Grant Proposal:
Misaligned Exploration–Exploitation Trade-Offs in Financial Regulation

Laura Blattner
Stanford University

Jann Spiess
Stanford University

April 9, 2021

As the use of machine learning (ML) and artificial intelligence (AI) in consumer lending grows, new questions arise on how the use of these algorithms should be regulated. A key challenge for designing policy is how to address potentially conflicting regulatory objectives. For example, regulators might want to promote the use of innovative lending products that expand access to finance for under-served populations. At the same time, (other) regulators want to ensure that lenders do not make excessively risky loans and discourage lending to potentially high-risk subgroups. These competing objectives lead to a trade-off between exploration and exploitation. Concerns about equal access to credit might push the regulator towards incentivizing exploration, e.g. the Community Reinvestment Act incentives loans to under-served communities. Concerns about risk, in contrast, dis-incentivize exploration, e.g., regulatory capital constraints, stress testing, and safety and soundness requirements are all designed to limit a lender's exposure to loan losses.

In this project, we propose to study optimal financial regulation as an exploration–exploitation trade-off, which allows us to leverage insights and techniques from the computer-science literature on contextual-bandit problems. We first develop a theoretical model to derive the regulators' optimal solution to the exploration–exploitation trade-off and compare this solution to current regulatory practices. We then explore empirically to what extent consumer lending markets in the US are characterized by under-exploration and quantify gains from allocating loans based on a contextual bandit algorithm that mirrors the regulators' optimal solution.

Theoretical framework. Our theoretical framework studies this problem through the lens of a principal-agent framework in which a regulator imposes restrictions on a lender who uses a credit-risk prediction algorithm to allocate loans. Lenders screen applicants who differ in the informativeness of their credit score (as in Blattner and Nelson, 2021). Lenders therefore face an exploration–exploitation trade-off. Lenders can focus on approving applicants for whom default risk can be well assessed (exploitation). Alternatively, they can choose to learn about the noisy applicants (exploration) either by approving them (and observing the default outcome) or by experimenting with novel data that might be informative about noisy applicants.

Lenders and regulators face *different* exploration–exploitation trade-offs, which motivates the need for regulatory intervention. First, this divergence arises because regulators might have a stronger preference than the lender for equal access to credit across groups and/or a stronger preference for mitigating risky lending. Differences in social preferences and lenders' preferences can lead to over- or under-exploration relative to first best – depending on the nature of the preference misalignment.[1] Second, lenders do not fully internalize the value of exploration for other lenders. For example, data reported to credit bureaus as a result of a loan the lender makes improves signal informativeness for other lenders. The latter friction leads to under-exploration since lenders do not take into account the positive externalities of information acquisition.

Our theoretical model aims to characterize optimal policy that navigates these various objectives and scenarios: (i) social preferences for equal access to credit across groups exceed that of the lender, (ii) social preferences for mitigating loan losses exceed that of the lender, (iii) preferences on both equality and risk dimensions diverge, (iv) there are separate regulatory entities maximizing each goal in isolation, (v) there is a single regulatory entity jointly optimizing both goals.

Empirical analysis. Our empirical analysis can offer insights into (i) the degree of under-exploration in US consumer lending markets, (ii) the gains from introducing a contextual bandit algorithm that increases exploration, (iii) whether these gains rely on exploring new applicant populations or on exploring new data sources, and (iv) how these gains vary across markets, product, and lenders differentially affected by under-exploration frictions.

Our empirical analysis builds on a novel dataset that combines bank account transaction data ("cash-flow data") with credit-bureau records. We proceed in two steps. Our first step is to build a decision algorithm that mimics the observed approval/rejection decisions in a (training) sample of loan applicants. We initially focus on unsecured personal loans and credit cards. We then train contextual bandit algorithms on the same training sample. The bandit algorithms are designed to reflect the preferences of an exploration-focused regulator, an exploitation-focused regulator, and a jointly-optimizing regulator. We plan to study both bandit algorithms that explore by extending loans to additional set of applicants (e.g. those with little credit report info or from sparse parts of the distribution) as well as the introduction of new data (cash-flow data). Cash-flow data is one of the most widely used sources of alternative data in credit underwriting but has been under-studied in economics due to data limitations.

Our approach allows us to compare outcomes across the two loan allocation rules on common test data. In particular, we can answer questions such as:

1. If we adapt a contextual bandit algorithm, what is the evolution of credit access across groups and realized default over time?

2. How do the results differ if we explore by lending to new types of applicants vs. using new types of data?

3. How do results differ across settings that vary in the severity of the under-exploration frictions?

[1] For example, if the regulator is only concerned about high realized default, the lender might be over-exploring relative to the regulator's preferred approach that focuses only on exploitation.

