Research Statement

Michael J. Dickstein

May 2024

In my research, I develop and apply methods from the field of industrial organization to study healthcare markets. The health industry is attractive to analyze in part due to its size—in the U.S., health care spending as a share of gross domestic product reached nearly 18% in 2021—but also because of the unique economics of the sector. Beginning with Arrow's (1963) analysis, economists have emphasized that various forms of uncertainty in medical care, including unknown treatment efficacy, adverse selection in insurance, and moral hazard among providers, make it unlikely that a competitive market will lead to socially desirable outcomes. Consequently, regulators play a role in designing key aspects of health markets.

I target my research contributions to inform regulators' market design decisions. For example, I consider how patient prices influence prescription drug adoption or how re-defining the boundaries of a health insurance pool may lead to greater enrollment, particularly among sicker households. Quantifying the effects of such large policy changes, however, requires distinct empirical tools. Experiments, such as subsidizing insurance for a randomized population, can be extremely costly to implement in healthcare settings, and the randomization itself may be viewed as morally repugnant. In addition, for-profit or non-profit firms in the market may respond to the policy change in ways not envisioned by the experiment. Structural models of behavior provide an alternative path; a through-line of my work is to tailor and extend modeling tools to match the complex economic features of health care markets, including uncertain product quality, adverse selection, and imperfect competition.

I group my research contributions into three related topic areas. First, I examine learning and information acquisition. Using the setting of prescription drug choice, I study how policy levers, like price and promotion, can alter how a physician experiments to match treatments to a patient's illness (*Efficient Provision*). I then expand the scope of this agenda to understand information acquisition more generally, studying firms forecasting future revenues (*What Do Exporters Know?*) as well as physicians facing uncertain prices (*Patient Costs*).

Second, I empirically study the role of payment incentives for agents acting on behalf of principals. I explore how payment systems operate in the health care sector and quantify how these systems influence the extent of physician agency—that is, the tendency of physicians to diverge from patient interests in response to their own financial incentives (*Industry Input in Policymaking; Physician vs. Patient Incentives*).

Third, I develop models of demand and supply in health insurance to study the design of "managed competition" markets, such as the markets regulated under the Affordable Care Act of 2010. I exploit novel data on household spending and enrollment decisions to study the consequences of market segmentation and to explore the effects of mandates for insurance coverage (*Market Segmentation; Insurance without Commitment; Market Size and Insurance Premiums*). In recent work in progress, I also study the role of insurance network breadth on waiting times and patient outcomes (*Wait Times in Health Insurance Design).*

My work highlighted above provides both insights into the regulation of healthcare markets as well methodological contributions. Reflecting their impact on the field, several of these papers appear in the syllabi of PhD courses in industrial organization at top economics departments, including Yale, Northwestern, and Harvard University. My work is also cited in health economics courses for undergraduates and graduate students in public policy and economics at the University of Pennsylvania, Emory University, and UNC, among other institutions.

Learning and Information Acquisition

My research on learning addresses two central questions in the literature on demand under uncertainty: (1) How do decision-makers incorporate new information in their product choices? And (2) what information do agents use in their decisions?

I tackle the first question by building on an empirical literature (Erdem and Keane 1996, Ackerberg 2003, Crawford and Shum 2005) that develops and estimates models of demand with experiential learning. My innovation is to design a dynamic model specific to treatment choice for depression. The depression market is important to study both due to its size—in survey data for the years 2015-2018, 13.2% of U.S. adults reported using antidepressants in the last 30 days—and because of the structure of the choice set. Antidepressant drugs fall into classes based on their biological mechanism of action, meaning that drugs within a class may present a similar benefit-risk profile for a given patient; the path of experimentation might optimally exploit these class effects. Similar correlation within class appears in demand for many products and services, including retail consumer goods.

In "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice," I develop a model of within- and across-class learning to understand physician choice in the selection of an antidepressant. More than twenty antidepressant drug treatments compete for a share of physicians' prescribing. The large choice set and potentially infinite horizon for care makes it difficult to apply standard dynamic discrete choice methods. Instead, I summarize the treatment selection by developing and estimating a new model that features an "index rule" as the physician's rule of thumb.

