Research Statement

Richard L. Sweeney

Fall 2022

My research lies at the intersection of energy and environmental economics (EEE) and industrial organization (IO). One strand of my research leverages the unique features and data of energy markets to study core questions from IO and public economics. A second strand uses theory and empirical methods from IO to inform and evaluate energy and environmental policy. Recently these two strands have coalesced around the topic of energy innovation. Going forward, an overarching theme of my research will be using IO insights to quickly and cost-effectively transition society away from fossil fuels.

IO in Energy Markets

The benefits of auctions

Asset owners often need to identify and choose between potential contracting partners to monetize their asset's value. While economic theory has characterized the relative performance of different highly structured *formal* assignment mechanisms, in the real world, many important assets are allocated via informal, unstructured processes, and little is known about how they perform relative to theoretical benchmarks. In "Relinquishing Riches: Auctions vs Informal Negotiations in Texas Oil and Gas Leasing" (with Thomas Covert, conditionally accepted at *The American Economic Review*) we directly measure the gains from using a centralized auction, relative to using informal, decentralized transactions, in the market for mineral leases in Texas. To make causal claims, we restrict our attention to a large class of lands initially set aside for public use under the Texas Constitution, on which legislative decisions made nearly one hundred years earlier determined whether leases signed during the recent shale boom transacted using an auction or an informal "negotiation".

Our empirical strategy compares auctioned and negotiated leases that lie in narrowly defined geographic areas, which transact at approximately the same time. We find that auctioned leases sell for 70% more than similar negotiated leases do, despite the fact that the two mechanism are equally likely to result in a parcel being leased. We also find that auctioned leases produce 55% more output than negotiated leases do. This is consistent with the idea that a well-designed auction should reliably allocate an asset to its highest value user, while a decentralized informal mechanism might not. Using detailed information on the identities of the firms and the bids that they submit to each auction, we show that these results are driven primarily by auctions better matching exploration and production companies (E&P's) to leases on which they are idiosyncratically more productive. These results suggest large potential gains from employing centralized, formal mechanisms in markets that traditionally allocate resources in an unstructured fashion, including the broader $3 trillion market for privately owned minerals in the United States.

In land markets, decisions made by one agent can affect the outcomes of spatially proximate agents. In "Auctions vs. Negotiations with Spillovers: Evidence from the U.S. Land Ordinance

and the Shale Boom," (with Thomas Covert and Ryan Kellogg)[1], we look for evidence of productivity spillovers across parcels allocated via different mechanisms. We focus on Colorado and North Dakota, where the US Land Ordinance of 1785 quasi-randomly allocated square-mile parcels to state land trusts in a "checkerboard" pattern, interspersed between privately owned and negotiated parcels. Due to differences in geology, a single square mile parcel is large enough to contain a shale well in Colorado, while two adjacent square miles are necessary in North Dakota. We show that in Colorado, the allocative efficiency gains from auctions on state land are comparable to findings from Covert and Sweeney (2022). But in North Dakota we find that the need to combine parcels leads to allocative process spillovers: the effect of auctions on state-owned land is attenuated, while production increases on private land adjacent to state land.

Pass-through

In imperfectly competitive settings, the price a firm sets depends directly on its own costs, but also indirectly on the costs of its competitors. In "Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom" (with Erich Muehlegger, forthcoming at *The Review of Economics and Statistics*), we demonstrate that this simple observation has two important implications for the estimation and interpretation of "pass-through" (or the extent to which prices change when costs change), a central policy parameter with wide-ranging economic implications (Weyl and Fabinger, 2013). First, when estimating a firm's response to a change in its own costs, it is important to account for rival responses as well. Beyond simple omitted variable bias from correlated cost shocks, the strategic response of seemingly untreated firms to rival cost changes invalidates them as a control group. Second, when using pass-through for policy prediction, the identifying variation used in estimation must match the policy application. For example, a pass-through rate estimated off of idiosyncratic cost shocks by comparing firms within the same time period may substantially underestimate the price impact of an industry-wide policy change.

