Discussion*

Kyle Herkenhoff

May 30, 2025

*Herkenhoff: University of Minnesota, Federal Reserve Bank of Minneapolis, NBER.

1 Discussion

Credit scores are a fundamental component of American life. They govern key decisions regarding shelter, access to entrepreneurial capital, access to cars and transportation, and in many cases, jobs (Herkenhoff (2019), Corbae and Glover (2024)). Despite the prevalence of credit scores in American life, they are pervasively absent in standard defaultable debt models (Athreya (2002), Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007), Livshits, MacGee, and Tertilt (2007) among others). The present paper summarizes the authors' past and present attempts to model credit scores, while also providing fruitful paths forward for future research.

The authors frame their current paper around an investigation the relationship between "reputational" inequality — as proxied by credit scores — and income inequality. The motivation stems from the observation that both income inequality and the dispersion in credit scores tend to increase over the life cycle, raising the question of whether and how these dynamics are linked. To address this, the authors develop and estimate a structural model in which the causal relationship between income and credit scores runs both ways. The framework is a consumer default model featuring private information about agents' patience and effort, and potentially unobserved medical debt. Lenders form beliefs about borrower types based on observable signals, including income, and price loans accordingly. Importantly, the model incorporates expectation errors (trembles) over default behavior and asset/debt choices to discipline lenders' off-equilibrium beliefs. The model is successful in broadly replicating key empirical patterns observed in the data, and they use the model to assess the welfare implications of information restrictions, including broad limitations (no tracking/scoring) and narrower limitations (no medical debt histories). Their main result from the quantitative calibration is that a full information regime is welfare improving. The intuition is simple: high-type households get better rates now that they are not lumped in with low-type households; and low-type households benefit from partially pooled credit terms, but they no longer need to take costly actions/effort to mimic high-type households.

The intellectual origins of this paper lie in the broader agenda of modeling and optimizing the institutions that govern market incompleteness, particularly in the context of consumer credit and default. Early foundational work, such as Athreya (2002), modeled pooled default risk, laying the groundwork for understanding the aggregate implications of credit markets with limited enforcement. A second generation of models, notably Chatterjee et al. (2007) and Livshits et al. (2007), introduced individual pricing, drawing on the insight from Eaton and Gersovitz (1981) that bond prices should depend on borrower-specific states. These models allowed interest rates and borrowing limits to vary with individual characteristics, such as income and loan size, resulting in highly state-contingent credit terms. Athreya, Tam, and Young (2009) highlighted that such finely tuned pricing delivered poor insurance properties – consumer credit access did little to smooth consumption in the face of income shocks. Introspection at the time, and subsequent work that I have been involved in writing has established that households receive significant insurance from unsecured credit markets.

One route to slowing down the pricing adjustments inherent in Chatterjee et al. (2007) is to rely on credit scoring. Credit scores serve as slow-moving reputational markers that can moderate the re-

sponsiveness of interest rates and credit limits to individual borrowing and shocks. Specifically, when borrowers take on small or moderate amounts of debt, the resulting updates to their credit terms are minimal, reflecting lenders' continued reliance on prior beliefs about borrower type. This intuition stems from the idea that modest borrowing doesn't reveal much new information, so credit pricing remains relatively stable. The current authors have been central to developing the agenda over market incompleteness and credit scoring over the past two decades, producing a body of work that includes five key papers. Chatterjee, Corbae, and Rios-Rull (2008) introduced a finite-horizon model with multiplicative preference shocks, while Chatterjee, Corbae, Dempsey, and Ríos-Rull (2023) extended the framework to an infinite horizon with trembles over asset choices. Corbae and Glover (2024) linked credit scores to labor market outcomes, and Briglia, Chatterjee, Corbae, Dempsey, and Ríos-Rull (2021) explored the broader economic consequences of introducing trembles. The latest contribution, here in the NBER Macroeconomics Annual 2025, integrates these elements into an richer finite-life cycle theory of inequality with costly effort choices and medical debt.

