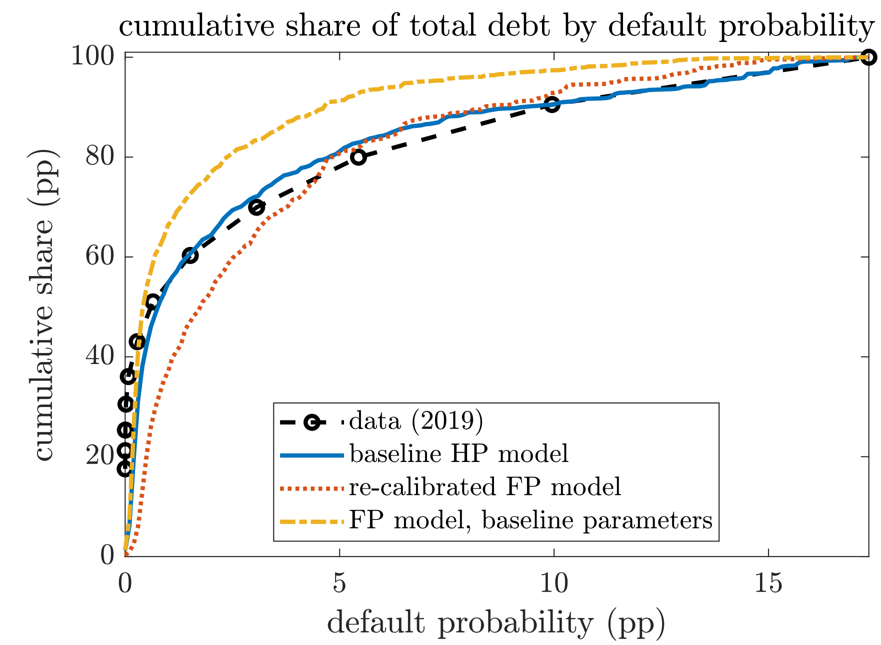

Figure 4: **Share of debt balances by default risk, models and data**

Notes: This figure plots the cumulative share of total debt balances (y-axis) up to each level of default probability (x-axis) in several model variants and in the data. The data are from our Y-14M sample, with default risk bins determined by credit score.

shifts in the composition of credit balances? Figure 4 provides insights into this by plotting the cumulative share of total debt balances by default probability for three versions of the model: (i) our baseline HP model; (ii) the re-calibrated FP model; and (iii) an FP economy solved under our baseline HP parameterization. This set of moments provides a straightforward way to summarize the joint distribution of credit quantities and riskiness in the model, just as Figure 3 does for credit prices and riskiness.

First, we compare the baseline HP model to the FP model with the same parameter values. This latter economy differs only due to the shape of loan price schedules and not parameter values. As such, this FP economy is sharply tilted towards low risk debt because households respond to the high costs of risky borrowing. The cumulative shares of debt with default probability less than 1%, 5%, and 10% are 21.2%, 12.6%, and 7.6% higher, respectively, in this "naive" (i.e. not re-calibrated) FP model compared to the baseline.[39]

Second, we compare the baseline HP and re-estimated FP models in order to account simultaneously for the shape of price schedules and estimated parameter differences. In this case, the CDF for the FP economy gathers mass more quickly than its counterpart for the baseline for all but very low risk levels, with much lower cumulative shares at low risk levels (32.4% lower at 1%

[39] Appendix Table B.3 reports the target moments of the FP model under HP parameters, as well as their deviations from the HP model. The differences are consistent with the results reported here: agents ration their borrowing primarily along the intensive margin, opting for lower risk, smaller loans than in the baseline model.

default probability), higher cumulative shares at high risk levels (2.7% higher at 15%), and very similar shares at intermediate levels (0.45% at 5%). The differences at low default probabilities stem from the lower discount factors and stigma costs in the FP economy: agents are riskier across the board. In the middle of the range of riskiness, though, the fundamental pricing force takes over, and agents begin to respond to incentives by staying away from high risk loans, leading to a thinner right tail in the FP model than in the baseline.

How do these models compare to the data with respect to this key metric? Even though the set of moments in Figure 4 is untargeted, our baseline model closely matches the empirical distribution from our merged Y-14M / Equifax data. By contrast, the FP model does not. This is intuitive: households' debt choices respond to prices, and our baseline model accurately captures how these prices respond to risk. Therefore, our baseline model also more accurately captures the distribution of debt balances with respect to risk. This offers an important additional reason for unsecured credit models to account for empirically correct borrowing premia.

4.2 Moments by income

Our unique data set allows us to construct key moments by conditioning directly on default risk. Since many widely-used data sets preclude this, most existing studies turn to a more readily available measure for cross-sectional analysis: borrower income. For comparability, Figure 5 plots key data moments by income quartile together with their counterparts in our baseline HP and FP models. Again, we find that our baseline model outperforms the FP model. First, panel (a) shows that our baseline delivers the observed modest decline in loan rates as income increases. The FP model, by contrast, delivers a much steeper profile, particularly overstating the interest rates paid by the bottom income quartile. Second, consistent with the share of debt by default risk from Figure 4, panel (b) demonstrates that, in line with the data, our baseline model features the lowest (highest) income quartiles taking on a larger (smaller) share of total debt than the FP model. Finally, panel (c) highlights that our baseline is also more consistent with the observed sharp decline in average individual leverage with respect to income, though neither model hits the high level of indebtedness for the lowest income quartile.

Summary of cross-sectional comparisons Motivated by our empirical analysis in Section 1, borrowing premia in our baseline HP model decline with borrower risk. The flat price schedules induced by these borrowing premia alter households' budget sets relative to the standard FP model, dampening their disincentives to take on high risk loans. While both the HP and FP models match key *aggregate* credit market features by construction, the HP model outperforms the main alternative in terms of matching untargeted *disaggregated* empirical patterns for key

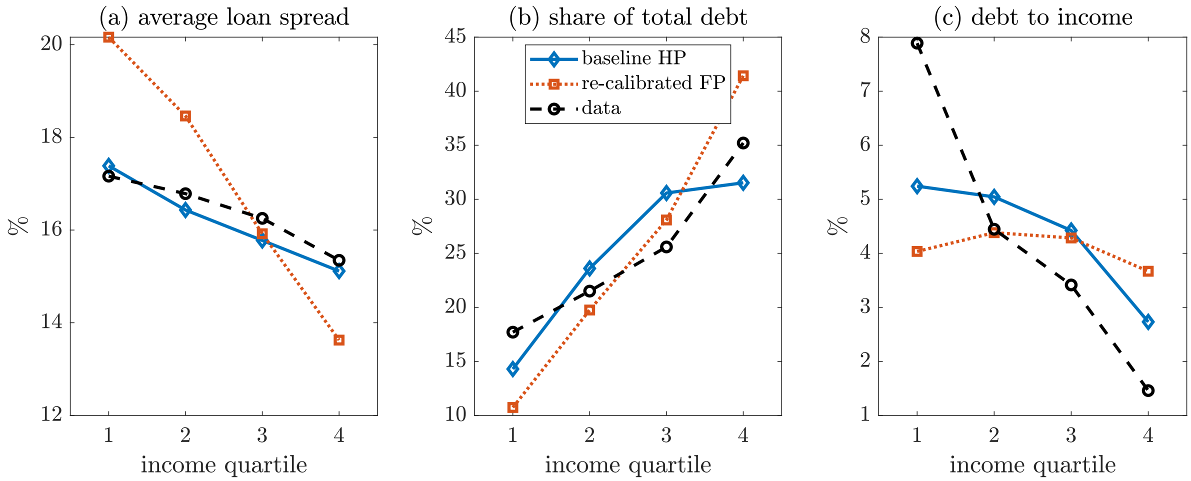

Figure 5: **Untargeted moments by income quartile: data vs HP and FP models**

Notes: Each figure plots the indicated moment by income quartile in several model variants and in the data. The data are from our merged Y-14M / Equifax sample, using income at origination as the income measure. Empirical rates are raw APRs scaled by the prime rate. The sums of squared errors between model and data for the average loan rate / share of total debt / debt to income ratio / total, evenly weighted across quartiles and moments, are 0.45 / 54.5 / 9.99 / 65.0 for the baseline model and 14.9 / 96.5 / 20.5 / 131.9 for the FP model.

credit variables by both borrower income and default risk. We view this as strong evidence in favor of incorporating empirically consistent borrowing premia in unsecured credit models.

5 Aggregate Dynamics and Borrowing Premia

Our analysis so far has examined the steady state of the HP and FP models to highlight cross-sectional effects of incorporating the variable incidence of borrowing premia we observe in the data. In this section, we use the model to study aggregate dynamics with a particular focus on two areas. First, we examine the differences between our baseline heterogeneous premia and fixed premium economies to highlight how the *incidence* of borrowing premia shape aggregate dynamics. Second, using data from our Y-14M / Equifax sample from the second quarter of 2020 – corresponding to the beginning of the Covid-19 pandemic –, we analyze how the *response* of borrowing premia to a negative shock affects aggregate dynamics. Our aggregate shock omits several salient features of the Covid-19 recession, but it is the only recession available in our Y-14M sample period, and so we use it for illustrative purposes.

5.1 Impulse response to a negative shock

The aggregate shock and borrowing premia We simulate a negative aggregate shock with two exogenous components: a 1% reduction in TFP and an increase in labor productivity risk.[40] This latter feature is important: income risk is critical for determining default risk, and therefore both default risk premia and borrowing premia. On impact, the variance of earnings increases from $\sigma_{\epsilon_2} = 0.129$ to $\sigma^R_{\epsilon_2} = 0.163$, following Storesletten et al. (2004). The persistence of the shock, $\rho_z = 1/3$, is consistent with a three-year recession.

In addition to these standard components, we consider two possible scenarios for the schedule of borrowing premia $b(p)$. In the first, there is no response: the schedule remains unchanged from its steady state level, which matches the average for the year 2019 as described in Section 3.1. In the second, the schedule shifts up on impact of the shock to the average observed on impact of the Covid-19 crisis (March to June 2020).[41] This upward shift raises the level of borrowing premia across the board, but does so more in relative terms for high risk borrowers. We term the first scenario "no response," and we term the second "tightening."

Results Table 3 presents our main aggregate results from simulations of this shock. The first three columns summarize the impulse responses by presenting the percentage difference relative to steady state for the indicated moment on impact of the shock.[42] We treat the results in column [1] as our benchmark, corresponding to the model with heterogeneous incidence and no response of premia on impact. Column [2] corresponds to the HP model in which premia shift up on impact, and column [3] corresponds to the FP model with no response.

Focusing first on commonalities across our simulations, the directional responses on impact for most variables reported in Table 3 are the same in each variant: the debt to income ratio, the bankruptcy rate, and average loan spreads all increase, while total debt and aggregate consumption decrease. The exceptions are: (i) the average borrowing premium, which declines in the benchmark model, increases in the version in which premia tighten, and remains flat (by construction) in the fixed premium model; and (ii) the fraction of households in debt, which increases in both HP models but decreases slightly in the FP model. In all cases, there is an increase in *demand* for credit coming from a decline in real incomes (via the equilibrium wage) and an increase in earnings risk. This increase in credit demand is offset somewhat by a surge in bankruptcies, since those who file cannot take on debt by construction. Offsetting the demand

[40] Formally, we assume $Y = zK^\alpha N^{1-\alpha}$, with $z = \overline{z} = 1$ in the steady state model. On impact of the shock, we have $z_1 = 0.99$, and then $z_{t+1} = \rho_z z_t + (1-\rho_z)\overline{z}$ thereafter over the course of the recovery.

[41] We use the exact same estimation procedure for the shock impact schedule as for the steady state schedule; polynomial estimates may be found in Table B.1, and a fit of the model to the empirical borrowing premium schedule may be found in Figure B.1.

[42] Full paths for these variables are presented in Appendix Figure B.3.

Incidence of premia	Heterogeneous		Fixed	Heterog.	Fixed
Response of premia on impact	None	Tighten	None	Tighten	None
Moment	% diff, impact v steady state			pp diff, $[j]-[1]$	
Column	[1]	[2]	[3]	$[2]-[1]$	$[3]-[1]$
total debt	-0.06	-0.30	-0.80	-0.24	-0.74
debt to income	0.77	0.54	0.15	-0.23	-0.63
fraction in debt	0.43	0.30	-0.05	-0.13	-0.47
bankruptcy rate	11.8	14.4	16.6	2.58	4.79
average loan rate spread	0.13	4.36	0.89	4.23	0.76
average borrowing premium	-0.55	3.75	0.00	4.30	0.55
aggregate consumption	-0.21	-0.22	-0.23	-0.01	-0.02

Table 3: **IRF: peak responses across model variants**

Notes: Columns [1] through [3] show the percentage difference in each row's moment at impact of the shock to the pre-shock steady state for the indicated model variant. The rightmost columns compare the indicated model variant to the version with variable premia and no response of premia to the impact of the shock, showing the percentage point difference in the max percentage difference moments from the leftmost columns. Full paths for the variables in this table are depicted in Figure B.3.

side, though, increases in individual-level default probabilities limit credit *supply*, shifting up interest rate schedules all else equal. Combining these forces, we see drops in both total debt and total consumption on impact.

We next extend the analysis from Section 3.4, comparing how our baseline HP and FP economies (columns [1] and [3]) differ in terms of aggregate dynamics. Outcomes are more severe across the board in the FP model, with the shifts in aggregate moments looking qualitatively similar to the case with tighter premia described below. As discussed above, a key difference between the HP and FP models is that increases in default premia associated with riskier loans are *not* offset by decreases in borrowing premia in the FP model. Therefore, marginal increases in risk lead to steep increases in spreads, and so generally borrowers choose loans with only small increases in risk relative to steady state.

How do these forces shape aggregate results? While the increase in loan spreads is only 0.76 pps larger in the FP model than in our benchmark, declines in credit quantities are much sharper: the debt to income ratio rises by 0.63 pps less, the fraction in debt rises by 0.47 pps less (actually *declining* relative to steady state), and total debt rises by 0.74 pps less (a drop over 10 times as large as the benchmark). Facing worse credit prospects, households default more and consume less: bankruptcies increase by nearly 5 pp more, and the drop in aggregate consumption is 9% larger in the FP model.

Finally, we analyze how the upward shift in the schedule of borrowing premia shapes the aggregate response of the credit market to the shock by comparing columns [1] and [2] in Table

3. In the benchmark model, individuals – particularly those with higher risk – face better terms of credit since borrowing premia do not increase; by adjusting their borrowing choices, they actually pay 0.55 pp *lower* borrowing premia on average than in steady state, which leaves average spreads virtually unchanged from their steady state level ($+0.13$ pps only) despite the increase in individual-level risk outlined above. By contrast, the tightening of the schedule of borrowing premia rules out such adjustments, and borrowing premia and spreads increase sharply on impact ($+3.75$ and $+4.36$ pps, respectively). These fundamental pricing shifts are reflected in all key credit market quantities. The average increase in borrower leverage and the share of debtors are both 30% lower when premia increase, and as a result the drop in total debt on impact is five times as large. Since tighter premia make future access to credit less valuable, we also observe a 22% larger initial surge in bankruptcy filings. The drop in aggregate consumption is just over 2% larger with tightening.