Identification challenge. We face two key identification challenges in this project. The first challenge is obtaining counterfactual default outcomes for rejected applicants whom the bandit algorithm approves. We solve this missing-outcome challenge by focusing either on applicants who were later approved for a similar loan product by a different lender or by using default outcomes on other loan products to impute the missing outcome (building on recent work by Lakkaraju et al., 2017, Coston et al., 2021, and Blattner and Nelson, 2021).

The second challenge is to obtain variation in the severity of the under-exploration frictions. We solve this challenge in several ways. First, we can compare repeat versus one-time loan products, e.g. unsecured personal loans vs mortgages, as repeat products allow lenders to capture more of the return to additional exploration. Second, we can compare bank and non-bank lenders since the latter are not subject to the the same safety and soundness regulation that govern risk management. This difference implies that non-bank lenders might be more exploration-minded relative to bank lenders. Alternatively, we can explore variation in enforcement actions on fair lending practices, e.g. the CFPB under the Trump vs. Biden administrations.

Contribution to the literature. We contribute to a literature that studies information in consumer lending.[2] Much of this literature has focused on exogenous changes in information, and the implications for efficiency and welfare. In contrast, we explicitly consider the lender's decision how much information to acquire and potential sources of under-acquisition of information. Moreover, we are focused on optimal regulatory policy in a setting where social and lender preferences for exploration potentially diverge.

We also contribute to a nascent literature on algorithmic decision-making (e.g., Athey et al., 2020) and how to regulate it (Blattner et al., 2021). Most of the existing work in this area has focused on questions of algorithmic fairness.[3] In contrast, we ask how to design optimal policy when there are competing policy objectives.

Finally, we contribute to a literature that uses tools from contextual-bandit problems in economic contexts, including hiring, job training, and health (Li et al., 2020, Currie and MacLeod, 2020, and Kasy and Sautmann, 2021). Relative to these approaches, we apply bandit algorithms on credit-bureau data to explore the role of exploration–exploitation trade-offs for financial regulation, including equity and risk concerns.

[2] Notable papers in this literature include Berg et al. (2020), Liberman et al. (2019), Bartik and Nelson (2020), Autor and Scarborough (2008), Agan and Starr (2018), Khandani et al. (2010), Corbae and Glover (2018), Fulford and Nagypál (2020), Bos et al. (2016), Kovbasyuk and Spagnolo (2018), Musto (2004), Einav et al. (2012), Avery et al. (2009), Avery et al. (2012), Keys et al. (2010), Keys et al. (2012), Rajan et al. (2015), Hurst et al. (2016), Agarwal et al. (2019), Bubb and Kaufman (2014), Laufer and Paciorek (2016)

[3] For example, Rambachan et al. (2020) study the regulation of algorithmic fairness in principal-agent framework.

Budget

- *Item*: Data acquisition: $12,000
- *Item justification*: This project requires merging credit-bureau data with account-level transaction data. We are in negotiations with Plaid, and two other account aggregators, to obtain a large-scale sample of their account-level transaction data for free. However, we need to pay for credit-bureau records that will then be merged to the transaction data. The credit bureau charges for both the data use as well as the work associated with the merge.
- *Item*: Cloud computing costs (AWS): $3,000.
- *Item justification*: We need a cloud computing set-up to store and pre-process large data sets and run computationally-intensive machine learning algorithms. Based on our AWS costs of existing projects, we estimate storage, SQL queries (Athena), and server time (EC2 instances) to cost approximately $3,000.

Data Availability

We anticipate to facilitate replication and data re-use in two ways. First, we are negotiating a clause in our data agreements that would allows us to build a synthetic version of the final data which can be made public. Second, we hope to build a contractual pipeline that other researchers can use to build a similar data merge in the future.

References

Agan, A. and S. Starr (2018): "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment," *The Quarterly Journal of Economics*, 133, 191–235.

Agarwal, S., S. Alok, P. Ghosh, and S. Gupta (2019): "Fintech and Credit Scoring for the Millennials: Evidence using Mobile and Social footprints," .

Athey, S. C., K. A. Bryan, and J. S. Gans (2020): "The Allocation of Decision Authority to Human and Artificial Intelligence," *AEA Papers and Proceedings*, 110, 80–84.

Autor, D. H. and D. Scarborough (2008): "Does Job Testing Harm Minority Workers? Evidence from Retail Establishments*," *The Quarterly Journal of Economics*, 123, 219–277.

Avery, R. B., K. P. Brevoort, and G. Canner (2012): "Does Credit Scoring Produce a Disparate Impact?" *Real Estate Economics*, 40, 65–114.

