Title: “Reference Pricing as a Deterrent to Entry: Evidence from the European Pharmaceutical Market."

Authors: Luca Maini and Fabio Pammolli

Journal: American Economic Journal: Microeconomics

Publisher: American Economic Association

Year: 2023

Doi: https://doi.org/10.1257/mic.20210053

American Economic Journal: Microeconomics 2023, 15(2): 345–383
https://doi.org/10.1257/mic.20210053

Reference Pricing as a Deterrent to Entry: Evidence from the European Pharmaceutical Market[†]

By Luca Maini and Fabio Pammolli[*]

External reference pricing (*ERP*), *the practice of benchmarking domestic drug prices to foreign prices*, *generates an incentive for firms to withhold products from low-income countries. Using a novel moment inequality approach*, *we estimate a structural model to measure how ERP policies affect access to innovative drugs across Europe. We find that ERP increases entry delays in eight low-income European countries by up to one year per drug. The European Union could remove these delays without replacing ERP by compensating firms through lump-sum transfers at the cost of around €18 million per drug.* (*JEL* L13, L51, L65)

External reference pricing policies (ERP) are a popular form of drug regulation in Europe. Governments use ERP to ensure they pay prices in line with other countries and to save money. Recent proposals have advocated for the introduction of ERP in the United States based on similar considerations.[1] However, these policies can affect firm behavior in ways that policymakers often fail to appreciate fully. ERP limits the ability to price discriminate across countries with different income levels. Economic theory suggests that firms could react to price discrimination constraints by selling only to consumers with a high willingness to pay. In this case, drug manufacturers could react to ERP policies by delaying the launch of new drugs in low-income countries where prices are referenced by high-income countries.[2] This logic suggests that ERP could have important global implications on access to innovative drugs.

[*]Maini: Harvard Medical School (email: maini@hcp.med.harvard.edu); Pammolli: Politecnico di Milano (email: fabio.pammolli@polimi.it). John Asker was coeditor for this article. Maini is extremely grateful to his advisers Ernst Berndt, David Cutler, Robin Lee, and Ariel Pakes, and his longtime collaborator Josh Feng. We would also like to thank Vincenzo Atella, Francesco Decarolis, Sharat Ganapati, Josh Gottlieb, Sonia Jaffe, Pranav Jindal, Myrto Kalouptsidi, Richard Manning, Brian McManus, Casey Mulligan, Tomas Philipson, Daniel Pollman, David Ridley, Mark Shepard, Ariel Stern, Elie Tamer, Pietro Tebaldi, Heidi Williams, Jon Williams, Annetta Zhou, and participants of various conferences and seminars where we presented early drafts of our work. Maini thanks the Becker Friedman Institute for their financial support through the Health Economics Fellowship. IMS Health provided access to the data.

[†]Go to https://doi.org/10.1257/mic.20210053 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1]The Department of Health and Human Services proposed the International Pricing Index Model for Medicare Part B in October 2018 (https://www.govinfo.gov/content/pkg/FR-2018-10-30/pdf/2018-23688.pdf, retrieved December 2019). House speaker Nancy Pelosi unveiled the Lower Drug Costs Now Act in September 2019 (https://www.speaker.gov/sites/speaker.house.gov/files/HR3%20Backgrounder%2010.2.19.pdf, retrieved December 2019).

[2]Firms can also react by selling their product at higher prices in low-income countries. We present some suggestive evidence that firms do so in Section C.3 of the online Appendix and argue that it does not affect the validity of our estimation or our counterfactuals.

This paper uses data from the European Economic Area (EEA) to show that firms react to ERP policies by strategically delaying entry in certain countries.[3] We develop a single-agent entry model that explicitly incorporates the price externality generated by ERP across markets and use it to isolate the effect of ERP from that of other policies and constraints. We then conduct a counterfactual exercise to calculate how access to innovative drugs would change without ERP.

The critical empirical challenge of this research question lies in separating *strategic* delays related to ERP from *idiosyncratic* delays that arise for unrelated reasons. In the EEA, new drugs are rarely launched everywhere right after receiving marketing approval, and therefore it is unlikely that observed delays are solely attributable to ERP. A primary source of delays, for example, is the requirement that before launching, manufacturers must send each country a separate pricing application. Unfortunately, we only observe launch dates and not application dates. Thus, in our data, an application undergoing review is observationally equivalent to an application temporarily withheld for strategic reasons.

Since our data do not distinguish between idiosyncratic and strategic delays, we rely on a model to separate them. Intuitively, our results are driven by differences in demand and prices across countries and drugs. The model interprets delays in small countries with low price levels as strategic if these countries are referenced by larger countries with higher price levels. Conversely, delays in large markets with high price levels are usually interpreted as idiosyncratic. Within-country variation in demand is also essential. Given two drugs with identical launch delays, the model assigns a greater probability of having incurred an idiosyncratic delay to the drug with higher predicted demand.

Not knowing the application date also presents a challenge for estimation, which we overcome by developing a novel moment inequality estimator. Firms solve a maximization problem by choosing when to send an application, but cannot control the exact launch date. Since the launch date is the only thing we observe in the data, we do not know the firm's strategy. We also cannot derive the optimal strategy as the solution to our model because the model is too complicated to solve, either analytically or numerically. This limitation rules out estimators that require knowing the firm's strategy, such as maximum likelihood or revealed-preference moment inequalities constructed as in, e.g., Pakes (2010) and Pakes et al. (2015). Instead, we build inequalities based on observed firm revenue. Our approach still relies on a revealed-preference argument but does not require observing the firm's strategy, though it imposes stricter assumptions on the error term. The estimator has potential applications beyond our empirical analysis. It can be used in settings where firms face unobserved strategy constraints, such as market entry regulation or capacity constraints, or when choice sets are unspecified (see, e.g., the problem in Barseghyan et al. 2020).[4]

[3] The EEA includes all European Union (EU) member states plus Norway and Iceland. We also include Switzerland, which has a series of bilateral trade agreements with the EU that allow the country to participate in its common market.

[4] We consider this estimator a valuable but not central contribution of the paper and include a derivation of its properties in Section D.5 of the online Appendix.

Our estimates suggest that replacing ERP with a pricing rule that does not link prices across countries would reduce delays in a set of lower-income Eastern European countries by up to one year per drug. The specific pricing rule does not matter for the results of the counterfactual, so our estimates hold across a broad range of alternative policies: from transitioning to a centralized European cost-effectiveness evaluation system (Drummond 2003) to two-part pricing systems with barriers preventing reference pricing and import-export of pharmaceutical products (Towse et al. 2015). While these policies would have a different impact on firm profits and consumer welfare, our model implies that they would have the same effect on access.

Our findings also suggest that recent calls for price transparency in Europe may increase access disparity.[5] While price transparency can reduce the cost of monitoring the affordability of pharmaceutical products, it can also tighten ERP constraints, which would exacerbate the problem of delays in low-income countries.

Finally, we use our counterfactual estimates to calculate the returns from managing the launch sequence. We find that, by engaging in strategic delays, firms earn an extra €18 million per drug on average, which—though a relatively small sum—is roughly equivalent to 40 percent of the lifetime earnings of the average drug in Eastern Europe. We suggest that the EU could improve drug access in these countries by compensating manufacturers for ERP-generated revenue losses with lump-sum transfers. Because lump-sum transfers have—to a first-order approximation—a small impact on welfare, the overall effect of this policy would almost certainly be positive.

This paper contributes to three main strands of economic literature. First, it belongs to a growing body of work studying how price regulation affects access to pharmaceuticals. The empirical side of this literature usually analyzes the impact of government policy on access using a reduced-form framework (Danzon et al. 2005; Kyle 2007; Danzon and Epstein 2012; Kyle and Qian 2013; Cockburn et al. 2016).[6] On the theory side, this literature has focused on simulating the impact of reference pricing on firm strategy (e.g., Stargardt and Schreyögg 2006; Toumi et al. 2013; Borja 2014; Houy and Jelovac 2015) or on establishing conditions under which regulation limiting price discrimination is beneficial or harmful to welfare (e.g., Brekke et al. 2007, 2015; Birg 2016; Brekke et al. 2016; Matteucci and Reverberi 2017). In contrast, our paper explicitly models the impact of reference pricing on firm incentives and develops an estimation strategy to isolate the effect of this policy on launch delays.

Second, our paper is related to a series of studies on the impact of tying prices to endogenous market benchmarks. While most of this literature focuses on price responses, our paper shows that firms can also respond along other margins (i.e., manipulating the entry strategy). Dubois et al. (2018) simulate the impact of ERP

[5] See, e.g., Wenzl and Chapman (2019), or a recent WHO resolution urging member states to improve the transparency of drug markets by disclosing information on actual prices paid: https://www.who.int/news-room/detail/28-05-2019-world-health-update-28-may-2019, retrieved March 2020. The Valletta Declaration of 2019 also mentions price transparency.

[6] A notable exception is Duso et al. (2014), who examines the welfare impact of parallel trade in Germany. Another methodologically related paper is Chaudhuri et al. (2006), which uses structural techniques to estimate the effect of patent policy on patient welfare in the Indian market for quinolones.

adoption in the United States. Duggan and Scott Morton (2006); Ridley and Lee (2020); and Feng et al. (2020) show that when US government programs tie drug reimbursements to the average private market price, the commercial market is affected in a variety of ways. Jaffe and Shepard (2017) and Decarolis (2015) show that linking subsidies to premiums affects equilibrium prices in health exchanges and Medicare Part D, respectively. More generally, price externalities across firms have been detected in the absence of government intervention. For example, Grennan (2013) and Grennan and Swanson (2016) show that knowing how much rival hospitals paid for medical devices can affect future prices.

Our third and final contribution is to the literature on partial identification. We extend the framework introduced by previous papers (e.g., Bajari et al. 2007; Pakes 2010; Pakes et al. 2015) to derive moment inequalities when firm strategies are unobserved. However, our estimator does not nest previous ones, as our setting requires the econometrician to recover the structural error term in the static estimation. This literature also includes several empirical papers (e.g., Katz 2007; Crawford and Yurukoglu 2012; Eizenberg 2014; Ho and Pakes 2014; Illanes 2016; Dickstein and Morales 2018; Wollmann 2018). The ones closest to us are Holmes (2011) and Morales et al. (2019), both of which also estimate single-agent models where entry decisions dynamically affect future profits.

The rest of the paper proceeds as follows. We begin by discussing the relevant features of the European pharmaceutical market in Section I. We then describe the data and present some reduced-form evidence in Section II. We outline our theoretical model in Section III and discuss estimation strategy and results in the following two sections: static estimation of demand and prices in Section IV and dynamic estimation in Section V. In Section VI we discuss our counterfactual result and its policy implications. Finally, in Section VII we provide some concluding remarks.

I. Overview of the European Drug Market

This section describes how new drugs receive marketing approval in the EEA and then provides a brief overview of drug price regulation across Europe, focusing on ERP.

New drugs can only be marketed after being reviewed for efficacy and safety. The European Medicines Agency (EMA), founded in 1995, oversees this review process in the EEA.[7] Drugs generally receive simultaneous approval in all EEA member states, though there are multiple ways to do so. The *centralized process*, administered by the EMA, grants marketing approval for the entire EEA. Alternatively, country-by-country approval is also an option. Firms can apply to a set of countries using the *decentralized procedure*, where one EEA member state acts as the primary reviewer. In these instances, firms can later extend the original approval to any other

[7] Switzerland is the only country in our data that does not automatically recognize EMA decisions on marketing approval. Additionally, Eastern European countries started recognizing EMA decisions automatically only upon joining the European Union. Estonia, Hungary, Lithuania, Latvia, Poland, and Slovenia joined in 2005, while Bulgaria and Romania joined in 2008.

country by applying for *mutual recognition.* Extensions are granted automatically within 180 days unless a member state raises additional safety concerns.

After receiving marketing approval, firms usually encounter additional delays because individual countries retain the ability to regulate prices independently from one another, and drugs cannot be sold before the firm and the government agree on pricing terms.[8] The time required to review an application and negotiate a price can vary significantly across countries. Data on turnaround times for applications are scarce, but survey evidence from the late 1990s and early 2000s indicates that the average varies substantially, from zero days in the United Kingdom and Germany to over two years in Poland (OECD 2008; PICTF 2006).

Firms petition for reimbursement status by submitting pricing and reimbursement applications to the government of each country. Requirements for these applications vary, though they must generally include evidence of the medical benefits of the drug as well as projected sales and a proposed price. The government then uses this information as an input into the pricing decision. The final price depends on various factors, including cost-effectiveness, internal reference pricing (which links prices of other molecules sold in the same country within a pre-specified equivalence class), and price-volume agreements to limit overall spending on high-volume drugs (Carone et al. 2012). Most European countries also use ERP as a criterion.

External reference pricing links the price of the same branded drug across countries.[9] The two most important aspects of ERP policies are the reference basket (i.e., the basket of countries whose prices are sampled) and the formula used to compute the reference price. For both, there is significant variation across countries (Figure 1).[10] Some governments (e.g., Austria, Belgium, Finland, Hungary, and Poland) require firms to submit prices from all other countries in the European Union. Others only reference similar countries in terms of geographical proximity, size, and income level—for example, Estonia references Hungary, Latvia, and Lithuania, while France references Germany, Italy, Spain, and the United Kingdom. With respect to the reference formula, most countries use the average across the reference basket, but a few (e.g., Latvia, Poland, and Romania) use the lowest price or other slight variations.

Countries may adhere to their stated ERP guidelines with varying stringency. Some governments (Belgium, Finland, France, Italy, Poland, and Spain) state that they only use ERP to "inform" the pricing decision, meaning that we might expect prices to be affected by ERP but not necessarily be perfectly aligned with the reference pricing benchmark. In other instances, governments may push for prices below the benchmark if they expect to sell higher volumes than the referenced countries.