The rule, a modification of Gittins' (1979) index, allows me to capture a physician's experimentation process with alternative choices in a way that is faster to compute than standard methods. I also define a Bayesian updating process that explicitly allows for correlation in learning: physicians may, for example, downgrade their expectations of an entire class of treatments from sampling just one member within the class.

I illustrate the value of my approach using commercial insurance claims from IBM's Marketscan Research Databases, which house claims from a set of large U.S. employers. I collect the relevant claims for patients with a depression diagnosis. In this empirical setting, I compare the fit of three competing models of physician behavior: a model that ignores the value of exploration; a model that incorporates exploration but ignores correlation in choices; and, my model, which features both an exploration incentive and product correlation. Applying the model testing approach of Vuong (1989), I find support in the data for my dynamic learning model with correlation in the choice set.

With the model estimates, I ask how the design of a health insurance plan, including the required patient out-of-pocket costs by drug, might interact with the physician's learning process to affect patient outcomes. In the data, patient costs largely correspond to the drug's wholesale cost; if a drug is cheaper to acquire, the patient pays less. In contrast, I design a new drug pricing schedule that lowers the patient cost for those drugs that the model suggests are best to sample early in the search process. By using these price incentives to redirect the search process, I find physicians identify the optimal treatment match faster. This design change produces gains for both insurers and patients: the health system experiences lower overall costs, and patients adhere to treatment more often, leading to better health outcomes.

The learning process in this work occurs *within patients*: physicians work to identify the best match between each patient and the available options. This model fits the antidepressant setting in which the choice set remains roughly stable during the period I analyze. In a more turbulent market—for example, in a market with new drug entry or warnings about a drug's effectiveness—learning *across patients* becomes more important. In such settings, a key goal is to understand decision

makers' methods of information acquisition and to characterize heterogeneity in their information about product characteristics, such as prices. Researchers rarely have direct data on the information used in product choice and so must rely on assumptions—often, perfect foresight or rational expectations—that are a function of a set of observed variables.

One path to relax the assumptions on decision makers' information is to develop a partially identified model that can handle unobserved expectations. In joint work with Eduardo Morales entitled "What do Exporters Know?," published in the *Quarterly Journal of Economics*, we follow this strategy. We build a novel moment inequality model that does not require the researcher to specify the full set of information that agents use in their decisions. We then apply our framework to study the discrete choice of manufacturing firms over whether to sell their output abroad when faced with uncertain returns in international markets. Using Chilean customs data, we show how to use our methodology to recover the parameters of a firm-level entry model.

Our framework also allows us to test whether firms differ in the information they possess about foreign markets. We find that larger firms possess better knowledge of market conditions in foreign countries, even when they have not exported in the past. To quantify the importance of this informational advantage, we conduct a counterfactual exercise in which we provide data on market conditions to all potential exporters. We predict that better informed firms choose to export to fewer markets: in the industrial sector we study, we observe a decrease of between 3.5% and 5.7% in the number of firm-to-destination export pairs. Interestingly, although the total number of firm-destination pairs decreases, the aggregate export revenue in the sector increases between 6.4% and 9.5%. These findings suggest a path for policymakers seeking to support a country's export sector: by providing better information to potential exporters, our model predicts fewer firms will make mistakes in their destination choices—that is, firms can avoid markets where they are likely to lose money. Consequently, the average firm's realized profits in a market would increase between 17.5% and 20.6%.

While we demonstrate the value of our new methodology using the exporter's problem, our goal in developing this estimation approach was to provide a tool that applies to a variety of choice settings where decisions depend on agents' forecasts of payoff-relevant variables. Agents can be firms, households, or other individuals facing choice under uncertainty. For example, Bombardini et al. (2023) employ our methodology to study the votes of members of the U.S. Congress to liberalize trade with China. They uncover a role for both constituent interests as well as political ideology in members' voting patterns. In a new paper, "Patient Costs and Physicians' Information," joint with Jihye Jeon and Eduardo Morales and under revision for the *American Economic Review*, I return to the prescription drug-choice setting, but now extend the partial identification toolkit I developed earlier to measure how physician treatment decisions respond to changes in patient costs.