There are, however, several alternative routes to generating realistic degrees market incompleteness in defaultable debt frameworks. These other contributions also document facts that raise important questions about the authors' modeling assumptions. Several papers that I was involved in writing argued that long-term relationships between creditors and consumers were important for attaining the right degree of market incompleteness observed in the data (Herkenhoff (2019)). Search frictions in the credit market (costly applications and potential denials) mimic several aspects of credit scoring. Credit access and post-default re-access rates are endogenous, depending on the agents' states, and movements in the terms of credit can be slowed down. In recent work with my coauthors (Braxton, Phillips, and Herkenhoff (2019) and Braxton, Herkenhoff, and M. Phillips (2024)), we find that credit limits and credit scores are remarkably *unresponsive* to events like job loss. We find that a 30% earnings loss translates to .05 standard deviation reduction in credit scores for job losers, which is extremely small, and a near zero response in credit limits. In the data, this commitment to the terms of credit (limits and rates) allows job losers to use credit markets to smooth consumption via unsecured credit markets. We find that roughly one third of job losers borrow more after job loss, one third do not respond, and one third default to self-insure (i.e. an equitable transfer of risk that is covered by the lending premia inherent in high credit card interest rates).

Given these findings, this begs the question, *if the current authors' goal is to model the degree of market incompleteness facing households properly, then why model lenders and scoring agencies as observing household income on a rolling basis?* This seems first order to modeling U.S. institutions and matching the degree of market incompleteness in the U.S. The modeling tools are flexible enough to handle these modifications, and so I expect future iterations of these projects will come to address the fundamental source of slow moving scores and what it means for insurance and market completeness.

The first empirical fact highlighted in the paper is that the variance of credit score ranks increases over the life cycle (Panel A of Figure 1). However, since rank is not a unitless measure, this raises interpretability issues. To better capture the relative dispersion of credit scores over time, the authors should consider transforming the data by either taking the log of ranks or by normalizing the standard devia-

tion by the mean, as in Panel B of Figure 1. Panel B of Figure 1 illustrates that uncertainty surrounding credit scores is highest early in life, peaking in individuals' early 30s. As individuals age, this uncertainty *declines*.

Figure 1: Uncertainty in scores over the lifecycle

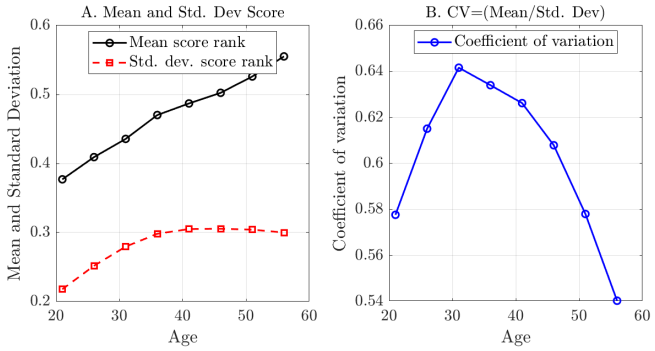

Why does the uncertainty over borrower types early in life shown in Panel B of Figure 1 matter for the economics of the problem? This uncertainty is particularly consequential because young households are often making costly, forward-looking investments such as in education, childcare, homeownership, and family formation. These investments have long-term implications for income trajectories and thus inequality, the main object of interest in the present article. Existing research, including Braxton, Chikhale, Herkenhoff, and Phillips (2024), shows that access to credit and the design of bankruptcy institutions can directly influence early childhood investments and upward mobility in the U.S. Similar themes are explored in work by Abbott, Gallipoli, Meghir, and Violante (2019), Daruich (2018), and Caucutt and Lochner (2020). However, in the current draft, investment effort shifts income realizations, which maps clearly into actions like human capital accumulation (e.g. schooling, occupation licensing, etc.). However, this effort exertion is modeled as a utility cost rather than a monetary one. The quasi-linear specification assumes that effort exertion e is at utility cost κ_e,

$$\max_{e\in\{0,1\}} -\kappa_e e + W_n(e;\omega,\tau) + \epsilon^e,$$

which effectively decouples financial constraints from lifetime income formation, potentially understating the economic role of early-life uncertainty over borrower types and thus borrowing frictions. This raises a broader theoretical question: *might a more compelling mechanism linking reputation and inequality involve early-life monetary investment costs (rather than utility costs), when uncertainty about types and credit constraints directly limit income-enhancing actions?*