Avery, R. B., K. P. Brevoort, and G. B. Canner (2009): "Credit Scoring and Its Effects on the Availability and Affordability of Credit," *Journal of Consumer Affairs*, 43, 516–537.

Bartik, A. and S. Nelson (2020): "Deleting a Signal: Evidence from Pre-Employment Credit Checks," *Chicago Booth Research Paper*.

Berg, T., V. Burg, A. Gombović, and M. Puri (2020): "On the Rise of FinTechs: Credit Scoring Using Digital Footprints," *The Review of Financial Studies*, 33, 2845–2897.

Blattner, L. and S. Nelson (2021): "How Costly is Noise? Data and Disparities in US Mortgage Market," .

Blattner, L., S. Nelson, and J. L. Spiess (2021): "Unpacking the Black Box: Regulating Algorithmic Lending," .

Bos, M., E. Breza, and A. Liberman (2016): "The Labor Market Effects of Credit Market Information," 47.

Bubb, R. and A. Kaufman (2014): "Securitization and moral hazard: Evidence from credit score cutoff rules," *Journal of Monetary Economics*, 63, 1–18.

Corbae, D. and A. Glover (2018): "Employer Credit Checks: Poverty Traps Versus Matching Efficiency," SSRN Scholarly Paper ID 3246828, Social Science Research Network, Rochester, NY.

Coston, A., A. Rambachan, and A. Chouldechova (2021): "Characterizing Fairness Over the Set of Good Models Under Selective Labels," *arXiv:2101.00352 [cs, stat]*.

Currie, J. M. and W. B. MacLeod (2020): "Understanding Doctor Decision Making: The Case of Depression Treatment," *Econometrica*, 88, 847–878.

Einav, L., M. Jenkins, and J. Levin (2012): "Contract Pricing in Consumer Credit Markets," *Econometrica*, 80, 1387–1432.

Jann Spiess

Stanford Graduate School of Business
655 Knight Way, Stanford, CA 94305

Phone: 857-284-9880
E-mail: jspiess@stanford.edu

Academic Positions

Stanford Graduate School of Business
Assistant Professor of Operations, Information & Technology, 2019–

Stanford Institute for Economic Policy Research
Faculty Fellow, 2021–

Microsoft Research New England
Post-Doctoral Researcher in Economics, 2018/19

Degrees

Harvard University
PhD, Economics, 2018
AM, Economics, 2015
MPP, Public Policy, 2013

University of Cambridge
MASt, Mathematics (Part III of the Mathematical Tripos), with distinction, 2011

Technical University of Munich
BSc, Mathematics (TopMath), with high distinction, 2010

Research

Publications

Abadie, A., & Spiess, J. (2021). Robust Post-Matching Inference. *Journal of the American Statistical Association*.

Gillis, T. B., & Spiess, J. L. (2019). Big Data and Discrimination. *The University of Chicago Law Review*, 86(2), 459–488.

Ludwig, J., Mullainathan, S., & Spiess, J. (2019). Augmenting Pre-Analysis Plans with Machine Learning. *AEA Papers and Proceedings*, 109, 71–76.

Mullainathan, S., & Spiess, J. (2017). Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives*, 31(2), 87–106.

Grants

A Welfare-Driven Approach to Explaining AI Decisions: The Case of Algorithmic Consumer Lending (with Laura Blattner), Stanford Human-Centered Artificial Intelligence (HAI) Seed Grant, 2021.

Advance the Science of Behavior Change through Machine Learning (with Behavior Change for Good Initiative and Katherine L. Milkman, Sendhil Mullainathan, Lyle Ungar), Schmidt Futures, 2019–2021.

Computational Applications for Behavioral Science (with ideas42 and Susan Athey, Guido Imbens, Jens Ludwig, Sendhil Mullainathan), Alfred P. Sloan Foundation and Schmidt Futures, 2019.