We motivate this work with a current policy debate in prescription drug markets. As drug spending nearly doubled in the past two decades—per-capita drug costs in the US rose from $668 to $1,126 between 2004 and 2019—health insurers introduced greater demand-side incentives in the form of higher patient out-of-pocket costs. The success of these incentives in steering demand toward lower cost options depends both on a physician's sensitivity to cost and on the physician and patient's awareness of price incentives. For policymakers evaluating such incentive programs, distinguishing information from preferences is a crucial first step in the market design problem. If an insurer raises the out-of-pocket cost for an expensive branded pharmaceutical drug, for example, and usage remains high relative to lower-cost options, is the lack of switching because the branded medication has higher effectiveness, because patients have high willingness-to-pay, or because physicians were unaware of the relative differences in out-of-pocket costs?

To answer this research question, we start by illustrating that estimates of the physician's responsiveness to price are sensitive to how the researcher specifies the physician's price forecast. In

the spirit of Manski (1991) and Manski (2004), we show that when agents' expectations differ from the researcher's assumptions about those expectations, this distinction generates bias in the measurement of price sensitivity. This bias propagates into counterfactual predictions about how, for example, welfare or market shares would change if insurers lowered the patient's out-of-pocket costs for a particular treatment.

To combat this potential for bias, we propose a solution using moment inequalities. In our framework, we can identify physician and patient preferences while placing only weak assumptions on the agents' information sets. We define two sets of moments. The first set, which we label "odds-based" inequalities, generalize our approach from Dickstein and Morales (2018). New to this paper, we modify our moments to handle settings in which consumers choose among more than two options. We derive the second set of moments, which we label "bounding" inequalities, using a linear approximation to our choice function, following the approach in Fujiwara et al. (2023). An important theoretical feature of these moments is that if the researcher includes agents' true information sets in the model, the moment inequalities will yield a point estimate equal to the true parameters.

We employ our estimation strategy to measure how health care use responds to higher out-of-pocket costs. Using data from the State of Oregon's All Payer All Claims (APAC) dataset for the years 2011 and 2016, we narrow our focus to patients seeking treatment for type 2 diabetes. We choose diabetes care as our market of interest both because of the size of the affected population—in the U.S. in 2019, 37.3 million people lived with diabetes—and because of the rapid growth in treatment costs.

Combining our model with data, we reject that physicians have perfect foresight about prices. Instead, we find physicians know only less detailed price averages, say at the drug-plan type-year level. In a counterfactual, we show that when we provide physicians perfect information on product prices in an office visit, say by introducing an automated pop-up in the patient's medical chart, we see a reduction in realized out-of-pocket costs of 11-22%, declining from roughly $45 per month to between $35 and $40 per month.

Physician Payments

In a second line of research, I study in more detail the incentive issues in physician reimbursement. The optimal design of these payments remains a key open question in health economics for both theoretical and empirical reasons. First, the patient-physician relationship is a canonical example of a principal-agent problem, where physicians choose treatments on behalf of their patients by balancing the patient's utility with their own financial incentives. I explore the effect of alternative payment designs on the extent of physician agency. Second, the payment levels themselves are high and often paid by government insurance. In 2019, payments for physician and clinical services in the U.S. equaled $772 billion or $2,356 per person per year. Roughly half of all insurance payments to physicians came from government payers, like Medicare and Medicaid. I explore how changes to the price-setting process in Medicare can change the allocation of funds to different types of physicians.