[8] Pricing restrictions typically only apply to drugs that the government pays for through the public health insurance system. However, since European citizens overwhelmingly access health care through government-funded programs, excluding a product from public formularies results in its de facto exclusion from the national market (European Commission 2012).

[9] Countries do not use the price of other products in their ERP functions. This is important because if the prices of different products marketed in different countries could affect each other, the launch strategies of all products would be closely interconnected. Modeling these interactions would require a multiple-agent model instead of the simpler single-agent model we utilize.

[10] See Section B.2 of the online Appendix for a more detailed discussion of how we collected the information in Figure 1, and for additional details on ERP policies.

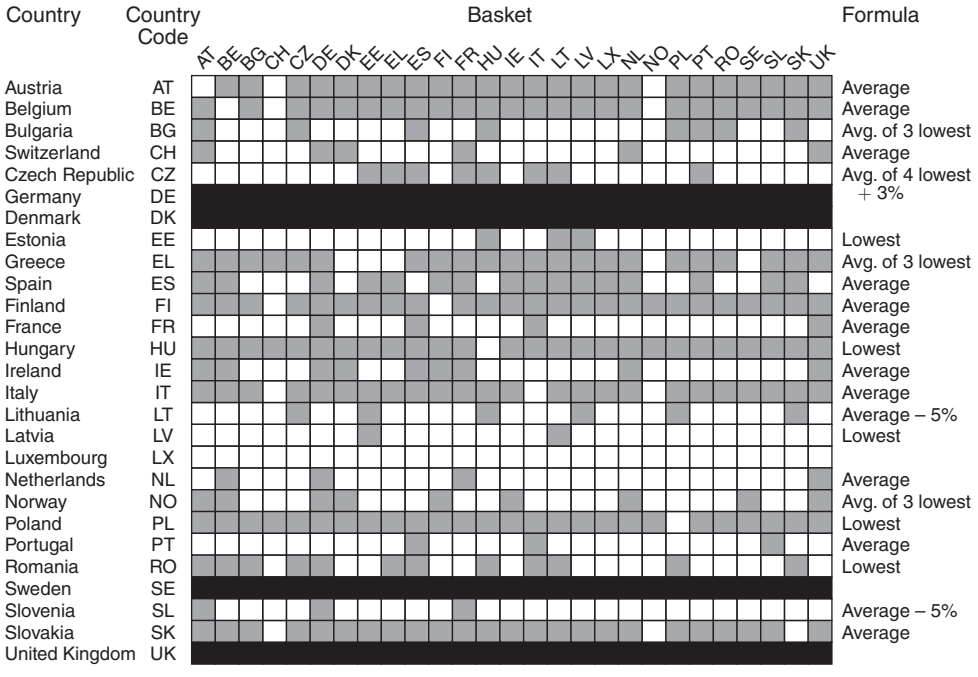

Figure 1. Reference Price Baskets and Formulas for Countries

Note: Luxembourg only references the drug's country of origin.

ERP affects incentives of firms by limiting their ability to price discriminate across countries. Another force that operates through the same channel is parallel trade. The EEA is a free trade area with no limits over the flow of goods across borders. This includes patent-protected products, such as prescription drugs. Parallel importers purchase drugs in countries with low prices and sell them abroad where prices are higher.[11] Without frictions, parallel trade would drive the price of all drugs to the minimum available price and act like an extreme version of reference pricing. In practice, this does not happen because firms can fight parallel trade by managing supply quotas and producing slightly altered versions of the same product (Kyle 2011). However, parallel trade can still undermine firm revenue. Using our data, it is impossible to completely separate the effect of reference pricing from that of parallel trade. However, our model generally applies to any policy that limits price discrimination, and our counterfactual analysis refers to a world where both reference pricing and parallel trade are unavailable.

II. Preliminary Evidence on Launch Delays

The primary data source for the empirical analysis is the MIDAS database (IMS Health 2012), maintained by IQVIA (formerly IMS Health). The dataset covers

[11] Crucially, ERP does not apply to versions of the drug sold by importers as parallel-traded products.

sales of all prescription drugs for European countries from 2002 to 2012. It consists of a quarterly panel of volume and revenue sales for each country in the EEA.[12] We calculate prices as the ratio of revenue to volume sales.[13]

We integrate the IMS data with a few additional sources. We collect approval dates for all EMA-approved medications from the EMA's website and the approval date of all mutual recognition applications from an internet database maintained by the Heads of Medicines Agencies (EMA 2012; HMA 2012). We also collect GDP, population, and exchange rate data from Eurostat and the European Central Bank and data on the incidence of diseases in each European country from the Global Burden of Disease (GBD) Study (ECB 2012; Eurostat 2012; IMHE 2012).

Our sample of analysis consists of patent-protected drugs with a potential market spanning multiple European countries. We select a subsample of patent-protected branded drugs that satisfies three criteria. First, the drug launched in Europe after January 1, 1995, ensuring that the marketing approval process was overseen by the EMA. Second, the drug had at least one new launch in a European country after January 1, 2002, meaning the drug's manufacturer still has some strategic choices to make that we can observe in the data.[14] Third, it was either matched to our approvals database with an approval date between 1995 and 2012 or is a patent-protected brand drug sold in at least ten countries by 2012. This last criterion ensures that the drug's market is wide enough to motivate a strategic approach to the launch sequence.

Our final selection consists of 481 drugs—defined as a molecule-firm-therapeutic class combination. Table 1 displays basic summary statistics for our sample. Unsurprisingly, the drugs in our main sample have much greater sales and diffusion relative to other drugs in the data, which are either generics or drugs sold in national markets only. Most of the products we include received approval directly from the EMA or applied for mutual recognition. Those that did not are drugs that we could not match with approval dates from any database but that we observe being sold in many European countries.[15]

From our main sample, we also selected a subsample of 87 drugs that lost exclusivity before December 31, 2012. This smaller group is used in the dynamic analysis because our methodology requires knowing the overall expected payoff of a drug over its entire lifecycle.[16]

[12] We are missing some information for specific countries and years. Please refer to Section A.4 of the online Appendix for a detailed description of the data cleaning and imputation process.

[13] Like virtually every other source of drug pricing data, the MIDAS database does not incorporate hidden discounts and rebates to payers. We discuss how the presence of these rebates may affect our analysis in Section A.1.3 of the online Appendix.

[14] Ideally, we would include all drugs that could *potentially* launch in an additional country. However, the assumptions of our model, detailed in Section IIIA, require us to observe sales in all countries included in the estimation. In practice, the set of drugs with potential launches and those with observed launches overlap almost completely.

[15] We calculate launch delays from the marketing authorization date. See section A.1.2 in the online Appendix for a description of how we estimated this latter variable.

[16] Patents generally expire at the same time in most countries since they are administered by the European Patent Office. However, some countries can choose to grant extensions to individual patents. In cases where the dates differ, we set period T as the latest expiration date among France, Italy, and Spain. These three countries are the three largest markets that use ERP. Therefore, when their patent protection expires, the strategic incentives to delay launches should all but disappear. Empirically, we observe only eleven new launches after period T.

Table 1—Summary Statistics

		Full sample	Main sample	Dynamic sample
Number of therapeutic classes		241	109	44
Number of firms		2,944	168	47
Number of molecules		6,354	475	86
Class-firm-mol combinations		55,131	481	87
Class-firm combinations		25,572	375	84
Number of mol. per firm	Mean	18.7	2.9	1.21
	Median	2	1	1
Number of mol. per firm-class	Mean	2.2	1.3	1.03
	Median	1	1	1
Diffusion	Mean	2.1	20.1	21.3
	Median	1	22	24
Yearly sales	Mean	€3,547,665	€115,427,475	€121,977,298
	Median	€92,489	€39,214,276	€46,340,438
Approval method	EMA		312	24
	MRP		127	46
	Other		42	17

While virtually all countries experience some launch delays relative to the approval date, the magnitude of these delays varies substantially across countries. Figure 2 shows the diffusion of drugs in our main sample across all European countries at various intervals. The maps show how some countries, particularly in Eastern Europe, lag behind the rest of the continent in terms of drug access. For example, more than one-third of all available products are unavailable in Bulgaria, Estonia, Latvia, and Romania six years after marketing approval.[17]

While the presence of delays is not enough by itself to conclude that firms are responding strategically to reference pricing, evidence from a 2004 survey of manufacturers suggests that application delays are not enough to explain the delays we see in the data (PICTF 2006). As discussed in the introduction, delays can arise because of the time required to obtain government approval for reimbursement. In Figure 3 we plot the average application delay reported in the survey against the average observed delay in our data for each country.[18]

If the application process caused all delays, we would expect the average application delay to match the delay observed in the data, but this is not the case. The plot shows that observed average delays tend to be greater than the turnaround times for pricing and reimbursement applications. Moreover, the relationship between these two measures is very different across high- and low-income countries. In Western Europe, where income is generally higher, the two measures are strongly correlated, suggesting that application delays can explain most of the variation. In lower-income Eastern European countries, however, we cannot detect any correlation, which we

[17] The Netherlands also appears to have low diffusion rates, but only because of missing data. Coverage for the Netherlands starts in 2007 and excludes hospital sales. Hence, our data do not include hospital-only products and products that exited before 2007.

[18] The survey did not report turnaround time information for Bulgaria, Latvia, Lithuania, Luxembourg, Romania, and Slovenia.

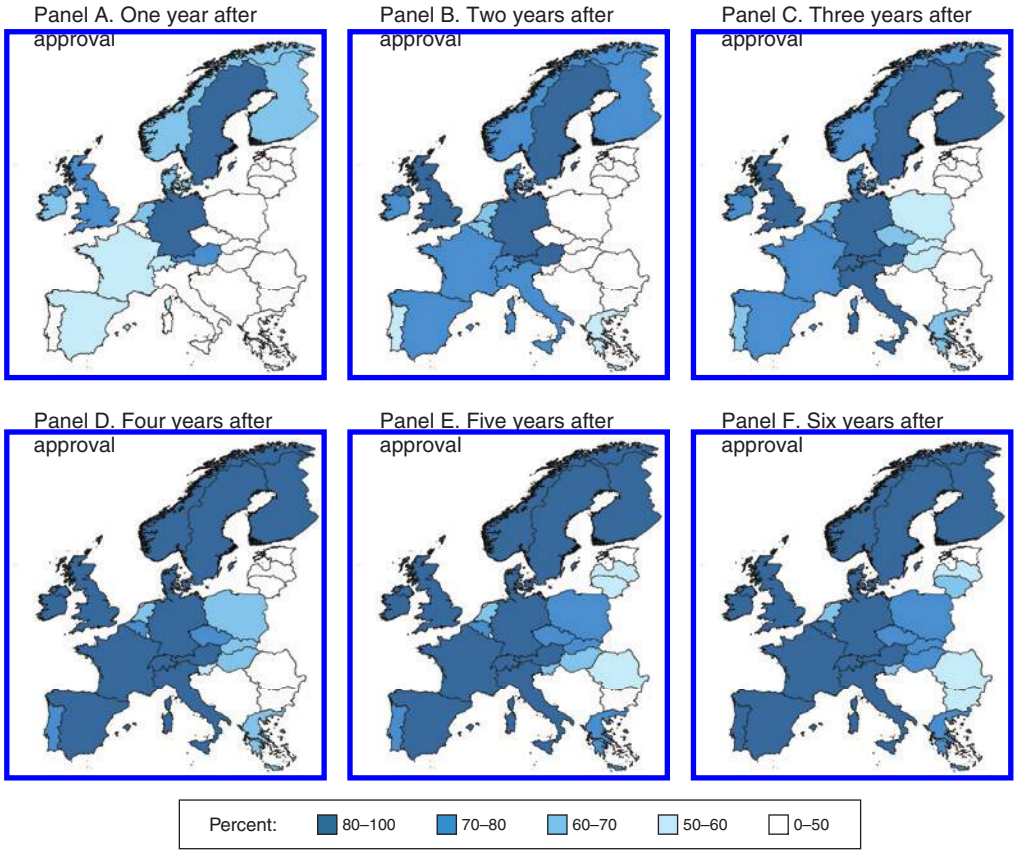

FIGURE 2. DIFFUSION OF EMA-APPROVED DRUGS IN EUROPEAN COUNTRIES

would expect if launch timing depended on strategic considerations rather than bureaucratic constraints.[19]

III. A Dynamic Entry Model with Price Externalities

We consider a single-agent model of entry across multiple markets. For simplicity, we consider single-product firms in our exposition of the model. We use the index j, which we use interchangeably to refer to both a firm and a product.[20] The firm's product has a marketing authorization for sale in a finite set $\mathcal{N}_j = \{1, \ldots, N_j\}$ of markets (European countries), indexed using the subscript k. The patent on each product has an expiration date, T_j periods into the future, at which point generic alternatives enter and profits fall to zero. The firm's objective is to maximize profits over the life cycle of its product.

[19] We present additional suggestive evidence that rejects alternative explanations for launch delays in online Appendix 3.1.

[20] As Table 1 shows, most firms in our data are single-product firms. The model can be easily extended to allow for multi-product firms, and our estimation strategy is valid for both single- and multiproduct firms.

