

Research Summary

January 2018

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

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 are *oligopolistic* and negotiate terms of trade with one another across different market segments. Understanding how firms compete and interact in these *vertical markets* is critical not only for regulatory, innovation and competition policy, but also for anticipating how macroeconomic shocks propagate through these industries and affect welfare. Yet, standard economic theory—generally tailored to settings where firms can unilaterally set prices or are price-takers—has struggled to provide sharp predictions in these more complicated environments where multiple parties have market power.

My research attempts to address these shortcomings. There are two primary goals of my work. The first is to develop formal analytical frameworks for studying firm competition in vertical markets characterized by bilateral oligopoly that can be “taken to data”: i.e., flexible enough to capture important institutional details of real world industries, and tractable enough for estimation and counterfactual evaluation. The second motivation is to adapt and apply these frameworks to a diverse set of economically significant industries, and synthesize both theory and empirics to address active policy-relevant issues and areas of debate.

I. Empirical Studies of Bargaining and Contracting in Vertical Markets and Bilateral Oligopoly

a. Health Care Markets

My most active area of research aims to deepen our understanding of 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 illuminate how competition among insurers and medical providers affects expenditures, access, and quality. Ultimately, I seek to inform 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 increase consumer welfare and constrain spending.

Bilateral negotiations between insurers and medical providers (e.g., hospitals and doctors) are central to health care provision in the U.S.: these negotiations determine insurers’ “networks”—i.e., 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 (induced, for example, by mergers or the introduction of insurance exchanges).

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 large California benefits manager, and uses estimates to examine how increased insurer competition for patients—e.g., induced by the formation of health insurance exchanges, entry by new insurers, or an increase in the variety of employer provided plans—affects negotiated payments, premiums, and welfare. Though increased insurer competition may lead to lower premiums holding fixed input prices (i.e., hospital negotiated rates), it may also lead to higher hospital prices as hospitals are better able to “play” one insurer off another. By simulating equilibrium outcomes upon the removal of an insurer from each market, we find that increased insurer competition—though often beneficial to consumers—can lead to both increases and decreases in negotiated hospital prices across markets, and redistribute rents across health providers.

The work in this previous paper addresses active areas of competition policy in health care. For example, the tradeoffs it measures are fundamental to evaluating market structure changes, such as the \$54 billion Anthem and Cigna insurance merger (challenged by the Department of Justice in 2016, and blocked in 2017). Additionally, the model it develops is used in ***The Price Effect of Cross-Market Hospital Mergers*** (with Leemore Dafny and Kate Ho, revise and resubmit, *RAND Journal of Economics*) to examine so-called “cross-market” hospitals 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 end consumer. However, as we note, such hospitals may still be viewed substitutes from the perspective of insurers that bundle hospital services for sale to multi-market employers or multi-person households. We pair this theoretical point with a new data set of cross-market hospital mergers over a 14-year period, and estimates sizable positive price effects following such mergers. Our results suggest that additional regulatory scrutiny is warranted.

Most recently, ***Equilibrium Insurer-Provider Networks: Bargaining and Exclusion in Health Care Markets*** (with Kate Ho, revised and resubmitted, *American Economic Review*) studies the efficacy of “narrow network” products that are increasingly being offered by insurers (whereby insurers selectively contract with a limited set of medical providers), and evaluates the potential effects of network adequacy standards being considered by the Centers for Medicare and Medicaid Services and several state regulators. In order to shed light on this question, we significantly extend our previous framework by incorporating both a model of how provider networks are formed in equilibrium, and allowing for an insurer to exclude hospitals in order to “play them off” one another (see further discussion below). 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 that private incentives to exclude hospitals on the part of insurers generally exceed social incentives—the insurer benefits from substantially lower negotiated hospital rates and consumers from lower premiums—and that regulation prohibiting exclusion generally increases rates and premiums without significantly influencing social surplus. Nevertheless, we note that such regulation has distributional consequences, and can prevent significant harm from consumers living close to otherwise excluded hospitals.