I began research on physician agency by examining incentives designed to limit moral hazard in health care. In "Physician vs. Patient Incentives in Prescription Drug Choice," I compare "demand-side" incentives, which impose costs on the patient to limit moral hazard, and "supply-side" incentives, which adjust the physician's compensation to discourage spending. Revisiting the depression treatment setting in this work, I find an important trade-off in drug treatment choice when physicians face capitation, a form of supply-side incentive in which physicians receive a lump-sum payment annually to treat a patient assigned to their practice. If the physician can keep a patient from returning to the office setting, she captures any savings relative to the lump-sum payment. I

find that with capitation of primary care office visits, physicians also change their prescription drug care even though drug care falls outside the capitation rate. In depression treatment, capitated physicians emphasize those treatments in the choice set that, in a large sample of claims data, can be shown empirically to require fewer follow-up visits. These antidepressants were typically cheaper generic drugs, which helped lower costs and improve adherence under capitation policies.

An alternative to supply- or demand-side incentives in insurance design is instead to regulate prices directly. In a joint paper with David C. Chan entitled "Industry Input in Policymaking: Evidence from Medicare" and published in the *Quarterly Journal of Economics*, we examine the operation of a regulated price system. Using rich internal data from the American Medical Association, we study the procedure for both establishing and updating regulated prices in Medicare's physician payment system. The system determines not only Medicare payments, which totaled $193 billion in 2019, but also influences the payment rates of roughly another $300 billion in private insurance payments that follow Medicare's system (Clemens and Gottlieb 2017).

To determine the underlying value of a physician service, Medicare solicits input from a committee of the American Medical Association. The committee, known as the Relative Value Scale Update Committee (RUC), comprises 25 representatives from different medical specialties who consider proposals for all changes to the relative values of services. Twenty-one of these members occupy permanent seats and the remaining four rotate. For example, a representative of the specialties of internal medicine, dermatology, and orthopedic surgery maintain permanent seats, while specialties including pediatric surgery and infectious disease rotate on and off the RUC.

We first ask whether the composition of the RUC leads to prices biased in favor of its members. To quantify this bias, we collect data from the American Medical Association on the details of proposals before the committee. Using this data, we develop a measure of affiliation between the committee members and the specialties proposing a price. Our affiliation measure reflects the alignment in preferences between the proposer and the specialties on the RUC who evaluate the proposal. Using quasi-experimental variation in affiliation, we find that increasing this measure of connectedness between the committee and proposers by one standard deviation increases prices by 10%. Our result provides new evidence of regulatory capture, here in the setting of health care markets.

We then turn to a central question of regulatory design: given the possibility of bias, what value does the government obtain from inviting industry input into policy making? In settings involving advisory committees, a key feature is the importance of policy-relevant knowledge (e.g., the safety and efficacy of a drug) held by industry participants. The government may form advisory committees that either contribute such knowledge directly or extract and synthesize information from outside special interests. We explore whether allowing some bias in these advisory committees may improve regulatory decisions by facilitating the communication of information needed for regulation. Specifically, we study whether Medicare can extract more information about physician services and set more appropriate prices by employing the RUC as an intermediary in decision-making.

To address this question, we develop a conceptual model, borrowing ideas from a large literature on the extraction of information from biased experts. Our model allows us to compare two types of information helpful for regulatory decisions. First, we consider information that is "hard," or verifiable. A committee adversarial to the specialty expert will encourage the expert to generate more of such evidence (Dewatripont and Tirole 1999; Hirsch and Shotts 2015). Second, information may be "soft," or unverifiable. Soft information must be credibly communicated (Crawford and Sobel 1982); a committee biased in favor of the specialty expert may improve such communication (Dessein 2002). The net effect of bias on information extraction thus depends on the nature of information relevant for decisions.

We use the Medicare setting to test the implications of our model. We find that a more biased committee leads to less hard information; proposals submitted to a RUC with greater affiliation feature lower quality data. To examine a policy-relevant metric of the overall level of information—both hard and soft—that Medicare collects through the RUC, we measure the degree to which Medicare price changes correlate with private insurance price changes. We interpret greater correlation as evidence that the committee's price reflects higher quality information. In our data, we find that price changes in private insurance track those changes in Medicare more closely when the Medicare price changes arise from proposals submitted to a more biased RUC. Thus, our evidence suggests a trade-off between bias and expertise in the committee's work; allowing more bias can improve the overall quality of information in Medicare pricing decisions.