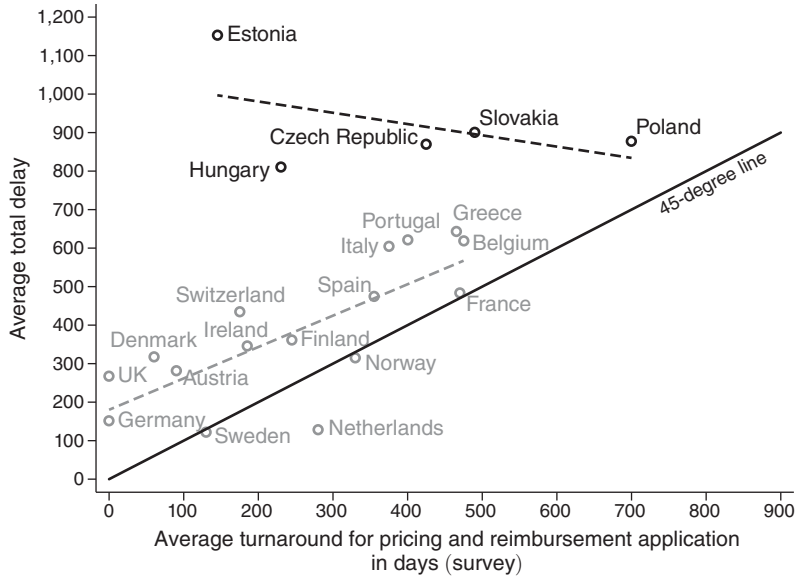

FIGURE 3. AVERAGE APPLICATION TURNAROUND VERSUS AVERAGE DELAY

In each period, the firm is solving a two-part problem: (i) in what countries should the product be launched?, and (ii) at what price? We are interested in understanding strategic launch delays, which are the outcome of the first part of the problem. The optimal launch strategy will depend on the equilibrium prices in each country. However, firms have limited agency in determining these prices because European governments strictly regulate drug prices. Therefore, we do not explicitly model the price-setting stage but instead, use a flexible parametric function to predict equilibrium prices.

We start by introducing some notation. Denote the *launch sequence* of firm j as $S_j = \{s_{jk}\}_{k \in \mathcal{N}_j}$, where s_{jk} denotes the period of entry of product j in country k. Furthermore, denote the launch sequence at the end of period t as $S_{jt} = \{s_{jkt}\}_{k \in \mathcal{N}_j}$, where

$$s_{jkt} = \begin{cases} s_{jk}, & \text{if } s_{jk} \leq t; \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

Once a product has entered, we assume that it cannot be voluntarily withdrawn.[21] We similarly denote the launch sequence of other firms as S_{-j}. We occasionally also use the shorthand S or S_t to indicate the launch sequences of all firms. Which sequence maximizes profits depends on demand and prices. The firm takes the

[21] We are aware of at least one case in which the firm voluntarily withdrew a drug due to a price disagreement (Tresiba, in Germany in 2015), but none in our data, which ends in 2012. All product exits we observe either occur following several years of falling sales numbers (suggesting the product is no longer economically or therapeutically viable) or can be linked to a suspension of the marketing authorization.

demand system and the price-setting equation as given when making entry decisions. We describe each in turn before specifying the dynamic entry model.

A. *Demand for Pharmaceutical Products*

We base demand on the logit random utility model. Our market definition is a country, year, and therapeutic class. We aggregate products within a therapeutic class at the molecule-brand status level. We define three possible brand statuses: originator (i.e., the brand sold by the patent holder or main manufacturer), non-originator brand (usually a parallel traded product), and generic.

The utility of consumer i in country k from consuming drug j in year t is given by

$$u_{ijt} = \delta_{jkt} + \nu_{ijt}. \tag{2}$$

To obtain more realistic substitution patterns, we also add a nesting structure at the molecule level. The error term ν_{ijt} is parametrized as

$$\nu_{ijt} = \zeta_{m(j)} + (1 - \sigma)\epsilon_{ijt}, \tag{3}$$

where $m(j)$ indicates the molecule of drug j, σ lies on the unit interval, ϵ_{ijt} is distributed according to a standard extreme value type 1 (EV 1) distribution, and $\zeta_{m(j)}$ is an error term whose distribution satisfies the property that ν_{ijt} is distributed according to an EV 1 distribution as long as ϵ_{ijt} is also EV 1 (Cardell 1997).[22]

We parametrize δ_{jkt} as

$$\delta_{jkt} = \alpha_{jk} + \beta_j age_{jt} + \eta_j NF_{jkt} + \xi_{jkt}. \tag{4}$$

Our specification incorporates two important empirical features of drug demand: heterogeneity in preferences across countries and growing demand over time.[23] α_{jk} captures a country-specific preference for each drug, which could reflect differences in prescribing guidelines or disease burden. β_j accounts for drug-specific time trends, which could be generated by physician learning or by the slow diffusion of information. We measure age starting with the drug's approval date. For non-originator products, we also keep track of the number of selling firms as a separate control variable NF_{jkt}.[24] Finally, we add a drug-country-year random shock, ξ_{jkt}. We do not include a coefficient for price since we do not observe the price that patients pay. Instead, we include realized demand as a control in the price function,

[22] For the nested logit model to make sense, σ must lie on the unit interval. We do not implement this restriction in the estimation. Instead, we switch to a simple logit whenever the parameter falls outside the unit interval.

[23] Demand for drugs can differ substantially across countries because of heterogeneity in prescribing guidelines, the incidence of disease, and patient preferences. Moreover, possibly because drugs are generally considered experience goods (Crawford and Shum 2005), demand for most products increases over the life cycle.

[24] All originator products are sold by a single firm in each country, though the firm is not necessarily the same across countries. However, most molecules face multiple brand and generic competitors, which we aggregate to avoid excessive entry and exit, and because many of these products have vanishingly small market shares.

implicitly assuming that the government mediates any relationship between price and volume.[25]

Including a drug-country fixed effect imposes some practical limits because we can only estimate this coefficient for launched products. As a result, we must exclude drug and country pairs where we do not observe sales from our analysis. We discuss the implications of this assumption for estimation in Section VB.

Inverting market shares (and normalizing the utility of the outside option to zero) yields the standard estimating equation for market shares MS_{jkt}:

$$\ln\left(\frac{MS_{jkt}}{MS_{0kt}}\right) = \alpha_{jk} + \beta_j age_{jt} + \eta_j NF_{jkt} + \sigma_m \ln\left(\frac{MS_{jkt}}{MS_{mkt}}\right) + \xi_{jkt}, \tag{5}$$

where MS_{0kt} is the share of the outside good and MS_{jkt} and MS_{mkt} are the market share of the product and the overall market share of the molecule nest, respectively.[26] We denote the demand function generated by this model as $D_{jkt}(S_t, \xi_{kt})$, where $\xi_{kt} = \{\xi_{jkt}\}$ is the vector of shocks for all products in country k and year t.

B. *Price-Setting Equation*

Drug prices are negotiated by firms and governments, and the exact form of the negotiation is hard to capture explicitly in a model. Governments are trying to reconcile several goals, such as providing access to valuable medications and rewarding innovation while at the same time facing a budget constraint. Since we do not have any information on the government's objective function, we opt for an agnostic approach and model prices using a flexible control function.[27]

Our price-setting equation includes two components. The first component is what we call *government price*, p^{gov}_{jkt}:

$$p^{gov}_{jkt}\big(D_{jkt}(\cdot)\big) = \theta_j \cdot \gamma_k \cdot \exp\Big[\beta_Z Z_{jkt} + \beta_D \ln\big(D_{jkt}(\cdot)\big)\Big]. \tag{6}$$

This is the price the government would set in the absence of reference pricing. The function includes product and country fixed effects θ_j and γ_k and three additional variables that try to capture the effect of other potential price-control policies. Z_{jkt} includes a flexible function of the number of other molecules available in the same market and—following Kyle (2006)—an indicator for whether the firm has headquarters in country k. The availability of alternatives should decrease the additional welfare generated by a drug, increase competitive pressure on prices, and provide additional benchmarks to governments that use internal reference pricing (the practice to set prices using the lowest price available within a group of substitutable drugs). Finally, we include log-realized demand for the drug

[25] Patients in European countries usually only pay a fraction of the cost of prescription drugs, so any degree of price elasticity that is picked up in the data is likely driven by the government.

[26] We provide a formal derivation of equation (5) in Section D.1 of the online Appendix.

[27] More classical forms of price competition, such as Bertrand Nash games, are not a good fit for this market and are rarely used in the literature. We present some evidence in this regard and additional evidence justifying various features of our price function in Section C.6 of the online Appendix.

$D_{jkt}(S_t, \xi_{kt})$. This variable should also have a negative sign. Governments make widespread use of price-volume agreements and can use soft nudges to steer patients away from expensive drugs to save money (Carone et al. 2012). Effectively, our approach approximates the price elasticity of revenue—which is the key element affecting dynamic entry decisions—through a combination of fixed demand and an inverse demand function (our pricing equation).

The second component of the price-setting equation is the *reference price*. We do not directly observe the reference price, but we know the reference functions F_{kt}^{ref} and baskets R_{kt}. We assume that governments see prices with a one-period lag and apply ERP before volume adjustments to calculate an estimate of the reference price.[28] The reference pricing function we implement empirically is

$$p_{jkt}^{ref}\big(S_t, D_{jkt}(\cdot)\big) \;=\; F_{kt}^{ref}\Big(\big\{p_{jkt-1}\big(S_t, D_{jkt}(\cdot)\big)\big\}_{k\in(R_{jt}\cap L_{jt-1})}\Big),$$

where L_{jt-1} is the set of countries where product j has launched as of time $t-1$.

We combine these two components by assuming that whenever the governments observe a reference price that is lower than the government price, the equilibrium price is set as a weighted average of the two. We let the weight be country-specific to capture heterogeneity in the application of reference pricing guidelines. The overall price-setting equation is given by

$$(7)\; p_{jkt}\big(S_t, D_{jkt}(\cdot)\big) = \begin{cases} p_{jkt}^{gov}\big(S_t, D_{jkt}(\cdot)\big), & \text{if } p_{jkt}^{ref}(\cdot) \;\geq\; p_{jkt}^{gov}(\cdot); \\ (1-\mu_k)p_{jkt}^{gov}\big(S_t, D_{jkt}(\cdot)\big) + \mu_k p_{jkt}^{ref}\big(S_t, D_{jkt}(\cdot)\big), & \text{if } p_{jkt}^{ref}(\cdot) \;<\; p_{jkt}^{gov}(\cdot); \end{cases}$$

where $\mu_k \in [0,1]$ represents the weight given to the reference price.[29]

Modeling prices using a control function has drawbacks compared to an explicit bargaining model between firm and country. A bargaining model could capture the underlying mechanism of ERP by potentially showing that firms charge artificially higher prices in low-income countries to minimize reference pricing spillovers—even after delaying.[30] This effect is implicitly incorporated in our control function through the country fixed effect.

[28] It seems natural to assume some form of delay in ERP because firms may not report new prices immediately and because governments may not update reference price right away. We are unable to be more precise because we cannot observe when governments update reference prices. We also think that applying volume adjustments at the end is natural since ERP sets the price at the beginning of the year, while volume discounts can only be applied at the end. This sequence implies that governments use initial prices for ERP (i.e., *before* volume adjustments), though we observe the final price (inclusive of eventual volume discounts). Therefore, when calculating reference prices, we exclude the volume component.

[29] We choose to include a kink in the price equation for two reasons. First, our logical interpretation of p_{jkt}^{gov} is that of a "reservation" price for the government. Hence, it would not make sense to assume that the government is willing to go above this price, even if other countries are doing so. Second, the price of drugs almost never increases in countries that use reference pricing, even though reference prices often do (this is particularly common in the first few years after approval, when the exact ordering of launches is not necessarily correlated with price levels). We also generally see that prices are more reactive to changes in reference prices when the reference price is lower. We provide some evidence of these patterns in Section C.6.2 of the online Appendix.

[30] We show some evidence that this is happening in Section C.3 of the online Appendix.

Ultimately, two features of our setting and data convinced us to avoid a fully specified model. First, not having application dates makes it difficult to compute credible threat points for negotiations. The standard way to calculate the firm's threat point is to simulate counterfactual profits under the assumption that no other negotiation fails (see, e.g., Fong and Lee 2013). Unfortunately, we do not know which other negotiations may be ongoing at any given time because we do not know when firms send their pricing applications. Second, we do not have any information about the objective function of the government. We could assume that the government acts as a perfect agent for its citizens, but we do not believe that our demand estimates represent consumer welfare accurately.[31] In Section D.2 of the online Appendix, we show that equation (8) can be derived as the solution to a simplified, static Nash Bargaining game where the government's bargaining power is a function of the reference price.

C. *Entry Dynamics*

The goal of the firm is to maximize profits over the finite life-cycle of its product by choosing the order and timing of entry in each country, conditional on the demand and price functions.

We assume that firms face stochastic shocks in the form of random entry delays. Formally, we model delay shocks as binary Bernoulli random variables ρ_{jkt} with country-specific mean ψ_k, independently distributed across countries, years, and drugs. If $\rho_{jkt} = 1$, then drug j cannot enter country k until period $t+1$, when a new shock is drawn. These shocks help capture variation in delays that cannot be explained through the reference pricing channel.

At the beginning of each period, the firm chooses a set of countries where to send entry applications. We represent this action as a binary vector $\{a_{jkt}\}_{k \in \mathcal{N}_j}$, where $a_{jkt} = 1$ whenever firm j chooses to send an entry application to country k. A *strategy* for firm j is a series of conditional actions that depend on the state variable of the problem, which consists of the firm's launch sequence at the end of the previous period, S_{jt-1}, as well as the launch sequence of its competitors, S_{-jt-1}.[32] Formally, we denote a strategy as a map

$$\mathcal{A}_{jt} : \mathcal{S}_{t-1} \rightarrow \{A_{j\tau}\}_{\tau=t}^{T_j},$$

[31] Our data contains aggregate quarterly sales, with only minimal drug characteristics to aid in adding structure to substitution patterns. We believe individual-level data is necessary to represent patient welfare accurately.

