

Research Summary

July 2018

Robin S. Lee (robinlee@fas.harvard.edu)

Research Fields: Industrial Organization, Applied Theory, Health Economics

Industries are rarely perfectly competitive. They also rarely comprise only firms who transact solely with consumers. Rather, in many industries—including health care, technology, media, and finance—firms wield significant market power and negotiate with one another over terms of trade. Standard economic theory—generally tailored to settings where agents either unilaterally choose prices or are price-takers—has often struggled to provide sharp predictions in these more realistic and complicated environments. Yet understanding how firms compete and interact in these markets is critical for regulatory, innovation and competition policy, and for anticipating how market-level and macroeconomic shocks propagate through industrial supply chains to affect welfare.

My research attempts to address these shortcomings and deepen our understanding of how firms bargain, contract and form supply relationships in imperfectly competitive markets. There are two primary goals of my work. The first is to develop formal analytical frameworks for modeling firm competition that can be “taken to data”: that is, flexible enough to capture important institutional details of real world industries, and tractable enough for estimation and the evaluation of counterfactual scenarios. The second is to adapt and apply these frameworks to a diverse set of economically significant industries, and synthesize both theoretical and empirical insights to inform policy-relevant areas of debate.

I. Empirical Studies of Bargaining, Integration, and Competition in Vertical Markets

a. Health Care Markets

My most active area of research focuses on supply-side responses to market structure and policy changes in the U.S. private and commercial health care market (representing over 60% of non-elderly consumers), and illuminates how competition among insurers and medical providers affects expenditures, access, and quality. Ultimately, I seek to improve the evaluation of merger policy and network adequacy regulation, enable more effective health marketplace design, and assist final payors for health care services—particularly self-insured employers and government agencies—in working with insurers to create more sophisticated plans that can constrain spending while improving consumer welfare.

Much of my work has focused on the bilateral negotiations between insurers and medical providers that are central to health care provision in the U.S.: these negotiations determine insurers’ “networks”—the set of doctors and hospitals that a consumer can utilize—and providers’ reimbursement rates. These rates, in turn, determine much of the annual \$1 trillion in spending in the private insurance sector (approximately 40% of total national health expenditures) and are projected to continue growing in the coming years. In general, however, little is known about their determination, and how they would be affected by policy or market structure changes.

In response to this need, *Insurer Competition in Health Care Markets* (with Kate Ho, *Econometrica*, 2017) develops a framework to analyze the complex supply-side interactions in employer-sponsored

health care markets. At its heart is a structural model that incorporates bilateral bargaining between hospitals and insurers, insurer premium negotiations with employers, household demand for insurance products, and individual demand for hospitals. The paper estimates the model using a unique and comprehensive data set comprising eligibility and claims information for all employees and their dependents of a major California benefits manager, and uses estimates to study how insurer competition for patients—affected by the formation of health insurance exchanges, entry or exit by insurers, or changes in the variety of employer provided plans—affects payments, premiums, and welfare. Though increased insurer competition may lead to lower premiums holding fixed hospital payments, it may also provide hospitals with greater bargaining leverage, thereby leading to higher negotiated rates and mitigating premium reductions. We simulate equilibrium outcomes upon the removal of an insurer from each market and find that reduced insurer competition—though often reducing consumer welfare—can nevertheless lead to potential decreases in hospital rates across markets, and redistribute rents across health providers. Critical determinants for whether such changes are ultimately harmful to consumers include the characteristics of the removed insurer, and the presence of effective constraints on premium levels such as those imposed by a large employer purchasing insurance for its employees.