Managed Competition in Health Insurance

In a third line of research, I study the design of "managed competition" health insurance markets. The term often refers to insurance systems like those of the Netherlands, Germany, and Switzerland, where private insurers compete in a marketplace in which the government sets rules about pricing, insurance plan differentiation, and other elements of the market design. Since the passage of the Affordable Care Act, the markets for individual insurance and small group insurance in the U.S. operate under a system akin to managed competition. My goal in this line of research is to inform policy by quantifying the effects of changes to specific market regulations that are under active consideration. I consider changes to the scope of segmentation in health insurance markets, to the design of individual mandate penalties for insurance, and to the breadth of insurance networks.

In joint research with Kate Ho and Nathaniel Mark, entitled "Market Segmentation and Competition in Health Insurance" and published in the *Journal of Political Economy*, I examine the regulatory question of insurance pooling. In the U.S., households obtain insurance from a diversity of payers, including government plans and private plans sponsored by employers or purchased directly in the individual market. For private plans, the premium a household faces depends on the expected health spending of all households pooled with the enrollee in the market segment and geographic region. For example, a household purchasing individual coverage in Portland, Oregon, will pay a premium that reflects the ages of the household's members and the expected costs of covering all households purchasing individual coverage in Portland. Similarly, the premiums for small employer coverage in Portland reflect the expected costs of all households in the small employer segment in Portland.

We measure the consequences of this segmentation on the level of public expenditures for insurance subsidies, on employer expenditures for coverage, on consumer surplus, and on a household's decision to obtain coverage. I previously described how geographic market segmentation under the Affordable Care Act could affect premiums and insurer participation in The Impact of Market Size on Premiums, joint with Mark Duggan, Joe Orsini, and Pietro Tebaldi, and published the *American Economic Review: Papers & Proceedings*; the innovation in my newer research is to build a supply and demand model to predict equilibrium prices, insurance enrollments, and patient spending under alternative market designs. With these elements, we can measure the gains and losses to consumers and payers from combining market segments.

We use the setting of small group insurance and individual market insurance in Oregon as our laboratory to examine the consequences of insurance segmentation. Nationally, the two insurance markets collectively covered 33.5 million households in 2016; we observe roughly 415,000 of those households in Oregon. We build our model to measure changes in consumer welfare when households who obtain insurance through a small employer must shift their coverage to the individual market. We choose these markets because many insurance regulations—such as

community rating and guaranteed issue of plans—are common to the two segments.[1] We can also study a key regulatory change from 2020. Beginning in that year, employers can offer individual coverage health reimbursement arrangements (ICHRAs). Through these accounts, employers no longer sponsor coverage directly. Instead, they simply seed a tax-exempt account with subsidies, which the employee can then use to purchase coverage in the individual market. If adopted widely, ICHRAs could transform the U.S. insurance system, eliminating the employer's role in plan choice.

We consider two economic mechanisms that influence household costs and surplus under alternative pooling designs. First, changing the pool of patients can reduce the extent of adverse selection. In our data, small group households prove, on average, healthier than individual market households. By shifting healthier enrollees to pool with existing households in the individual market, regulators can reduce adverse selection in the market and lower premiums for individual market participants; healthier small group enrollees, however, might face higher premiums in the new pool. Second, we consider evidence that various frictions in the employer market—including added costs from broker intermediation—lead to higher markups in the small group market. We find small group households in Oregon, despite pooling with sicker individual-market households, nonetheless gain surplus from shifting to individual coverage. The benefit from reduced markups in the individual market outweighs the costs of being pooled with a less healthy population.

The measurement exercise we conduct requires building a demand model that features multiple behavioral elements, including risk aversion, adverse selection, moral hazard, and selection on moral hazard, as in Einav et al. (2013). We also need a supply model that incorporates expected health costs and administrative costs and outputs equilibrium premiums under new pooling environments. We contribute to the empirical literature by showing how to build a model with these features that is straightforward to estimate.