[32] The definition of the state space does not track submitted applications awaiting approval, which indirectly implies that firms can withdraw applications that have been submitted if they incur a delay. Not having data on applications means that we do not observe applications awaiting approval, so we cannot include this variable in the state space. The impact of this assumption on the model should be minimal. First, even though firms probably cannot withdraw applications explicitly, they can almost certainly prolong the process on their end, should they believe that to be beneficial. Second, situations where sending an application has a positive NPV in one period but not in the next, while theoretically possible, are empirically unlikely. Third, this hypothetical ability to withdraw an application does not play a role in the estimation procedure: one side of the inequality relies only on actual launch data, while for the side that uses simulations, we only use strategies where firms keep applying until approved.

where $\mathcal{S}_{t-1}$ is the set of all possible realizations of the launch sequence of all firms at the end of period $t-1$. $\mathcal{A}_{jt}$ generates a set of functions $A_{j\tau}$ for each period τ from t until T_j. Each function $A_{j\tau}$ maps each possible value of $\left(S_{j\tau}, S_{-j\tau}\right)$ to a vector $\left\{a_{jk\tau}\right\}_{k\in\mathcal{N}_j}$.

After the firm sends the applications the vector of binary shocks for the current period $\left\{\rho_{jkt}\right\}$ is realized, determining where entry is feasible. Once the shocks are realized, the state variable updates according to the following rule:

$$s_{jkt} = \begin{cases} t, & \text{if } s_{jkt-1} = 0,\ a_{jkt} = 1, \text{ and } \rho_{jkt} = 0; \\ s_{jkt-1}, & \text{otherwise.} \end{cases} \tag{9}$$

Finally, governments set prices, products are sold, and profits are realized.

The firm's value function at time t is given by

$$V_t\left(S_{jt-1}, S_{-j,t-1}\right) = \max_{\mathcal{A}_{jt}=\left\{A_{j\tau}\left(S_{\tau-1}\right)\right\}_{\tau=t}^{T}} \sum_{S_j} \left\{ \sum_{S_{-j}} \left[\sum_{\tau=t}^{T} \beta^{\tau-t} R_\tau\left(S_{j\tau}, S_{-j\tau}\right) \right] \cdot \Pr\left(S_{-j} \middle| S_{-j,t-1}, \mathcal{A}_{-jt}\right) \right\} \cdot \Pr\left(S_j \middle| S_{jt-1}, \mathcal{A}_{jt}\right),$$

where β is the discount factor; $R_\tau\left(S_{j\tau}, S_{-j\tau}\right)$ is the expected period revenue of the firm for a given realization of the all launch sequences; and $\Pr\left(S_j \middle| S_{jt-1}, \mathcal{A}_{jt}\right)$ and $\Pr\left(S_{-j} \middle| S_{-j,t-1}, \mathcal{A}_{-jt}\right)$ are the probabilities of S_j and S_{-j} conditional on S_{jt-1} and S_{jt-1}, for given strategies $\mathcal{A}_{jt}$ and $\mathcal{A}_{-jt}$ of the firm and its competitors.

The expected period revenue is defined as

$$R_\tau\left(S_{j\tau}, S_{-j\tau}\right) = E\left[\sum_{k\in L_{j\tau}} p_{jk\tau}\left(S_\tau, D_{jk\tau}(\cdot)\right) D_{jk\tau}(\cdot)\right]. \tag{11}$$

The expectation is taken over the possible realizations of the stochastic error $\xi_{k\tau}$ in the demand system.

Discussion of Assumptions.—Before proceeding to the estimation, we note two crucial simplifying assumptions of this model. The first simplifying assumption is that firms operate as single agents. In other words, while the actions of other firms can affect expected revenue (by stealing market share and potentially affecting equilibrium prices), we assume that firms do not anticipate competitors' reactions to strategy deviations.[33] This assumption, while undesirable, is necessary to construct the moment inequalities we use in the estimation, which compare the payoff of observed firm behavior to the predicted payoff of off-equilibrium deviations. When firms compete strategically in a dynamic environment, the only way to predict

[33] We note that this implies some degree of internal inconsistency in the model since firms take the entry sequences of competitors into account but do not realize that competitors could be doing the same.

other firms' reactions to off-equilibrium deviations is to use the model as a guide. However, this requires a full solution—something we do not have.[34] The second simplifying assumption is that the marginal cost of production and the fixed cost of entry in each country is zero, effectively equating revenue and profit. While there are costs involved in the production and distribution of drugs, estimating marginal costs in markets where prices are constrained by regulation is complicated, and we think this is a reasonable approximation for two reasons.[35] First, brand drugs enjoy significant markups over production costs. Second, most fixed costs are sunk by the time drugs receive marketing approval.[36]

We discuss the implications of these assumptions on estimation in Section VB, when we present our moment inequality approach.

IV. Estimation of Drug Demand and Prices

In this section we describe our statistical model and estimation procedure for drug demand and prices. Before presenting our results, we outline the assumptions on the error term in each estimating equation and discuss the conditions under which identification is valid. We conclude by showing that our estimated demand and price primitives justify delays from a revenue perspective.

A. *Demand Estimation*

We estimate demand from equation (5). All variables come from IMS data, except for age, which is calculated using the approval date from the European regulatory authorities. To measure market size, we borrow a map from the four-digit Anatomical Therapeutic Classification (ATC4, available in the IMS data) to GBD indications from Costinot et al. (2019) and use it to calculate the number of patients that might use drugs in a given therapeutic class.[37] We then use the number of patients to estimate of market size in terms of standard units, which we use to construct market shares from data on sales volumes.[38]

[34] An alternative solution to this problem could be to assume that competitors will react to off-equilibrium behavior by taking whatever action minimizes the profits of the firm that deviated. It is unclear whether this strategy is feasible in our setting. The likely profit-minimizing strategy is to immediately send applications to all countries whenever off-equilibrium play is detected. However, this response is almost certainly unreasonable and would make the moment inequality approach much less effective. More realistic, less punishing strategies would require additional work to show that they constitute a lower bound on expected profits of the deviation. We leave a more thorough study of this alternative to future work.

[35] Dubois and Lasio (2018) show how to estimate profit margins with regulation constraints using anti-ulcerant drugs in France as an empirical case study. In theory, we could use their methodology to estimate costs. However, this would mean separately replicating their work for 25 countries and about 100 therapeutic classes. Proper identification would also require a policy change, which may not exist in all countries we need to include in the analysis.

[36] We are not aware of papers that estimate fixed costs of entry for pharmaceutical products in Europe, but several industry insiders have confirmed to us in conversation that these costs are low.

[37] We thank the authors of the paper for sharing the map with us ahead of publication.

[38] For details on the construction of the market size variable, see Section A.3 of the online Appendix.

Identification of Demand System Parameters.—Two potential identification issues arise. The first is that $\ln\left(\frac{MS_{jkt}}{MS_{mkt}}\right)$ (i.e., the within-molecule market share of product i) is correlated with the error term ξ_{jkt}, so we need instruments to recover a consistent value for σ. We use three instruments. The first one is the total number of other firms that are selling the same product. The idea is that in a logit model, the within-molecule share will be mechanically related to the number of alternative options. The second one is years since the patent on molecule m expired. This instrument exploits the gradual shift of market shares to generic manufacturers after loss of exclusivity. The third instrument is the average within-molecule market share of parallel traded products for other molecules in the same country. This instrument captures the average propensity of a government to nudge patients toward cheaper parallel traded products.

The second potential identification issue is that firms might be able to observe ξ_{jkt} before entry, leading to a classic selection problem common to many settings: countries where entry is recorded would have unobservably high values of ξ_{jkt}, leading to a biased estimator. Allowing for drug-country-specific preferences helps attenuate these concerns, but our model may still be misspecified if firms have more information about year-to-year fluctuations in demand. In practice, however, we expect selection to be a second-order concern in this case. Firms never exit voluntarily, so we do not need to worry about exit selection. The remaining concern is entry selection: demand in years prior to entry could be unobservably low, causing firms to wait before entering. This effect would generate an upward bias in our drug-country coefficient. Section VB argues that this effect leads to a more conservative estimate of ERP's impact on strategic delays.

B. *Price and External Reference Pricing Parameters*

Since our prices are yearly averages, our data almost certainly contain some degree of measurement error. We include a measurement error term η_{jkt} that is independent and identically distributed across countries, drugs, and years but do not include any other source of error for simplicity. Since our price function is multiplicative, we also assume that η_{jkt} is multiplicative (i.e., additive in logs). Denoting p_{jkt} as the model-predicted price, and p^{o}_{jkt} as the observed price, the estimation equation becomes

$$\ln\left(p^{o}_{jkt}\right) = \begin{cases} \ln\left(p^{gov}_{jkt}(\cdot)\right) + \eta_{jkt}, & \text{if } p^{ref}_{jkt}(\cdot) \geq p^{gov}_{jkt}(\cdot); \\ \ln\left((1-\mu_k)p^{gov}_{jkt}(\cdot) + \mu_k p^{ref}_{jkt}(\cdot)\right) + \eta_{jkt}, & \text{if } p^{ref}_{jkt}(\cdot) < p^{gov}_{jkt}(\cdot). \end{cases} \tag{12}$$

Our estimation routine searches the vector of parameters that minimizes the difference between the model prediction and the data. To improve the speed and

efficiency of the procedure, we match log differences in price, which do not depend on the product fixed effect θ_j.[39] The estimating equation in differences is

$$\ln\left(\frac{p^o_{jkt}}{p^o_{jk't+1}}\right) = \ln\left(\frac{p_{jkt}(\cdot)}{p_{jk't+1}(\cdot)}\right) + \eta_{jkt} - \eta_{jk't+1}, \tag{13}$$

and our routine minimizes the sum of squares of the two error terms, subject to the constraint that μ_k lies on the unit interval.[40]

$$O(\gamma_k, \mu_k, \beta_Z, \beta_D) = \sum_{j,k,k',t}\left[\ln\left(\frac{p_{jkt}(\cdot)}{p_{jk't+1}(\cdot)}\right) - \ln\left(\frac{p^o_{jkt}}{p^o_{jk't+1}}\right)\right]^2. \tag{14}$$

Identification of the Pricing Equation.—The main threat to identification is the possibility that the price shocks η_{jkt} might be correlated across countries.[41] For example, a negative cost shock affecting all of Europe might result in lower prices everywhere, which our model could erroneously interpret as a consequence of reference pricing. To minimize this possibility, we calculate reference prices using *predicted* prices instead of observed ones.[42] This approach ensures that μ_k is identified through co-movements between observed prices and predicted reference prices, limiting the chance that common shocks could generate spurious correlations between prices of different countries.

Variation in predicted reference prices comes from two primary sources: (i) changes in reference functions and (ii) new launches. Variation in reference functions almost always comes from the entry of new countries in the EU or the Eurozone, which is exogenous.[43] Variation in launch timing is at least partly exogenous due to the randomness of the application process. Under the assumptions of our model, idiosyncratic delays generated by the application process are orthogonal to prices and strategic considerations.[44]

Variation in entry sequences also helps us identify other components of the pricing equation. Nearly all countries in our data show up as the first entry in the launch sequence of at least a few drugs. Since governments cannot observe a reference

[39] We prove this in the Appendix. Intuitively, our estimator is conceptually similar to a first-difference estimator. To avoid differencing out the country fixed effect however, we compare prices of different countries—instead of the price of the same country in consecutive years. We do not believe there is an advantage to selecting a specific sequence of country differences, so we determine the sequence randomly. We also include $\frac{p_{jk2012}}{p_{jk'2002}}$ as a moment. By doing so we retain all but one observation per drug (the remaining observation is then used to pin down θ_j).

[40] A negative value of μ_k does not make sense; a value of μ_k greater than one—while theoretically possible—raises the possibility of negative prices in counterfactual predictions, which is undesirable.

[41] Demand shocks that are correlated across countries may also be a problem, though our pricing model excludes the volume component in calculating the reference price. Hence, demand shocks do not affect the reference price in our estimation (though they do affect the price).

[42] To construct reference prices we use a simple loop that combines equations (6), (7), and (8). The loop proceeds as follows. In period 1 there are no reference prices, and prices are given by the reservation prices in equation (6). In period 2 reference prices can be built from available reservation prices using equation (7). Combining reference and reservation prices using equation (8) yields predicted observed prices in period 2. In period 3 one can build reference prices from prices observed in period 2 using equation (7), and then construct observed prices from equation (8). Repeat this process until the end of a drug's life cycle. See online Appendix D.3 for more details.

[43] We provide a full list of reference function changes in Section B.2 of the online Appendix.

[44] To be sure, not all delays are random. Strategic delays are clearly linked to prices. These include delays that may be part of government negotiations but arise because of disagreements over price.