5.2 Summary of business cycles

In this section, we explore aggregate dynamics further and include business cycle moments in our empirical comparisons. Specifically, we augment our baseline model to consider stochastic equilibria, featuring persistent shocks to aggregate TFP. We calibrate these shocks to match the frequency and duration of recessions and the cyclical volatility of aggregate output, and link both the riskiness of the persistent component of labor productivity (following Storesletten et al. (2004)) and the schedule of borrowing premia (following the estimation described above) to the level of TFP.[43] Given our rich household heterogeneity and the endogenous aggregate evolution of the distribution $\mu(x)$, we solve the model using the method of Krusell and Smith (1998). Details of the extension, calibration, and computation of the model, as well as a full set of the results referenced in this subsection, are available in Appendix B.4.

While all three variants of the model from Table 3 are broadly consistent with key business cycle moments of the credit market, our heterogeneous premia models: (i) more closely match the data than the FP version; and (ii) demonstrate that including the response of borrowing premia in a downturn is crucial for matching the volatility of interest rate spreads. Along the former dimension, the HP model performs particularly better than the FP model with regards to credit quantity moments such as total debt, the average debt to income ratio, and bankruptcy filings. In particular, total debt in the FP model overpredicts the countercyclicality we observe in the data and underpredicts volatility. Our HP models correct this by more accurately capturing the prices faced by the crucial high-risk segment of the credit market. Along the latter dimension, the

[43] Given that we only have reliable Y-14M data dating back to 2014, we are unable to measure borrowing premia in historical recessions outside the current one, which limits our ability to perform a richer estimation of the $b(\cdot)$ function over the cycle.

HP model with borrowing premia that do not shift over the cycle yields implausibly low volatility in average spreads. This suggests that movements in borrowing premia, rather than pure risk premia, drive cyclical movements in unsecured credit prices.

6 Lending Standards and Borrowing Premia

Our focus so far has been on capturing the empirically accurate incidence and response of borrowing premia to measure their cross-sectional and dynamic impacts on the unsecured credit market. This analysis has been intentionally agnostic about the *sources* of borrowing premia. While an exhaustive study of the drivers of these borrowing premia is out of scope for the present paper, they are necessarily linked the commonly cited but less well-understood concept of lending standards. These are related to such factors as the economic outlook, competitive landscape of the credit market, banks' risk tolerances, and many more. In this section we extend our model to include a notion of lending standards and endogenize one likely driver: credit market tightness.

We preserve virtually all of the structure from our baseline model, but enrich it in one key dimension: we allow the *level* of borrowing premia to be endogenous. This level effect is determined by the endogenous aggregate demand for credit and supply of loanable funds in the economy. Meanwhile, the *incidence* of borrowing premia – i.e. the *shape* of the schedule – is exogenous. We argue that this exogenous component of borrowing premia offers a useful, quantitative formalization of the portion of bank lending standards that is supply-driven. This component is little understood but quantitatively significant, as documented in Bassett et al. (2014) and **?** for example. We proceed by first describing existing measurements of lending standards and their shortcomings. We then present our extended model and conduct our main quantitative experiment: an estimate of shifts in lending standards.

6.1 Overview of lending standards

Each quarter, the Senior Loan Officer Opinion Survey (SLOOS) conducted by the Federal Reserve surveys a representative sample of U.S. banks about changes in their lending standards along many dimensions (e.g. approvals, limits, spreads, etc.). Typically, these qualitative responses are aggregated into diffusion indices intended to measure the extent to which lending standards have tightened or eased in the aggregate. These indices show that standards tighten sharply during recessions, ease gradually during recoveries, and remain largely unchanged during expansions.[44]

[44] See Appendix A.4 for additional details on the SLOOS and the corresponding construction of diffusion indices. Broadly, this index takes the responses of banks on changes in lending standards and weights them by banks' shares of total credit card lending. The patterns we document hold broadly for all categories of bank loans. See Figure A.6a for trends in standards and terms diffusion indices for the U.S. credit card market.

Notably, measures of lending standards constructed in this way necessarily describe quarterly changes as opposed to levels, making comparisons across time difficult.[45]

While many theories use tightness in lending standards measured by these diffusion indices as justification for modeling a restricted supply of credit, these papers do not readily point to a common type of change in credit supply.[46] Why do banks tighten standards, and how are tightened standards reflected in interest rate spreads, credit limits, or other terms of credit? Our goal in this section is to combine our model, our empirical estimates of borrowing premia, and survey-based evidence to provide a more concrete understanding of lending standards.

It is worth noting that several existing empirical studies shed light on some of the nuances of lending standards.[47] In addition, certain editions of the SLOOS ask special questions about these issues. For example, in the October 2019 SLOOS, banks cited as main reasons for reduced willingness to approve credit card applications: (i) concerns related to borrowers' ability to consistently make payments; (ii) expected deterioration in portfolio quality; (iii) a less favorable or more uncertain economic outlook; and (iv) a reduced tolerance for risk. These responses support two key facets of the framework we now outline: that lending standards are closely related to borrower risk, and that they respond to aggregate conditions.

6.2 Extending the model

Overview We augment our baseline model with a constraint on lenders' loan supply which dictates the maximum fraction of loanable funds (savings) in the current period that lenders can allocate to credit provision (loans): that is, for a given standards weighting function $\lambda(\cdot)$,

$$\text{total standards-weighted funds borrowed today} \ \leq \ \text{total funds saved today} \tag{14}$$

Consistent with our empirical approach, we assume that the weight on a given loan varies with both the riskiness of the loan and time (or the aggregate state). Equation (17) below adapts (14) into model notation. While this specification is admittedly *ad hoc*, as demonstrated below it has two features of first order importance to our exercise: (i) it allows borrowing premia to have an endogenous component which depends on overall credit market tightness; and (ii) it preserves the pricing functions from our baseline model.

[45] While the July SLOOS inquires about the levels of standards relative to the midpoint of the range since 2005, levels are still difficult to compare across time in general.

[46] See, for example, Dell'Ariccia and Marquez (2006), Darst et al. (2020), and Fishman et al. (2019). Most related in spirit to our paper is **?**, who build and estimate a DSGE model with credit fictions. The authors here argue that a shock to credit supply accounts for a significant portion of cyclical changes in lending standards. Shocks to credit supply in their framework originate from moral hazard; in our framework, there are no direct shocks to credit supply, just shifts in the quantity of loanable funds relative to loan demand.

[47] See, for example, Bassett et al. (2014), Castro et al. (2020), and **?**.

Since households in the model behave competitively – taking the savings price $\overline{q}$ and loan price schedules $q(\ell; x)$ as given –, this change in the lender problem has no direct impact on the specification of the household block of the model. In particular, equations (5) through (12) are unchanged. All the differences in this extended model come from how prices are determined on the lender side, and so we turn our attention there.

Modified lender problem As in our baseline model, competitive lenders maximize expected discounted flow profits, but now this profit maximization problem is subject to an additional constraint of the form (14). Flow profits π for a lender are the sum of net capital returns and repayments on old contracts, net of issuances for new contracts:

$$\begin{aligned} \pi(K', \mathcal{M}'; K, \mathcal{M}) &= \underbrace{(1 + r - \delta)K - K'}_{\text{net return on capital}} + \underbrace{\int_{\mathcal{X}\times\mathcal{L}} q(\ell; x)\ell dm'(\ell; x)}_{\text{issuances, new contracts}} \\ &\quad - \underbrace{\int_{\mathcal{X}_{-1},\mathcal{X},\mathcal{L}} \left(1 - g_{BK}(\ell; x) + \xi g_{BK}(\ell; x)\right) \ell dm(\ell; x_{-1}) d\mathbb{P}(x|x_{-1})}_{\text{repayments, last period contracts}} \qquad (15) \end{aligned}$$

The first term in this profit equation is the return on aggregate capital (rK), net of investment ($K' - (1 - \delta)K$). The second term reflects payments from households for contracts issued today, $\mathcal{M}' \equiv \{m'(\ell; x)\}$, at discount prices $\{q(\ell; x)\}$. The third term reflects repayments from households on last period's contracts, $\mathcal{M} \equiv \{m(\ell; x)\}$. The $g_{BK}(\cdot)$ term in this expression accounts for bankruptcies declared on loans made in the previous period, while $\mathbb{P}(x|x_{-1})$ accounts for exogenous transitions of individual states between periods. The ex ante repayment probability for a loan of size ℓ to a borrower in state x, $p(\ell; x)$, is calculated exactly as in equation (12). Therefore, the lender problem is

$$W(K, \mathcal{M}) = \max_{K', \mathcal{M}'} \pi(K', \mathcal{M}'; K, \mathcal{M}) + \frac{1}{1+i}\mathbb{E}\left[W(K', \mathcal{M}')\right] \qquad (16)$$

$$\text{subject to} \quad -\int_{\mathcal{X}\times\mathcal{L}_-} \lambda(p(\ell; x))q(\ell; x)\ell dm'(\ell; x) \leq \int_{\mathcal{X}\times\mathcal{L}_+} q(\ell; x)\ell dm'(\ell; x) \qquad (17)$$

The objective function (16) is the recursive formulation of the discounted sum of the flow profits defined in (15), with lenders discounting at the equilibrium interest rate. The expectation operator in this expression accounts for all individual-level transitions. The key addition in our model is constraint (17), which implements the generic version (14). The summation on both sides of (17) is over all idiosyncratic states, but the summation on the left (right) side of the equation is only over loans, $\mathcal{L}_- = \{\ell | \ell < 0\}$ (savings, $\mathcal{L}_+ = \{\ell | \ell > 0\}$).

Analyzing this problem using standard techniques and placing a multiplier $\eta \geq 0$ on constraint

(17) delivers the following two pricing equations:[48]

$$\overline{q} = \frac{1}{(1+i)(1+\eta)} \tag{18}$$

$$q(\ell;x) = \frac{p(\ell;x)+\xi(1-p(\ell;x))}{(1+i)\,(1+\eta\lambda(p(\ell;x)))} \text{ if } \ell<0 \tag{19}$$

Equation (18) defines the common price on all savings contracts. Since there is no default risk, individual states x do not affect this price, but lending standards still have an indirect, general equilibrium effect via the tightness of the loan supply constraint. If the constraint binds so that $\eta > 0$, then $\overline{q} < (1+i)^{-1}$, and savers earn a premium relative to our baseline model, easing the constraint by promoting savings among households. Equation (19) is the price of a loan. The numerator is the expected repayment per unit, and the denominator reflects that borrowers pay an additional premium above and beyond the base cost of funds, both exactly as in the loan pricing equation (13) from the baseline model.

The key difference between equations (13) and (19), then, is that the borrowing premium now has an endogenous component (via η) in addition to the exogenous one (via $\lambda(\cdot)$). Applying the definition of the borrowing premium as in our empirical analysis, we obtain

$$b(p) = \eta\lambda(p) \tag{20}$$

When the supply of savings is low relative to the demand for credit, the multiplier η increases to ease constraint (17) by raising loan and savings rates to discourage borrowing on the margin. Importantly, this increase in borrowing costs is accomplished purely through borrowing premia (separately from risk premia), and this effect is more pronounced for loans which have high λ-weights in constraint (17). Note that the *shape* of the borrowing premium schedule with respect to default probability is entirely determined by the exogenously specified $\lambda(\cdot)$ function, while its *level* is endogenously determined via the equilibrium value of the multiplier η. This multiplier summarizes the endogenous effect of credit market tightness, that is, the demand for debt relative to the supply of savings.

6.3 Estimating lending standards

Approach Since lending standards are inherently dynamic, our approach to inferring them combines the steady state of the extended model with the endogenous response to a negative shock. The former element delivers a baseline lending standards function, $\lambda_0(p)$, and equilibrium multiplier, η_0, while the latter yields the shifted versions, $\lambda_1(p)$ and η_1. Comparing these two $\lambda(\cdot)$

[48] Details and derivations may be found in Appendix C. Note that $q(\ell;x) = \overline{q}$ when $\ell \geq 0$.

functions offers a well-defined notion of "tightening" or "easing" standards, both in the aggregate (summing across *all* risk levels) and in the cross-section (considering *each* risk level). We require that our endogenous borrowing premia schedules (20) match exactly those in the data, using the FY2019 schedule as our baseline schedule, $b_0(p)$, and the March - June 2020 schedule as our on-impact schedule, $b_1(p)$.[49] In our modified environment, this implies that the parameters of the lending standards function $\lambda_t(\cdot)$ are identical to those of the borrowing premia function $b_t(\cdot)$ from Table B.1, *up to an endogenous level shift from the multiplier* η_t. We use the same parameterization as our baseline (heterogeneous premia) model from Table 2.[50]

In order to bring our model environment closer to the empirical setting of the Covid shock observed in early 2020, we use a slightly different aggregate shock relative to the one from Section 5.1, whose primary purpose was to illustrate aggregate dynamics with borrowing premia in a general sense. We maintain the assumptions that the shock has an expected duration of 3 years and that earnings risk increases (consistent with Storesletten et al. (2004)) on impact, but we assume a larger initial drop in TFP of 9.8%. This generates the drop in output observed in the first quarter of the pandemic recession.[51]

From the perspective of our inference on lending standards, the impact of the shock has two main consequences: (i) demand for credit increases, driving up η; and (ii) borrowing premia shift out, changing $b(p)$. Recall that our calibration ensures that $\eta_t \lambda_t(p) = b_t(p)$ for $t = 0, 1$. The estimated proportional shift in lending standards, then, adjusts observed shifts in borrowing premia for shifts in credit demand:

$$\underbrace{\lambda_1(p)/\lambda_0(p)}_{\text{estimated shift in standards}} = \underbrace{b_1(p)/b_0(p)}_{\text{observed shift in premia}} \Big/ \underbrace{\eta_1/\eta_0}_{\text{endogenous change in multiplier}} \tag{21}$$

Equation (21) offers a clean summation of how our approach links standards, credit demand, and borrowing premia. Our estimated proportional shift in lending standards scales observed shifts in borrowing premia at each level of risk by the endogenous response of aggregate credit demand, summarized by the change in the multiplier on constraint (17). An important feature of equation (21) is that while it holds p-by-p, the scale factor for aggregate demand is independent of p. As a result, even if borrowing premia increase for all risk levels, our model reveals that this does not automatically imply that lending standards increase for all risk levels.

[49] This is the extended model's analog of the approach from the impulse response analysis in Section 5.1.

[50] The only differences between the steady states of our baseline and extended models come from the additional return on savings associated with a positive η per equation (18). Since the value of η is small (it increases the return on savings by less than 6 bps), the differences between these economies are negligible.

[51] With better earnings data from the Covid-19 crisis it could be possible to match this component of the shock more closely to the empirical episode as well. Absent this data, though, and recognizing that other policy interventions played a role in shaping the response to the crisis, we maintain our initial assumption for transparency.