Professional Services

Refereeing

Management Science, Manufacturing & Service Operations Management, Quarterly Journal of Economics, Econometrica, American Economic Review, Review of Economic Studies, Economica, Economic Journal, Journal of Econometrics, Journal of Public Economics, Biometrika, Proceedings of the National Academy of Sciences

Program Committee

ML for Economic Policy, NeurIPS 2020

21th ACM Conference on Economics and Computation (EC'20)

20th ACM Conference on Economics and Computation (EC'19)

Session Organization

Welfare-Driven Optimal Experimentation and Targeting, INFORMS 2020

Leveraging Disparate Sources of Data and Machine Learning to Improve Causal Inference (with Johann Gagnon-Bartsch), 2019 Joint Statistical Meetings

Selected Honors

David A. Wells Prize for best dissertation, Department of Economics, Harvard University, 2018

Restud Tour, 2018

Harvard University Certificate of Distinction in Teaching, 2012, 2014, 2016

McCloy Fellow, John F. Kennedy School of Government, Harvard University, 2011–2013

Bateman Scholar, Trinity Hall, University of Cambridge, 2011

TopMath Study Award & TopMath Presentation Award, Technical University of Munich, 2010

Last updated: January 29, 2021

LAURA BLATTNER

lblattner@stanford.edu
Stanford University ⋄ Graduate School of Business
655 Knight Way ⋄ Stanford, CA 94305

ACADEMIC APPOINTMENTS

Stanford Graduate School of Business, *since 2018*
Assistant Professor of Finance

Louise and Claude N. Rosenberg Jr. Faculty Scholar *since 2019*

EDUCATION

Harvard University *2018*
Ph.D. Political Economy & Government (Economics Track)

Oxford University, Nuffield College *2012*
M.Phil. Economics, Distinction

Oxford University, Merton College *2010*
B.A. Philosophy, Politics & Economics, First Class

RESEARCH INTERESTS

Household Finance, Financial Regulation, AI and FinTech, Banking

WORKING PAPERS

When Losses Turn Into Loans:
The Cost of Weak Banks with Luisa Farinha and Francisca Rebelo
2020 - Conditional Accept, American Economic Review

We provide evidence that banks distort the composition of credit supply in order to comply with ratio-based capital requirements in times of economic distress. An unexpected intervention by the European Banking Authority provides a natural experiment to study how banks respond to falling below minimum required capital ratios during an economic downturn. We show that affected banks respond by cutting lending but also by reallocating credit to distressed firms with underreported loan losses. We develop a method to detect underreported losses using loan-level data. The credit reallocation leads to a reallocation of inputs across firms. We calculate that the resulting increase in input misallocation accounts for about 13% of the decline in productivity in Portugal in 2012.

How Costly is Noise?
Data and Disparities in the US Mortgage Market? with Scott Nelson
2021- Submitted

With the adoption of credit scores and, more recently, the growth of machine learning, concerns about equality and efficiency in credit markets are shifting from human frictions toward "algorithmic frictions" that affect how algorithms allocate credit. Focusing on widely used credit scores, we document that these scores are statistically noisier indicators of default risk for historically under-served groups in the US mortgage market. This noise emerges primarily through the explanatory power of the underlying

credit report data (e.g., thin credit files), not through issues with model fit (e.g., the inability to include protected class in the scoring model). Estimating a structural model of the US mortgage market that builds on a classic Becker test of marginal borrowers' characteristics, we show that reducing algorithmic frictions by equalizing the precision of credit scores can shrink disparities in approval rates and in credit misallocation for disadvantaged groups by roughly half.

Unpacking the Black Box:
Regulating Algorithmic Lending with Scott Nelson and Jann Spiess
2021

We characterize optimal financial regulation in a world where lending decisions are made by complex credit scoring models but regulators are limited in the amount of information they can learn about these models. We show that limiting lenders to algorithms that are simple enough to be fully transparent is inefficient as long as the bias induced by misalignment between social preferences and lenders' preferences is small relative to the uncertainty about the true state of the world. Ex-post algorithmic audits can improve welfare, but the gains depend on the design of the audit tools. Tools that focus on minimizing overall information loss, the focus of most commercial 'explainer' software tools, will generally be inefficient since they focus on explaining the average behavior of the model rather than sources of mis-prediction, which matter for welfare-relevant outcomes. Targeted tools that focus on the source of incentive misalignment, e.g., risk preferences or racial disparities, can provide first-best solutions. We provide empirical support for our theoretical findings using a large-scale credit bureau data set.

What Did 1 Trillion Euros Buy Us? Evidence on
the Real Effects of Quantitative Easing in Europe with Luisa Farinha and Gil Nogueira
2020 - Awarded Lamfalussy Fellowship in 2019

This paper studies the transmission channels of the European Central Bank's (ECB) asset purchase programs via the banking sector using proprietary data from the Bank of Portugal. Banks that hold larger amounts of assets eligible for ECB purchase prior to announcement of the programs realize trading profits from selling these assets following the announcement. Banks use most of these gains to increase their cash holdings. We find a moderate positive effect on loan approval rates for new corporate borrowers, which is stronger for riskier borrowers. However, banks do not offer significantly lower interest rates or provide additional loans to existing customers. We investigate an additional origination channel. The ECB's purchase of asset-backed securities (ABS) and covered bonds does not lead banks with pre-existing issuance technology to originate more loans. Our results suggest that the pass-through of asset purchase programs to lending conditions may occur through channels other than bank balance sheets.