TABLE 2—PRICE ESTIMATION RESULTS

Country	$\ln(\gamma_k)$		μ_k	
Austria	−0.105	(0.021)	0.330	(0.207)
Belgium	−0.123	(0.021)	0.195	(0.240)
Bulgaria	−0.203	(0.106)	1.000	(0.003)
Denmark[b]	−0.084	(0.016)	0	
Estonia	−0.185	(0.060)	1.000	(0.150)
Finland	−0.135	(0.021)	0.262	(0.322)
France	−0.095	(0.019)	0.000	(0.267)
Germany[a,b]	0		0	
Greece	−0.094	(0.042)	1.000	(0.052)
Hungary	−0.263	(0.080)	0.984	(0.216)
Ireland	−0.078	(0.068)	0.592	(0.352)
Italy	−0.177	(0.036)	1.000	(0.151)
Latvia	−0.244	(0.043)	0.875	(0.273)
Lithuania	−0.244	(0.048)	1.000	(0.100)
Luxembourg[c]	−0.239	(0.024)	0	
Netherlands	−0.207	(0.021)	0.001	(0.182)
Norway	−0.169	(0.020)	1.000	(0.310)
Poland	−0.061	(0.089)	0.914	(0.145)
Portugal	−0.194	(0.041)	0.999	(0.271)
Romania	−0.281	(0.148)	1.000	(0.081)
Slovenia	−0.247	(0.020)	0.999	(0.101)
Spain	−0.160	(0.023)	1.000	(0.204)
Sweden[b]	−0.108	(0.017)	0	
Switzerland	−0.004	(0.015)	0.000	(0.010)
United Kingdom[b]	−0.193	(0.016)	0	
Controls				
log quantity sold	−0.025		(0.003)	
Home firm indicator	0.051		(0.019)	
At least 1 other molecule in class	−0.009		(0.044)	
At least 2 other molecules in class	0.005		(0.026)	
At least 5 other molecules in class	−0.024		(0.023)	
At least 10 other molecules in class	−0.005		(0.015)	

Notes:
[a] The price level is normalized to Germany's.
[b] Denmark, Germany, Sweden, and the United Kingdom do not use ERP during 2002–2012.
[c] Luxembourg references the price of the country of origin of the drug. Since we do not know country of origin, we assume that μ_j equals zero.

price at the beginning of the launch sequence, this variation helps us identify the components of the government price function.

To further check our model, we conduct a placebo test where we assign hypothetical ERP functions to Denmark, Germany, Sweden, and the United Kingdom—countries that do not use ERP. If cross-country correlation in η_{jkt} were driving our estimates, we would probably pick up some spurious effects of ERP in these countries as well. Reassuringly, we find that μ_k coefficients for all four countries are almost exactly zero.[45]

Results.—We report results for the vector of parameters $\left(\hat{\gamma}_k, \hat{\mu}_k, \hat{\beta}_Z, \hat{\beta}_D\right)$ in Table 2. Standard errors—calculated using nonparametric bootstrap with sampling at the drug level—are in parentheses. In the first column, we show the coefficients

[45] We report the results of this test in Section C.5 of the online Appendix.

for $\ln(\gamma_k)$, which can be interpreted as a percentage difference in price relative to a benchmark (in this case, the omitted coefficient used as the benchmark is the one for Germany).

The point estimates roughly match our intuition: lower-income countries tend to pay lower prices.[46] We note, however, that one should not interpret country fixed effects as capturing the absolute price level of their respective country but instead reflect the equilibrium played in the data, which includes the externality of ERP. For example, we estimate a relatively low coefficient for Norway—one of the wealthiest countries in Europe on a per-capita basis. Norway is only referenced by Finland and therefore generates negligible ERP-related spillovers relative to, say, France, which is referenced by fourteen countries. In the absence of ERP, we expect that the price levels of all countries would adjust to reflect a new equilibrium. Our model cannot predict what the adjustment would be, but this does not affect our counterfactual.

The second column shows estimates for μ_k, which measures how strictly each country adheres to its own ERP guidelines. We observe significant heterogeneity across countries in this respect. Thirteen countries have coefficients above 0.85, meaning that they closely follow reference pricing guidelines. However, five countries have coefficients below one-third, which suggests that they either do not follow their guidelines closely or apply them only to selected drugs. In particular, France, the Netherlands, and Switzerland do not appear to use reference pricing at all, with coefficients estimated to be almost exactly zero.[47]

The coefficients on the control variables generally behave as expected. Higher quantity is associated with lower prices, and prices tend to be approximately 5 percent higher in countries where the firm has headquarters. In addition, having a higher number of competitors in the same class is associated with slightly lower prices, though the relationship appears to be nonlinear and noisy, probably because prices are not determined through a competitive process.

C. *Simulation-Based Evidence of Optimality of Delays*

Our results suggest that ERP affects equilibrium prices, but is the implied externality strong enough to generate delays? To answer this question, we simulate expected firm revenue from various entry sequences and compare it to the expected revenue of a naive entry sequence where products are launched immediately in every country. The advantage of these simulations is that they rely solely on our static price and demand estimates. Hence, they do not depend on parametric assumptions about the distribution of idiosyncratic delays.[48] The downside is that they cannot identify the extent to which firms engage in strategic delays.

[46] There are a couple of exceptions. For example, Poland has a higher coefficient than many other countries with higher income. However, it uses the minimum price in Europe as its reference, and its coefficient on μ_k is close to one. In this case, the government may be willing to grant higher prices, knowing that reference rules will bring them down quickly.

[47] In the case of France, this is consistent with the way reference pricing is implemented. As Dubois and Lasio (2018) point out, ERP in France does not apply to all brand drugs but only to a small group of innovative products.

[48] The expected value of revenue is taken over possible realizations of the error term in the demand system. The idiosyncratic delay component does not factor in these simulations, which also hold fixed the entry sequences of all other products in the market.

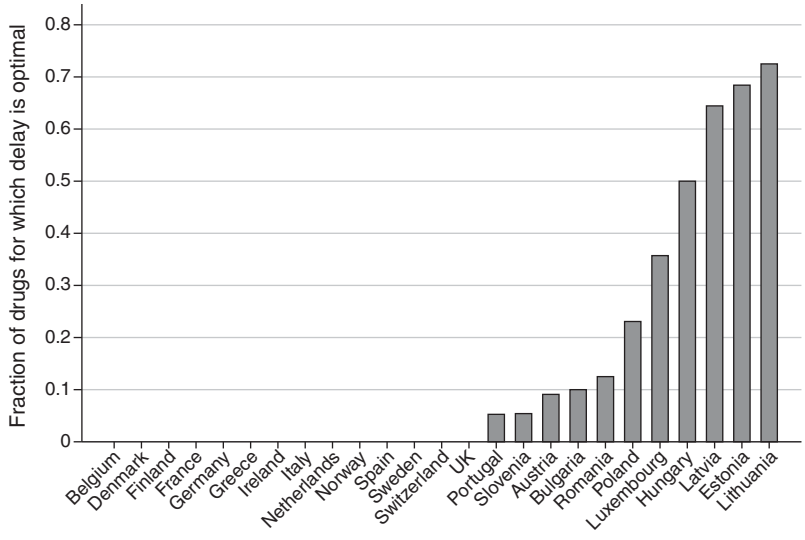

FIGURE 4. OPTIMALITY OF STRATEGIC DELAYS BY COUNTRY

For each of the 87 drugs in our dynamic product sample, we simulate entry sequences that consist of immediately launching in all but one country, country k (we test all possible delays for the country where entry does not occur immediately). We then calculate the fraction of drugs for which delaying in country k is optimal. Figure 4 plots the results. We find that delays are often optimal in countries with lower income but rarely elsewhere.[49]

The delay patterns in Figure 4 are remarkably consistent with the entry patterns from Figure 2. This result confirms the intuition that firm behavior follows the incentives laid out by the regulatory environment and supports the idea that our model captures the relevant features of this market.

We also test several other potential strategies, including (i) delays in all but *two* countries, (ii) symmetric delays in all Eastern European countries, and (iii) head-to-head delays where the product is launched in at most two countries. All simulations provide broad support to the idea that delays in Eastern Europe are often optimal, but delays in Western Europe rarely are. We report detailed output from these additional simulations in Section C.7 of the online Appendix.

V. Dynamic Analysis

With estimates of demand and price in hand, we turn our attention to the parameters governing idiosyncratic delays. The traditional way to estimate these parameters is to predict launch patterns as a function of the delay parameters ψ_k and

[49] Only two higher-income countries would experience delays according to this simulation: Austria and Luxembourg. Luxembourg is a small market, so it is not surprising that for some drugs it would be optimal to exclude it. The model also predicts a delay for one drug in Austria. In this case, the drug in question was indeed launched after a long delay and earned low revenue. Even though Austria tends to have high price levels relative to most other countries, it can affect the prices of countries with higher price levels (e.g., Ireland).

then estimate them by matching predicted patterns to the data. Unfortunately, we cannot adopt this strategy because our model does not have a closed-form solution, and the number of strategies available makes numerical approaches unfeasible.

This problem is not unprecedented in the literature, and previous papers have often solved it by using revealed-preference moment inequalities. This method does not require a full solution to the model but instead relies on the assumption that the strategy observed in the data represents the solution to the maximization problem modeled by the econometrician. Under these conditions, it is relatively straightforward to derive restrictions that compare the firm's expected payoff from an arbitrary strategy to the expected payoff of the strategy chosen in the data.

This approach also does not work in our case because we do not observe the firm's strategy. What we observe is a launch sequence, which is the result of a strategic choice of the firm, plus an idiosyncratic component (represented by the delay parameter ψ_k). In other words, when we observe a delay, we do not know whether it was a strategic choice of the firm or the result of a random shock.

Our solution is to extend the revealed-preference inequality framework to allow for unobserved strategies. The intuition behind our extension is simple. When the strategy is observed, one can recover the expected payoff of the observed strategy (call it $\mathcal{A}^o$) for any arbitrary value ψ' of the unknown parameter (i.e., $E\left[\tilde{V}\left(\mathcal{A}^o, \psi', \cdot\right)\right]$, where $\tilde{V}(\cdot)$ represents the expected payoff conditional on playing a specific strategy). The traditional revealed preference moment inequalities compare $E\left[\tilde{V}\left(\mathcal{A}^o, \psi', \cdot\right)\right]$ to $E\left[\tilde{V}\left(\mathcal{A}', \psi', \cdot\right)\right]$ for alternative strategies $\mathcal{A}'$. When the strategy is unobserved, $E\left[\tilde{V}\left(\mathcal{A}^o, \psi', \cdot\right)\right]$ cannot be recovered. However, we show that under some additional assumptions we can use revenue data to recover the expected payoff of the firm's strategy for the true value of the parameter ψ_0 (i.e., $E\left[\tilde{V}\left(\mathcal{A}^o, \psi_0, \cdot\right)\right]$).[50] We then use $E\left[\tilde{V}\left(\mathcal{A}^o, \psi_0, \cdot\right)\right]$—instead of $E\left[\tilde{V}\left(\mathcal{A}^o, \psi', \cdot\right)\right]$)—to generate restrictions.

Since our inequalities effectively use less information, our extension comes with two drawbacks. First, it requires stricter assumptions on error terms. Any structural error term in the revenue function must be recovered during static estimation. In our setting, this rules out drug-country-year specific revenue shocks or stochastic fixed costs of entry known to the firm. Second, our inequalities generally provide a one-directional bound—which in our case is the lower bound.[51] To complete the identified set, we construct an upper bound using data on approval and entry dates, noting that firms can only apply after receiving approval. The average probability of an idiosyncratic delay is then bounded above by the average probability of overall delay.

We describe our estimation approach starting from the upper bound estimation, which is more straightforward, and follow that with the derivation of the moment inequalities based on revenue data.

[50] We denote the strategy played by the firm as $\mathcal{A}^o$ in both expressions even though the strategy is technically unobserved in the second case. This slight abuse of notation nonetheless underlines that the strategy is the same, regardless of whether it was observed or not.

[51] This property holds under broad assumptions. We provide a rigorous proof in Section D.5 of the online Appendix and discuss strategies to obtain two-directional bounds in Section D.6. Unfortunately, these strategies do not yield meaningful bounds in our empirical setting.

A. *Inequalities Based on Entry Data*

To recover an upper bound on the delay parameter ψ, we use entry and approval data. Our distributional assumption on the random delay shocks is that the probability of an application for entry in country k being delayed in any given period is ψ_k. Let $\left(1-\bar{\psi}_k\right)$ be the overall probability that product j will enter country k in any given year. $\left(1-\bar{\psi}_k\right)$ is the combination of the probability that the firm will apply times the probability of the application being accepted. Hence, for all j, $\bar{\psi}_k \geq \psi_k$.

To estimate $\bar{\psi}_k$, we simply calculate the probability of a delay by using data on approval dates and launch dates. Suppose that a product approved in year 0 enters France in year 2. Then the expected probability of entry is one-third: the product had three opportunities to enter—in years 0, 1, and 2—and registered one success—in year 2.

B. *Inequalities Based on Revenue Data*

In this section we derive our novel moment inequality estimator. We start from a general statement of the problem, which maps to the standard presentation of revealed preference inequality problems in the literature (see, e.g., Pakes 2010; Pakes et al. 2015).

To begin, let $\mathcal{I}_{jt}$ be the information set of firm j at time t, and let $\mathcal{A}^{\star}_{jt}\left(\mathcal{I}_{jt}\right)$ denote the optimal strategy of firm j starting in period t, conditional on its information set. Under standard revealed-preference assumptions, $\mathcal{A}^{\star}_{jt}\left(\mathcal{I}_{jt}\right)$ satisfies

$$\sup_{\mathcal{A}'_{jt}} E\left[\tilde{V}_t\left(\mathcal{A}'_{jt}, Y_j, \psi_0\right)\middle|\mathcal{I}_{jt}\right] \leq E\left[\tilde{V}_t\left(\mathcal{A}^{\star}_{jt}\left(\mathcal{I}_{jt}\right), Y_j, \psi_0\right)\middle|\mathcal{I}_{jt}\right], \tag{15}$$

where ψ_0 is the true value of the parameter, Y_j is any variable that affects payoffs other than the decision variable, and expectations are taken over likely values of Y_j.[52]

To transform 15 into an empirical condition for estimation, we need two steps. First, we must specify a model for how Y_j changes as a function of $\mathcal{A}_{jt}$. In our setting, Y_j contains the entry sequence of product j and all other products in the market, plus the variables included in the estimation of demand and prices, which we denote collectively as Z_j. Hence, $Y_j = \left\{S_j, S_{-j}, X_j\right\}$. The firm's strategy $\mathcal{A}'_{jt}$ only affects the entry sequence S_j and does so—together with the idiosyncratic delay shocks—according to the rule specified in equation (9). All other interactions are ruled out. $\mathcal{A}'_{jt}$ does not affect S_{-j} because we have assumed that firms operate as single agents. Moreover, other payoff-relevant random variables in X_j (such as age, the home-country indicator, and fixed-effects) are specified to be independent of the entry sequence, and therefore are unaffected by the firm's strategy.