The second key empirical fact in the paper is that income and credit scores are highly but imperfectly correlated. However, the authors rely on income data imputed by Equifax, which is itself derived from credit variables. Since the imputation is based on the same information set used to construct credit scores, a sufficiently flexible mapping could, in principle, generate a near-perfect correlation between the two.

While it is plausible that the true relationship is indeed strong but not exact, this particular empirical fact would be more convincing if validated using external income data, such as Y-14 regulatory filings or other administrative sources. Doing so would help ensure that the observed correlation reflects true economic co-movement rather than mechanical artifact.

The main methodological innovation is that the paper incorporates taste shocks over asset and default choices as a way to discipline off-equilibrium-path beliefs, building on previous work by the authors Chatterjee et al. (2023). However, this modeling choice raises important conceptual and computational questions. Are these trembles primarily a numerical device to ensure convergence and smooth policy functions, or do they capture meaningful economic behavior so that we should not take the limit of the trembles to zero? The fact that they do not take the limit of trembles to zero suggests that the authors believe there is some economic content to the shocks. This has direct implications for estimation of parameters (such as the discount factor), savings behavior, welfare calculations, counterfactual interpretations. The authors put forth a thorough and clear appendix that puts forth solutions to several of these issues.

While the authors provide a detailed discussion of the economics of the taste shocks, I want to devote part of this discussion to help clarify what the authors do and why, and to issue words of caution to future potential users of this framework. Consider a simple 2-period model economy. The agent is born with assets a_0 and income y and then makes a saving decision and consumes assets. The utility function $u(\cdot)$ is concave and satisfies the usual Inada conditions. Assume that there is an extreme value taste shock ϵ^a over all assets and no discounting:

$$V(a_0, y) = E_{\epsilon^a} \max_{a \in A} \{U(c_1) + U(a) + \epsilon^a\} \quad s.t. \quad c_1 = a_0 + y - a$$

With no taste shocks, the solution is simple – split labor and non-labor income evenly, $a = \frac{a_0+y}{2}$.

What happens to saving behavior with taste shocks? I illustrate the resulting "soft-max" choice in Figure 2. The even-split is optimal without taste shocks, and so the agent chooses that option most often, but the taste shocks make agents tremble to the surrounding savings level, with the choice probability diminishing in the distance from the no-tremble optimum. In Panel A, we see individual get 7 draws of the taste shocks since their feasible set is bound above by y.

Consider now restarting the economy in which $a_0 \approx y$ (so consider someone who trembled to the top of the asset grid). That person now draws 9 taste shocks and thus has a higher expected mean taste shock draw. This puts upward pressure on savings, over and above the discount factor: the more you save, the larger your feasible set of asset grid points, and the more taste shocks you draw. The authors discuss these points in their article and how it pushes agents toward patience and they also discuss adjustments for grid density (which alters the number of draws), citing their companion paper Briglia et al. (2021). These adjustments are non-trivial but essential for those seeking to adopt this framework for future research. These solutions require additional modifications when the problem is extended to correlated taste shock draws or researchers want to adopt non-grid based solution methods. Future users of these tools need to pay careful attention to the way taste shocks alter the economics of the problem.

The presence of trembles also raises important questions about welfare and the class of counterfactuals that can be studied. For instance, some of the gains that researchers would attribute to a looser credit limit counterfactual will stem from more draws of taste shocks as opposed to the typical source of gains from greater consumption smoothing. The authors attempt to address this by using the resulting policy functions and turning off the taste shocks. Future users of this framework should conduct similar checks.

