Research Statement for Brad Larsen, Stanford University Dept. of Economics, 2021 (latest version here)

I am an empirical industrial organization (IO) economist who studies intertwined topics: bargaining, occupational licensing, and online markets. My work explores how efficient/equitable bargaining is, how well licensing works, and how frictions affect online markets. My research has been funded by 4 NSF grants and grants from the Sage, Arnold, and Hellman Foundations, as well as year-long fellowship awards at Hoover, IRiSS, and Stanford Impact Labs.

1) Bargaining: My work examines efficiency and equity in real-world bargaining. I have long been fascinated by bargaining. Eleven years ago at MIT, I acquired unique data: every sequential offer from hundreds of thousands of business-to-business negotiations at used-car auctions—a scale and detail unheard of in empirical bargaining work. This sparked more than a decade of investigation into the *efficiency* and *equity* of bargaining.

1.A) Efficiency: The first strand of my bargaining work shows that, when a negotiating buyer and seller disagree, it is often inefficient; they are missing out on gains from trade. In my used-car data, over 30% of negotiations fail. This impasse is *inefficient* if the buyer truly values the car more than the seller (and the impasse is *efficient* otherwise). So how inefficient is real-world bargaining? This fundamental question was unanswered, despite decades of previous bargaining theory and lab experiments, because buyer and seller values are *private*—they aren't recorded anywhere in field data. The methods I designed, and my novel data, made me uniquely positioned to answer this question.

I developed a new method to infer *bounds* on private values in bargaining. For example, when a seller accepts an offer, she must prefer that offer over her private value. And when a seller walks away from an offer on the table, she must prefer keeping the car, so her private value exceeds the offer. These intuitive inequalities, applied to data, yield bounds on the *distributions* of private values, which can be used to quantify inefficiency. In "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used Auto Auctions," (*Review of Economic Studies*, 2021), I found that at least 21% of negotiations end in disagreement even though trade gains exist—the buyer's value exceeds the seller's—a large inefficiency indeed. I found twice that level of inefficient impasse (43%) in eBay bargaining ("How Well Does Bargaining Work in Consumer Markets? A Robust Bounds Approach," working paper, with Freyberger). For this latter analysis, we constructed even more general bounds on private values that allow for the possibility that *consumers* (who comprise our eBay negotiators) may be less sophisticated than the *professional* negotiators in my used-car data.

I investigated a potential cause of this inefficiency in used-car markets: auction house employees mediating negotiations. With Lu and Zhang, in "Intermediaries in Bargaining: Evidence from Business-to-Business Used-Car Inventory Negotiations" (working paper), we showed that human mediators can be thought of varying in their ability to implement the theoretically efficient mechanism, and this skill variation explains much of the inefficient impasse: the 75th percentile mediator is 22 percentage points more likely to seal a deal than the 25th. We also found that better mediators know when to stop bargaining; they do not push disagreeing parties to persist.

These three studies, and those described in 1.B below, represent some of the first empirical bargaining research incorporating incomplete/asymmetric information. This stands in stark contrast to most empirical IO work, which assumes complete information (e.g., Nash bargaining), where agents perfectly know opponents' values and hence no impasse occurs. Allowing for incomplete information in empirical work is nontrivial, and little guidance is offered by existing game-theoretical models, which either ignore the possibility of impasse completely or only permit one or two offers, clearly violating the real-world data I work with. For these reasons, I designed robust bounding tools to quantify inefficiency. Motivated by questions that theory or standard empirics alone cannot address, my work might be characterized as "empirical theory" or "empirical mechanism design"—empirical analysis in the spirit of mechanism design theory (e.g., the Myerson-Satterthwaite Theorem).

1.B) Equity: The second strand of my bargaining work looks at fairness in real-world bargaining. In used-car markets, a sea of state-level laws govern manufacturer-dealer relationships, which previous work argues raises the *bargaining power* of dealers (buyers in my data). Unlike complete-information, Nash bargaining, no standard notion of bargaining power or equity exists for incomplete-information. Zhang and I proposed one in "Quantifying Bargaining Power Under Incomplete Information: A Supply-Side Analysis of the Used-Car Industry" (working paper). We developed a method relying on mechanism-design arguments to identify (rather than bound) agents' private values. We found that manufacturers (e.g., Ford, GM) have more bargaining power than dealers and more than other sellers (e.g., rental companies). Our earlier methodological working paper offers extensions to other settings ("A Mechanism Design Approach to Identification and Estimation").