The tradeoffs this paper identifies and measures are central issues in active areas of health care competition policy, including the evaluation of the \$54 billion Anthem and Cigna insurance merger that was challenged by the Department of Justice in 2016 and blocked in 2017. Additionally, the theoretical model it develops is used in *The Price Effect of Cross-Market Hospital Mergers* (with Leemore Dafny and Kate Ho, revised and resubmitted, *RAND Journal of Economics*) to examine “cross-market” hospital mergers that occur across distinct product or geographic markets. Most existing hospital merger models employed by the courts and antitrust agencies have had difficulty reconciling price increases from hospitals that do not compete with one another at the point of service for the same consumer. However, as we articulate in this paper, such hospitals may still be viewed as substitutes from the perspective of insurers that bundle hospital services for sale to multi-market employers or multi-person households; consequently, a merger may enable the parties to obtain higher negotiated rates. We pair this theoretical point with a new data set of cross-market hospital mergers over a 12-year period; we focus on the prices of hospitals in an acquiring system that are not located in the same local geographic market as hospitals of the acquired system, and estimate sizable positive price effects following such mergers. Our results suggest that additional scrutiny is warranted.

Most recently, *Equilibrium Insurer-Provider Networks: Bargaining and Exclusion in Health Care Markets* (with Kate Ho, forthcoming, *American Economic Review*) studies the efficacy of increasingly prevalent “narrow network” products—whereby insurers selectively contract with a limited set of medical providers—and evaluates the effects of network adequacy standards, such as those considered by the Centers for Medicare and Medicaid Services and several state regulators. We significantly extend existing models of bargaining in bilateral oligopoly by developing a new bargaining solution concept that allows for an insurer to exclude hospitals from its network in order to “play them off” against one another, and predicts which provider networks are formed in equilibrium (see further discussion below). We incorporate this into our earlier framework (Ho and Lee, 2017), and use previously obtained estimates to simulate equilibrium outcomes in California had existing regulations on network breadth been relaxed. We find that our new solution can rationalize levels of exclusion that we observe in our data, whereas previously used bargaining solutions cannot. In our setting, we establish that bargaining motives introduce an economically meaningful incentive to distort network breadth and quality away from the social optimum; furthermore, we show an insurer often benefits from exclusion through substantially lower negotiated hospital rates, and that regulation prohibiting exclusion generally increases rates and premiums without significantly influencing social surplus. However, such regulation

has distributional consequences, and can prevent significant harm from consumers living close to otherwise excluded hospitals. We also find that although an insurer's incentives to exclude may be greater than those faced by a social planner, they may be relatively well-aligned with consumer and employer preferences.

This last point suggests that insurers and final payors of health care services may be able to work together to constrain health care spending by designing customized networks. Ongoing work, **Health Insurance Menu Design: Managing the Spending-Coverage Tradeoff** (with Kate Ho, in preparation), examines whether these parties can realize savings along alternative dimensions as well. Using a new comprehensive dataset covering insurance enrollment decisions and health care claims for a large employer over a three year period—where this employer introduced a new set of high-deductible health plans in the middle of our sample period—we aim to estimate the distribution of health risks and preferences for the covered population using a model of plan choice and health care utilization that controls for selection and moral hazard. Our goal is to develop a framework that can be used by employers seeking to optimize the menu and financial characteristics of plans offered to employees. Our project will shed light on the optimal set of plans that balances cost-control with coverage and risk protection motives, and how this depends on employee attributes.

Finally, much of my work in health care markets examines changes in reimbursement rates and premiums; these changes may reallocate consumers across insurers and medical providers, and have implications for entry, exit, and investment behavior by firms. It is also the case that such changes may also influence the nature of health care delivery itself. **Physician Prices, Hospital Prices and Treatment Choice in Labor and Delivery** (with Kyna Fong and Patricia Foo, *American Journal of Health Economics*, 2017) provides evidence (and contributes to a broader literature) linking changes in provider reimbursement rates to changes in health care utilization. This paper uses within-provider variation in payments from different insurers across time to establish that physicians appear to respond to financial incentives when making certain procedure choices.

b. Technology and Media Markets

Another strain of my research explores the effects of vertical integration in technology and media markets. Although there is a substantial theoretical literature on the pro- and anti-competitive consequences of integration, the empirical evidence on magnitudes and overall welfare effects is still limited. Two of my papers contribute to this literature.