This framework is valuable for a range of applications outside Oregon. For example, one can apply our approach to study the effect of expanding eligibility for public insurance under Medicare to populations younger than 65 years old. The key insight from our research is that to judge the full costs and benefits of changes to market segmentation, policymakers need to consider not only the health of enrollees in the two pools, but also the efficiency of insurance provision. The degree of competition, the role of brokers, and the administrative costs in each market segment can determine the desirability of pooling.

In subsequent work, I tackle a related regulatory challenge in managed competition insurance markets: when consumers can exit coverage in any month of an enrollment year, does the presence of dropouts lead to higher equilibrium premiums for full-year enrollees? Is there scope for mandated insurance in such an environment? In joint work with Rebecca Diamond, Timothy McQuade, and Petra Persson entitled "Insurance without Commitment: Evidence from the ACA Marketplaces," prepared for re-submission to the *Quarterly Journal of Economics* (under a reject and resubmit status), we develop and estimate a dynamic structural model of household insurance enrollment to quantify the costs and benefits of attrition.

We report several facts new to the health economics literature. First, we document significant rates of insurance drop-out: in the 2014 and 2015 open enrollment periods in California, about 30% of enrollees exit within nine months of signing up. We observe this pattern of exit both in administrative data from Covered California, the state's health insurance marketplace, and from a large commercial bank transactions dataset. In the high-frequency bank data, we can identify households who purchase individual insurance coverage through credit and debit card transactions.

[1] Community rating rules prevent insurers from setting premiums based on health status, including past illness history. Guaranteed issue rules require insurers to enroll households regardless of health status; households cannot be rejected at the time of application because of a history of pre-existing illness.

Furthermore, before households exit, they can choose to concentrate their health spending into the few months they maintain coverage. In subsequent months, they can exit coverage and save the remaining premium payments for the year. In our transactions data, we indeed observe sharp increases in a household's health consumption upon enrollment, with more transactions and greater out-of-pocket health care and drug spending during the period of coverage. We observe this pattern particularly among younger and single households.

To analyze the effect of drop-out on welfare, we begin with a conceptual model. We show that when a market has both traditional adverse selection and drop-out—that is, when some households can re-time their spending and then exit coverage—the attrition need not always harm welfare in the insurance market. In fact, one-sided commitment contracts can *raise* welfare for all households. We illustrate this result using a two-type special case of our model. Suppose, for example, that enrollees who would choose to drop coverage mid-year spend less on health care on an annual basis than do full-year enrollees. Without the option to drop coverage mid-year, healthier dropouts might choose not to enroll in insurance at all, as the annual premium may be too high given their low level of health care needs. In contrast, with drop-out, insurers can effectively price discriminate. The high-cost enrollees pay higher prices in total over the year than low-cost enrollees, who exit early and pay premiums for only a fraction of the year. The option to drop out can encourage low-cost households to enroll, bringing down the average monthly cost in the insurance pool.

To quantify the welfare effect of attrition in our empirical setting, we estimate a structural model that features both endogenous household exit from insurance coverage and endogenous timing of health care consumption. With estimates of the model parameters, we predict how drop-out rates and premiums would change were regulators to impose penalties for early drop-out or ban early exit entirely.

Finally, in work-in-progress entitled "Incorporating Wait Times in Health Insurance Design", joint with Pierre Bodéré and Guillaume Fréchette, we explore the consequences of network breadth. Insurers offering plans in managed competition marketplaces often limit the set of physicians a patient can visit with full insurance coverage; we measure how narrowing the network a patient can access might affect waiting times, and subsequently patient health outcomes and insurer costs.

In this agenda, we first propose a measure of wait times—detection to treatment (DTT)—which we compute solely using data from medical claims rather than using administrative measures. We describe patients' high-dimensional medical trajectories using machine learning tools to represent their risk profile over time. Using IBM's Marketscan data for our analysis, we find preliminary evidence that our DTT measure is higher for patients with insurance plans that restrict patient access to some physicians. Further, exploiting quasi-exogenous variation in physicians' capacity to treat patients in more restrictive plans, we show that longer wait times result in adverse outcomes post-surgery: patients are more likely to be readmitted to a hospital, pay higher total out-of-pocket costs, and use addictive drugs, such as opioids. We can use this measurement to quantify an important dynamic trade-off that arises from restrictive provider networks. These restrictions may lower utilization and healthcare costs in the short run, coming at the expense of higher illness costs and worsened health outcomes in future periods.