In my analysis of used-car and eBay data, I came across a curious (and previously undocumented) fact regarding *fairness* in real-world bargaining: agents favor offers halfway between the two most recent offers. Previous lab

experiments documented a preference for fairness, but only in ultimatum games with *complete* information (where agents know the pie). In the settings I studied, I observed clear evidence instead of *incomplete* information*;* I documented this evidence (and other aspects of fairness) on eBay with Backus, Blake, and Tadelis, "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions" (*Quarterly Journal of Economics*, 2020). For that analysis, we introduced a massive new dataset—90 million eBay listings with 25 million distinct bargaining sequences—that we were permitted to release publicly for other researchers. In "Fairness in Incomplete Information Bargaining: Theory and Widespread Evidence from the Field" (working paper), with Keniston, Li, Prescott, Silveira, and Yu, we showed that this propensity to favor *split-the-difference* offers extends far beyond eBay and used-car markets, occurring in international trade, housing, pre-trial settlement, a game show, and transportation. We constructed a formal theory of this behavior (and documented consistent evidence) for how this can be *fair*, showing it can be rationalized by agents' robust inferences about the size of the pie.

My bargaining work takes an approach to analysis that can be considered *auction*-like, in that I take seriously *incomplete information*. Indeed, the bounds approaches described in 1.A above are in the spirit of English auction work by Haile and Tamer (2003), while the mechanism design approach in 1.B is a generalization of Guerre, Perrigne, and Vuong (2000) from first-price auctions to bargaining. We developed these tools because we needed them; no off-the-shelf tools existed for studying incomplete-information bargaining. A similar need to develop tools arose in studying used-car *auctioneers* in "Bid Takers or Market Makers? The Effect of Auctioneers on Auction Outomes" (2016, *American Economic Journal: Microeconomics*) with Lacetera, Pope, and Sydnor. There we showed that human auctioneers exploit bidders' behavioral biases to improve the sale probability. I wanted to know how important the auctioneer was relative to traditional auction-design instruments (e.g., reserve prices). No tool existed to address this question, so Coey, Sweeney, and I developed one: a bound on revenue gains from auction design that can be computed as the average gap between the second- and third-highest bids scaled by 2/n (the number of bidders). See "The Bidder Exclusion Effect," *RAND Journal of Economics*, 2019). We explored other empirical auction bounds with Waisman in "Ascending Auctions with Bidder Asymmetries" (*Quantitative Economics*, 2017).

2) Occupational Licensing: My work studies whether these laws make sense. A second strand of my research takes an IO perspective on occupational licensing—one of the most common forms of entry regulation, affecting nearly 25% of U.S. workers by barring them from the market if they do not satisfy certain requirements. While clearly of interest for labor economics, I consider vital the *IO perspective* on the topic, which is infused with questions of asymmetric information, entry regulation, reputation, and antitrust. I see my work as bringing in core questions that, while at the center of occupational licensing debates, are understudied, including, "Does the left tail of quality (not just the average) improve with stricter licensing? Could reputation systems reduce a need for licensing? Are licensing restrictions consistent with safety concerns? Are entry barriers imposed by licensing laws an antitrust concern?"

In "The Effect of Occupational Licensing Stringency on the Teacher Quality Distribution" (working paper), with Ju, Kapor, and Yu, we showed that licensing stringency raises the left tail of public school teacher quality, contributing to a decades-long debate over teacher certification. To study these tail effects, we exploited a quantile regression approach that I developed with Chetverikov and Palmer, "IV Quantile Regression for Group-Level Treatments, with an Application to the Distributional Effects of Trade" (*Econometrica,* 2016). Licensing laws vary widely by state/year, constituting a group-level treatment, which our method was designed to handle. Several previous studies analyzed effects on *average* quality; we focus instead on the main pro-licensing argument—raising the left tail.

For decades, opponents have argued that licensing is unnecessary when adequate reputation mechanisms exist. This argument may be even more important given the rise of online markets. In work with Farronato, Fradkin, and Brynjolffsen, "Consumer Protection in an Online World: An Analysis of Occupational Licensing" (working paper), we studied new data from a large online labor platform where consumers hire professionals via procurement auctions (many aspects of my work are intertwined with auctions and with online markets). We showed that consumers value reputation more than licensing, and stricter licensing reduces competition and raises prices without improving quality.

The requirements for licensure vary widely across states and professions, sometimes in counterintuitive ways, mandating that barber training be ten times as long as EMT training. In "Regulatory Inconsistency: Occupational Licensing and the Consequence of Error" (ongoing work), I propose a method to aggregate pairwise occupation comparisons to quantify how *inconsistent* regulation is with safety arguments. Barbers, cosmetologists, and opticians come out particularly high on the inconsistency ranking. With fellowship funding from Stanford Impact Labs, I have spent the past year designing a policy lab that will take these insights to state legislatures and agencies and help them identify and change regulations that make the least sense. Through this effort, I have engaged with the Council of State Governments, the Department of Consumer Affairs and State Contractors Board of California, and other groups.