As a second step, we must specify a measurement function for $\tilde{V}_t(\cdot)$. In our case, under the assumption of no entry or production costs, the payoff of the firm is simply

[52] This equation is equivalent to C1 in Pakes (2010) for a single-agent model. Also notice that $\tilde{V}_t\left(\mathcal{A}^{\star}_{jt}\left(\mathcal{I}_{jt}\right), Y_j, \psi_0\right)$ is equivalent to $V_t\left(S_{jt-1}, S_{-j,t-1}\right)$ as defined in equation (10).

the revenue earned from selling the product across various European countries. To measure it, we use the demand and price functions recovered in the previous estimation stage, which represent unbiased estimates of the firm's expectation under the assumption that there are no structural errors after controlling for all variables in X_j. This is equivalent to assuming that the econometrician observes the firm's information set $\mathcal{I}_{jt}$.

Using demand and price estimates, the revenue that firm j earns starting in period t, given entry sequences S_{jt} and S_{-jt}, is given by

$$(16)\; R_t\left(S_j, S_{-j}, X_j\right) \;=\; E\left[\sum_{\tau=t}^{T} \beta^{\tau-t} \sum_{k \in L_{j\tau}} p_{jkt}\left(S_j, S_{-j}, X_j, D_{jkt}\left(S_j, S_{-j}, X_j\right)\right) D_{jkt}\left(S_j, S_{-j}, X_j\right)\right].$$

Since we assumed that there are no structural errors, the only difference between $R_t\left(S_j, S_{-j}, X_j\right)$ and $\tilde{V}_t\left(\mathcal{A}'_{jt}, Y_j, \psi_0\right)$ comes from the country-year-specific vector of shocks ξ_{kt}, which is unknown to the firm prior to entry.[53] Under these assumptions, given an entry sequence $\left(S_j, S_{-j}\right)$, we can simulate a consistent estimate for $R_t\left(S_j, S_{-j}, X_j\right)$ by following a three-step procedure:

(1) Simulate N matrices of demand shocks $\xi^n \;=\; \left\{\xi^n_{jkt}\right\}$, where ξ_{jkt} is the shock for firm j, country k, and period t.

(2) For each matrix ξ^n, calculate

$$R^n_t\left(S_j, S_{-j}, X_j, \xi^n\right) \;=\; \sum_{\tau=t}^{T} \beta^{\tau-t} \sum_{k \in L_{j\tau}} p_{jkt}\left(S_j, S_{-j}, X_j, D_{jkt}\left(S_j, S_{-j}, X_j, \xi^n\right)\right) D_{jkt}\left(S_j, S_{-j}, X_j, \xi^n\right).$$

(3) Estimate $\tilde{R}_t\left(S_j, S_{-j}, X_j\right)$ as

$$\tilde{R}_t\left(S_j, S_{-j}, X_j\right) \;=\; \frac{1}{N} \sum_{n=1}^{N} R^n_t\left(S_j, S_{-j}, X_j, \xi^n\right).$$

This procedure returns consistent estimates for the expected revenue of a given entry sequence. If we knew the application strategy of each firm, we could use the procedure to calculate its expected value as a weighted average of the revenue of all possible realizations of the entry sequence given the strategy and an arbitrary value of the delay parameter vector ψ. While we do not observe the application strategy however, the observed entry sequence $\left(S^o_j, S^o_{-j}\right)$ is the result of optimal play (i.e., the optimal strategy $\mathcal{A}^{\star}_{jt}\left(\mathcal{I}_{jt}\right)$) plus the idiosyncratic delay shock. As a result, $\tilde{R}_t\left(S^o_j, S^o_{-j}, X_j\right)$ is a draw from the distribution of the revenue of firm j. This suggests that aggregating $\tilde{R}_t\left(S^o_j, S^o_{-j}, X_j\right)$ across firms will yield a consistent estimate for the average expected revenue across firms.

[53] The difference between estimated prices and observed prices is assumed to be measurement error, which means that our estimated prices are the true prices.

THEOREM 1: *For any* $\epsilon > 0$, *we can find* M' *such that*

$$\frac{1}{M}\left|\sum_{j=1}^{M}\left(\tilde{R}_t\left(S_j^o, S_{-j}^o, X_j\right) - E\left[\tilde{V}_t\left(\mathcal{A}_{jt}^{\star}\left(\mathcal{I}_{jt}\right), Y_j, \psi_0\right)\middle|\mathcal{I}_{jt}\right]\right)\right| < \epsilon,$$

for all $M > M'$.

Theorem 1 indicates that, even without knowing $\mathcal{A}_{jt}(\psi_0)$, we can recover a consistent estimate of $\frac{1}{M}\sum_{l=1}^{M} E\left[\tilde{V}_t\left(\mathcal{A}_{jt}^{\star}\left(\mathcal{I}_{jt}\right), Y_j, \psi_0\right)\middle|\mathcal{I}_{jt}\right]$ (i.e., the average expected payoff across firms). A formal proof is available in Appendix B. The result relies on a generalized law of large numbers that requires draws from independent, but not necessarily identical distributions.[54] In our empirical application the distributions differ because they depend both on the initial state S_{jt-1}^o and on the optimal strategy $\mathcal{A}_{jt}^{\star}(\mathcal{I}_{jt})$, both of which will differ across firms.

The second term is the expected revenue from playing an arbitrary strategy $\mathcal{A}_{jt}'$. To compute this term, we add an additional simulation layer to the previous procedure. For a given guess ψ' of the parameter vector, and for each firm j, we draw $\{\omega_j^r\}_{r=1}^{nsim}$, where $\omega_j^r = \{\omega_{jkt}^r\}$ is a matrix of shocks for each period t, and country k. Combining the shocks and $\mathcal{A}_{jt}'$, we calculate the simulated entry paths $S_j(\omega_j^r, \psi')$ and denote

$$(17)\quad R_t^{sim}\left(\mathcal{A}_{jt}', S_{jt-1}^o, S_{-j,t-1}^o, X_j; \psi'\right) = \frac{1}{nsim}\sum_{r=1}^{nsim}\left[\sum_{S}\tilde{R}_t\left(S_j, S_{-j}, X_j\right)\cdot \Pr\left(S\middle|S_{t-1}, \mathcal{A}_{jt}'\right)\right],$$

where *nsim* is the number of simulations.[55]

THEOREM 2: *For any* $\epsilon > 0$, *we can find* M' *such that*

$$\frac{1}{M}\left|\sum_{j=1}^{M}\left(R_t^{sim}\left(\mathcal{A}_{jt}', S_{jt-1}^o, S_{-j,t-1}^o, X_j; \psi'\right) - E\left[\tilde{V}_t\left(\mathcal{A}_{jt}', Y_j, \psi'\right)\middle|\mathcal{I}_{jt}\right]\right)\right| < \epsilon.$$

Theorem 2 indicates that through simulation we can recover a consistent estimate of $\frac{1}{M}\sum_{j=1}^{M} E\left[\tilde{V}_t\left(\mathcal{A}_{jt}', Y_j, \psi'\right)\middle|\mathcal{I}_{jt}\right]$ (i.e., the average expected revenue of playing $\mathcal{A}_{jt}$ across firms for a given value ψ' of the unknown parameter). The formal proof is similar to the Proof of Theorem 1 and can be also found in Appendix C. Intuitively, the result depends on the same independent draws assumption on S_{-j} as Theorem 1, as well as the assumption that we can recover a consistent estimate of the revenue function.

[54] Our proof assumes independence, but this is a sufficient, and not a necessary condition. More general versions of the law of large numbers can accommodate correlation across draws.

[55] Please refer to Section D.4 of the online Appendix for a detailed description of the simulation procedure.

Theorems 1 and 2 prove that we can build aggregated moment inequalities. Summing up equation (15) across firms $j \in \{1,2,\ldots,M\}$, we get

$$\frac{1}{M}\sum_{j=1}^{M}\left(E\left[\tilde{V}_t\left(\mathcal{A}^{\star}_{jt}(\mathcal{I}_{jt}),Y_j,\psi_0\right)\middle|\mathcal{I}_{jt}\right]-E\left[\tilde{V}_t\left(\mathcal{A}'_{jt},Y_j,\psi_0\right)\middle|\mathcal{I}_{jt}\right]\right)\geq 0. \tag{18}$$

for all $\mathcal{A}'_{jt}$. Combining the results of the two theorems, the sample analog of equation (18) is

$$\frac{1}{M}\sum_{j=1}^{M}\left[\tilde{R}_t\left(S^o_j,S^o_{-j},X_j\right)-R^{sim}_t\left(\mathcal{A}'_{jt},S^o_{jt-1},S^o_{-j,t-1},X_j;\psi'\right)\right]\geq 0. \tag{19}$$

Notice that equation (19) is an *unconditional* moment inequality, which averages across all observations to eliminate the source of error (in our case, application delays). In theory, one could construct a series of additional *conditional* moment inequalities, which would be averages across firms with specific observable characteristics, as long as these characteristics are not correlated with application delays. In practice, our limited sample size (84 firm-therapeutic class combinations) means that splitting observations in smaller groups to create conditional inequalities would run the risk of invalidating the law of large number argument that is necessary for the inequalities to hold under Theorem 2.

Any strategy $\mathcal{A}'_{jt}$ can then be used to create an inequality restriction $IR\left(\mathcal{A}'_{jt},\psi'\right)$ on a guess ψ' of the unknown parameter

$$IR\left(\mathcal{A}'_{jt},\psi'\right)=\min\left\{0,\frac{1}{M}\sum_{j=1}^{M}\tilde{R}_t\left(S^o_j,S^o_{-j},X_j\right)-R^{sim}_t\left(\mathcal{A}'_{jt},S^o_{jt-1},S^o_{-j,t-1},X_j;\psi'\right)\right\}, \tag{20}$$

with a corresponding identified set Ψ^I that satisfies

$$\Psi^I=\left\{\psi':\sum_{\mathcal{A}'_t}IR\left(\mathcal{A}'_{jt},\psi'\right)=0\right\}. \tag{21}$$

Impact of Assumptions on Empirical Estimates.—Before we look at the empirical implementation of our methodology, we discuss how various model and estimation assumptions might affect our results.

The most consequential model assumption is that firms behave as if they were single agents. This restriction has two effects. First, it rules out more sophisticated equilibria where firms optimally split markets to avoid competition. In practice, two aspects of the pharmaceutical market imply that separating equilibria are unlikely to arise. The first aspect is that branded drugs effectively have a finite life-cycle due to generic entry following patent expiration. At the end of the life-cycle, any separating equilibrium will unravel because the ERP externality stops mattering, and all firms find it optimal to launch in all countries.[56] The second aspect is that arrival of products is staggered, sometimes by several years. As a result, most firms that arrive

[56] ERP becomes increasingly less relevant as drugs approach the end of their patent life because countries need some time to observe new prices and incorporate them in their ERP functions.

on the market face established competitors whose products are already available in many countries. The presence of established competitors makes separating equilibria difficult both because (i) coordination is more challenging if one firm has already entered most markets, and (ii) because older products are closer to the end of their life cycle, and to the point when a separating equilibrium would unravel. In the data, we see suggestive evidence confirming this intuition. Drugs enter most countries by the end of their life cycle. Moreover, the order of entry usually matches the order of marketing approval (this rules out separating equilibria because a separating equilibrium requires different firms to be the first to enter in separate sets of countries).[57]

The second effect of the single-agent assumption is that the inability of firms to anticipate future reactions to potential deviations affects our calculation of counterfactual revenues under different strategy profiles. The effect does not have a clear sign. In some situations, we may expect revenue to be biased downward. For example, a strategy that anticipates entry might yield unrealistically high expected revenue if competitors react by also anticipating entry. In other situations, however, the bias could have the opposite sign—for example, if anticipating entry discourages a competitor from entering at all.

In terms of the estimation, the three most important assumptions we make are (i) zero costs of entry or production, (ii) excluding countries where we do not observe entry from the strategy space, and (iii) the absence of a structural error in the demand and price estimation. We believe all three assumptions are conservative and, if violated, would imply longer strategic delays.

Assuming zero costs means equating revenue and profits. If high, fixed entry costs and marginal costs of production are likely approximately constant across countries. As a result, costs make up a greater fraction of profits in countries with lower prices and smaller market sizes. Hence, the incentive to launch in these countries will be lower than what our model predicts. In terms of our estimation approach, removing costs implies that strategies that anticipate entry in small-market countries earn a higher payoff, which translates to a smaller likelihood of predicting strategic delays.

Excluding countries where we do not observe entry from the strategy space also leads to underestimation of delays. If we were to drop this assumption, we would have to replace country-drug interaction fixed effects with separate sets of country and drug fixed effects. That combination would likely overestimate the popularity of drugs in countries where we do not observe entry. As a result, deviations that anticipate entry in these countries would appear profitable in the moment inequality estimation, and any delays would be attributed to idiosyncratic shocks.[58]

[57] Notice that this does not mean that there are no drug-specific delays. For example, the monoclonal antibodies Rituximab, Trastuzumab, and Alemtuzumab were approved in December 1997, November 1998, and August 2001, respectively. They entered in that order in every single European country except Estonia and Lithuania, where Rituximab and Trastuzumab entered in reverse order, but still within six months of each other. However, their respective delays in each country are quite heterogeneous. For example, Rituximab entered Lithuania with a five-year delay, but Trastuzumab and Alemtuzumab entered Lithuania with only a two- and three-year delay, respectively. Similarly, Rituximab and Alemtuzumab entered Slovenia with about a three-year delay, but Trastuzumab entered less than one year after receiving marketing approval.