First, in **Vertical Integration and Exclusivity in Platform and Two-Sided Markets** (*American Economic Review*, 2013), I explore whether a dominant hardware platform's use of exclusive and integrated software was anti-competitive, allowing it to unfairly maintain its early advantage, in the U.S. videogame industry from 2000-05. During this period, the industry was dominated by a small number of hardware (platform) manufacturers and a concentrated market of software developers. I develop and estimate a dynamic model of consumer demand for hardware and software, and software developer "demand" for which platforms to support, to simulate market outcomes had hardware providers been unable to integrate or acquire exclusive software titles. Whereas previous empirical literature in hardware-software markets often assumes consumers care only about the number of software titles onboard a platform, this paper crucially allows the benefit from each title to differ; given the skewed sales and hit-driven nature of the industry, this flexibility is important for identifying the marginal benefit of a single software product to each hardware platform's sales. The paper makes a number of additional methodological contributions to account for dynamics (products are durable, and consumers anticipate future price changes and software releases) and preference heterogeneity (consumers select onto and

across platforms over time), features that are pervasive in most technology industries. Results indicate that exclusivity and integration may have enhanced platform competition and aided smaller, entrant platforms: without exclusive arrangements, high quality software would typically be released on all platforms; lower quality titles, constrained by the costs of developing for multiple systems, would likely have developed first for the larger incumbent platform due to its larger user base, and only later, if at all, developed a version for others. As a result, neither entrant platform would have been able to offer consumers any benefit over the incumbent; with exclusivity, however, entrants could create a competitive advantage, and was in fact leveraged by them to gain traction in this networked industry. These results suggest that integration and exclusivity—while perhaps affecting the entry of new software products differently (as in *U.S. v. Microsoft*)—may foster platform competition.

Second, in ***The Welfare Effects of Vertical Integration in the Multichannel Television Market*** (with Greg Crawford, Michael Whinston, and Ali Yurukoglu, *Econometrica*, 2018), we examine the role of integration by content distribution into content provision in the multichannel television industry. Such issues are at the heart of the recently proposed and challenged \$85 billion merger between AT&T and Time Warner, as well as other approved (Comcast/NBC in 2011) and abandoned (Comcast/TimeWarner in 2015) mergers in the past. In the US, over 110 million households subscribe to a multichannel television operator; households are limited to receiving only channels with which their operator has contracted. Though integrated channels generally are not allowed to be withheld from rival distributors, a loophole in legislation allowed for the exclusion of rivals in certain cases. We construct a comprehensive data set on the U.S. multichannel television industry, collected and synthesized from numerous sources, that contains information on viewership and subscription patterns, channel ownership and integration status, and prices, quantities, and channel carriage “lineups” for cable and satellite bundles for the years 2000 to 2010. Leveraging substantial cross-market variation in ownership and regulatory policy facing regional sports networks in our data, we develop and estimate a structural model of viewership, subscription, distributor pricing, bargaining, and carriage decisions. Our applied framework, adaptable to other content distribution settings (and indeed referenced in the recent AT&T/Time Warner merger case), is one of the first to capture and incorporate the potential for double marginalization, incentives to engage in foreclosure, and—also novel to this literature—the possibility of imperfect coordination and internalization within an integrated firm. We use our model and estimates to comprehensively measure the potential anti-competitive and efficiency enhancing effects of integration by simulating a series of vertical mergers and divestitures between content and distribution. In our setting, we find that integration, even absent regulations ensuring access to integrated content by rival distributors, increases consumer and total welfare on average.