We demonstrate these points empirically by studying the responsiveness of prices to cost changes in the U.S. oil refining industry. Leveraging rich variation in input costs to competing firms stemming from the fracking boom, we estimate pass-through rates that vary from near zero, for firm-specific shocks, to one, for industry-wide shocks. Viewed as a continuum, these estimates reconcile seemingly conflicting pass-through estimates from the recent fuel pass-through literature. Using regional variation in energy prices, Ganapati et al. (2020) find that refineries are largely unable to pass-through costs; while using national variation in the cost of renewable fuel credits, Knittel et al. (2017) find that the incidence falls fully on consumers. We interpret the disconnect in these estimates as reflective of the distinction between idiosyncratic and shared cost variation, and show analytically that price changes will naturally vary for these two shocks, even after conditioning on a firm's own costs.

To demonstrate the relevance of these results for policy analysis, we consider a hypothetical carbon tax on refineries, the second-largest industrial point-source of emissions in the United States. When paired with a border adjustment tax, a carbon tax represents an industry-wide cost shock. Properly accounting for the indirect effects of competitors, our results suggest consumers would bear virtually all of the burden of the tax. Yet, an estimate of tax pass-through based on

[1]This paper and "Frack Time versus Slack Time" (discussed below) are supported NSF Grant #1919499, "Productivity, Efficiency, and Distortions in the U.S. Shale Boom."

within market-time variation in costs would suggest that prices would rise by a mere five cents per dollar of tax imposed. This exercise highlights the importance of matching the cost variation used to estimate pass-through with the scope of the policy to which the estimate is applied.

In earlier work, "The Incidence of US Climate Policy: Alternative Uses of Revenues from a Cap-and-Trade Auction," (with Dallas Burtraw and Margaret Walls, *National Tax Journal*, 2009), we studied the equity implications of a national carbon price in the United States. The price changes resulting from the introduction of a fee on carbon emissions would be regressive, but could be outweighed by the distributional effects of reallocating the revenue raised. We evaluate five alternatives for returning carbon revenues to households, and find that some alternatives, such as expanding the Earned Income Tax Credit or cap-and-dividend are net progressive, while others, like a reduction in the income tax, amplify regressivity.

Environmental Policy Design

Input vs output subsidies

The first best solution to the problem of environmental externalities is a Pigouvian tax. However, such taxes are politically unpopular. In lieu of this, efforts to reduce environmental externalities from fossil fuels have typically taken the form of subsidies for green substitutes instead.

In "Investment versus Output Subsidies: Implications of Alternative Incentives for Wind Energy", (with Joseph Aldy and Todd Gerarden, conditionally accepted at *The Journal of the Association of Agricultural and Resource Economists*) we ask whether it is more cost-effective for the government to subsidize renewable energy production or renewable energy capital. Parish and McLaren (1982) demonstrate that the optimal margin to target is theoretically ambiguous, and both types of subsidy can be observed in the real world. Estimating how well targeted these subsidies are in practice is complicated by the fact that policy design is endogenous, and observed variation in subsidy margin may be correlated with other factors that affect policy outcomes.

We circumvent this problem by leveraging a unique natural experiment in the U.S. wind industry. Wind power in the United States was historically subsidized by the production tax credit (PTC), an output subsidy. An unexpected policy innovation stemming from the 2008 financial crisis temporarily gave developers the option to forego the PTC in favor of an up-front investment subsidy. Using instrumental variable and matching estimators, we find that investment subsidy claimants produced at least 10% less power than they would have under the output subsidy. Accounting for extensive margin effects, we show that output subsidies are more cost-effective than investment subsidies over a large range of wind energy output targets.

The role of retail

A long literature dating back to the 1970's has documented an apparent energy efficiency "gap": when purchasing energy-using durables (like cars or water heaters), consumers appear to underinvest in energy saving technology that is privately net beneficial. In response to this belief that consumers are making mistakes, governments around the world have enacted a wide range of policies and programs aimed at informing consumers about energy use, and nudging them to

make better decisions.[2]

In "The Role of Sales Agents in Information Disclosure: Evidence from a Field Experiment" (with Hunt Allcott, *Management Science*, 2017), we consider the role of sales agents in implementing these policies. Sales agents are consumers' primary source of information in many markets, and an integral part of the delivery of informational consumer policy. To study the way in which sales agents might help (or hinder) information acquisition about energy costs, we implement a large natural field experiment at the call center of a major water heater retailer. Using randomized assignment of subsidies and treatment scripts, combined with an extensive auditing of sales calls, we begin by simply documenting that sales agents often fail to deliver energy-use information, even when instructed to. Rather than simple shirking, we show that sales agents observe latent consumer interest in energy efficiency on the sales call, and strategically withhold the information from customers who are unlikely to be interested. Even after accounting for non-compliance, we find that strong informational treatments and generous subsidies have relatively small impacts on energy efficiency takeup (consistent with work in other settings by Hunt). However, we find that offering subsidies to both consumers *and* sales agents significantly increases takeup, relative to generous consumer subsidies alone. Overall, this research demonstrates that sales agents observe important idiosyncratic information during the sales process, and suggests that retail could be leveraged to more cost effectively increase the adoption of energy saving technology.