References

Ackerberg, Daniel A., "Advertising, learning, and consumer choice in experience good markets: an empirical examination," *International Economic Review*, 2003, 44 (3), 1007–1040.

Arrow, Kenneth J. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review*, 1963, 53 (5), 941-973.

Bombardini, Matilde, Bingjing Li, and Francesco Trebbi, "Did US Politicians Expect the China Shock?," *American Economic Review*, 2023, 113 (1), 174-209.

Chan, David C., and **Michael J. Dickstein**, "Industry Input in Policymaking: Evidence from Medicare," *Quarterly Journal of Economics*, 2019, 134 (3), 1299–1342.

Clemens, Jeffrey P., and Joshua D. Gottlieb, "In the Shadow of a Giant: Medicare's Influence on Private Physician Payments," *Journal of Political Economy*, 2017, 125, 1–39.

Crawford, Gregory S. and Matthew Shum, "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 2005, 73 (4), 1137–1173.

Crawford, Vincent P., and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, 50, 1431–1451.

Dessein, Wouter, "Authority and Communication in Organizations," *Review of Economic Studies*, 2002, 69, 811–838.

Dewatripont, Mathias, and Jean Tirole, "Advocates," *Journal of Political Economy*, 1999, 107, 1–39.

Diamond, Rebecca, **Michael J. Dickstein**, Timothy McQuade, and Petra Persson, "Insurance without Commitment: Evidence from the ACA Marketplaces" Working Paper, 2023.

Dickstein, Michael J., "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice," Working Paper, 2021.

Dickstein, Michael J., "Physician vs. Patient Incentives in Prescription Drug Choice," Working Paper 2017.

Dickstein, Michael J., Mark Duggan, Joseph Orsini, and Pietro Tebaldi, "The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the ACA," *American Economic Review: Papers & Proceedings,* 2015, 105 (5), 120-125.

Dickstein, Michael J., Kate Ho, and Nathaniel Mark, "Market Segmentation and Competition in Health Insurance" *Journal of Political Economy*, 2024, 132 (1), 96-148.

Dickstein, Michael J., Jihye Jeon and Eduardo Morales, "Patient Costs and Physicians' Information," Working Paper, 2023.

Dickstein, Michael J., and Eduardo Morales, "What Do Exporters Know?," *Quarterly Journal of Economics*, 2018, 133 (4), 1753–1801.

Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen, "Selection on Moral Hazard in Health Insurance," *American Economic Review*, 2013, 103 (1), 178–219.

Erdem, Tulin and Michael P. Keane, "Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets," *Marketing Science*, 1996, 15 (1), 1–20.

Fujiwara, Thomas, Eduardo Morales, and Charly Porcher, "A Revealed-Preference Approach to Measuring Information Frictions in Migration Decisions", Princeton Univ. Working Paper, 2023.

Gittins, John C., "Bandit processes and dynamic allocation indices," *Journal of the Royal Statistical Society - Series B (Methodological)*, 1979, 148–177.

Hirsch, Alexander V., and Kenneth W. Shotts, "Competitive Policy Development," *American Economic Review*, 2015, 105, 1646–1664.

Manski, Charles F., "Nonparametric Estimation of Expectations in the Analysis of Discrete Choice Under Uncertainty," in William Barnett, James Powell, and George Tauchen, eds., Nonparametric and Semiparametric Methods in Econometrics and Statistics, Cambridge: Cambridge University Press, 1991.

Manski, Charles F., "Measuring Expectations," *Econometrica*, 2004, 5 (9), 1329–1376.

Vuong, Quang H., "Likelihood ratio tests for model selection and non-nested hypotheses," *Econometrica*, 1989, pp. 307–333.