c. Financial Exchanges

I have recently begun work that synthesizes methods and results from a variety of fields, including market microstructure and market design, to better understand the industrial organization of financial markets. In the U.S., approximately a dozen stock exchanges bring together investors and trading firms, and handle more than 1.5 trillion shares (representing \$60 trillion in value) annually. These exchanges use variants of the continuous limit order book market design, a design that Budish, Cramton and Shim (2015, *QJE*) show enables latency arbitrage, results in rents being extracted by high frequency trading firms (HFTs) from investors, and leads to a socially wasteful high-frequency arms race. In ***Will the Market Fix the Market? A Theory of Stock Exchange Competition and Innovation*** (with Eric Budish and John Shim, in preparation), we build and analyze a model of competition among stock exchanges to understand whether market participants have incentives to adopt new market designs that ameliorate these inefficiencies. Our formal results are derived from a multi-stage game of competition among

exchanges, sophisticated HFTs, and investors. One of our theoretical contributions is the development of a new solution concept that, in a manner analogous to the “reactive” informational equilibrium notion of Riley (1979, ECMA) used to restore existence in models of insurance, enables us to analyze informed trading behavior across multiple exchanges. We elucidate the relationship between exchanges and HFTs, and show that exchanges, through the levying of exchange-specific technology fees for co-location and proprietary data fees, are able to capture some of the rents that HFTs extract from investors; critically, these rents are not dissipated through standard price competition. We use detailed trades and quotes data paired with financial filings to confirm that the theoretical predictions of our model conform to empirical realities. We thus argue that private incentives to adopt new market designs that eliminate latency arbitrage rents are socially inefficient, and suggest potential policy responses.

II. Network Formation and Bargaining in Bilateral Oligopoly

What links together the industries presented in the previous section is that they are vertical markets characterized by bilateral oligopoly: in all cases, a small number of intermediaries (health insurers, hardware manufacturers, content distributors, financial exchanges) negotiate with other firms (medical providers, software developers, content providers, trading firms) over prices and terms of trade.

Theoretical models of multilateral bargaining are often either too stylized (for example, they rule out contracting externalities that are prevalent in these settings) or too complicated for use in applied work. One exception, however, is the “Nash-in-Nash” solution proposed in Horn and Wolinsky (1988, *RAND*). This solution, adopted by recent empirical work in industrial organization and in several of the papers described in this research summary, is a set of agreements (specifying, for example, prices) for all pairs of contracting agents such that each agreement is the solution to a bilateral Nash bargaining problem for each pair—that is, a “Nash equilibrium in Nash bargains.” However, as Nash bargaining is a cooperative game theory concept, there is a question of whether the more general Nash-in-Nash solution concept is consistent with a non-cooperative model.

“Nash-in-Nash” Bargaining: A Microfoundation for Applied Work (with Allan Collard-Wexler and Gautam Gowrisankaran, *Journal of Political Economy*, forthcoming) makes progress towards answering this question. We provide a non-cooperative foundation for the Nash-in-Nash solution that generalizes the Rubinstein (1982, *ECMA*) alternating offers bargaining game to a setting with multiple firms. We provide necessary and sufficient conditions under which there exists an equilibrium of this game that generates the Nash-in-Nash solution, and provide stronger sufficient conditions under which the equilibrium outcome is unique. Importantly, we allow firms to engage in multilateral deviations across multiple bargains, and do not restrict attention to stationary strategies. Thus, the paper establishes that this solution—fast becoming a workhorse bargaining model in applied work—may be reasonable for determining the division of surplus in certain settings.

Nevertheless, in spite of its advantages, the Nash-in-Nash solution has certain limitations. First, it is often employed as a surplus division rule for a *given* set of contracting partners or network; as such, it typically provides limited guidance as to which network(s) will form in equilibrium. Second, standard implementations often implicitly assume disagreement points when bargaining are “static,” and firms cannot not adjust other contracting decisions when there is disagreement. These issues are not innocuous, and can lead to erroneous predictions with economically substantive implications.