Hidden disamenities

Individuals concerned about exposure to local disamenties can avoid them by moving away from them (Tiebout, 1956). Housing markets will adjust to this behavior, and the price gradient observed across markets with respect to disamenities can be used to recover consumers' willingness to pay to avoid them (Rosen, 1974). These two observations, which underlie rich empirical literatures on sorting and hedonic property estimation, rely on the assumption that homebuyers are perfectly informed about the prevalence of disamenities in all markets. Yet, in many cases, it is clear that homeowners are *not* particularly well-informed. In these situations, it is possible that house prices understate the public's true willingness to pay.

In "Housing Market Capitalization of Pipeline Risk: Evidence from a Shock to Salience and Awareness" (with Evan Herrnstadt, R&R at *Land Economics*), we study the case of natural gas pipelines, which are buried underground and difficult to observe, and about which it has been documented that people have little awareness. Recent pipeline expansions stemming from the shale gas boom have faced considerable local opposition over (stated) community safety concerns; yet, in the cross-section, pipelines do not appear to meaningfully affect house prices. To determine whether this reflects true indifference or just poor information, we study the fallout from one of the largest pipeline explosions in U.S. history, in San Bruno, California in 2010. We document that the explosion, which killed 8 people and injured 53, generated a large increase in salience (news coverage) and information acquisition (Google search behavior). Six months after the explosion, the operator of the pipeline sent an informational letter to all households within 2000 feet of any natural gas pipeline, alerting them of their proximity.

[2]With some notable exceptions, a large body of recent, rigorous economics research has largely failed to find any evidence that consumers are, on average, making large systematic mistakes when it comes to energy efficiency. See Allcott (2016) for a review.

Using a difference-in-differences strategy, we find that properties close to pipelines in the San Francisco area declined by 2% following the explosion. The response is larger in the immediate aftermath of the explosion, and among properties closest to the pipeline, exceeding 5% for properties right on top of a pipeline in first quarter after the event. However, we see no response among properties similarly exposed to pipeline risk in other media markets within the state. We also fail to detect any response to the informational letter sent to households the following year. These results suggest that homebuyers are willing to pay to avoid pipeline risk when the issue has their attention, but that this attention hard to capture and fleeting.

Regulatory heterogeneity and spillovers

Environmental regulation often incompletely spans the set of products that firms supply or markets that they serve. When this happens, regulation in one market can affect outcomes in unregulated adjacent markets. In "Environmental Regulation, Imperfect Competition and Market Spillovers: The Impact of the 1990 Clean Air Act Amendments on the US Oil Refining Industry," I study the implications of this in the U.S. oil refining industry.[3]

The 1990 Clean Air Act (CAA) took the two most important refined petroleum products, gasoline and distillate, and divided their end markets up. Areas of the country with severe ozone problems were required to use a new clean grade of gasoline, called reformulated gasoline (RFG), and highway diesel consumers were required to purchase a new low sulfur grade of distillate, called low sulfur diesel (LSD). An important feature of this industry is that all refineries are multiproduct firms, producing regulated products *jointly* along with products not subject to regulation. Given this, it is imperative that any benefit-cost analysis of the industry measure changes in all markets, not just those directly affected. For example, a refinery that finds it costly to produce low sulfur diesel might convert less crude to diesel, and more to jet fuel. This would raise the price of diesel, but it would also reduce the price of jet fuel, and both margins need to be accounted for when computing welfare.

The fact that products and markets which are unaffected by the CAA are interlinked with regulated markets precludes simple difference-in-difference style estimation. Instead, I first specify a novel multiproduct production function, and directly estimate the costs associated with producing all products, including those newly created by the CAA. I then simulate counterfactual prices and quantities in all markets, not just regulated markets, to compute welfare changes. I find that spillovers into non-regulated markets are important: while consumers in regulated markets experienced welfare losses on the order of $34 billion during this period, this loss was partially offset by gains of $14 billion dollars among consumers in markets not subject to regulation.