[58] We also believe that using separate drug and country fixed effects would lead to significant misspecification. In section C.4 of the online Appendix, we show that delays are significantly correlated with the residual from a regression of volume sales on separate drug and country fixed effects. This correlation suggests that firms know which products will be more successful in certain countries and prioritize them in the launch sequence.

Finally, removing sources of structural errors from demand estimation also leads to a more conservative estimate of ERP's impact on strategic delays. Structural error in the form of a known country-year-drug specific shock (say, ν_{jkt}) would lead to a classic selection problem common to many settings: sales prior to entry are unobservably low. The presence of ν_{jkt}, if unaccounted, would imply that we overestimate the payoff of counterfactual strategies that send early applications to countries with lengthy delays. Hence, these strategies will look more attractive in our counterfactual calculations, leading the model to interpret those delays as idiosyncratic rather than strategic.

Empirical Implementation.—In the empirical implementation, we aggregate moments at the firm and therapeutic class level, treating the class-specific revenue of each firm as a separate observation. This aggregation is consistent with our assumption that there are no cross-elasticities between drugs in different therapeutic classes. We also assume that ψ_k takes on the functional form $\psi_k = \bar{\psi}_k \times \psi_g$, where $\bar{\psi}_k$ is the overall probability of delay (calculated using approval and entry data), g indexes a group of countries, and $\psi_g \in [0,1]$ is the parameter to be estimated, which we assume is identical across countries in the same group.[59] We do this because the computational demands of moment inequality estimation grow exponentially with the number of parameters; therefore, the computational cost to estimate bounds for each country k would be high.

We calculate ψ_g separately for countries in Western Europe and countries in Eastern Europe. We use this grouping because, following our simulation results from Section IVC, we believe strategic delay incentives are negligible in Western European countries but significant in Eastern European countries.

We estimate the identified set by building moment conditions based on equation (19) and by checking the range of parameters that satisfies it for a series of possible strategies $\mathcal{A}_{jt}$. Calculating $\tilde{R}_t\left(S_j^o, S_{-j}^o, X_j\right)$ and $R_t^{sim}\left(\mathcal{A}'_{jt}, S_{jt-1}^o, S_{-j,t-1}^o, X_j; \psi'\right)$ does not require observing drugs from their original launch, but it does require observing them until period T_j. Hence, we perform this analysis on the dynamic sample (see Table 1 for details).

Since the strategy space is vast, there are many possible strategies that we can use in our empirical implementation. In practice, however, only strategies that closely resemble the optimal strategy will deliver tight bounds.

We consider three groups of possible strategies. The first group consists of strategies based on the observed launch sequence. To do so, we take the observed entry periods s_{jk}^o and assume firms apply τ periods prior to observed entry, with $\tau \in (0,1,2,3)$. The second group of strategies uses a simple rule: apply in all Western European countries right away, and apply in Eastern Europe with a delay of τ periods, with $\tau \in (0,1,\ldots,T)$. Finally, the third and last group of strategies attempts to approximate the best-performing launch sequences using our simulation exercise

[59] We choose this specific functional form because overall delays are strongly correlated with the average turnaround time for a pricing and reimbursement application—especially for Western European countries (see Figure 3). This approach maintains some heterogeneity across Western European countries, which we think is important given that the probability of delay varies significantly across Western Europe. The adjustment is less important for Eastern Europe, which has more homogeneous delay probability rates.

in Section IVC. Since in that exercise we only test sequences where entry is delayed in at most one country, we approximate the best entry sequence by creating a profile that combines the optimal delay of each country. For example, suppose that after testing single-country delays for drug j in the simulations, we find that revenue increases when delaying entry in countries A, B, and C by τ_A, τ_B, and τ_C periods, respectively. In this case, we would build strategies that try to achieve an entry sequence where entry occurs as quickly as possible in all countries, and with τ_k delay in countries A, B, and C. To do so, the firm would immediately send applications for entry in all countries and then would wait to send applications to countries A, B, and C until some number of periods prior to τ_k. As in the previous cases, we test strategies that send applications τ periods before the targeted entry period, with $\tau \in (0,1,2,3)$.[60]

The strategies that deliver the tightest bound are going to be the ones that, for each drug, maximize revenue in simulations. Unsurprisingly, we consistently find that strategies in this last group are the ones that achieve this effect.

Figure 5 shows our results. The darker area is the identified set, whose two extremes are $\psi_{EE} = 0.69$ and $\psi_{WE} = 0.90$. Using these two pairs, we can construct identified sets by country. The results are shown in Figure 6.

VI. Implications of ERP on Access and Revenue

A. *Delays in the Absence of ERP*

Repealing ERP implies that firms no longer have an incentive to delay entry in any country. The specific pricing rule adopted to replace ERP does not matter, as long as it eliminates externalities across countries (crucially, this also requires imposing limits on other policies that could potentially generate externalities, such as parallel trade).

Because the pricing rule does not matter, we do not need to rely on our pricing model to work out the implications of this counterfactual. This strengthens the external validity of our estimates because our result will not depend on the out-of-sample accuracy of the pricing model. Instead, delays only depend on our estimates of ψ_k.

ψ_k is set-identified, so our estimated counterfactual will yield a range of possible outcomes. Since all our results suggest that the relevant delays are in Eastern Europe, we fix $\psi_{WE} = 1$ to simplify the exposition. Our measure of interest is the average number of country-years of delay per drug; that is, the total number of years of delays across all drugs (e.g., four country-years of delay means a total of four years of delays across the eight Eastern European countries in our data, an average delay of six months in each country).

The identified set for ψ_{EE} implies a delay reduction of up to 8.5 country-years (i.e., approximately one year in each of the Eastern European countries in our

[60] Strategies built this way may miss instances where delays are profitable only when combined. For example, I might find that delaying in Estonia is pointless if I'm launching in Lithuania, and vice versa, but *jointly* delaying in both countries increases revenue. We partially test for these interactions when we simulate delays of country pairs, and we find limited evidence of these interactions (in our testing, we find only four cases where delaying in a pair of countries is optimal, but delaying in either one alone is not).

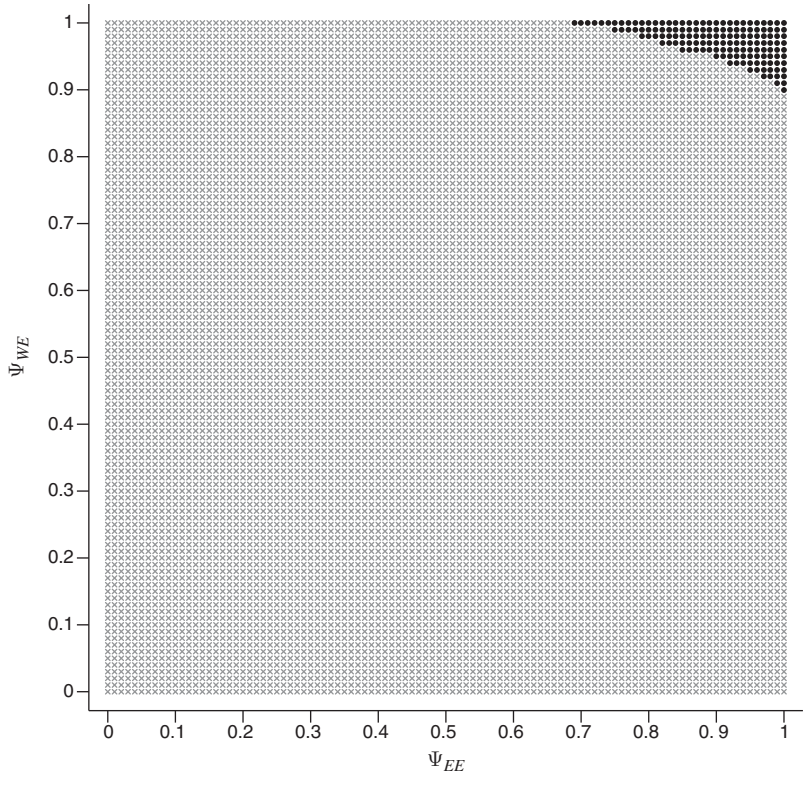

FIGURE 5. IDENTIFIED SET OF ψ_{WE} AND ψ_{EE}

sample). To obtain this number, we simulate delays for all possible values of ψ_{EE} in our identified set. The upper bound of our identified set implies that all observed delays are idiosyncratic. The implied values of ψ_k lead to a total simulated delay of 15.3 country-years, just shy of an average delay of 2 years in each country. At the lower bound, total simulated delays are 6.8 country-years, a reduction of 8.5 country-years.

An implication of this simulation exercise is that the total amount of idiosyncratic delays is larger than the amount of strategic delays, as our results rule out the presence of significant strategic delays in Western Europe. Under the most conservative scenario, idiosyncratic delays lead to approximately 11.6 years of delay across 17 Western European countries (about eight months in each country, on average).

B. *Revenue Implications*

Calculating how much firms gain from engaging in strategic delays can help us understand whether they have a strong incentive to respond to the incentives generated by ERP. We simulate the expected revenue of the naive strategy of sending entry applications everywhere right away and compare it to observed revenue. As

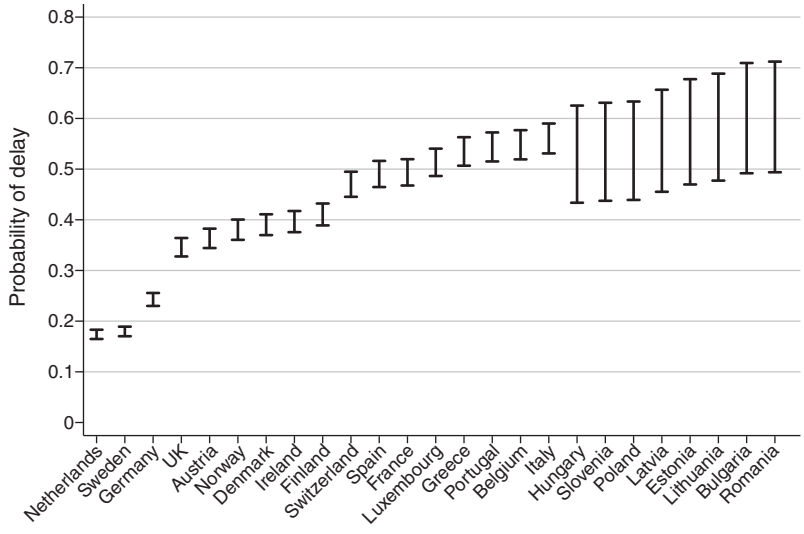

FIGURE 6. IDENTIFIED SET OF ψ_k BY COUNTRY

in the construction of the moment inequalities, we need to average across firms to eliminate the noise generated by the realization of the delay shocks. What we recover is a consistent estimate of the average revenue loss across all firms.[61]

We find that firms gain approximately €18 million per drug by engaging in strategic delays.[62] The average drug earns €46 million in Eastern Europe over its life cycle. Hence, strategic delays have a significant impact on marginal revenues in Eastern Europe. At the same time, €18 million is only a tiny fraction of the average lifetime expected revenue in the EEA for drugs in our sample—around €1.8 billion. There are two reasons why this number is low. First, in several Western European countries, prices are only marginally higher than in Eastern Europe, reducing the impact of reference pricing when it is applied. Second, our estimates suggest that countries with lower prices also have a higher probability of stochastic delays. As a result, even when firms apply everywhere simultaneously, drugs tend to enter later in these countries, which reduces the impact of ERP.

Since the revenue loss is not large, the EU could compensate firms with lump-sum transfers in exchange for forgoing strategic delays. This solution does not require EU member states to give up the prerogative to manage drug pricing independently. A centralized European agency could hand out the subsidies upon confirmation that an entry application has been sent and approved in all European countries. The overall budget impact of this policy would be negligible according to our estimates. On average, between 1995 and 2017, around 27 new drugs received approval in the EEA. Hence, the overall impact of this subsidy would be less than €500 million per year. For context, the overall budget of the EU was around €150 billion in 2016. Since

[61] We describe the exact procedure in more detail in Section D.7 of the online Appendix.

[62] This estimate assumes our lower bound estimate of $\psi_{EE} = 0.69$. Higher values of ψ_{EE} would generate a lower number.

lump sums constitute transfers, any gains from early access would improve overall welfare.[63] We leave whether such a mechanism satisfies incentive-compatibility constraints to future research.

VII. Conclusion

This paper studies the extent to which ERP policies contribute to the disparity in access to innovative drugs across countries. ERP generates complex incentives for firms that might benefit from strategically delaying entry in low-income countries. Using a novel moment inequality approach, we characterize the impact of these policies on launch delays. Our methodology allows us to obtain identification even though the firm's actions are unobserved, thus contributing to a growing body of literature showing how moment inequalities can make even the most complicated models tractable.

Our results contribute to two ongoing discussions among European policymakers. First, our estimate of ERP's impact can guide a data-driven assessment of the policy and its alternatives. Several proposals to replace reference pricing have been suggested (Kanavos et al. 2011; OECD Health Policy Studies 2008; Towse et al. 2015; Vogler et al. 2015). If reform proves difficult, our paper suggests an alternative solution that does not require changes to the current pricing system. Firms could be induced to forgo strategic delays through a system of lump-sum transfers. The overall budget impact of this policy would be around half a billion euros per year.

Second, our results suggest caution in the wake of recent calls for increased price transparency. Proponents of price transparency claim that it can streamline the monitoring of affordability and availability of pharmaceutical products. While this is undoubtedly true, we argue that it can also exacerbate the access problem of low-income countries by making ERP constraints more stringent.