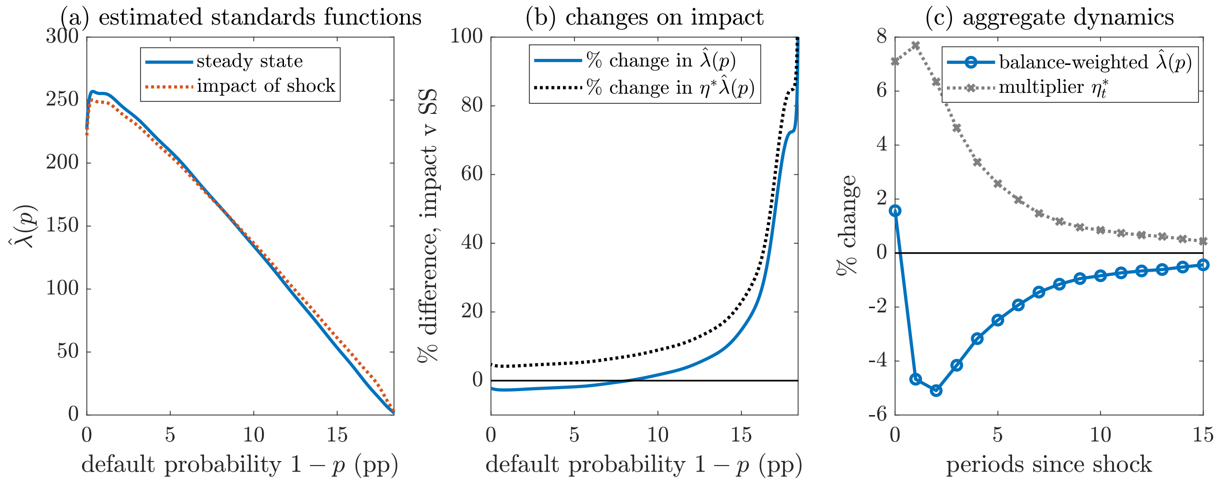

Figure 6: **Estimated lending standards**

Notes: This figure shows the inferred lending standards schedules for pre-Covid and Covid periods, respectively. Percentage differences are computed with 2019 ("pre-Covid") as the base year. The equilibrium values of the loan constraint multiplier are $\eta_0^* = 5.68\text{e-}4$ in steady state and $\eta_1^* = 6.08\text{e-}4$ on impact of the shock. In panel (c), the balance-weighted average lending standards is computed using the distribution of borrowing choices within the period as in Figure 4.

Results Figure 6 presents the results of this analysis. Panel (a) plots the estimated standards functions in steady state and on impact of the shock. By construction, the shapes of these schedules and their corresponding borrowing premia schedules match exactly, but their levels vary because the equilibrium value of the multiplier on impact, $\eta_1^* = 6.08\text{e-}4$, is 7.1% higher than its steady state level, $\eta_0^* = 5.68\text{e-}4$. Panel (b) plots the percentage difference in the standards function as in equation (21), as well as the observed shift in borrowing premia. The latter change confirms that borrowing premia shifted up across all levels of risk, but particularly so for high risk borrowers; for example, no borrower with a default probability lower than 10% saw an increase in borrowing premium of more than 8.9%, while a borrower with a default probability of 15% saw their borrowing premium increase by 23.0%.

The estimated shifts represented by the blue line in panel (b) imply that standards moved against higher risk borrowers. In particular, we observe "tightening" – a positive shift in the $\lambda(\cdot)$ function – for all default probabilities above 8.2%. Using our empirical mapping from default probabilities to credit scores, this corresponds to borrowers with credit scores of approximately 670-690 or below. For very high risk borrowers – for example, those with default probability in excess of 14.1% or a credit score lower than about 640 – the effect is especially stark, with estimated $\lambda(\cdot)$-weights increasing by more than 10%. By contrast, intermediate and low risk borrowers faced little change in standards with estimated shifts in $\lambda(\cdot)$-weights in the range of -2.7% to -0.1%. The modest increases in borrowing premia these borrowers faced, then, were entirely attributable to the general equilibrium effect of increased demand for credit.

Validation: how do standards in our model compare to survey evidence? We now compare our results to survey evidence from the Senior Loan Officer Opinion Survey (SLOOS). As discussed above, it is difficult to incorporate qualitative survey data directly into a quantitative model, and quantitative diffusion indices have no direct analog in the model. However, by delving deeper into the details of these surveys we can gain insight into how the model's interpretation of standards compares to real world outcomes.

Broadly, SLOOS data suggest that lending standards and key terms for credit card loans are tighter in recessions than expansions.[52] Banks start tightening standards in the lead-up to downturns and subsequently ease gradually after recessions. Panel (c) of Figure 6 shows that our model reflects this pattern. On impact of the shock, the balance-weighted average of lending standards increases 1.57%. This is consistent with a modest tightening of standards on average. Subsequently this average declines into negative (easing) territory, gradually returning to steady state thereafter. This pattern is consistent with the dynamics of standards during the pandemic presented in Appendix Figure A.6a. Why does our model deliver a short-lived tightening followed by a persistent easing of standards, despite the fact that borrowing premia remain elevated over the entire course of the recovery? The dynamics of the multiplier η_t in panel (c) highlight that net demand for credit remains persistently elevated. This raises the endogenous component of borrowing premia, and so inferred lending standards must ease to offset this effect.

What about borrower-specific factors? In our model, *premia* tighten for all borrowers in response to the shock, but *standards* tighten only for high risk borrowers and loans. This suggests that, in a sense, credit supply has shifted to the low risk segment of the market from the high risk. We find evidence for this differential shift in the October 2019 SLOOS as well.[53] Banks reported that they were less likely to approve loans for borrowers with FICO scores of 620 in comparison with the beginning of the year, while they were about as likely to approve such loans for borrowers with FICO scores of 720 over this same period. While this partition of scores is relatively coarse, the pattern is well within the ranges reported in Figure 6.

7 Conclusion

This paper uses administrative data from Y-14M and Equifax to document borrowing premia in excess of default risk premia in the U.S. credit card market. These premia are large, varying systematically and tending to decline with default probability. This pattern holds conditional on borrower income and across time. Borrowing premia rise across the board in crises, but rise the

[52] See Figure A.6 for details.

[53] This vintage of the survey included a set of special questions asking banks to assess the likelihood of approving credit card applications by borrower FICO score in comparison with the beginning of the year.

most for high risk loans. Our findings reveal that borrower risk is a sufficient statistic for these premia, while other borrower characteristics and bank-specific factors, are secondary.

Notably, these premia have been absent from standard unsecured credit models. This is a crucial omission because these premia sharply alter the relationship between loan prices and loan risk, implying a much flatter slope of interest rate spreads with respect to default probability than in existing models. This flat schedule dampens households' disincentives to take on riskier loans. Therefore, embedding this simple feature changes the quantitative implications of standard credit models in a first order way. In the cross-section, capturing the empirical incidence of borrowing premia allows the model to match the joint distribution of debt and credit risk very closely, whereas a model with a fixed premium independent of default risk cannot. Furthermore, empirically correct borrowing premia play a crucial role in shaping aggregate credit market dynamics. In response to a negative aggregate shock, the relatively low cost of high risk loans facilitates borrowing for very constrained agents, though the tightening of borrowing premia dampens this effect. Extending our model to endogenize borrowing premia allows us to relate them to lending standards, supporting survey evidence that credit supply tightens only for the riskiest portion of the borrower population.

Our analysis presents three main limitations, which we view as avenues worth pursuing in future research. First, we interpret additional costs of borrowing in excess of those implied by measurable default risk as arising from supply frictions which we model exogenously. While sufficient for the measurement purposes in this paper, this approach is silent about the origin of these costs. Future research accounting for regulatory constraints, search frictions, or price discrimination strategies, for example, would provide further insights on the mechanics underlying lending policies.[54] Second, non-price terms of credit, such as credit limits, are also an important part of lending policies in the unsecured credit market. These terms also vary widely across borrower risk, akin to what we emphasize for spreads.[55] While our model can deliver some measures of credit limits, it is mostly silent on non-price terms of credit. Third, the limited time span of Y-14M hinders our ability to consider recession episodes before the 2020 recession, constraining our business cycle analysis and estimation of standards. Related, the particularities of the Covid crisis – government support via stimulus checks, expanded unemployment benefits, specific income shocks, etc. – are likely to change households' borrowing and saving behavior. We hope to tackle these facets of our analysis as more data become available.

[54] In Castro et al. (2020), we empirically analyze these mechanics and drivers of changes in standards across banks using matched survey data on banks' lending policies with detailed data on banks' loan portfolios.

[55] See Appendix Figure A.6. Furthermore, banks' SLOOS responses suggest that non-price terms are a large margin of adjustment over the cycle.

References

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, "Regulating Consumer Financial Products: Evidence from Credit Cards," *Quarterly Journal of Economics*, 2015, pp. 111–164.

Allen, Jason, Robert Clark, and Jean FranÃ§ois Houde, "The effect of mergers in search markets: Evidence from the Canadian mortgage industry," *American Economic Review*, 2014, *104*, 3365–3396.

_ , _ , and _ , "Search frictions and market power in negotiated-price markets," *Journal of Political Economy*, 2019, *127*, 1550–1598.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro, "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *Quarterly Journal of Economics*, 2017, *132*, 715–753.

Athreya, Kartik B., "Welfare implications of the Bankruptcy Reform Act of 1999," *Journal of Monetary Economics*, 11 2002, *49*, 1567–1595.

Athreya, Kartik, Jose Mustre del Rio, and Juan M. Sanchez, "The Persistence of Financial Distress," *Review of Financial Studies*, 2018.

_ , Xuan S. Tam, and Eric R. Young, "Unsecured credit markets are not insurance markets," *Journal of Monetary Economics*, 2009, *56*, 83–103.

Bassett, William F., Mary Beth Chosak, John C. Driscoll, and Egon Zakrajsek, "Changes in bank lending standards and the macroeconomy," *Journal of Monetary Economics*, 2014, *62*, 23–40.

Castro, Andrew, Kyle Dempsey, David Glancy, and Felicia Ionescu, "What are Bank Lending Standards?," *Working Paper, Ohio State University*, 2020.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan, "Business cycle accounting," *Econometrica*, 2007, *75*, 781–836.

Chatterjee, Satyajit, Dean Corbae, Kyle Dempsey, and José-Víctor Ríos-Rull, "A Quantitative Theory of the Credit Score," *Working Paper, Ohio State University*, 2020.

_ , _ , Makoto Nakajima, and José-Víctor Ríos-Rull, "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 2007, *75*, 1525–1589.

Chen, Kaiji, Patrick Higgins, and Tao Zha, "Cyclical lending standards: A structural analysis," *Review of Economic Dynamics*, 2021, *42*, 283–306.

Darst, R. Matthew, Ehraz Refayet, and Alexandros Vardoulakis, "Banks, Non Banks, and Lending Standards," *Federal Reserve Board Finance and Economics Discussion Series*, 10 2020, *2020*, 1–52.

Dell'Ariccia, Giovanni and Robert Marquez, "Lending booms and lending standards," *Journal of Finance*, 10 2006, *61*, 2511–2546.

Drozd, Lukasz A. and Michal Kowalik, "Credit Cards and the Great Recession: the Collapse of Teasers," *Working Paper, Federal Reserve Bank of Philadelphia*, 2019, pp. 1–55.

Dvorkin, Maximiliano, Juan Sanchez, Horacio Sapriza, and Emircan Yurdagul, "Sovereign Debt Restructurings," 2019.

Fishman, Michael J, Jonathan A Parker, and Ludwig Straub, "A Dynamic Theory of Lending Standards," *Working Paper, MIT*, 2019, pp. 1–39.

Galenianos, Manolis and Alessandro Gavazza, "Regulatory Interventions in Consumer Financial Markets: The Case of Credit Cards," 2020.

Gilchrist, Simon and Egon Zakrajsek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 2012, *102*, 1692–1720.

Greenwald, Daniel L., John Krainer, and Pascal Paul, "The Credit Line Channel," *Federal Reserve Bank of San Francisco, Working Paper Series*, 2020.

Gross, David B. and Nicholas S. Souleles, "Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data," *Quarterly Journal of Economics*, 2002, *117*, 149–185.

Herkenhoff, Kyle, "The Impact of Consumer Credit Access on Unemployment," *Review of Economic Studies*, 2019, *86*, 2605–42.

_ and Gajendran Raveendranathan, "Who Bears the Welfare Costs of Monopoly? The Case of the Credit Card Industry," *Working Paper, Federal Reserve Bank of Minneapolis*, 2020.

Herkenhoff, Kyle F., Gordon Phillips, and Ethan Cohen-Cole, "The Impact of Consumer Credit Access on Employment, Earnings, and Entrepreneurship," *NBER Working Paper*, 2016, pp. 1–78.

Krusell, Per and Anthony A Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 1998, *106*, 867–896.

Livshits, Igor, James C. MacGee, and Michèle Tertilt, "Accounting for the Rise in Consumer Bankruptcies," *American Economic Journal: Macroeconomics*, 4 2010, *2*, 165–193.

_ , _ , and MichÃšle Tertilt, "Consumer Bankruptcy: A Fresh Start," *The American Economic Review*, 2007, *97*, 402–418.

Lown, Cara and Donald P Morgan, "The Credit Cycle and the Business Cycle: New Findings Using the Loan Officer Opinion Survey," *Journal of Money, Credit and Banking*, 2006, *38*, 1575–1597.

Nakajima, Makoto and José-Víctor Ríos-Rull, "Credit, Bankruptcy, and Aggregate Fluctuations," *Working Paper, Federal Reserve Bank of Philadelphia*, 2019, pp. 1–61.

Raveendranathan, Gajendran, "Revolving credit lines and targeted search," *Journal of Economic Dynamics and Control*, 2020, *118*, 103964.

Schreft, Stacey L. and Raymond E. Owens, "Survey evidence of tighter credit conditions: What does it mean?," *Economic Review*, 1991, *77.*

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, "Cyclical dynamics in idiosyncratic labor market risk," *Journal of Political Economy*, 2004, *112*, 695–717.

Sullivan, James X., "Borrowing during unemployment: Unsecured debt as a safety net," *Journal of Human Resources*, 2008, *43*, 383–412.

FOR ONLINE PUBLICATION: Appendix for "Borrowing Premia in Unsecured Credit Markets"

A Data Appendix

A.1 Main data sources, sample construction, and additional facts

As described in Section 1, we combine two data sources to gain insights into credit terms and usage: (i) Y-14M, a detailed loan-level panel data set built from the portfolios of large bank holding companies in the United States, collected by the Federal Reserve Board as part of the Comprehensive Capital Analysis and Review; and (ii) the FRBNY Consumer Credit Panel (Equifax), a nationally representative five percent sample of all credit files for U.S. borrowers. This is feasible given that the population of borrowers in the two data sets is the same. Given the dominance of the largest banks in the credit card market, almost the entire universe of credit card loans and borrowers (as represented by Equifax) is in Y-14M.

Equifax These data are a nationally representative anonymous random sample from Equifax credit files. The data set is a quarterly panel, beginning in 1999, with snapshots of consumers' credit profiles at the end of each quarter. Starting in January 2020, the data is available at a monthly frequency. We drop observations missing either Risk Score or raw delinquency data. Given the low frequency of transactions for the three lowest Risk Score bins, we trim these bins from our sample.

For robustness, we compute three default measures for borrowers with credit cards as follows: (i) **baseline**, which includes bankruptcy or severe derogatory; (ii) **narrow**, which includes only bankruptcy; and (iii) **broad**, which includes bankruptcy, severe derogatory, or 120 days past due. Equifax contains additional information for account payment or delinquency status including (in increasing order of severity): 30 days past due (DPD); 60 DPD; 90 DPD; 120+ DPD; in collections; severe derogatory; and bankruptcy. The patterns we find are the same with only slight differences in levels.[56]

We use all four quarters in 2019 and compute the average probability of default for each credit score bin for this period. The baseline and broad definitions are quite similar, while the narrow measure preserves the pattern with a downward level shift (Figure A.1, right panel). When recomputing data moments for the Covid period, we keep our measure of probability of default unchanged since there is little difference in default probabilities across credit scores of borrowers (Equifax Risk Score) between 2019 and March-June 2020 (Figure A.1, left panel) for all but the three lowest Risk Score bins, which we drop as described above.

[56] In principle, one could construct the same exact narrow definition of default using the bankruptcy flag variable in Y-14M data or the broader definition using information on DPD. However, severe derogatory is not reported in Y-14M, and default information is generally unreliable in Y-14M.