A series of papers attempts to address these limitations. In ***Equilibrium Insurer-Provider Networks: Bargaining and Exclusion in Health Care Markets*** (discussed above), we develop a new bargaining solution concept, “Nash-in-Nash with Threat of Replacement,” that builds on Nash-in-Nash. The solution, which we prove emerges from another formal non-cooperative alternating offers bargaining

game, endogenizes firms' outside options to include the possibility of swapping one trading partner for another. This is an important extension in health care environments: for example, if an insurer negotiates with only one of two children's hospitals in a market, its disagreement point under Nash-in-Nash bargaining typically involves having no children's hospital in its network; instead, it may be more plausible that the insurer is able to form a contract with the other hospital upon disagreement. We show that Nash-in-Nash cannot rationalize observed levels of exclusion, whereas our new concept can. Though the implementation is tailored to our health care setting, the framework can be applied more widely to environments where firms can commit to a set of contracting partners, and credibly use excluded partners as bargaining leverage.

In ***Markov Perfect Network Formation: An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks*** (with Kyna Fong, revise and resubmit, *Review of Economic Studies*), we adopt an alternative approach. We explicitly move away from static analysis, and develop a flexible dynamic model of network formation and bargaining that admits multiple upstream and downstream firms and allows each firm's contracting partners to change over time. Reflecting real world bargaining disputes and breakdowns, these changes are anticipated by firms when contracting and bargaining, and affect both negotiated payments and the equilibrium set of trading partners. We show how the framework can be used to estimate underlying structural parameters that govern static payoffs. We also use this framework to simulate hospital mergers in health insurance markets, and demonstrate the impact of dynamic re-contracting and bargaining on merger effects.

III. Theoretical Work on Exclusive and Selective Contracting in Platform Markets

Earlier papers, begun in graduate school (though published later), represent initial forays into examining vertical markets. Using stylized theoretical models, this line of research analyzes the intense competition among a small number of downstream firms, or *platforms*, to get upstream suppliers onboard. Though these theoretical exercises prove useful explorations of what is possible and influenced my later research, they do not on their own identify which mechanisms are empirically relevant, nor necessarily predict market outcomes without further assumptions on underlying economic primitives. (Motivated by a desire to resolve such ambiguities, my later work, discussed above, relies more heavily on empirical and econometric analyses, and dives into specific industries.) I summarize these papers here.

The first paper in this earlier line of work, ***Competing Platforms*** (*Journal of Economics and Management Strategy*, 2014), explores the question of why multiple platforms exist even if it would be efficient to only have one. Such is often the case when network effects and coordination benefits are large. Using a stylized model, I show how strategic contracting between suppliers and intermediaries can lead to an inefficient "market-splitting" outcome in which multiple platforms are active. It thus provides an alternative explanation for platform multiplicity that does not rely on standard arguments, which include coordination failure. The second paper, ***Exclusivity and Control*** (with Andrei Hagiu, *Journal of Economics and Management Strategy*, 2011), studies when a supplier might choose to exclusively provide its services through a single platform as opposed through multiple ones. The paper analyzes the tradeoff faced by firms: although a supplier forgoes reaching a larger audience by not joining multiple platforms, a platform benefits when a supplier joins exclusively as it is able to extract greater rents from consumers through higher prices. If the second effect is larger, then the platform can compensate the supplier for exclusivity. We show which effect dominates depends on measurable properties of consumer demand and the intensity of competition. The third paper, a policy oriented piece entitled ***Subsidizing Creativity: Zero Pricing and Network Neutrality*** (with Tim Wu, *Journal of Economic Perspectives*, 2009), applies insights from the theoretical literature on two-sided markets and the aforementioned work to the "net neutrality" debate. We articulate why certain aspects of net

neutrality, specifically a ban on fees levied by internet service providers (ISPs) on large content providers, may protect incentives to innovate, and prevent potentially harmful fragmentation of the Internet (whereby certain content providers are only accessible on certain ISPs).