Debt or Demand: Which Holds Investment Back?
Evidence from an Investment Tax Credit with Luisa Farinha and Francisca Rebelo
2018

We study how debt frictions and demand affect corporate investment using administrative data from a large temporary investment tax credit in Portugal. We obtain exogenous variation in demand for exporting firms from product-destination-level changes in foreign demand. We proxy debt frictions by an index of different debt-earnings ratios. Debt has a strong, non-linear effect on the likelihood that a firm invests in response to the tax credit. Firms in the lower two quartiles of our debt-earnings index have roughly equal predicted take-up probabilities. For firms in the third quartile, predicted take-up drops by 50% while firms in the worst debt-earnings quartile have a predicted take-up rate close to zero. We show that the effect of demand is mediated by the size of a firm's debt burden. While demand has a strong positive effect for the bottom debt quartiles, demand ceases to affect take-up in the highest debt quartile. These results highlight that the distribution of debt, rather than the absolute stock of debt, matters for understanding post-crisis investment dynamics.

WORK IN PROGRESS

Quantifying the Financial Accelerator in the Eurozone? with Moritz Lenel
2021

This project aims to quantify the importance of the financial accelerator mechanism by using cross-sectional empirical estimates of the 'bank lending channel' to discipline a quantitative general equilibrium macroeconomic model that features a financial sector. First, we build a meta-database of cross-sectional estimates of the bank lending channel, converted into common units to make estimates comparable across studies. We document the extent and potential sources of heterogeneity in these cross-sectional estimates. Second, we build a quantitative general equilibrium macroeconomic model that features heterogeneous banks and intermediary asset pricing. While our model nests workhorse financial accelerator models in the tradition of Gertler 2010 and Gertler and Karadi 2011, the additional features we introduce allow us to replicate cross-sectional regression designs in our model and to calibrate our model to cross-sectional estimates. The model can also account for key sources of heterogeneity in cross-sectional estimates.

Fiscal Multipliers and Inflation: New Evidence from Tax Rebates with Susan Cherry, Xavier Jaravel, Hala Moussawi
2021

SCHOLARSHIP, HONORS, AND AWARDS

Human-Centered Artificial Intelligence Seed Grant	*2020*
AWS Cloud Computing Grant	*2020*
AQR Top Finance Graduate Award	*2018*
Young Innovators Award, Vordenker Forum, Plansecur	*2018*
BlackRock Applied Research Award (Winner)	*2017*
Lamfalussy Fellowship (For: When Losses Turn Into Loans), European Central Bank	*2017*
Research Grant, Harvard Institute for Quantitative Social Science	*2017*
Harvard University Certificate of Distinction in Teaching	*'16, '17*
Macro Financial Modeling Fellowship, Becker Friedman Institute	*2016*
Research Grant, Lab for Economic Applications and Policy	*2015*
GSAS Summer Pre-dissertation Fellowship, Harvard	*2015*
Research Grant, Hirtle Callaghan Fund	*2015*
Ph.D. Fellowship, German National Merit Foundation	*2014*
ERP Fellowship, German Federal Ministry of Economics and Technology	*2013*

PROFESSIONAL ACTIVITIES

Referee *American Economic Review, Journal of Finance, Quarterly Journal of Economics, Journal of Financial Economics, Journal of International Economic, Journal of the European Economic Association, Review of Financial Studies Review of Finance, Management Science, Journal of Monetary Economics*

PRESENTATIONS

UCSF, RCSF, Ohio State, Chicago Fed	*2021*
UC Davis, NBER SI household finance, NBER Economics of AI, Stanford SITE	*2020*
Berkeley-Stanford seminar, St Louis Fed	
LBS, CMU, Brookings, BYU, UCSC, Banque de France	*2019*
Copenhagen Business School, Chicago (Booth), Northwestern (Kellogg),	*2018*
MIT Sloan (Finance, Applied Economics), Penn (Wharton), Yale (SOM),	
Michigan (Ross), NYU (Stern), Duke (Fuqua), Princeton, Berkeley (Haas),	
Stanford (GSB), Young Scholars Finance Consortium (Texas A&M)	
LSE, Columbia (GSB), NYU Stern WAPFIN	
Bank of Portugal, Bank of England, Harvard University	*2017*

TEACHING EXPERIENCE

Innovative Financial Inclusion (MBA), Stanford GSB	*2020-21*
Empirical Corporate Finance (PhD), Stanford GSB	*2020-21*