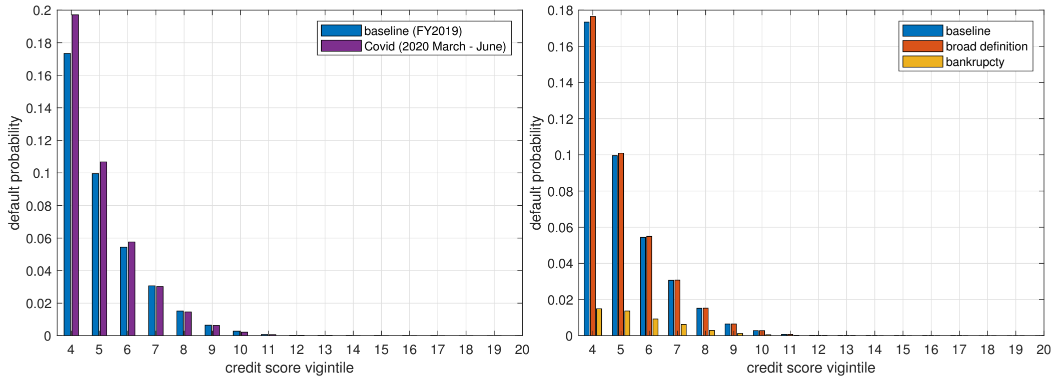

Figure A.1: **Probability of default across periods and alternative measures**

Notes: The left panel plots the probability of default by 5% credit score bin using our baseline (severe derogatory plus bankruptcy) measure for 2019 and 2020. The right panel plots the three measures (baseline, broad, and narrow) for the year 2019.

Y-14M We construct our sample as follows. First, we take a 5 percent sub-sample and restrict attention to general purpose credit cards that are not securitized, secured or under promotion for accounts that are open (active or inactive) in each month. We focus on consumer bank cards which are general purpose credit cards, and we exclude non-consumer / business cards and private label / propriety cards, since these can only be used in the stores of the retailer issuing the card. We drop observations missing FICO score, APR, or end-of-cycle balance. Additionally, we drop observations for which: (i) the Vantage score is reported; (ii) FICO scores are below 300 or over 850; (iii) APR is below 5 percent. All these are very rare. We use the most recently updated credit score available for the primary account holder using a commercially available credit score unless it is missing, in which case we replace this with the credit score at origination. About 10% of the sample is replaced by the score at origination. Furthermore, we truncate the sample to eliminate the highest risk borrowers in the three lowest FICO score vigintiles. This keeps consumers with probabilities of default up to just under 20% on average in our sample, which we take as a high upper bound for our measure of default and representative of the economic behavior our model captures. This truncated sample still covers more than 90% of credit card balances.

Using this data set we compute for the 17 groups of borrowers: (i) average interest rates (unconditional on debt and conditional on the population-wide median level of debt, which we proxy by the middle quintile); (ii) share of total debt by income level and default risk; (iii) average debt (unconditional and conditional on having a balance); and (iv) debt to income. Again, we average over all 12 months of 2019. We compute the share of debt held by each bin as the ratio of the total balances in that bin to the total balances across all accounts. The construction of model moments is described in Section A.5. For robustness, we also compute these moments by income quartiles. Lastly, we recompute these objects for the March-June 2020 period.

For interest rates, we use the purchase APR for accounts that are not in default / workout, for which the default / workout APR is used. To average, we weight this APR by the total outstanding balance for the account at the end of each month's cycle. We also recompute this average interest rate conditional on observations with balances near the median level of debt (40th to 60th percentile). For robustness, we compute interest rates controlling for smaller debt levels with more borrowing: patterns look similar. Aggregation is over FICO score bins and then

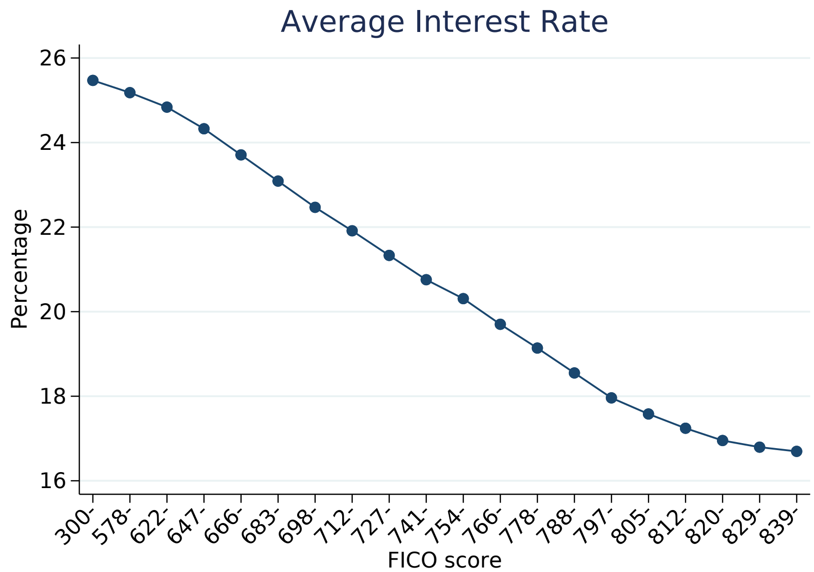

Figure A.2: **Interest rates by FICO scores**

Notes: This figure shows the interest rate schedule across bins. Note that figures here are reported in raw percentage points, not spreads as in the main text.

over time. Figure A.2 shows the interest rate schedule across bins. Note that figures here are reported in raw percentage points, not spreads as in the main text.

Borrower credit risk We use as the common identifier borrower's credit risk type, as measured by the j^{th} vigintile (5% bin) of credit scores, $j = 1, ...20$. For our analysis, we use the proprietary credit score developed by Equifax ("Risk Score") and the FICO score available in Y-14M. The Risk Score has been used extensively in household finance research by academics and policymakers given its inclusion in the Equifax data set. The FICO score is the most commonly used credit score by lenders when providing and setting terms of credit. We map each vigintile of Risk Score in Equifax into its corresponding vigintile of FICO score in Y-14M. As argued in Section 1, this is reasonable because the two measures are equivalent in assessing likelihood of default. We provide further details on this here.

Both credit ratings integrate several common sources of information about consumers into a single score. Debt payment history is the most important determinant of one's credit score, but other factors matter as well: levels of indebtedness, length of credit history, and credit limit utilization, among others. This information is collected by Credit Reporting Agencies (CRA) such as TransUnion, Experian, and Equifax in order to create detailed credit reports. Second, while numerous credit scoring models have been constructed in the past decades, they all use the CRA data to generate credit scores that can be queried by lenders and borrowers. Previous research has documented that these different credit scoring models provide largely similar information about the creditworthiness of consumers. For example, as the 2012 Consumer Financial Protection Bureau (CFPB) report shows, correlations across different credit scores are high (over 0.9), especially for consumers with lower credit scores.

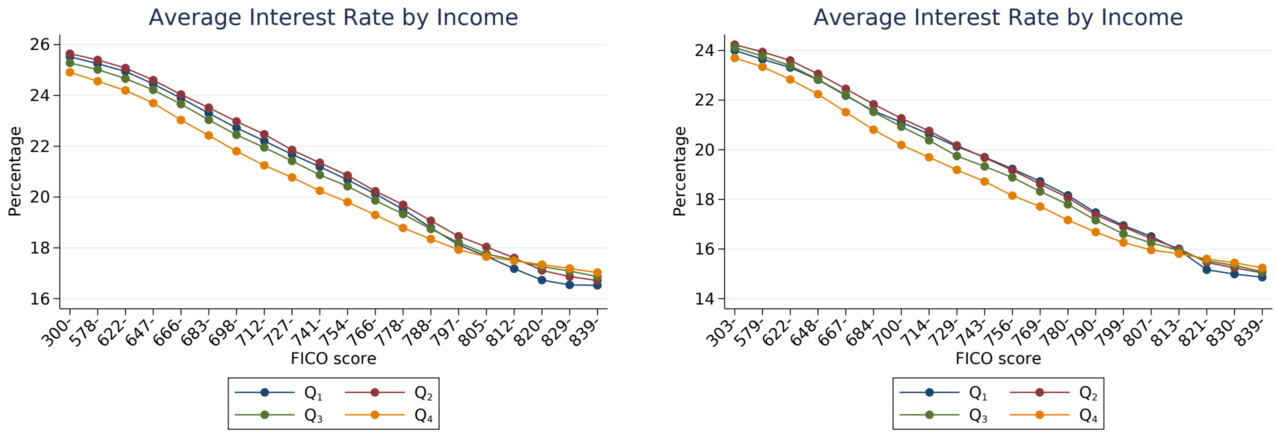

(a) **Full year 2019**

(b) **Covid crisis: March - June 2020**

Figure A.3: **Interest rates conditional on income**

Moments conditional on income Additionally, we recompute interest rates across credit scores also controlling for income and report findings in Figure A.3. The patterns in Figure A.2 are preserved when controlling for income. While borrowers with higher earnings generally face lower rates, the differences across income groups are small. Unlike credit variables, income is self-reported. Y-14M includes both income at origination and updated income at the end of the cycle for each month. While the latter variable is more desirable to use all else equal, it is not well populated and so we use the former. We restrict attention to accounts originated after 2005 to keep a reasonable level of informational value for borrowers' reported income. This trims only 7% of the sample and covers 90% percent of balances. We winsorize the income variable at 2% at the bottom and 1% at the top, and use CPI to convert income into 2019 dollars. For the remaining variables used for model validation, we proceed in the same way. Lastly, we turn to the Covid crisis and recompute patterns using March-June 2020 in Y14M. Figure A.3b reports the schedules of interest rates and credit limits in the crisis period conditional on income. In general, interest rate schedules shifts upward and credit limit schedule downward across credit score and income groups, albeit by small margins.

A.2 Alternative regression specifications

We run three alternative sets of regressions to complement Section 1.3. First, we run the specification of column [1] from Table 1 using income as the only independent variable to drive home the point that income itself has little explanatory power; the R^2 of this regression is 0.007. Adding FICO to this specification returns us to the specification in column [2] in Table 1.[57] Second, in Panel A of Table A.1, we run a specification in which we include a quadratic FICO term to capture the curvature in the borrowing premium schedule in Figure 2. While this increases the explanatory power by 14.6% (R^2 increases from 60.7% to 69.6% in our baseline specification from column [1]), the key findings from the main text are unchanged.

[57] Full results for this specification are available upon request; we omit them here for brevity.

Panel A: dependent variable: borrowing premium (quadratic)

	[1]	[2]	[3]	[4]	[5]	[6]
FICO score	0.642	0.645	0.629	0.632	0.632	0.634
	(2e-4)	(2e-4)	(2e-4)	(2e-4)	(2e-4)	(2e-4)
FICO^2	-5.2e-6	-5.2e-6	-5.3e-6	-05.2e-6	-5.2e-6	-5.2e-6
	(2e-9)	(2e-9)	(2e-9)	(2e-9)	(2e-9)	(2e-9)
income		-0.002		-0.003	-0.003	-0.004
		(2e-5)		(3e-5)	(3e-5)	(3e-5)
revolver			X	X	X	X
multiple relationships			X	X	X	X
new account			X	X	X	X
quarter FE					X	X
bank FE						X
R^2	0.696	0.697	0.699	0.699	0.699	0.710

Panel B: dependent variable: interest rate

	[1]	[2]	[3]	[4]	[5]	[6]
FICO score	-0.141	-0.139	-0.145	-0.143	-0.143	-0.137
	(1e-4)	(1e-4)	(1e-4)	(1e-4)	(1e-4)	(1e-4)
income		-0.001		-0.002	-0.002	-0.003
		(2e-5)		(2e-5)	(2e-5)	(2e-5)
revolver			X	X	X	X
multiple relationships			X	X	X	X
new account			X	X	X	X
quarter FE					X	X
bank FE						X
R^2	0.146	0.147	0.168	0.170	0.170	0.283

Table A.1: **Additional regression specifications**

Notes: Number of observations is 14,426,760. Results are significant at 1 percent level. All specifications are variations on estimating equation (4). Column [1] imposes $\beta_2 = \gamma_1 = \gamma_2 = \gamma_3 = 0$; [2] imposes $\gamma_1 = \gamma_2 = \gamma_3 = 0$; [3] imposes $\beta_2 = \gamma_2 = \gamma_3 = 0$; [4] imposes $\gamma_2 = \gamma_3 = 0$; and [5] estimates all coefficients.

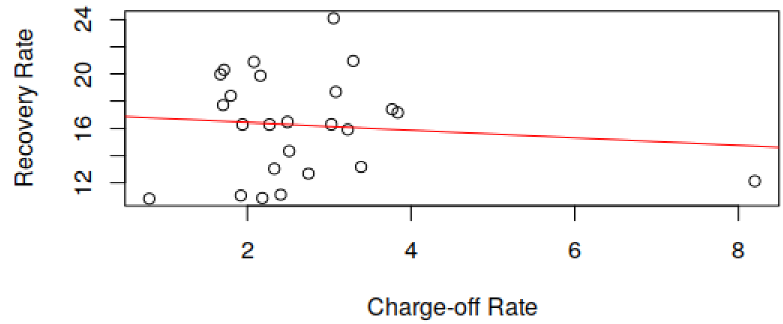

Figure A.4: **Recovery and charge-off rates across banks**

Notes: This figure plots recovery rates against charge-off rates for the 25 banks comprising our sample from the Call Reports which we use to estimate recovery rates.

Third, in Panel B of Table A.1 we repeat our analysis replacing the borrowing premium b_{it} with the interest rate R_{it} in our regression equation (4) to provide a more comprehensive view of evidence regarding credit card pricing. Our primary takeaways go through in this modified analysis with three caveats. First, overall explanatory power is much lower, with R^2 declining by about 75% overall relative to the analogous specifications in Table 1. Second, adding borrower and account controls has a slightly larger contribution to overall explanatory power, increasing the R^2 by over 2 pp as opposed to under 0.1 pp. Third, bank fixed effects explain a lot of the residual variation in interest rates unexplained by the other variables of interest.

A.3 Recovery rates

Empirical measure and data We compute recovery rates for each bank-quarter pair for the top 25 banks by assets (98 percent of credit card assets in 2019) as the ratio of total recoveries to gross charge-offs using Call Report data. We then weight across banks using the bank's share of credit card assets in the quarter, and then average across all quarters in the sample. Since recoveries may lag charge-offs, we also compute the recovery rate in two additional ways for robustness. First, we use an annual rather than quarterly basis, comparing the sum of recoveries to the sum of charge-offs for each year. Second, we use a 3 period right-aligned rolling quarterly sum for recoveries and a 3 period left-aligned rolling sum for charge offs. In all these exercises we find values ranging from 15.7% to 16.3%, and our chosen ξ parameter is the midpoint of this range.

Our borrowing premia calculations assume that the recovery rate does not vary with borrower risk. While borrower-level recovery rate data is not available, we use the Call Reports to document that there is no correlation between recovery and charge-off rates at the bank level (see Figure A.4). This lack of correlation suggests that banks that might have credit card portfolios with systematically different risk levels do not experience systematically different levels of recoveries, which is consistent with our (stronger) borrower- or bin-level assumption.