[3]This was my job market paper in 2015, and it relied on confidential data from the Energy Information Administration. Under the Trump Administration, I abruptly and unexpectedly lost access to the data, which precluded any revision. It is possible that access could be restored when a new EIA Administrator is confirmed.

Energy Innovation

Measuring dynamic policy benefits

The past two decades have seen a significant expansion of renewable energy worldwide, fueled by generous subsidies. While these subsidies have undoubtedly been effective at increasing renewable energy adoption, there is a growing body of literature which compares their costs with careful estimates of the contemporaneous emissions avoided, and concludes that the costs exceed the *static* benefits (Van Benthem et al., 2008; Abrell et al., 2019; Greenstone and Nath, 2020). However, generous renewable subsidies were at least partially motivated by the belief that accelerating the deployment of renewable technology today would lower the costs of adoption tomorrow, and these dynamic policy benefits need to be accounted for. To this end, in "Winds of Change: Estimating Learning by Doing without Cost or Input Data" (with Thomas Covert), we estimate the size of dynamic cost reductions from policy-induced manufacturing experience in the global wind industry.

Estimating the extent to which production in the past has reduced manufacturing costs today is challenging because costs are not observed. Prices are not a substitute because the industry is highly concentrated, and because firms will incorporate any expected future learning benefits into their current prices (Benkard, 2004). We overcome these problems by modeling the confidential procurement process as a second-price menu auction, and show that this allows us to infer latent manufacturer opportunity costs from observed developer turbine choices. To isolate manufacturing costs from dynamic "markdowns", we borrow insights from Berry and Pakes (2000) that allow us to control for dynamic pricing incentives without computing the equilibrium of a dynamic game.

We estimate this model using data on the universe of commercial wind farm projects worldwide since 2000. We find that a doubling of manufacturing experience reduces manufacturing costs by approximately 30%. Only 1-2% of experience spills over to other turbine models produced by the same firm, and spillovers to turbines produced by other firms are on the order 0.1-0.2%. Nevertheless, even though inter-firm spillovers are small in relative terms, we show that in aggregate they rationalize a large share of observed cost reductions over time, consistent with policymaker motivation for generously subsidizing the industry.

Secrecy policy

The rapid maturation of hydraulic fracking and horizontal drilling (i.e. "fracking") at the turn of the century transformed the oil and gas industry. But, in order to take advantage of this new technology, exploration and production companies (E&P's) still had to figure out *where* to apply it. Mineral exploration is a classic example of a noncooperative environment with information externalities. Given the large cost of drilling, if possible, firms prefer to wait for their rivals to explore so that they can make a riskless decision afterwards. Such "waiting games" have been shown to generate costly delay, suboptimal information acquisition and inefficient sequencing (Hendricks and Kovenock, 1989). A common regulatory response to the problem of free riding on innovative efforts is to allow firms to keep secrets (Friedman et al., 1991). In the case of oil and gas exploration, many governments keep the information they collect for tax purposes confidential for a number of years before disclosing it to the public. By delaying or eliminating the

possibility of observing rival outcomes, strict secrecy reduces free riding. However, secrecy can exacerbate losses from incomplete information aggregation. During a secrecy period, some wells which shouldn't be drilled (based on all available information) will be, while other opportunities which are profitable will go unexploited.

In "Secrecy Rules and Exploratory Investment: Theory and Evidence from the Shale Boom" (with Thomas Covert), we study the net effect of these forces and ask whether confidentiality laws improve investment efficiency. Using a two-firm, two-period model, we characterize equilibrium behavior under policies which disclose whether investment outcomes exceed a predefined level. These policies include complete secrecy, in which players only observe rival actions, as well as full disclosure, in which players also perfectly observe rival outcomes. We provide a sufficient condition which guarantees that full disclosure generates more welfare than complete secrecy, and show that it depends on the extent of firm patience. We also characterize the optimal amount of partial disclosure, and the level which improves information aggregation but does not increase firms' incentive to free ride.