IV. Other Papers

a. Matching

The theory of two-sided matching generally assumes that agents know their true preferences over potential partners prior to engaging in a match. However, in many matching markets, information acquisition plays an important role with costly interviews, dates, and meeting. Since these interviews affect the formation of preferences, their assignment can dramatically affect the outcome of the eventual matching process. In ***Interviewing in Two-Sided Matching Markets*** (with Michael Schwarz, *RAND Journal of Economics*, 2017), we generalize the one-to-one matching model of Gale Shapley (1962) to include a prior stage in which firms first choose a subset of workers to interview; we show that even if two equilibria involve the same number of interviews for all firms and workers, outcomes can be substantively different: the number of agents that remain unmatched depends not only on the number of interviews that each agent receives but also on the number of common interviewing partners among agents. We introduce the concept of *overlap* that captures this notion, and prove the probability of being matched is maximized with what we refer to as *perfect overlap*—if two firms interview a common worker, then those firms interview the exact same set of workers.

b. Learning

In ***Multiple Equilibria and Selection by Learning in an Applied Setting*** (with Ariel Pakes, *Economic Letters*, 2009), we explore the use of learning algorithms to assist with counterfactual analysis in the presence of multiple equilibria. We study the counterfactual reallocation of ATMs among banks following a hypothetical merger accompanied by an unexpected cost shock, and show: (i) simple enumeration of all potential equilibria is informative as the number of total equilibria is small compared to the number of potential allocations, and (ii) there may be a distribution of predicted equilibria given the variance of cost shocks, and this distribution has a notable dependence on both the cost specification and on the learning process. Thus, to use learning algorithms as a means of guiding equilibrium selection, it would be helpful to obtain evidence on which processes are more relevant for the setting being considered.

V. References

Except for preliminary work, all papers are available at <http://www.robinslee.com/>

a. Refereed Publications

Crawford, G., R. S. Lee, M. D. Whinston, and A. Yurukoglu, "The Welfare Effects of Vertical Integration in Multichannel Television Markets," 2018, *Econometrica*, 86(3): 891-954.

Foo, P., R. S. Lee, and K. Fong, "Hospital and Physician Prices and Treatment Choice in Labor and Delivery," 2017, *American Journal of Health Economics*, 3(3): 422-453.

Gowrisankaran, G., A. Collard-Wexler, and R. S. Lee, "'Nash-in-Nash' Bargaining: A Microfoundation for Applied Work," forthcoming, *Journal of Political Economy*.

Hagiu, A. and R. S. Lee, "Exclusivity and Control," 2011, *Journal of Economics and Management Strategy*, 20:3, 679-708.

Ho, K. and R. S. Lee, "Insurer Competition in Health Care Markets," 2017, *Econometrica*, 85(2), 379-417.

Ho, K. and R. S. Lee, "Equilibrium Insurer-Provider Networks: Bargaining and Exclusion in Health Care Markets," forthcoming, *American Economic Review*.

Lee, R. S., "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," 2013, *American Economic Review*, 103(7): 2960-3000.

Lee, R. S., "Competing Platforms," 2014, *Journal of Economics and Management Strategy*, 23(3), 507-526.

Lee, R. S. and M. Schwarz, "Interviewing in Two-Sided Matching Markets," 2017, *Rand Journal of Economics*, 48(3): 835-855.

Lee, R. S. and A. Pakes, "Multiple Equilibria and Selection by Learning in an Applied Setting", 2009, *Economic Letters*, 104:1, 13-16.

b. Other Publications

Lee, R. S. and T. Wu, "Subsidizing Creativity through Network Design: Zero Pricing and Net Neutrality," Summer 2009, *Journal of Economic Perspectives*, 23:3, 61-76.

c. Working Papers

Dafny, L., K. Ho, and R. S. Lee, "The Price Effects of Cross-Market Hospital Mergers," 2018, revised and resubmitted, *RAND Journal of Economics*.

Lee, R. S. and K. Fong, "Markov-Perfect Network Formation: An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks" 2013, revise and resubmit at *Review of Economic Studies*.