In the context of default problems, extra caution is necessary. In the present paper, taste shocks over default/not default are only drawn with negative net worth. The resulting asset choice probabilities, with no adjustments to the mean of those default taste shocks, is sketched in panel C. Since only those indebted can default and draw taste shocks over default, households will always want to hold ϵ debt rather than have zero net worth. In the data, there is a significant mass of zero net worth households. The authors again adjust the mean of the taste shocks to adjust for this problem (see their Appendix); however, there are simpler more intuitive solutions such as allowing all households (with positive or negative net worth) tremble over default and mapping positive-asset defaults into "unmodeled, wealth-destroying shocks" that lead to default such as divorce, medical shocks, etc. This is the approach I take in my joint work Herkenhoff and Raveendranathan (2024). All of this discussion, however, points toward the caution one must take when extending or using these tools in richer settings with richer delinquency settings. For instance, when studying the recent Chapter 7/13 streamlining proposals in the U.S., how should one adjust taste shocks over Chapter 7 and Chapter 13 default after the option to file Chapter 13 is removed or more options are added?

In one of their main counterfactual exercises, the authors analyze the effects of moving from a credit scoring regime to one of "no tracking, " where lenders observe no information about borrower types. This exercise retains the assumption of atomistic, one-period lenders, which may overstate the information loss from eliminating credit scores. A strong case can be made—drawing on classic arguments such as Bernanke (1983)—that in the absence of formal scores, lenders would engage in relationship lending or develop alternative screening mechanisms, thereby partially mitigating the informational loss. This raises the question of whether we should interpret their findings as an upper bound on the potential effects of removing credit score inputs, rather than a realistic estimate of the actual consequences.

Im summary, while I raise some concerns about modeling assumptions in this discussion, the current authors have made major inroads on an important topic and have provided the profession with the tools to answer my questions and many more. These papers have advanced the science of economics, and I look forward to their future development and applications.

