

I am an economic theorist studying how information and algorithms shape incentives, allocation, and welfare in markets, organizations, and digital platforms. My research asks how information should be generated, disclosed, sold, and used, and how those choices shape the design of markets and institutions. My work contributes to three areas:

Information Design Much of my work treats information as a policy instrument. Across settings, the central question is the same: what should agents learn, and when? In [3], I frame performance evaluation as a dynamic persuasion problem: organizations choose how much to learn about workers and when to reveal it, taking into account its effects on incentives and retention. The results explain wage compression, tenure-based wage growth, and up-or-out careers as consequences of optimal evaluation design. In [2], I study a dynamic matching market in which capacity-constrained providers decide whether to accept incoming requests. Full transparency induces cream skimming and lowers welfare by reducing trade and thinning the market. The paper shows that simple information policies can restore efficient trade when sellers are homogeneous and improve welfare more generally. [7] develops a welfare framework for information policy. It characterizes welfare outcomes achievable through disclosure rules and shows that Pareto-efficient welfare distributions can be recovered through persuasion problems. The analysis recasts algorithmic fairness as an economic question about who gains and who loses under alternative information policies. Finally, [9] provides a tractable method for information design in classic economic games. A duality-based approach yields certificates of optimality and robustness and explains why simple policies, such as targeted disclosure, often suffice across settings ranging from polarization games to competitive pricing. Across these papers, the broader message is that more transparency is not always better: well-designed information can improve incentives, reduce externalities, and outperform full revelation.

Mechanism Design, Screening, and Pricing A second line of research studies how screening, disclosure, and product design interact when contracting parties have private information. The applications include data markets, platforms, and complex product pricing. [1] develops a framework for data markets. A monopolist sells data to a buyer with private beliefs. Screening works through data design as well as prices, and the paper characterizes simple optimal menus of directional tests. [5] studies a seller who jointly chooses what to disclose about a multi-attribute product and how to price it. I show that optimal disclosure often takes a linear form and identify when no disclosure is optimal. The results provide a tractable theory of selective disclosure and shrouding. In [4], I ask when upgrade pricing, nested bundles with increasing quality, is optimal and develop tools showing when multidimensional screening collapses to a simple tiered menu. In [6], mechanism design moves upstream: an agent under moral hazard strategically chooses the production technology on which she will later be screened. The paper shows how coarse technologies can be used to engineer agency rents. [8] studies data governance when a privately informed seller chooses what consumer data to access and shows that effective screening is often impossible: many outcomes from individualized data access can already be replicated by a common public rule. Most recently, [11] studies mechanism design when the same institution is used repeatedly and agents learn from the outcomes they observe. The paper shows how this learning constrains implementable outcomes and provides a tractable characterization of optimal designs by combining information design with mechanism design. Together, these papers show that private information shapes what can be screened, what can be versioned, and why simple pricing rules are often exactly optimal.

The Economics of Artificial Intelligence My recent work applies these tools to the design, pricing, and governance of AI systems. [10] studies product recommendation algorithms when sellers respond by adjusting prices. I show that recommendation systems should sometimes distort what users see to counteract seller pricing power and neutralize the welfare losses from finer market segmentation. [12] develops a framework for markets for large language models. I show that a high-dimensional pricing problem can be reduced to one dimension, yielding predictions about token allocation and rationalizing simple menus such as committed-spend contracts, model versioning, and linear token pricing. [13] studies how a decision-maker should use advice from an informed but potentially misaligned AI system. The optimal robust policy creates a zone of calibrated trust: moderate recommendations are used at face value, while extreme ones are clipped to safer interpretations, rationalizing guardrails such as second reads and corroboration. I see AI as the next major area in which economic theory can shape institutions, and several related projects are underway on AI-mediated markets, human-AI interaction, and the governance of data and models.

REFERENCES

- [1] Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin (2018) “The Design and Price of Information,” *American Economic Review*, 108(1), 1–48.
- [2] Romanyuk, Gleb and Alex Smolin (2019) “Cream Skimming and Information Design in Matching Markets,” *American Economic Journal: Microeconomics*, 11(2), 250–276.
- [3] Smolin, Alex (2021) “Dynamic Evaluation Design,” *American Economic Journal: Microeconomics*, 13(4), 300–331.
- [4] Bergemann, Dirk, Alessandro Bonatti, Andreas Haupt, and Alex Smolin (2022) “The Optimality of Upgrade Pricing,” in *Web and Internet Economics*, 41–58.
- [5] Smolin, Alex (2023) “Disclosure and Pricing of Attributes,” *RAND Journal of Economics*, 54(4), 570–597.
- [6] Garrett, Daniel F., George Georgiadis, Alex Smolin, and Balázs Szentes (2023) “Optimal Technology Design,” *Journal of Economic Theory*, 209, 105621.
- [7] Doval, Laura and Alex Smolin (2024) “Persuasion and Welfare,” *Journal of Political Economy*, 132(7), 2451–2487.
- [8] Ichihashi, Shota and Alex Smolin (2025) “Data Provision to an Informed Seller,” *Games and Economic Behavior*, 153, 131–144.
- [9] Smolin, Alex and Takuro Yamashita (2025) “Information Design in Smooth Games,” conditionally accepted at *Theoretical Economics*.
- [10] Ichihashi, Shota and Alex Smolin (2025) “Buyer-Optimal Algorithmic Recommendations,” working paper, revision requested at *Review of Economic Studies*.
- [11] Doval, Laura and Alex Smolin (2026) “Calibrated Mechanism Design,” working paper.
- [12] Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin (2026) “Menu Pricing of Large Language Models,” working paper.
- [13] Dworczak, Piotr and Alex Smolin (2026) “Robust Trust,” working paper.