In the second part of the paper, we apply this model to rich oil and gas exploration data in the Appalachian shale basin where, at the start of the recent shale boom, Pennsylvania had a policy of complete secrecy, while neighboring West Virginia had a policy of full disclosure. We fit our theory model of each disclosure regime to the data, and estimate the structural primitives underpinning exploration in this setting. We find that the model corresponding to the disclosure policy on the books in each state rationalizes the data better than a model imposing the alternative, empirically validating our theoretical framework. With these primitives in hand, we then simulate counterfactual outcomes under the full continuum of disclosure thresholds. We find that full disclosure is the surplus maximizing policy in each state, generating at least 50% more expected value than complete secrecy.

Process innovation and slack time

Rather than a stable process, the U.S. onshore oil and gas industry has repeatedly discovered new and improved ways to profitably extract hydrocarbons from shale formations over the past 20 years. One curious "fact" that has been extensively documented in non-academic coverage of the industry is the observation that when oil prices decline, firms figure out how to become more productive in order to protect profits. This is curious because a simple model of productivity enhancing investment suggests the opposite relationship: the returns to drilling more productively are *increasing* in the output price. In "Frack Time versus Slack Time" (with Thomas Covert and Ryan Kellogg), we hypothesize that this pattern could be rationalized by the classical "slack time" theory of innovation from the management literature (Nohria and Gulati, 1996). According to this theory, a firm might allocate all available resources to maximizing short-run output during a boom, but then reallocate slack labor towards long-run innovative effort during a bust.

Testing this hypothesis empirically is challenging due to the fact that latent mineral quality varies considerably across space. If firms at least partially observe this, then when prices are high they will drill lower quality locations, on average, than when prices are low. If we do not account for this selection, then, even if productivity remains constant over time, it would look like firms are getting more productive when prices are low. To account for endogenous unobserved well quality, we introduce a new non-parametric productivity estimator which relies solely on the order in which investments are made. Crucially, this estimator does not require us to model firms' beliefs

about the evolution of productivity. We use this estimator to recover firm-by-time-by-location productivities for every major shale play in the United States over the last two decades. We then leverage variation in the relative price of natural gas and oil over time to rationalize changes in productivity as a function of slack time.

Future work

Renewable energy procurement and deployment

The International Renewable Energy Agency estimates that fully decarbonizing the electricity sector will require over $130 trillion of investment in renewable energy. A core theme of my future research agenda will be understanding how to achieve decarbonization at the lowest possible social cost.

My current, early-stage, work in this area concerns renewable energy auctions. The rapid growth in renewable energy starting in the 2000s was fueled primarily by "feed-in tariffs" – fixed subsidy payments for renewable energy output. While these subsidies were undoubtedly effective at encouraging renewable energy development, they were also very expensive. In response to this, recently numerous governments around the world have replaced these fixed tariffs with renewable energy auctions. Motivated by this policy shift, as well as my earlier work on the benefits of auctions in the oil and gas sector, I aim to answer two questions: How much have auctions reduced renewable procurement costs to date? Are there changes to the renewable auction process which will reduce public costs further?

To answer these questions, I focus on Germany, the third largest renewable energy market in the world. Germany's feed-in tariff regime was uniquely formulaic and site-specific. I use this empirical variation, along with the detailed engineering and economic data employed in my "'Winds of Change" paper, to estimate the distribution of renewable plant construction costs. I then develop an empirical model of wind developer bidding in a multi-year, pay-as-bid procurement auction. With these estimates and model in hand, I am in the process of estimating the effect of Germany's introduction of auctions in 2017 two ways. First, I will simulate the level of renewable entry and the cost to the German government post-2017 in a counterfactual world where the auctions were never enacted. Second, I will simulate outcomes and government costs prior to 2017 in a counterfactual world where auctions were always in place from the advent of the industry. These two complementary strategies, will provide estimates of the gains from this widely debated renewable energy policy, in an environment where simple pre-post comparisons are difficult to interpret causally due to trending confounders. In future work, I plan to simulate counterfactual outcomes under an alternative auction design recently proposed by the European Commission, which replaces a bid ceiling with a rule that awards contracts to the lowest 80% of bidders.

My Work

Aldy, Joseph E, Todd D Gerarden, and Richard L Sweeney, "Investment versus Output Subsidies: Implications of Alternative Incentives for Wind Energy," *accepted at The Journal of the Association of Agricultural and Resource Economists*, 2022.

Allcott, Hunt and Richard L Sweeney, "The Role of Sales Agents in Information Disclosure: Evidence from a Field Experiment," *Management Science*, 2017, *63* (1), 21–39.