Figure 2: The economics of taste shocks

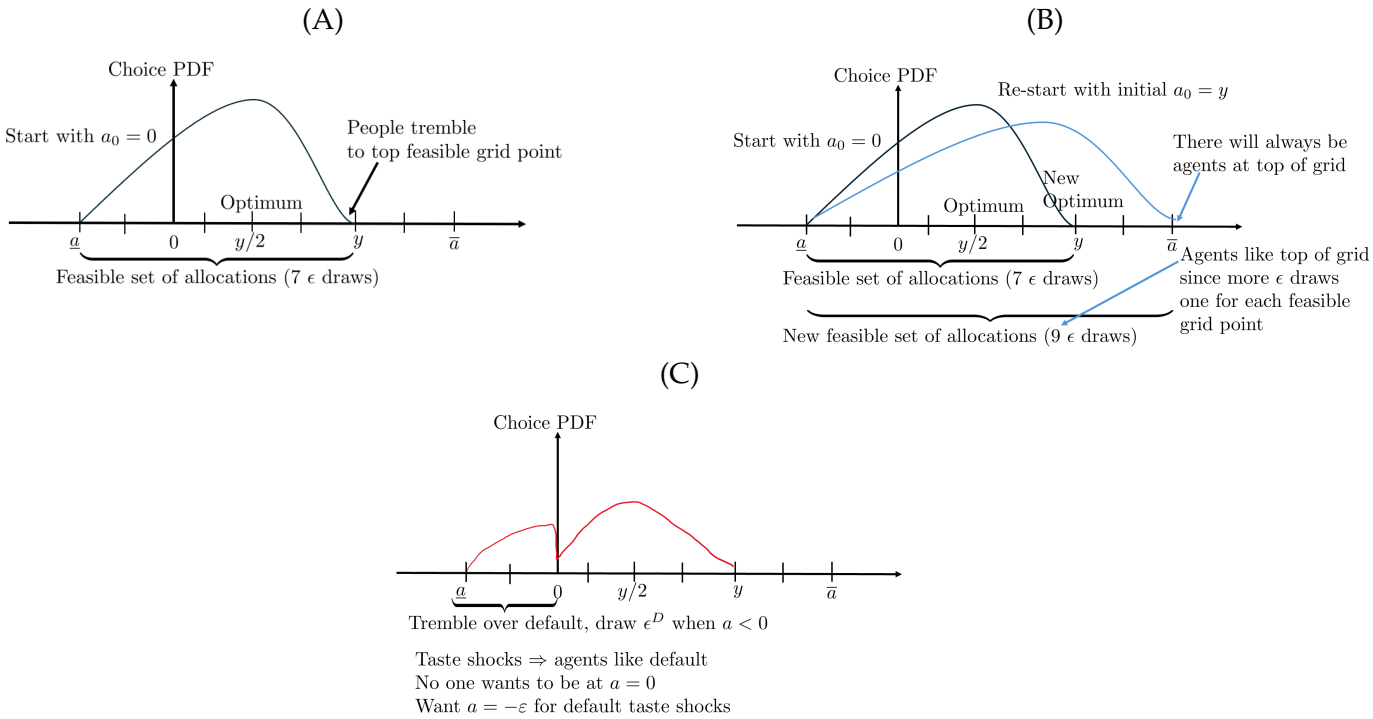

Taste shocks $\Rightarrow$ agents like default
No one wants to be at $a = 0$
Want $a = -\varepsilon$ for default taste shocks

References

Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2019). Education policy and intergenerational transfers in equilibrium. *Journal of Political Economy 127*(6), 2569–2624.

Athreya, K., X. S. Tam, and E. R. Young (2009). Unsecured credit markets are not insurance markets. *Journal of monetary Economics 56*(1), 83–103.

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

Bernanke, B. S. (1983). Non-monetary effects of the financial crisis in the propagation of the great depression. Technical report, National Bureau of Economic Research.

Braxton, J. C., N. Chikhale, K. F. Herkenhoff, and G. M. Phillips (2024). Intergenerational mobility and credit. Technical report, National Bureau of Economic Research.

Braxton, J. C., K. Herkenhoff, and G. M. Phillips (2024). Can the unemployed borrow? implications for public insurance. *Journal of Political Economy 132*(9), 3025–3076.

Braxton, J. C., G. Phillips, and K. Herkenhoff (2019). Can the unemployed borrow? implications for public insurance. In *2019 Meeting Papers*, Number 323. Society for Economic Dynamics.

Briglia, L., S. Chatterjee, D. Corbae, K. Dempsey, and J.-V. Ríos-Rull (2021). Saving for a sunny day: An alternative theory of precautionary savings. *Unpublished.[12]*.

Caucutt, E. M. and L. Lochner (2020). Early and late human capital investments, borrowing constraints, and the family. *Journal of Political Economy 128*(3), 1065–1147.

Chatterjee, S., D. Corbae, K. Dempsey, and J.-V. Ríos-Rull (2023). A quantitative theory of the credit score. *Econometrica 91*(5), 1803–1840.

Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Ríos-Rull (2007). A quantitative theory of unsecured consumer credit with a risk of default. *Econometrica 75(6)*, 1525–1589.

Chatterjee, S., D. Corbae, and J.-V. Rios-Rull (2008). A finite-life private-information theory of unsecured consumer debt. *Journal of Economic Theory 142*(1), 149–177.

Corbae, D. and A. Glover (2024). Employer credit checks: Poverty traps versus matching efficiency. *Review of Economic Studies*, rdae095.

Daruich, D. (2018). The macroeconomic consequences of early childhood development policies. *FRB St. Louis Working Paper* (2018-29).

Eaton, J. and M. Gersovitz (1981). Debt with potential repudiation: theoretical and empirical analysis. *Review of Economic Studies 48(2)*, 289–309.

Herkenhoff, K. F. (2019). The impact of consumer credit access on unemployment. *The Review of Economic Studies 86*(6), 2605–2642.

Herkenhoff, K. F. and G. Raveendranathan (2024). Who bears the welfare costs of monopoly? the case of the credit card industry. *Review of Economic Studies*, rdae098.

Livshits, I., J. MacGee, and M. Tertilt (2007). Consumer bankruptcy: A fresh start. *American Economic Review 97(1)*, 402–418.