The FTC is currently keenly interested in occupational licensing. The Supreme Court recently ruled, in *NC State Board of Dental Examiners v. FTC*, that state licensing board actions can violate the Sherman Act. In ongoing work, I am investigating the impacts of state licensing board makeup on market outcomes in each licensed profession.

3) Online Markets: My research considers efficiencies and frictions in online markets, how they are impacted by market design and regulation, and who the winners and losers are. I am particularly interested in consumer search, price discrimination, and copyright infringement. In "Discounts and Deadlines in Consumer Search" (*American Economic Review*, 2020), with Coey and Platt, we introduced a notion of a *deadline* into search. Deadlines lead consumers to grow increasingly impatient and willing to pay more during a search. Using eBay auction data, which allowed us to track individuals over time, including at each failed attempt to secure a product, we found robust evidence consistent with deadlines. A major theme throughout my work is an interest in *inefficiency* in real-world markets. Here we quantified dynamic inefficiencies introduced into search markets by deadline-motivated behavior.

A current policy debate for tech companies revolves around the regulation of and access to consumer data. In "Dynamic Competition in the Era of Big Data" (working paper), with Kehoe and Pastorino, we designed a model of dynamic competition between firms accessing personal data. Our model describes *personalized pricing* of experienced goods, where firms may manipulate consumers' tastes through prices and product offerings. We estimated dynamic demand from online sales of Apple/Samsung products and simulated counterfactual personalized pricing. We found that personalized pricing can be efficient, with the average consumers unharmed when there is sufficient competition and symmetric firm access to personal data. This idea of how firms may exploit Big Data also appears in my work with Coey, Sweeney, and Waisman ("Scalable Optimal Online Auctions," Marketing Science, 2021, lead article), where we showed how optimal design of online auctions can be greatly simplified with large datasets.

I am also interested in digital copyright law. In "Identification in Ascending Auctions, with an Application to Digital Rights Management" (*Quantitative Economics,* forthcoming), Freyberger and I analyzed whether consumers react to digital copyright law. We estimated willingness to pay for unlocked smartphones before and after a 2013 change to the Digital Millennium Copyright Act (DMCA) banning unlocking. We found that buyer values decreased by 39%. In separate ongoing analysis of copyright in digital markets ("The Effect of Legitimizing Arbitrage Through Parallel Trade: A Textbook Case"), Dana, Moshary, and I exploit a recent Supreme Court case (*Kirtsaeng v. John Wiley*) in international textbook markets to analyze how online arbitrage erodes producers' price discrimination. We have also acquired data on millions of DMCA copyright infringement notices sent to Google, allowing us to study how the legalization of arbitrage affects piracy in this market.

During the pandemic, I have pivoted some of my online markets work to run a massive RCT (currently in the field) on Facebook using content my team (various researchers and professional videographers) designed to appeal to vaccine-hesitant Republicans, including military messengers and rare footage of Trump endorsing vaccines. We measure effects on vaccine uptake (from CDC) and attitudes (from Facebook ad auction metrics). Earlier in the pandemic, we ran a similar experiment encouraging masks.

Teaching and Service. I teach two half-quarter graduate IO courses (Econ 257 and 258), covering broad IO subjects and incorporating my work on bargaining and other topics. These courses attract students from economics, marketing, finance, political science, and engineering. The content includes material from previous lecturers and content I have developed and continue to modify. I have taught undergraduate IO (Econ 157) for several years, largely with content I have prepared. The course consistently receives very high ratings from undergraduates, who enjoy learning practical applications of economics. I also taught a full-quarter economics course to Stanford undergrads studying abroad in Madrid in 2019 (OSP 40, "Pirates, Soccer, and Dons: A Sampler of Economics and Data Science in Spain"). The course introduced students (many who were not economics majors) to microeconomics, econometrics, and statistical programming using real-world examples from historic and modern Spain. 24 of the 32 Stanford-in-Madrid students chose to enroll in the course. I will teach a similar course I designed for Stanford's Florence program in 2022. Students in any of my classes will be quick to say that I am an enthusiastic teacher—whether talking about bargaining, licensing, IO, or learning in general. I thoroughly enjoy mentoring undergraduate and graduate research assistants (typically 2-5 at any point in time). Six of my RAs or former advisees have gone on to PhDs. I also enjoy coauthoring with students: Ju, Lu, Yu, Waisman, and Zhang are examples. I have advised or am advising 11 PhD dissertations and 5 undergraduate honors theses.

During my time at Stanford, I have co-organized several conferences, including a SITE conference on negotiation with researchers from economics, organizational behavior, and social psychology (canceled due to COVID-19); a SITE conference on occupational licensing; and the NBER Summer Institute IO meetings.