Could different recovery rates deliver a fixed borrowing premium? Our measurement approach in Section 1.1 assumes a fixed recovery rate across all loans, and as a result we obtain

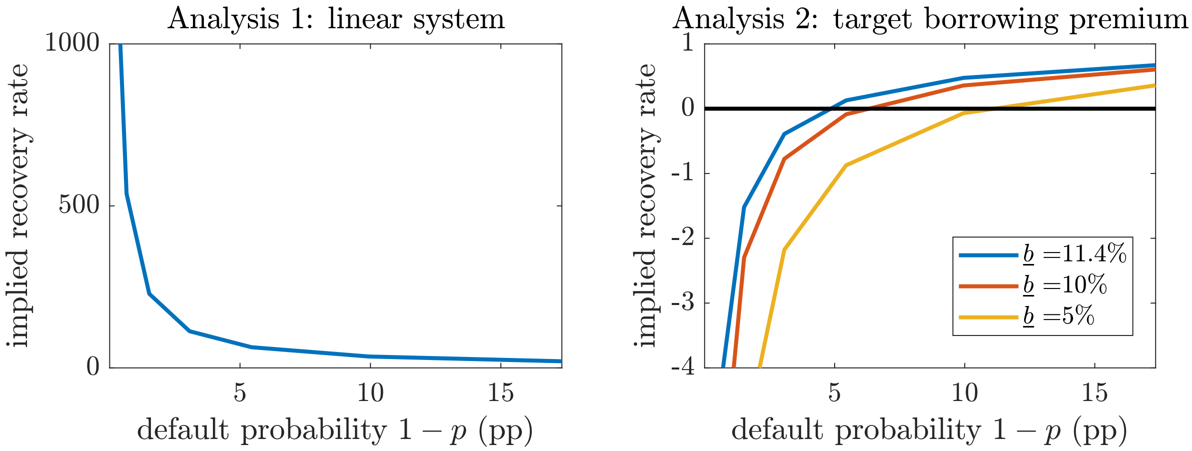

Figure A.5: **Recovery rates and borrowing premia**

Notes: This figure is constructed according to the analyses described in Appendix A.3.

the heterogeneous schedule of borrowing premia depicted in Figure 2. While we have argued using data available to us that the assumption of a fixed recovery rate is a reasonable one, we would like to examine the question further given its centrality to our paper. To this end, in this section we present two analyses which help address this question.

Our first analysis asks: are our borrowing premia results explainable by omitted differences in recovery rates across credit score bins? We can formalize this question in a slightly different way by asking if there is a set of recovery rates (ξ_i for all bins i) for which our measurements would imply a fixed borrowing premium ($b_i = \bar{b}$ for all bins i). For completeness, we require that this borrowing premium match the balance-weighted average we obtain in the data, $\bar{b} = \sum_i s_i b_i$, where s_i is bin i's share of total credit balances. Using equations (2) and (3), we can summarize these restrictions by the system of linear equations

$$\begin{aligned} b_i &= R_i \left[p_i + (1-p_i)\xi_i\right] - 1 \text{ for all } i \\ \bar{b} &= \sum_i s_i b_i \\ b_i = \bar{b} \text{ for all } i \implies 0 &= \sum_i (1-s_i) R_i \left[p_i + (1-p_i)\xi_i\right] \text{ for all } i \end{aligned}$$

where N is the number of bins. This is a system of N equations in N unknowns, $\{\xi_i\}_{i=1}^N$ which eliminates b_i. Therefore, we can use data on interest rates, repayment probabilities, and shares $\{R_i, p_i, s_i\}_{i=1}^N$ to solve this system. The results of this exercise are plotted in the left panel of Figure A.5. Immediately, we see that the implied recovery rates are completely implausible, far exceeding the natural upper bound of 1 for all but the highest risk loans. We can further test this possibility by comparing the balance-weighted average borrowing premium implied by this inferred set of recovery rates to its empirical counterpart; the former is 416%, while the latter is 11.3%. Both findings lead us to reject the possibility that all variation in borrowing premia is explainable by recovery rates.

Second, we can ask a more specific question targeting the especially puzzling portion of our borrowing premium schedule: does there exist a schedule of recovery rates that explains the low borrowing premia for high risk borrowers? To address this question, we can consider the following exercise. First, assume the borrowing premium for the lowest risk bin is effectively independent of the recovery rate, since $p_i \approx 1$ for this group. To be as conservative as possible in this exercise, assume further that these agents have full recovery in default so that their loans are riskless. Then, their borrowing premium is $\underline{b} \equiv R_i - 1$ by equation (3), since $\tilde{R}(p_i) = 1$ for them by equation (2). Second, we can ask what ξ_i would rationalize $b_i = \underline{b}$ for each bin i? Rearranging equation (3), it must be

$$\xi_i = \frac{1 + \underline{b} - R_i p_i}{R_i(1 - p_i)}$$

The results of this analysis are presented in the right panel of Figure A.5. Since the first portion of this algorithm is in principle optional, we report a baseline set of estimates corresponding to the empirical $\underline{b} = 11.4\%$, and two additional sets corresponding to lower alternatives, $\underline{b} = 10\%$ and $\underline{b} = 5\%$. For default probabilities below 5%, the implied recovery rate would have to be negative to rationalize observed borrowing premia. Furthermore, for higher risk loans, we'd need implausibly high recovery rates, on the order of 30-50%, which our Call Report analysis rules out. For these reasons, we also reject that recovery rates can explain our estimates of borrowing premia.

A.4 Senior Loan Officer Opinion Survey (SLOOS)

A.4.1 Construction of diffusion indices

To construct an index for changes in standards and terms of credit we follow the methodology in Bassett et al. (2014) and use questions that ask participating banks to report whether they have changed their standards, or changed terms during the survey period.[58] Specifically, questions about changes in standards follow the general pattern of "Over the past three months, how have your bank's credit standards for approving credit card loans changed?" The possible answers, on a 1-5 scale, are: (i) eased considerably; (ii) eased somewhat; (iii) about unchanged; (iv) tightened somewhat; and (v) tightened considerably.

Similar questions are asked for credit terms (interest rate spreads, credit limits), and the approach we describe below can be replicated for these specific terms as opposed to standards more broadly. Historically, SLOOS respondents rarely characterize their standards as having changed "considerably." Therefore, we collapse the scale as follows: $S_t^i = -1$, if bank i reported easing standards in quarter t, $S_t^i = 0$, if bank i left standards unchanged in quarter t, and $S_t^i = +1$ if bank i reported tightening standards in quarter t.

[58] In constructing our indexes, we revisit the method used in Bassett et al. (2014) in two directions: first, we use more granular data by loan types when we compute the share for each loan category on banks balance sheet; second, when computing weights associated with each type of loan, we expand the universe of banks beyond respondents in the SLOOS in line with the Call Reports data.

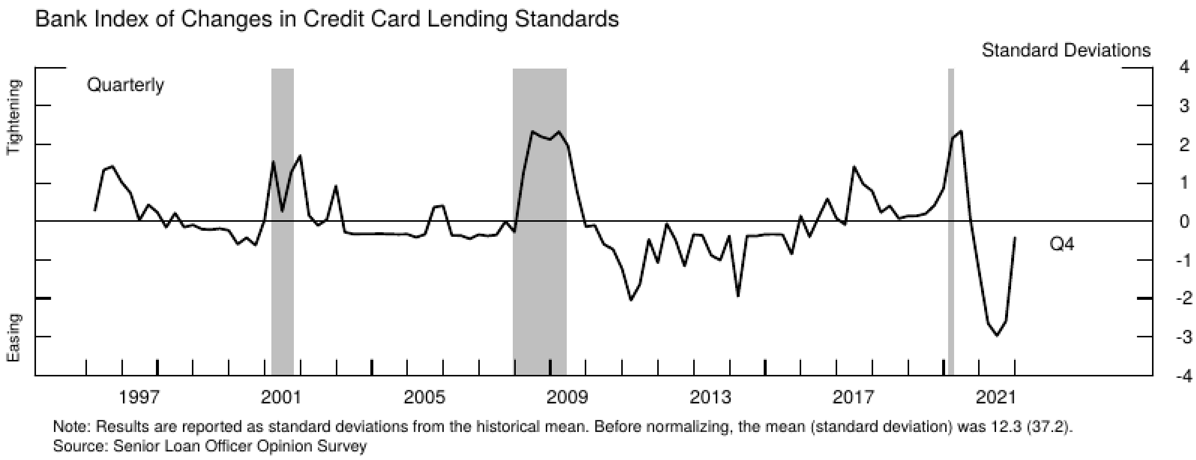

(a) **Lending standards**

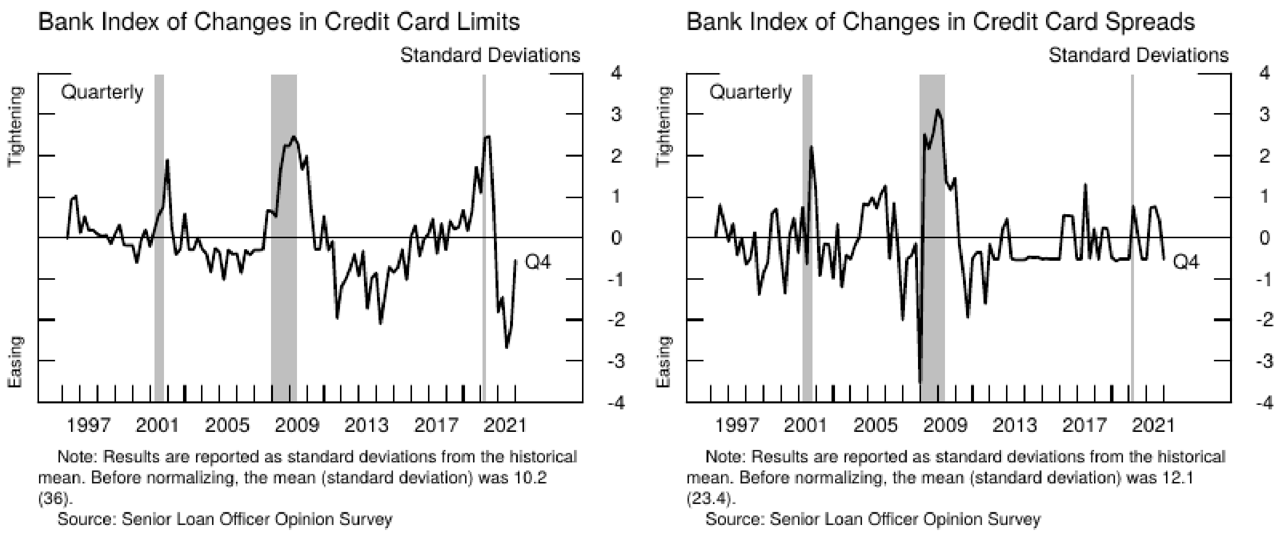

(b) **Terms (limits and spreads)**

Figure A.6: **Diffusion indices for credit card loans**

We weight a bank's responses by that bank's share of total consumer loans in the previous period ($w_{i,t-1}$) to obtain the aggregate measure of changes in lending standards, $\Delta S_t = \sum_i w_{i,t-1} S_t^i$. These weights are computed using the Call Reports. These indices reflect the net percent of consumer loans subject to tightened standards: positive (negative) values indicate eased (tightened) standards. We normalize these aggregate measures by their historical average to create an overall lending standards index, I_t^S, which measures standard deviations in each quarter t from its historical average. Figure A.6a in the main text shows the index for lending standards for consumer loans, and Figure A.6b shows the indices for spreads and limits as well. The changes in standards for consumer loans and for credit card loans are quite similar, since credit card loans represent the largest category of loans included.

A.5 Moment definitions

In this section, we define key moments and metrics used throughout the paper, mainly model moments used in our model versus data comparison, that are not described in the main text. To ease notation, we express the individual state vector as x and describe a specific element of the state vector as a function of x. Except where necessary, we suppress explicit dependence on the time period t or the aggregate state s.

Macro moments The macro moments we use throughout the paper are standard. For business cycle moments, we take logs of the data and apply the HP filter with an annual smoothing parameter of 6.25. We use Equifax for all credit quantities for which we present cyclical properties, except for debt to income for which we use Y-14M. For interest rates on all credit card loans we use schedule G19 from the Flow of Funds and for interest rates on risk free loans we use Y-14M. We use 1999-2019 for series constructed from Equifax and G19, given that 1999 is the first year available in Equifax, whereas for series constructed from the Y14M data, we start only in 2014, the first year available.

Fraction in debt, volume of debt, and share of total debt The total volume of debt is: $-\sum_x \mu(x) g_a(x)(1-g_d(x))\mathbf{1}[g_a(x)<0]$. The share of total debt accounted for by a subgroup applies the above metric to the subgroup, then weights by the total. In the data, we define credit as the total credit card debt balances in Equifax, excluding bankruptcy, the exact counterpart of our model. Recall that in the event of a bankruptcy, the household cannot borrow or save in the current period.[59] We define the fraction of agents in debt to be the fraction of agents choosing debt, $a'<0$. For a given period t, this fraction is $\sum_x \mu(x)(1-g_d(x))\mathbf{1}[g_a(x)<0]$. In the data, we compute this fraction in Equifax as the fraction of consumers that hold positive credit card debt.

Debt to income ratio We compute the debt to income ratio conditional on borrowing: the ratio of an agent's total debt ($-g_a(x)\mathbf{1}[g_a(x)<0]$) divided by total labor income ($w\epsilon_1\epsilon_2\epsilon_3$). We normalize by the fraction in debt, defined above. In the data, we use Y-14M at the account level on credit card debt outstanding at the end of balance and income at origination (recall that income at the end of balance is an optional reporting variable in the data and quite sparse and so we are limited to using income at origination). Although imperfect, given that income is highly persistent we believe that the properties of this series represent a reasonable proxy for current debt to income measure.

[59] Alternatively, one can use Total Real Revolving Consumer Credit Outstanding from the Flow of Funds/G.19 data, Federal Reserve Board (FRB)), case in which credit cannot be adjusted for bankruptcy at the individual level as Equifax data allow us to do. If the stock of revolving credit is adjusted using the credit card bankruptcy or charge-off rate, the correlation with GDP is negative, otherwise it is positive. Furthermore, while the large volatility is robust relative to the time periods used, this is not the case for the correlation. For instance, in the subperiod 1980-2018 the correlation is 0.33 as noted by Nakajima and Ríos-Rull (2019). Fieldhouse et al. (2016) point out that over the 1993-2006 period, such correlation was negative. In line with our paper, they also point out that this correlation is even more negative when adjusted for charge-offs.

Bankruptcy and charge-off rates The bankruptcy rate is $\sum_x g_d(x)\mu(x)$; the charge-off rate is $\sum_x -a g_d(x)\mu(x)$, normalized by the total debt in the economy. To compute a conditional metric (either on being in debt or by income group), we normalize the distribution by the desired population. In the data, we use the exact same definition of default rate from Equifax used for default probability in the model as described in Section 1 (bankruptcy and severe derogatory) as well as using only the bankruptcy flag in these data to compute cyclical properties for default and bankruptcy rates in Section 6. Constructed variables in this way are based on the universe of consumers modeled in our paper.[60] The share of suboptimal bankruptcy is the total mass of bankruptcy coming from agents for whom bankruptcy is not the modal action, $\sum_x g_d(x)\mu(x)\mathbf{1}[g_d(x) < 0.5]$, divided by the total bankruptcy rate.