Lastly, much of my work in health care markets focuses on changes in reimbursement rates and premiums; in turn, 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

In many vertical markets, firms engage not only in selective contracting, but also vertically integrate along the supply chain. Although there is a substantial theoretical literature on vertical integration exploring pro- and anti-competitive consequences, 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 use a dynamic model of consumer demand for hardware and software and a dynamic model of how software developers choose 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 proves integral when identifying the marginal benefit of a single software product to each hardware platform's sales. In specifying a rich consumer demand system for hardware and software products, the paper also highlights the importance of accounting for dynamics (e.g., products are durable, and consumers anticipate future price changes and software releases) and preference heterogeneity, features pervasive in most technology industries. Results from the exercise indicate that exclusivity and integration may have actually enhanced platform competition and aided the 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.¹

Second, in ***The Welfare Effects of Vertical Integration in the Multichannel Television Market*** (with Greg Crawford, Michael Whinston, and Ali Yurukoglu, *Econometrica*, forthcoming), 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 debate over recently approved (Comcast/NBC in 2011), abandoned (Comcast/TimeWarner in 2015), and proposed (AT&T/Time Warner in 2016) mergers. In the US, over 110M 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

¹ In ***Home Videogame Platforms*** (a chapter in the *Oxford Handbook of the Digital Economy*, 2012), I survey recent economic research on the videogame industry and discuss how insights from this work can be used to answer other related questions.

in certain cases. Leveraging substantial cross-market variation in ownership and regulatory policy facing regional sports networks, we estimate a structural model of viewership, subscription, distributor pricing, bargaining, and carriage decisions. Our model incorporates the potential for double marginalization, incentives to engage in foreclosure, and 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

In the U.S., over a dozen stock exchanges function as trading platforms, bringing together investors and trading firms; on these platforms, more than 1.5 trillion shares—representing \$60 trillion in value—are exchanged annually. These exchanges use variants of a market design that Budish, Cramton and Shim (2015, *QJE*) have shown to cause latency arbitrage and 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 develop a model of platform competition among stock exchanges to understand whether market participants have incentives to adopt new market designs that ameliorate inefficiencies arising from latency arbitrage. We use detailed trades and quotes data paired with financial filings to confirm that the theoretical predictions of our model conform to empirical realities. Central to our analysis is the relationship between exchanges and high frequency trading firms (HFTs): our model indicates that exchanges, through the levying of exchange-specific technology fees to HFTs for co-location services and proprietary data feeds, share in the rents that HFTs extract from investors. Thus, we argue that private incentives for market design innovation are socially inefficient, and use our model to evaluate potential policy and market responses.

II. Network Formation and Bargaining in Bilateral Oligopoly

What links together the industries presented in the previous section is that they are all 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 (e.g., 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, e.g., prices) for all pairs of contacting agents such that each agreement is the solution to a bilateral Nash bargaining problem for each pair—i.e., 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) answers 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 their 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 attempt to address these limitations. In ***Equilibrium Provider Networks*** (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 show emerges from a 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 contracting 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, breakdowns, and changes to provider networks. Importantly, firms anticipate future changes when contracting and bargaining, affecting both the “network” of firms that contract and negotiated payments. We show how the framework can be used to estimate underlying structural primitives. We also use this framework to simulate hospital mergers in health insurance markets, and demonstrate the impact of dynamic network and rate changes on merger effects.

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

Earlier work, 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 (described 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 strain of research, ***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. The paper decomposes 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 argue that certain aspects of network neutrality, specifically a ban on fees levied by internet service providers (ISPs) on large content providers, may be beneficial by protecting incentives to innovate, and by preventing 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 crucially 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*—i.e., 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.

References

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

Refereed Publications

Crawford, G., R. S. Lee, M. D. Whinston, and A. Yurukoglu, "The Welfare Effects of Vertical Integration in Multichannel Television Markets," forthcoming, *Econometrica*.

Foo, P., R. S. Lee, and K. Fong, "Hospital and Physician Prices and Treatment Choice in Labor and Delivery," Summer 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.

Lee, R. S., "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," December 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.

Other Publications

Lee, R. S., "Home Videogame Platforms," 2012, in *Oxford Handbook of the Digital Economy*, eds. Martin Peitz and Joel Waldfogel.

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.

Working Papers

Dafny, L., K. Ho, and R. S. Lee, "The Price Effects of Cross-Market Hospital Mergers," 2017, revise and resubmit, *RAND Journal of Economics*.

Ho, K. and R. S. Lee, "Equilibrium Insurer-Provider Networks: Bargaining and Exclusion in Health Care Markets," 2018, revised and resubmitted, *American Economic Review*.

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*.