Burtraw, Dallas, Richard Sweeney, and Margaret Walls, "The Incidence of US Climate Policy: Alternative Uses of Revenues from a Cap-and-Trade Auction," *National Tax Journal*, 2009, *62* (3), 497–518.

Chan, Gabriel, Robert Stavins, Robert Stowe, and Richard Sweeney, "The SO2 Allowance Trading System and the Clean Air Act Amendments of 1990: Reflections on 20 Years of Policy Innovation," *National Tax Journal*, 2012, *65* (2), 419–452.

Covert, Thomas R and Richard L Sweeney, "Relinquishing Riches: Auctions vs Informal Negotiations in Texas Oil and Gas Leasing," *conditionally accepted at The American Economic Review*, 2022.

_ **and** _ , "Secrecy Rules and Exploratory Investment: Theory and Evidence from the Shale Boom," submitted to The Review of Economic Studies, 2022.

_ **and** _ , "Winds of Change: Estimating Learning by Doing without Cost or Input Data," working paper, 2022.

_ **, Ryan Kellogg, and Richard L Sweeney**, "Auctions vs. Negotiations with Spillovers: Evidence from the U.S. Land Ordinance and the Shale Boom," in progress, 2022.

_ **,** _ **, and** _ , "Frack Time versus Slack Time," in progress, 2022.

Herrnstadt, Evan and Richard L Sweeney, "Housing Market Capitalization of Pipeline Risk: Evidence from a Shock to Salience and Awareness," *revision resubmitted to Land Economics*, 2022.

Muehlegger, Erich and Richard L Sweeney, "Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom," *The Review of Economics and Statistics*, forthcoming.

Sweeney, Richard L, "Environmental Regulation, Imperfect Competition and Market Spillovers: The Impact of the 1990 Clean Air Act Amendments on the US Oil Refining Industry," working paper, 2015.

Other References

Abrell, Jan, Mirjam Kosch, and Sebastian Rausch, "Carbon abatement with renewables: Evaluating wind and solar subsidies in Germany and Spain," *Journal of Public Economics*, 2019, *169*, 172–202.

Allcott, Hunt, "Paternalism and Energy Efficiency: An Overview," *Annual Review of Economics*, 2016, *8* (1), 145–176.

Benkard, C Lanier, "A dynamic analysis of the market for wide-bodied commercial aircraft," *The Review of Economic Studies*, 2004, *71* (3), 581–611.

Benthem, Arthur Van, Kenneth Gillingham, and James Sweeney, "Learning-by-doing and the optimal solar policy in California," *The Energy Journal*, 2008, *29* (3).

Berry, Steve and Ariel Pakes, "Estimation from the Optimality Conditions for Dynamic Controls," December 2000.

Friedman, David D, William M Landes, and Richard A Posner, "Some economics of trade secret law," *Journal of Economic Perspectives*, 1991, *5* (1), 61–72.

Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker, "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes," *American Economic Journal: Applied Economics*, April 2020, *12* (2), 303–42.

Greenstone, Michael and Ishan Nath, "Do renewable portfolio standards deliver cost-effective carbon abatement?," *Becker-Friedman Institute Working Paper*, 2020.

Hendricks, Kenneth and Dan Kovenock, "Asymmetric information, information externalities, and efficiency: the case of oil exploration," *The RAND Journal of Economics*, 1989, pp. 164–182.

Knittel, Christopher R., Ben S. Meiselman, and James H. Stock, "The Pass-Through of RIN Prices to Wholesale and Retail Fuels under the Renewable Fuel Standard," *Journal of the Association of Environmental and Resource Economists*, 2017, *4* (4), 1081–1119.

Nohria, Nitin and Ranjay Gulati, "Is slack good or bad for innovation?," *Academy of Management Journal*, 1996, *39* (5), 1245–1264.

Parish, Ross M. and Keith Robert McLaren, "Relative Cost-Effectiveness of Input and Output Subsidies," *Australian Journal of Agricultural Economics*, April 1982, *26* (1), 1–13.

Rosen, Sherwin, "Hedonic prices and implicit markets: product differentiation in pure competition," *Journal of Political Economy*, 1974, *82* (1), 34–55.

Tiebout, Charles M, "A Pure Theory of Local Expenditures," *Journal of Political Economy*, 1956, *64* (5), 416–424.

Weyl, E. Glen and Michal Fabinger, "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, June 2013, *121* (3), 528–583.