Average interest rate and borrowing premium We compute the loan-weighted average interest rate spread conditional on borrowing as

$$\frac{\sum_x g_a(x)(1-g_d(x))\mathbf{1}[g_a(x)<0]\left(\frac{\overline{q}}{q(g_a(x);x)}-1\right)\mu(x)}{\sum_x g_a(x)(1-g_d(x))\mathbf{1}[g_a(x)<0]\mu(x)}$$

In steady state, the denominator is equal to the (constant) total debt, but this is not in general true when aggregate shocks are present in the environment. For business cycle calculations, we use G19 data from FRB for interest rates of all credit card accounts which we adjust for the federal funds rate. As in the case of the interest rate on all loans, we adjust for the prime rate. The loan-weighted average borrowing premium is computed exactly as in the expression above, with the spread term $\overline{q}/q(g_a(x);x)$ replaced by the borrowing premium $b(p(g_a(x);x))$.

B Quantitative Appendix

B.1 Estimating the schedules of borrowing premia

Figure B.1 depicts the fit of our model borrowing premia to the data for our baseline steady state model (data corresponding to FY 2019) and for our dynamic analysis (impact of shock corresponding to 2020 March to June). As described in the main text, the fit is nearly exact for both schedules for a 20th order polynomial. We choose 20th order because lower orders struggle to match the curvature for very low levels of default risk, and higher orders do not increase the fit by much. The polynomial specification for $N = 20$ is

$$b(p) = \begin{cases} \sum_{n=0}^{N} x_n \left(\frac{p-m_0}{m_1}\right)^n & \text{if } p \geq \underline{p} \\ 0 & \text{if } p < \underline{p} \end{cases} \tag{B.1}$$

The coefficients x_i (as well as normalization constants m_0 and m_1 and the threshold $\underline{p}$) for these polynomials are presented in Table B.1. The truncation at $\underline{p}$ accounts for scarcity of data at high

[60] Alternatively, one can use the number of bankruptcies (obtained from U.S. Courts) normalized by the number of households (from the Census), as in Nakajima and Ríos-Rull (2019), with properties of these objects presenting similar patterns. In particular, bankruptcy filings are also highly volatile and countercyclical. Specifically, they find that for all bankruptcies relative volatility is 8.5 and correlation with GDP is -0.47.

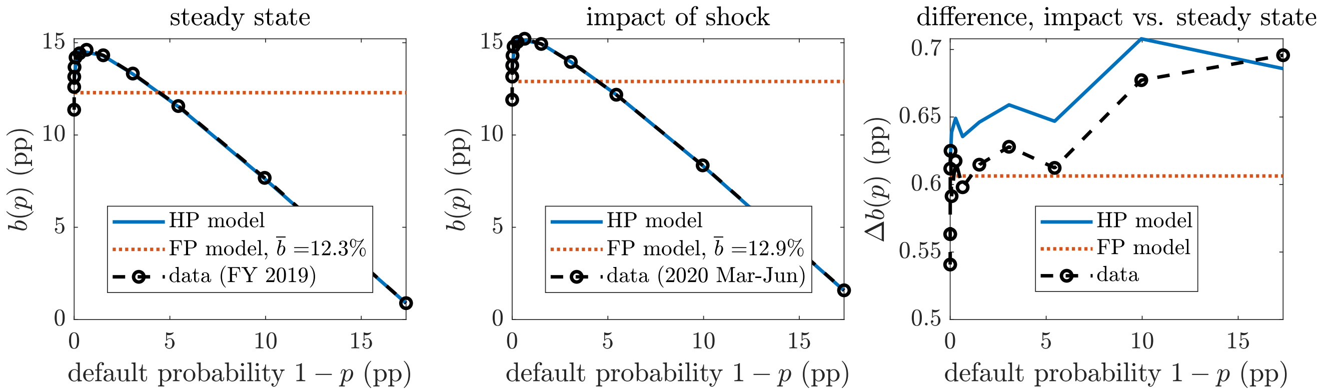

Figure B.1: **Model fit to observed borrowing premia schedules**

coef.	**20**	**19**	**18**	**17**	**16**	**15**	**14**	**13**	**12**	**11**
FY2019	-0.002	-0.002	0.030	0.028	-0.176	-0.153	0.565	0.459	-1.088	-0.811
2020Q2	-0.003	-0.002	0.035	0.023	-0.204	-0.126	0.653	0.375	-1.251	-0.657

coef.	**10**	**9**	**8**	**7**	**6**	**5**	**4**	**3**	**2**	**1**
FY2019	1.279	0.855	-0.895	-0.517	0.346	0.160	-0.063	-0.019	0.002	0.048
2020Q2	1.464	0.686	-1.022	-0.408	0.395	0.121	-0.072	-0.013	0.002	0.047

coef.	**intcpt.**	m_0	m_1	$\underline{p}$
FY2019	0.083	0.907	0.053	0.815
2020Q2	0.089	0.907	0.053	0.815

Table B.1: **Polynomial fits to borrowing premia data**

Notes: The polynomial specification is given by equation (B.1).

levels of default risk. The m coefficients allow for higher precision in the polynomial fitting by rescaling the probabilities of default, which cluster near 0. As described in the main text, the fixed premium is chosen to be equal to the balance-weighted average from the variable premium model, which is 12.3% for the steady state and 12.9% on impact of the shock.

B.2 Sensitivity and Jacobian Analysis

This subsection complements Section 3.3 by providing quantitative results in Table B.2 and a fuller discussion of the mappings: (i) from model parameters to model moments (Jacobian analysis); and (ii) from empirical moments to estimated parameters (sensitivity analysis). We conduct these analyses for both the baseline (HP) and fixed premium (FP) economies, focusing on the former and highlighting key differences of the latter with respect to the former. We begin with the Jacobian analysis and then move to sensitivity.

B.2.1 Jacobian analysis

Overview Each column of panels A.2 and B.2 shows the elasticity of the indicated model moment (column) with respect to the indicated model parameter (row), computed as the ratio of log differences between moment and parameter. Each elasticity is computed by perturbing the indicated model moment by 0.1% from its estimated value in Table 2.

By column: which model parameters affect each model moment? First, consider panel A.2 of Table B.2 column-by-column to assess which model moments are most affected by which parameters in the baseline model. The bankruptcy rate (BKR) is most sensitive to the average discount factor: a 1% increase in $\overline{\beta}$ yields nearly a 7.5% lower bankruptcy rate. The default preference parameters χ and ζ and the share of high beta types μ_{β_H} also impact this moment. The fraction in debt (FID) is sensitive to parameters describing the extent of differences between types and the switching process between them, with the largest elasticities being with respect to μ_{β_H} and the low-to-high transition or "churn" parameter, $\Gamma^{\beta}_{LH'}$. This follows from the GE nature of the model as described in the main text. A bigger gap between the types increase FID, while a lower share of β_H or more churn in types decreases it. The debt to income ratio (DTI) is quite sensitive to $\overline{\beta}$ (more patience means borrowers can credibly take on more leverage with similar repayment probabilities) and $\Gamma^{\beta}_{LH'}$ (more churn means current type is a worse signal of future type, leading to rationing). Average spreads depend negatively on $\overline{\beta}$ since less default implies lower rates and positively on μ_{β_H} since more high types means more borrowing is accounted for by low types. The capital-output ratio is driven largely by β process in intuitive ways. The charge-off rate (COR) responds negatively to increases in overall type quality ($\overline{\beta}$ and μ_{β_H}) and the default preference parameters; an increase in the utility cost of default χ lowers COR by limiting filings on the margin, while an increase in the extreme value scale parameter ζ increases COR by increasing the likelihood that agents with relatively low values attached to default will choose to do so anyway. The share of suboptimal bankruptcies (SoBKR) responds in the same way to χ and ζ, while increases in the average level of β strongly increase the SoBKR. This last effect arises because it increases the weight on continuation values in individual value functions, which are in general the part of $V_0^{BK}(\cdot)$ that depresses $V_0^{BK}(\cdot)$ relative to the repayment value $V_0^{R}(\cdot)$.

By row: which model moments are most affected by each model parameter? Next, consider panel A.2 once more, this time looking row-by-row to highlight which moments are most affected by a given parameter. The average discount factor plays a central role in determining nearly every model moment. Beyond that, μ_{β_H} and the gap between types $\Delta\beta$ have big impacts on all three default moments and the capital-output ratio. The churn parameter $\Gamma^{\beta}_{LH'}$ has the biggest impacts on the credit quantity moments FID and DTI. The stigma and extreme value scale, unsurprisingly, exert their largest influences on all three default moments.

Moment		BKR	FID	DTI	IRS	K/Y	COR	SoBKR
Panel A: Heterogeneous Premia (Baseline, HP) Model								
A.1: Sensitivity Matrix (AGS 2017)								
avg discount factor	$\overline{\beta}$	-0.00	-0.03	0.02	0.09	-0.02	-0.02	-0.01
diff in disc. factor	$\Delta\beta$	-0.01	1.91	-1.33	-0.40	0.44	0.83	0.09
high β share	μ_{β_H}	0.02	-1.88	1.20	-0.20	-0.31	-0.52	-0.17
low to high β prob	$\Gamma^{\beta}_{LH'}$	0.01	1.54	-0.05	0.19	0.36	0.78	0.08
utility cost of def.	χ	-0.03	2.77	-1.79	-0.09	0.05	0.05	0.17
EV scale	ζ	-0.05	2.47	-1.76	-0.51	-0.48	-1.31	-0.02
A.2: Jacobian Matrix								
avg discount factor	$\overline{\beta}$	-7.42	-0.02	4.13	-7.16	9.43	-7.42	37.8
diff in disc. factor	$\Delta\beta$	-0.72	0.11	0.64	0.36	-0.88	-0.80	5.40
high β share	μ_{β_H}	-2.29	-0.35	-0.20	0.61	-1.26	-2.50	3.97
low to high β prob	$\Gamma^{\beta}_{LH'}$	-0.34	-0.70	-1.06	0.08	0.00	0.27	0.53
utility cost of def.	χ	-0.99	-0.09	0.24	-0.12	-0.06	-2.61	-2.17
EV scale	ζ	1.04	-0.21	-0.23	0.18	-0.05	1.41	1.36
Panel B: Fixed Premium (FP) Model								
B.1: Sensitivity Matrix (AGS 2017)								
avg discount factor	$\overline{\beta}$	-0.03	0.14	-0.06	0.29	0.10	-0.05	-0.01
diff in disc. factor	$\Delta\beta$	-0.27	0.97	-0.33	-1.61	-1.31	0.02	0.37
high β share	μ_{β_H}	-0.09	0.30	0.09	-0.50	-0.38	0.13	0.08
low to high β prob	$\Gamma^{\beta}_{LH'}$	-0.46	-1.01	1.53	-2.25	-2.05	0.82	0.21
utility cost of def.	χ	0.60	-0.26	-1.05	2.69	2.82	-0.18	-0.53
EV scale	ζ	0.66	-0.07	-1.14	3.22	3.00	-0.67	-0.57
B.2: Jacobian Matrix								
avg discount factor	$\overline{\beta}$	-14.8	-0.53	5.29	-6.02	8.46	-4.43	18.3
diff in disc. factor	$\Delta\beta$	20.0	0.32	4.34	-0.39	1.03	0.75	17.1
high β share	μ^{β}_{H}	-4.08	-3.70	-5.03	2.35	-3.61	0.16	-0.25
low to high β prob	$\Gamma^{\beta}_{LH'}$	19.3	0.06	3.06	-0.42	0.98	0.63	18.5
utility cost of def.	χ	-1.86	0.47	0.87	-0.51	-0.03	-2.55	-4.03
EV scale	ζ	21.1	-0.70	2.27	0.54	0.26	3.24	22.7

Table B.2: **Sensitivity and Jacobian for Baseline and Fixed Premium Estimations**

Notes: Notation: BKR = bankruptcy rate; FID = fraction in debt; DTI = debt to income ratio; IRS = average interest rate spread; $\frac{K}{Y}$ = capital-output ratio; COR = charge-off rate; and SoBKR = share of suboptimal bankruptcy. The sensitivity measure is a transformation of the Λ matrix from Andrews et al. (2017). All figures are numerical elasticities, $d\log x/d\log y$. For panels A.1 and B.1, each cell is the elasticity of the estimated parameter (row) with respect to a 1% increase in the indicated empirical moment (column), filtered through the Λ matrix. For panels A.2 and B.2, each cell is the elasticity of the indicated model moment (column) with respect to the indicated parameter (row).

Comparing the FP model to the HP model Finally, compare panel A.2 (for the HP model) to panel B.2 (for the FP model). For this analysis, it is useful to recall that the HP and FP calibrations differ in important ways described in Section 3.3. The average discount factor plays the same centrally important role in both models. Furthermore, the two default preference parameters and the share of high β types affect model moments in much the same way across the two models. There are, however, two important differences. First, all default moments in the model are far more sensitive to the scale parameter ζ in the HP model than the FP model because (i) the scale parameter is already 15% lower in the FP model, so a 0.1% increase in ζ is a larger increase in actual "noisiness" and (ii) the types are worse in the FP model, leading them to downweight continuation values relative to the baseline model and therefore be more "susceptible" to noise from the extreme value shocks. Second, the gap between types, share of high types, and low-to-high transition probability have slightly different and generally much larger affects on model moments in the FP model than the baseline. For example, while the elasticity of DTI and FID with respect to $\Delta\beta$ is positive in both models, the analogous elasticities of IRS are positive for the HP model and negative for the HP model. This is because predominant effect of making low types worse in the baseline model is that the pay higher spreads in equilibrium, while in the FP model the predominant effect is that these types are priced out on the margin, improving the pool of borrowers.

B.2.2 Sensitivity analysis

Overview Given a J-vector of target moments m and a P-vector of parameters θ to estimate, Andrews et al. (2017) show that: (i) the $P \times J$ sensitivity matrix $\Lambda = -(G'WG)^{-1}G'W$ may be computed using only the $J \times P$ Jacobian matrix G and the $J \times J$ weighting matrix W from the initial estimation; (ii) given Λ and a $J \times J$ empirical variance-covariance matrix Ω, one may compute standard errors as the diagonals of the $P \times P$ matrix $V = \Lambda\Omega\Lambda'$; and (iii) for a J-vector of constants c corresponding to possible "mismeasurements" in m, the first-order asymptotic bias in the estimator of θ is $B(c, \Lambda) = \Lambda c$. For a given Ω, then, small values of Λ correspond to more precise parameter estimates. A small value of Λ_{ij} indicates that even if empirical moment j (m_j) were quite different, the model's estimate of parameter i (θ_i) would change little. Therefore, if the (appropriately-weighted, squared) sum across all m_j of Λ_{ij} for θ_i is small, then θ_i is precisely estimated given the set of targets $\{m_j\}$. By contrast, a large value of Λ_{ij} suggests that the estimate of θ_i is quite sensitive to a change in m_j. This implies that: (i) variation in m_j is critical for determining the value of θ_i; and (ii) the empirical variance of m_j and its covariances with other moments must be relatively low in order for θ_i to be estimated precisely.

Panels A.1 and B.1 of Table B.2 present a transformation of the Λ matrix described above for the baseline and FP models, respectively. Specifically, we posit a 1% mismeasurement in each empirical moment (i.e. $c = 0.01 \times m$). Then, using the Jacobian matrix G (whose computation is described above), we compute the bias matrix $B(c, \Lambda)$. Then for each moment m_j and parameter θ_i, we compute the numerical elasticity $[\log(\hat{\theta}_i + B_{ij}) - \log \hat{\theta}_i]/[\log(1.01 m_j) - \log(m_j)]$.

By column: how do parameter estimates change with each empirical moment? First, consider panel A.1 column-by-column. For this analysis, it is useful to consider all parameter changes induced by the change in measured empirical moment *simultaneously*. For example, a 1% increase in interest rate spreads would lead us to estimate a slightly higher $\overline{\beta}$, which is perhaps counterintuitive since households are on average more patient. Tracing through the calculations, though, this would be achieved by modestly lowering both β_L and β_H (β_L more so than β_H) while increasing the share of β_H types, the persistence of the β_L type, and the likelihood of transition from β_H to β_L. This last effect is quite large and makes high types fundamentally riskier, increasing the spreads they pay when they borrow. Importantly, this analysis reflects how parameters would adjust to the mismeasurement in spreads *while still matching the other moments at their targeted levels*.

All bankruptcy rate elasticities are 5 bps or below, with the largest magnitudes unsurprisingly for the default preference parameters χ and ζ. Increasing FID would suggest a wider gap between the types and also a lower share of high types, as well as revising up both default parameters. Increasing DTI would have the exact opposite directional effects, suggesting that targeting both moments is useful in pinning down this set of parameters. One of the most important things learned comes from comparing the effects of changes in the charge-off and suboptimal bankruptcy rates on χ and ζ; an increase in COR has a large downward effect on the EV scale parameter and little effect on stigma, while and increase in SoBKR has a large positive effect on stigma and little effect on ζ. These competing forces underscore how these moments combine to separately pin down these two crucial parameters.

By row: which empirical moments drive the largest parameter changes? The average discount factor is not very sensitive to small perturbations of any model moment. Most parameters that we calibrate are more sensitive to the fraction in debt and the debt to income ratio than any other model moments. This is consistent with how these moments pin down the balance between borrowers and savers in our general equilibrium context. Notably, while all our parameters are not very sensitive to the bankruptcy rate – a frequency measure of default –, many are sensitive to the charge-off rate – a dollar-weighted measure. This suggests that there is important additional information from the set of loans on which borrowers default above and beyond the raw share.

Comparing the FP model to the HP model We now compare panels A.1 and B.1, the sensitivity matrices for the baseline and FP models, respectively. While some of the same principles apply (e.g. the average level of β is the least sensitive parameter to virtually all moments), we see considerable differences across these matrices owing to the different parameter estimates and the different economic mechanisms at play in the two models. For example, an increase in FID in the baseline model leads to: a lower share of high β types, more frequent β_L to β_H transitions, and higher stigma and EV scale parameters. By contrast, in the FP model an analogous increase leads to an increase in the high β share, less frequent transitions, and modestly lower stigma and scale parameters. While the FID moment offers a particularly stark example, there are other, smaller differences across the models as well.

B.3 Additional quantitative results

Equilibrium distributions of wealth and income We compare our models' distributions of wealth (panel (a)) and labor income (panel (b)) to those in the data (Survey of Consumer

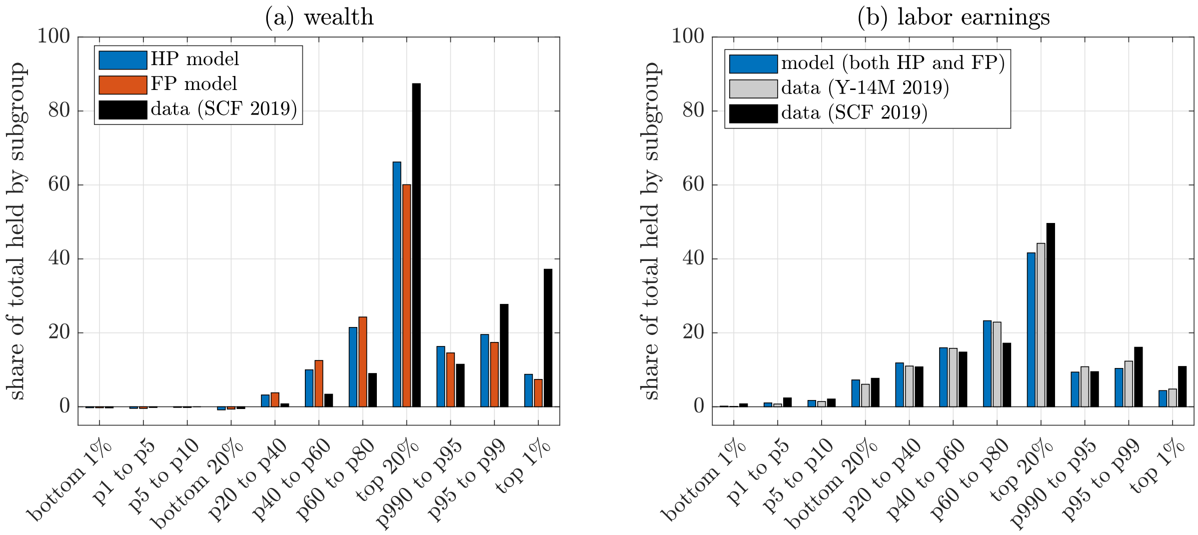

Figure B.2: **Distribution of wealth and labor earnings**

Notes: This figure reports the steady distributions of wealth and labor income in the baseline (HP) and fixed premium (FP) models and in the data (2019 SCF). We also report the empirical distribution from Y-14M in 2019 for labor earnings. Since labor productivity is exogenous in the model and the HP and FP model have the same equilibrium wage by construction, the distribution of labor earnings in these economies is identical. The SSE of the HP (FP) model relative to the data for the wealth moments shown in panel (a) is 0.155 (0.208). The wealth Gini coefficient for the HP (FP) model is 0.655 (0.616). The SSE of the model relative to the Y-14M (SCF) data for the labor earnings moments shown in panel (b) is 0.002 (0.018).

Finances (SCF) 2019 for wealth and income and Y-14M for income) in Figure B.2.[61] While broadly replicating the shapes of these distributions, our model does not capture the extreme skewness or concentration at the top of the distribution in either metric, particularly wealth. Notably, the model-implied distribution of income – which is identical across the HP and FP models by construction – matches quite closely with that implied by Y-14M data, which is less skewed than the SCF, with a sum of squared errors for the depicted moments of 0.002. Furthermore, the distribution of wealth in our baseline model matches the distribution of wealth from the SCF more closely than the FP model (SSE 0.155 vs 0.208), another set of untargeted moments in favor of the baseline model to complement those presented in Section 3.4.

Additional cross-sectional moments and the "naive" FP model Table B.3 presents additional cross-sectional results to augment the analysis in Section 4. In particular, we report our empirical target moments not only for the HP and FP models considered in the main text, but also for the FP model evaluated under the parameters of the HP model. This exercise highlights how the aggregate credit market changes under a different incidence of borrowing premia *without re-estimating* the parameters of the model. To complement Figure 4, we also provide in Panel B a selection of moments describing the cumulative share of total debt balances with respect to default risk.

The differences between the "naive" FP economy and the HP and FP economies in the

[61]SCF wealth and earnings calculations come from Moritz Kuhn and José-Víctor Ríos-Rull, https://sites.google.com/site/kuhnecon/home/us-inequality?authuser=0.

Incidence of premia	Heterog.	Fixed			
Parameterization	HP	FP	HP	FP	HP
Moment		level		% diff, $[j]-[1]$	
Column	[1]	[2]	[3]	$[2]-[1]$	$[3]-[1]$
Panel A: empirical targets (pp)					
bankruptcy rate	0.374	0.390	0.213	+4.53	-43.0
fraction in debt	12.3	12.2	12.1	-1.36	-1.73
debt to income	4.34	4.32	1.87	-0.52	-57.0
average loan rate spread	21.1	21.1	19.1	+0.01	-9.25
capital-output ratio	3.09	3.09	3.09	-0.10	+0.01
charge-off rate	3.79	3.69	1.90	-2.69	-50.0
suboptimal bankrupt share	45.5	45.7	53.1	+0.37	+16.8
Panel B: other moments (pp)					
average borrowing premium	12.2	12.2	12.2	0.00	0.00
cumulative share of total debt, def. prob $\leq$					
1%	54.8	37.1	66.4	-32.4	+21.2
2%	64.4	53.0	76.9	-17.7	+19.5
5%	81.2	80.8	91.4	-0.45	+12.6
10%	90.5	92.9	97.4	+2.59	+7.61
15%	96.9	99.5	99.8	+2.68	+2.97

Table B.3: **Additional cross-sectional moments**

Notes: The entries in columns [1] and [2] of Panel A correspond exactly to Panel C of Table 2. Column [3] reports moments for the FP model solved for the parameters of the HP model. The two rightmost columns present the percentage difference for each version of the FP model relative to the baseline HP model.

aggregate are quite stark. Agents in this economy self-ration their borrowing choices sharply along the intensive margin: the average debt to income ratio is 57% lower than in our baseline economy. Notably, the extensive margin is largely unchanged, with the share of debtors declining less than 2%. This result is intuitive: the cost of a riskless loan is quite similar in all these economies, but the FP economies feature a sharp increase in spreads for high levels of default risk. Without adjusting the parameters of the model as described in Section 3.4, agents shy away from these riskier loans and borrow only small amounts when they borrow. As a result, the bankruptcy and charge-off rates decline by 43% and 50%, respectively, relatively to the baseline economy, and the level of spreads declines.

Full paths for IRF analysis Figure B.3 presents the full paths from impact of the shock to recovery from the shock for the variables included in Table 3. This analysis places the "on-impact" results in context of the full recovery from the shock.

There are two striking points beyond those observable in Table 3. First, panel (a) shows that while total debt drops slightly on impact in our baseline model, it actually remains elevated relative to steady state for 5 periods, before dipping down again and gradually returning to steady

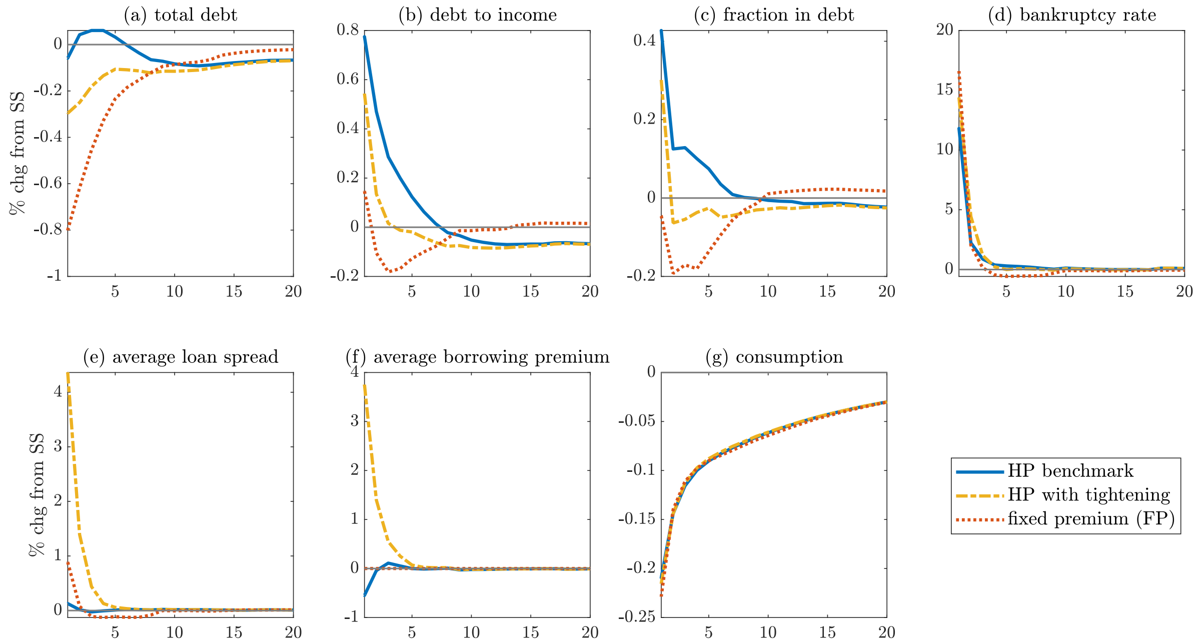

Figure B.3: **Full IRF paths for simulations and variables in Table 3**

Notes: This figure reports the full impulse response path to the shock described in Section 5.1 for the variables highlighted in Table 3.

state. This makes the implied cyclicality of debt in this model ambiguous *ex ante*. This is not the case for either alternative model, for which the level of debt remains below its steady state level for the entirety of the recovery. Second, panels (b), (c), and (e) show that the rates of recovery in debt to income, fraction in debt, and the average loan spread differ sharply across our model economies. For example, while the debt to income ratio increases on impact in all three economies, it remains elevated for 8 periods in our baseline, 4 in the case when premia tighten, and only 1 in the FP case. A similar pattern holds for the fraction of households in debt. For loan spreads, average spreads return to their steady state levels in no more than 3 periods in either our baseline, no tightening economy or the FP economy. By contrast, spreads remain above steady state levels for at least 6 periods in the case with tightening borrowing premia.

B.4 Business cycle model

Extending the model We make three main assumptions. First, the production function now has an aggregate TFP term, z, which follows an AR(1) process; that is, $Y = zK^{\alpha}N^{1-\alpha}$, with z following a transition process Γ^{z}. TFP is the fundamental exogenous shock to the economy which drives aggregate dynamics. Second, consistent with empirical evidence on the cyclical properties of earnings risk, we assume that the transition matrix which governs the persistent component of individuals' labor productivity varies with aggregate TFP; that is, Γ^{ϵ} is now $\Gamma^{\epsilon}(z')$. Third, the schedule of borrowing premia may be linked to the level of TFP; that is, $b(p)$ is now $b(p;z)$.

We summarize the aggregate state of the economy by $s = (z, \mu)$, which includes the exogenous level of TFP and the endogenous distribution of households over idiosyncratic states. All equilibrium value functions, decision rules, and pricing functions, now depend not only individual states x but also the aggregate state s. Furthermore, the equilibrium distribution $\mu(x)$ is no longer stationary, but evolves in the aggregate in a manner still consistent with households' decision rules. Aside from these changes, everything else from our baseline model is unchanged.

Computation The distribution $\mu(x)$ has a very high dimensionality, which makes solving the model using finite state space approximation difficult. We tackle this issue using the well-known approach of **?**. Specifically, competitive households must be able to correctly forecast future aggregate equilibrium prices – r', w', and i' – state-by-state given currently available information. Given the pricing equations implied by firm optimality and the aggregate resource constraint, households need only know aggregate capital, K', and aggregate labor productivity in each possible future state $N'(z')$.

The results of these forecasts for the calibrated model are presented in Table B.4. Panel A is for our baseline model with heterogeneous incidence of borrowing premia and no tightening of borrowing premia in recessions; Panel B is for the version with heterogeneous premia and tightening in recessions; and Panel C is for the fixed premium model. The results in Table B.4 show that households are able to (almost) perfectly forecast the parts of the approximate aggregate state required to compute equilibrium prices at each date and in each state. The overall standard errors of each forecasting regression are vanishingly small, and the (adjusted) R^2 of each regression exceeds 0.999.

Calibration The TFP process is discretized into two states: z_E (expansion) and z_R (recession), with transition matrix Γ^z. We set $z_R = 0.99$ so that recessions in the full business cycle model are of a similar magnitude (-1%) as our impulse response experiment from Section 5.1. We parameterize Γ^z by setting the share of recession periods, μ_{z_R}, equal to 0.211, the share of years with a recession in the postwar U.S. The probability of exiting a recession, $\Gamma^z_{RE'}$, is set to $2/3$ to match an average duration of a recession of 1.5 years. We then calibrate the expansion TFP level, z_E, to match the standard deviation of (logged, HP-filtered) aggregate output in U.S. data, 1.20%. Under our assumption that average TFP is equal to one ($\overline{z} = \mu^z_E z_E + \mu^z_R z_R = 1$), maintained for comparability with our steady state model, this pins down the last remaining parameter, the probability of entering a recession, $\Gamma^z_{ER'}$. This is achieved with $z_E =$ 1.0027, which implies $\Gamma^z_{ER'} =$ 0.178.

Given the TFP process, we link the labor productivity and borrowing premia processes as follows. For labor productivity, we assume that the transition matrix for the persistent component, ϵ_2, has standard deviation $\sigma^E_{\epsilon_2} =$ 0.094 in expansions and $\sigma^R_{\epsilon_2} =$ 0.163 in recessions, taking estimates directly from Storesletten et al. (2004). The persistence of ϵ_2 and the other parameters of the ϵ process are unchanged relative to the baseline. For borrowing premia, given data scarcity considerations described in the main text, we assume that the schedule of borrowing premia for expansions is the same as our baseline (2019) schedule, and that the analogous schedule in recessions is the same as our shock impact (2020) schedule. Estimates of these processes are discussed in Section B.1.

Results Table B.5 presents our business cycle results for the data (cols. [1] and [2]) and the three model variants (baseline: cols. [3] and [4]; HP with tightening: [5] and [6]; and FP: [7] and [8]) presented in the IRF analysis in Section 5.1.

forecast var.	agg. TFP z	z'	intercept	$\log K$	$\log N$	adj. R^2	s.e.	N
Panel A: Heterogeneous premia (HP) without tightening (Benchmark)								
$\log K'$	z_H	-	0.204	0.895	0.000	0.999	1.9E-04	1564
$\log K'$	z_L	-	0.090	0.939	0.125	1.000	1.8E-05	436
$\log N'$	z_H	z_H	0.039	-	0.783	1.000	2.0E-05	1280
$\log N'$	z_L	z_H	0.040	-	0.776	1.000	7.4E-06	284
$\log N'$	z_H	z_L	0.074	-	0.610	1.000	2.3E-05	284
$\log N'$	z_L	z_L	0.076	-	0.602	1.000	8.2E-06	152
Panel B: Heterogeneous premia (HP) with tightening								
$\log K'$	z_H	-	0.205	0.895	0.000	0.999	2.0E-04	1564
$\log K'$	z_L	-	0.088	0.940	0.127	1.000	1.8E-05	436
$\log N'$	z_H	z_H	0.039	-	0.783	1.000	2.0E-05	1280
$\log N'$	z_L	z_H	0.040	-	0.776	1.000	7.4E-06	284
$\log N'$	z_H	z_L	0.074	-	0.610	1.000	2.3E-05	284
$\log N'$	z_L	z_L	0.076	-	0.602	1.000	8.2E-06	152
Panel C: Fixed premium (FP)								
$\log K'$	z_H	-	0.185	0.905	0.000	0.999	1.9E-04	1564
$\log K'$	z_L	-	0.079	0.946	0.118	1.000	2.2E-05	436
$\log N'$	z_H	z_H	0.039	-	0.783	1.000	2.0E-05	1280
$\log N'$	z_L	z_H	0.040	-	0.776	1.000	7.4E-06	284
$\log N'$	z_H	z_L	0.074	-	0.610	1.000	2.3E-05	284
$\log N'$	z_L	z_L	0.076	-	0.602	1.000	8.2E-06	152

Table B.4: **Forecasting rules for model with aggregate uncertainty**

Notes: This table reports estimates and regression statistics for the forecasting equations for the business cycle model, $\log K' = \alpha_0(z) + \alpha_1(z)\log K + \alpha_2(z)\log N$ and $\log N' = \beta_0(z,z') + \beta_1(z,z')\log N$. "s.e." denotes the standard error of the regression. The last column presents the number of observations for each subsample of the total simulation of $T = T_1 + T_0 - T_0 = 2500 - 500 = 2000$ sample periods.

Panel A highlights two key macro aggregates: output and consumption. By construction, the standard deviation of output is identical across all three model economies and matches the data exactly. In all cases of the model and in the data, consumption is highly procyclical and less volatile than output. One of the shortcomings of our model is that the relative volatility of consumption in our model is only about 1/6 of what it is in the data. Comparing across model variants, though, we find corroboration for the aggregate consumption results in our IRF analysis: the relative volatility is lowest in our baseline model, followed by the HP model with tightening of premia in recessions, followed last by the fixed premium model. Our (steady state) calibration insures that agents' ability to smooth consumption *on average* (or, in the cross-section) across all model variants is the same, and so these results pick up purely cyclical differences.

Panel B presents key credit quantity moments. Perhaps surprisingly given our IRF analysis in Section 5.1, we find that total debt is countercyclical across all our models and in the data,

Premia Incidence	[Data]		Heterogeneous				Fixed	
Tighten in Rec.?			N		Y		-	
Moment	σ_X/σ_Y	ρ_{XY}	σ_X/σ_Y	ρ_{XY}	σ_X/σ_Y	ρ_{XY}	σ_X/σ_Y	ρ_{XY}
Column	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Panel A: Macro aggregates								
output	*1.20%*	*1.00*	*1.20%*	*1.00*	*1.20%*	*1.00*	*1.20%*	*1.00*
consumption	0.81	0.92	0.14	0.78	0.15	0.79	0.16	0.81
Panel B: Credit quantities								
total debt	3.20	-0.29	0.83	-0.78	0.82	-0.80	0.34	-0.92
bankruptcy filings	17.9	-0.11	1.51	-0.94	1.52	-0.95	1.14	-0.99
debt to income	5.23	-0.27	8.51	-0.99	8.85	-1.00	10.7	-1.00
fraction in debt	6.62	0.48	0.94	-0.94	0.96	-0.94	0.15	-0.82
Panel C: Interest rates								
avg IR, all loans	0.90	-0.88	0.10	-0.98	1.18	-1.00	0.55	-1.00
avg BP, all loans	0.19	-0.56	0.29	0.98	0.90	-1.00	0.00	0.00

Table B.5: **Business cycles: data, baseline, HP with tightening, and FP economies**

Notes: Business cycle moments are computed using the cyclical component of Hodrick-Prescott filtered logs of the data, with annual smoothing parameter of 6.25. σ_X/σ_Y is the standard deviation of variable X normalized by the standard deviation of GDP, and ρ_{XY} is the correlation of X with GDP. Data are for 1999-2019 except for risk free loan rate and debt to income for which we use Y-14M data. Details of data sources are provided in the Appendix. For the measure of default used in Section 1, the relative volatility and cyclicality are 5.11 and -0.29.

though the magnitude of this effect is larger in our model than the data (about -0.8 vs. -0.3). This "flip" comes from the larger assumed swing in the earnings risk in our business cycle model.[62] Total debt is significantly less volatile in our models than the data (about 3.2 vs. 0.8 / 0.3 HP / FP). Importantly, the variable premia case brings both the relative volatility and cyclicality of total more in line with their empirical values. Similarly, all our models capture that bankruptcies and the debt to income ratio are more volatile than output and countercyclical as in the data, and our baseline model outperforms the FP case across the board. All our models fail to capture the procyclicality and excess volatility of the fraction of households in debt, the former effect likely driving the gaps between models and data for total debt.

Finally, panel C presents two key credit pricing metrics: the average interest rate spread and average borrowing premium on all loans. Of course, the FP model (with no tightening) has no volatility or cyclicality in borrowing premia. All models capture the strong countercyclicality of loan spreads that we observe in the data. Notably, the HP model with tightening in recessions comes closest to matching the relative volatility of spreads we observe in the data, suggesting that the cyclical response of borrowing premia is important for matching this moment.

[62]Specifically, our IRFs assume that on impact σ_{ϵ_2} goes from 0.129 to 0.163, while our business cycle model assumes that this same variable goes from 0.094 in expansions to 0.163 in recessions. The reduction in risk in expansions dampens credit demand in expansions, driving the countercyclicality we observe here.

C Details of the Extended Model with Lending Standards

C.1 Theory

Pricing conditions The first order conditions of maximizing the bank objective (16) subject to (17) are

$$
\begin{aligned}
1 &= \frac{1}{1+i}\mathbb{E}\left[W_K(K',\mathcal{M}')\right] && \text{(C.1)}\\
q(\ell;x)\ell\left(1+\eta\mathbf{1}[\ell>0]+\eta\lambda(p(\ell;x))\mathbf{1}[\ell<0]\right) &= \frac{1}{1+i}\mathbb{E}\left[W_{m(\ell;x)}(K',\mathcal{M}')\right] && \text{(C.2)}
\end{aligned}
$$

and the corresponding envelope conditions are

$$
\begin{aligned}
W_K(K,\mathcal{M}) &= 1+r-\delta && \text{(C.3)}\\
W_{m(\ell;x_{-1})}(K,\mathcal{M}) &= -\int_{\mathcal{B}\times\mathcal{E}}\left(1-g_{BK}(\ell;dx)+\xi g_{BK}(\ell;dx)\right)\ell\mathbb{P}(dx|x_{-1}). && \text{(C.4)}
\end{aligned}
$$

Combining expressions (C.1) and (C.3) yields $i=r-\delta$ as in the baseline model. Applying the definition of expected repayment probability from equation (12) and combining equations (C.2) and (C.4) delivers equations (18) and (19) from the main text.

Modified equilibrium definition Most of the equilibrium definition in Section 2.4 for the baseline model with exogenous borrowing premia extends to the model with endogenous borrowing premia. In particular, household optimization (i), firm optimization (ii), distribution consistency (iv), and the market clearing conditions (v) remain the same. However, the lender optimality condition (iii) is subject to three changes. First, the loan price schedule must be given by (19) rather than (13). Second, the savings price must be given by (18) rather than $(1+i)^{-1}$. Third, the multiplier η must be consistent with the aggregate loan supply constraint (17); that is, if $\eta=0$, then the right hand side of this constraint must be strictly greater than the left hand side, and if $\eta>0$, then the right and left hand sides of the constraint must be exactly equal.

C.2 Computation

The algorithm for solving the baseline model with exogenous borrowing premia is completely standard, and so we do not include it here. The same applies to the impulse response and business cycle analyses for this version of the model. The former is solved using a perfect foresight transition, back-solving the household decision problem given a vector of equilibrium prices, then forward-solving the sequence of endogenous distributions over idiosyncratic states, then updating the equilibrium price vector until convergence. The latter is solved as described in Appendix B.4.

The version of the model with endogenous premia in Section 6, though, requires modifications to this algorithm, and so we describe our computational method here. The key insight of our algorithm is that we require the converged equilibrium to match a given schedule of borrowing premia, and so we do not need to assume a particular $\lambda(\cdot)$ in our solution; rather, we can solve the model until the value of the multiplier η has converged, then infer $\lambda(\cdot)$ from the borrowing premia schedule. We go through the details for the steady state, then describe the adaptation to solving an impulse response to an aggregate shock.

1. Make a guess of the equilibrium capital stock K_0 and the multiplier on the loan supply constraint η_0.

2. Compute the implied set of factor prices $r(K_0)$ and $w(K_0)$, interest rate $i(K_0)$, and return on savings $\overline{q}(K_0, \eta_0) = [(1 + i(K_0))(1 + \eta_0)]^{-1}$.

3. Make a guess of the bankruptcy decision rule, $g_{BK}(x)$, for all x. Use this to define the repayment probability function $p_0(\ell; x, K_0, \eta_0)$ according to (12).

4. Impose the desired borrowing premium schedule on the loan price schedule. That is, given an empirical schedule $\hat{b}(p)$, specify the loan price schedule

$$q(\ell; x, p_0, K_0, \eta_0) = \frac{p_0(\ell; x, K_0, \eta_0) + \xi\,(1 - p_0(\ell; x, K_0, \eta_0))}{(1 + i(K_0))\left(1 + \hat{b}(p_0(\ell; x, K_0, \eta_0))\right)} \tag{C.5}$$

   Note that the lending standards function $\lambda(\cdot)$ does not appear in equation (C.5).

5. Solve for household decision rules $g_a(x; p_0, K_0, \eta_0)$ and $g_{BK}(x; p_0, K_0, \eta_0)$ given the equilibrium prices assumed in the preceding steps.

6. Use equation (12) and the new bankruptcy decision rule to compute an updated guess of the repayment probability, $p_1(\ell; x, K_0, \eta_0)$. If $\max_{\ell,x} |p_1(\ell; x, K_0, \eta_0) - p_0(\ell; x, K_0, \eta_0)| < \varepsilon_p$, where ε_p is a tolerance level, proceed to step 7. Otherwise, set $p_0(\ell; x, K_0, \eta_0) = \chi_p p_1(\ell; x, K_0, \eta_0) + (1 - \chi_p)p_0(\ell; x, K_0, \eta_0)$, where $\chi_p \in (0, 1]$ is a relaxation parameter, and return to step 4.

7. **Assess aggregate market clearing.**

   (a) Using pricing conditions (18), (19), and (20), rewrite constraint (17) as

$$\underbrace{-\int_{\mathcal{X}\times\mathcal{L}_-} b(p(\ell; x))\frac{p(\ell; x) + \xi(1 - p(\ell; x))}{1 + b(p(\ell; x))}\ell dm(\ell; x)}_{\equiv A_-} \leq \frac{\eta}{1 + \eta}\underbrace{\int_{\mathcal{X}\times\mathcal{L}_+} \ell dm(\ell; x)}_{\equiv A_+} \tag{C.6}$$

   Use equation (C.6) to compute $A_-(K_0, \eta_0)$ and $A_+(K_0, \eta_0)$.

   (b) Compute the implied aggregate capital $K_1(K_0, \eta_0) = \int_{\mathcal{X}} a d\mu(x)$.

   (c) Define the following convergence metrics for η and K:

$$\begin{aligned} \Delta_\eta &= \left|\frac{\eta_0}{1 + \eta_0}A_+(K_0, \eta_0) - A_-(K_0, \eta_0)\right| \\ \Delta_K &= \left|K_1(K_0, \eta_0) - K_1\right| \end{aligned}$$

(d) If $\Delta \equiv \max\{\Delta_\eta, \Delta_K\} < \varepsilon$, where ε is a tolerance parameter, proceed to step 8. Otherwise, define

$$\eta_1(K_0, \eta_0) = \frac{A_-(K_0, \eta_0)}{A_+(K_0, \eta_0) - A_-(K_0, \eta_0)},$$

update $\eta_0 = \chi_\eta \eta_1(K_0, \eta_0) + (1 - \chi_\eta)\eta_0$ and $K_0 = \chi_K K_1(K_0, \eta_0) + (1 - \chi_K) K_0$, where $\chi_\eta \in (0, 1]$ and $\chi_K \in (0, 1]$ are relaxation parameters, and return to step 2.

8. Use equation (20) to back out an implied lending standards function, $\hat{\lambda}(p) = \hat{b}(p)/\eta_0$.

Solving for impulse responses simply adapts this algorithm to a perfect foresight transition path. The main departure is that rather than assuming a single empirical borrowing premia target $\hat{b}(p)$ as in the steady state, we must assume a series of targets, $\{\hat{b}_t(p)\}_{t=1}^{T}$, where T is the terminal date at which we assume the economy has returned to steady state. Given a steady state schedule $\hat{b}_0(p)$ and an impact schedule $\hat{b}_1(p)$, we define this sequence recursively via

$$\hat{b}_{t+1}(p) = \rho_z \hat{b}_t(p) + (1 - \rho_z)\hat{b}_0(p) \text{ for all } t = 1, ..., T - 1 \text{ and all } p$$

Given this, we can nest the algorithm above neatly into a standard structure in which the decision rules are back-solved, distributions are forward-solved, and equilibrium price sequences are gradually updated to convergence.