Finally, our paper also contains lessons for US policy. Current savings estimates of the introduction of ERP in the United States do not usually consider the strategic response of firms.[64] Our results suggest that this omission leads to an upward bias. Firms would react to US adoption of ERP by delaying entry in other markets. These delays would have two negative effects. First, they would limit the effectiveness of ERP and reduce the savings it generates. Second, they would decrease access in all other countries whose price the United States references.

[63] This conclusion might not hold in a model where additional frictions or dynamic considerations exist (e.g., a shadow cost of raising governments funds or dynamic implications on the incentives to invest in research and development). However, the size of the subsidy is small enough that these additional considerations will likely be second order.

[64] See, e.g., the analysis of the Ways and Means Committee Staff: https://waysandmeans.house.gov/sites/democrats.waysandmeans.house.gov/files/documents/US%20versus%20International%20Prescription%20Drug%20Prices_0.pdf, retrieved December 2019.

Appendix. Theoretical Derivations

A. *Derivation of the Price Equation*

The derivation hinges on showing that the reference price is a linear multiple of the drug fixed effect. Intuitively, this result arises because the reference price is a weighted average of government prices, which are all linear multiples of the drug fixed effect. We first prove the following lemma:

LEMMA 3: *Let* $\lambda_{jkt} = \gamma_k \exp(\beta_Z Z_{jkt-1})$, *and let* $\lambda_{jt} = \{\lambda_{jkt}\}_{k \in \mathcal{N}_i}$. *Then there exists a set of weights* $\phi_{jklt}\left(S_{t-1}, \{\lambda_{j\tau}\}_{\tau=1}^{t}\right)$ *such that for any drug* j, *country* k, *and year* t,

$$p_{jkt}^{ref} = \sum_{\tau=1}^{t} \sum_{l \in R_{k\tau}} \phi_{jklt}\left(S_{t-1}, \{\lambda_{j\tau}\}_{\tau=1}^{t}\right) p_{jl\tau-1}^{gov}\left(D_{jlt}(\xi_{lt})\right).$$

PROOF:

This lemma states that we can write p_{jkt}^{ref} as a linear function of past government prices whose weights depend on the entry sequences of all firms and structural parameters of the model (except drug fixed effects). We use proof by induction.

Start from $t = 1$. We want to show that for all k,

$$p_{jkt}^{ref} = \sum_{l \in R_{k1}} \phi_{jkl1}(S_0, \lambda_{j0}) p_{jl0}^{gov}\left(D_{jl1}(\xi_{l1})\right). \tag{A.1}$$

The definition of the reference price is

$$p_{jk1}^{ref}\left(L_{j0}, D_{jk1}(\xi_{k1})\right) = F_{k1}^{ref}\left(\left\{p_{jl0}\left(D_{jk1}(\xi_{k1})\right)\right\}_{l \in (R_{k1} \cap L_{j0})}\right),$$

where $L_{j0} = \{k : s_{j0} \neq 0\}$ is the set of countries where product j is available in period 0. Since reference pricing cannot be applied at time $t = 0$, the prices that can be referenced are the initial government prices: $p_{jl0}\left(D_{jk1}(\xi_{k1})\right) = p_{jl0}^{gov}\left(D_{jk1}(\xi_{k1})\right)$. F_{k1}^{ref} can be the average function, the minimum function, or the average of the three lowest prices.[65] If F_{k1}^{ref} is the average function, the proof is complete because average is a linear function. If F_{k1}^{ref} is the minimum function, or the average of the three lowest prices, we can construct weights as follows. First note that L_{j0} depends on S_0 only. Assume without loss of generality that L_{j0} is not empty, and let n_l denote the rank of $l \in L_{j0}$ in increasing order of λ_{j1}. In other words, if $n_l = 1$, then $\lambda_{jl1} = \min\{\lambda_{j\ell 1} : \ell \in (R_{k1} \cap L_{j0})\}$, and, more generally, $\lambda_{jl1} = \min\{\lambda_{j\ell 1} : \ell \in (R_{k1} \cap L_{j0}) \wedge n_\ell \geq n_l\}$. In other words, n_l is a ranking of countries in decreasing order of government price. Finally, let $m_{jk1} = \min\{|L_{j0}|, 3\}$, where the operator $|\cdot|$ indicates the cardinality of a set. If F_{k1}^{ref} is the average of the three lowest prices, write the weights as

$$\phi_{jkl1}(S_0, \lambda_{j0}) = \begin{cases} 1, & \text{if } n_l = 1; \\ 0, & \text{otherwise.} \end{cases}$$

[65] See Figure 1 and the series of ERP figures in the online Appendix.

If F_{j1}^{ref} is the average of the three lowest prices instead, construct the weights as

$$\phi_{jkl1}(S_0, \lambda_{j0}) = \begin{cases} \frac{1}{m_{jk0}}, & \text{if } n_l \leq m_{jk0}; \\ 0, & \text{otherwise.} \end{cases}$$

These weights are written as a function of S_0 and λ_{j0}; hence they satisfy the premise of the proposition.

Now suppose the assertion of the proposition is true for $\tau \in \{1, \ldots, t-1\}$. We can walk through the same exact steps as we did for $t = 1$ but substitute $p_{jlt-1}\big(L_{jt-1}, D_{jkt}(\xi_{kt})\big)$ for λ_{j1}. Doing this will give us weights for p_{jkt}^{ref} as a linear function of the prices in the previous period. By construction, prices in the previous period are a weighted average of government prices and reference prices. By the inductive assumption, the reference prices are linear functions of adjusted government prices. Since the sum of linear functions is also linear, the proposition holds for period t as well.

The lemma gives us a way to write the reference price as a weighted average of government prices.

PROPOSITION 4: *Let $p_{jkt}(\cdot)$ be as in equation (8). Then, for any $k, k' \in \mathcal{N}_j$ we have*

$$\ln\left[\frac{p_{jkt}(\cdot)}{p_{jk't+1}(\cdot)}\right] =$$

$$\begin{cases} \ln\left[\dfrac{\gamma_k \cdot \exp\big(\beta_Z Z_{jkt} + \beta_D \ln(D_{jkt})\big)}{\gamma_{k'} \cdot \exp\big(\beta_Z Z_{jk't+1} + \beta_D \ln(D_{jk't+1})\big)}\right], & \text{if } p_{jkt}^{ref}(\cdot) \geq p_{jkt}^{gov}(\cdot) \wedge \\ & p_{ikt+1}^{ref}(\cdot) \geq p_{ikt+1}^{gov}(\cdot); \\ \ln\left[\dfrac{\gamma_k \cdot \exp\big(\beta_Z Z_{jkt} + \beta_D \ln(D_{jkt})\big)}{(1-\mu_{k'})\gamma_{k'} \cdot \exp\big(\beta_Z Z_{jk't+1} + \beta_D \ln(D_{jk't+1})\big) + \mu_{k'} \tilde{p}_{jk't+1}^{ref}(\cdot)}\right], & \text{if } p_{jkt}^{ref}(\cdot) \geq p_{jkt}^{gov}(\cdot) \wedge \\ & p_{ikt+1}^{ref}(\cdot) < p_{ikt+1}^{gov}(\cdot); \\ \ln\left[\dfrac{(1-\mu_k)\gamma_k \cdot \exp\big(\beta_Z Z_{jkt} + \beta_D \ln(D_{jkt})\big) + \mu_k \tilde{p}_{ijt}^{ref}(\cdot)}{\gamma_{k'} \cdot \exp\big(\beta_Z Z_{jk't+1} + \beta_D \ln(D_{jk't+1})\big)}\right], & \text{if } p_{jkt}^{ref}(\cdot) < p_{jkt}^{gov}(\cdot) \wedge \\ & p_{ikt+1}^{ref}(\cdot) \geq p_{ikt+1}^{gov}(\cdot); \\ \ln\left[\dfrac{(1-\mu_k)\gamma_k \cdot \exp\big(\beta_Z Z_{jkt} + \beta_D \ln(D_{jkt})\big) + \mu_k \tilde{p}_{ijt}^{ref}(\cdot)}{(1-\mu_{k'})\gamma_{k'} \cdot \exp\big(\beta_Z Z_{jk't+1} + \beta_D \ln(D_{jk't+1})\big) + \mu_{k'} \tilde{p}_{jk't+1}^{ref}(\cdot)}\right], & \text{if } p_{jkt}^{ref}(\cdot) < p_{jkt}^{gov}(\cdot) \wedge \\ & p_{ikt+1}^{ref}(\cdot) < p_{ikt+1}^{gov}(\cdot); \end{cases}$$

where $\tilde{p}_{jkt}^{ref}(\cdot)$ is such that $p_{jkt}^{ref}(\cdot) = \tilde{p}_{jkt}^{ref}(\cdot)\theta_j$ and $\tilde{p}_{jkt}^{ref}(\cdot)$ is not a function of θ_j.

PROOF:

By Lemma 3, government prices are a multiplicative function of θ_j and therefore so are reference prices. Hence we can write $p^{ref}_{jkt}(\cdot) = \tilde{p}^{ref}_{jkt}(\cdot) \cdot \theta_j$, where $\tilde{p}^{ref}_{jkt}$ is not a function of θ_j. This yields

$$p_{jkt}(\cdot) = \begin{cases} \theta_j \gamma_k \cdot \exp\big(\beta_Z Z_{jkt} + \beta_D \ln(D_{jkt})\big), & \text{if } p^{ref}_{jkt}(\cdot) \geq p^{gov}_{jkt}(\cdot); \\ \theta_j \gamma_k \cdot \exp\big(\beta_Z Z_{jkt} + \beta_D \ln(D_{jkt})\big) + \mu_k \tilde{p}^{ref}_{jkt}(\cdot)\theta_j, & \text{if } p^{ref}_{jkt}(\cdot) < p^{gov}_{jkt}(\cdot); \end{cases}$$

which shows that $p_{jkt}(\cdot)$ is a multiplicative function of θ_j. Hence, when we consider $p_{jkt}/p_{jk't+1}$, θ_j will appear in both the denominator and the numerator and will drop out.

B. *Proof of Theorem 1*

To prove the theorem, we rely on the strong law of large numbers applied to non-identical, independent random variables.

Let $\mathbf{\Pi}_{jt}(\mathcal{A}'_{jt}, Y_j, \psi_0)$ be a random variable denoting the payoff of firm j from period t onward, conditional on firm l following strategy $\mathcal{A}_{jt}$, and on the firm's information set $\mathcal{I}_{jt}$. The conditional expectation of $\mathbf{\Pi}_{jt}(\mathcal{A}'_{jt}, Y_j, \psi_0)$ is $E\big[\mathbf{\Pi}_{jt}(\mathcal{A}'_{jt}, Y_j, \psi_0)\,\big|\,\mathcal{I}_{jt}\big]$. We start by proving a useful Lemma.

LEMMA 5: *Let* $R_\tau(S_{j\tau}, S_{-j\tau})$ *be defined as in equation* (11). *Then* $R_\tau(S_{j\tau}, S_{-j\tau}) < \infty$.

PROOF:

The lemma states that period payoffs are bounded. The realization of period payoffs depends on ξ_{kt}. Define

$$R_\tau(S_{j\tau}, S_{-j\tau}, \xi_{k\tau}) = \sum_{k \in L_{j\tau}} p_{jk\tau}\big(S_{j\tau-1}, S_{-j\tau}, D_{jk\tau}(S_{-j\tau}, \xi_{k\tau})\big) \cdot D_{jk\tau}(S_{-j\tau}, \xi_{k\tau}).$$

For any given product j and country k, we can write demand as

$$D_{jk\tau}(S_{-j\tau}, \xi_{k\tau}) = M_{k\tau} \cdot \frac{\exp(\alpha_{jk} + \beta_j age_{j\tau} + \eta_j NF_{jk\tau} + \xi_{k\tau})}{1 + \sum_{\ell \in E_{-j\tau}} \exp(\alpha_{\ell k} + \beta_\ell age_{\ell\tau} + \eta_\ell NF_{\ell k\tau} + \xi_{k\tau})},$$

where $M_{k\tau}$ is the market size in country k in period τ. Hence, $D_{jk\tau}(S_{-j\tau}, \xi_{kt}) \in [0, M_{k\tau}]$, which means demand is bounded. Price is also bounded by the government price, which is finite:

$$p_{jk\tau}\big(S_{j\tau-1}, S_{-j\tau}, D_{jk\tau}(S_{-j\tau}, \xi_{k\tau})\big) \leq p^{gov}_{jk\tau}\big(D_{jk\tau}(S_{-j\tau}, \xi_{k\tau})\big) < \infty.$$

Moreover, using the definition of government price in equation (6), we can rewrite

$$p^{gov}_{jk\tau}\left(D_{jk\tau}(\xi_{k\tau})\right) \cdot D_{jk\tau}(\xi_{k\tau}) \;=\; p^{gov}_{jk\tau}(1) \cdot \left(D_{jk\tau}(\xi_{k\tau})\right)^{1+\beta_D}.$$

Hence, the period payoff of a single drug in any given country is bounded above by $p^{gov}_{jk\tau}(1) \cdot \left(M_{j\tau}\right)^{1+\beta_D}$ and is bounded below by zero. This implies that the period payoff in any given country and period is finite, and therefore $R_\tau\left(S_{j\tau}, S_{-j\tau}\right)$ is also finite.

COROLLARY 6: $\mathbf{\Pi}_{jt}\left(\mathcal{A}'_{jt}, Y_j, \psi_0\right)$ *has finite variance.*

PROOF:

This corollary follows directly from Lemma 5. $\mathbf{\Pi}_{jt}\left(\mathcal{A}'_{jt}, Y_j, \psi_0\right)$ is defined as the discounted sum of the expected period payoffs. By Lemma 5, the expected period payoffs are finite. In particular, given that period payoffs in each country are bounded above by $p^{gov}_{jkt}(1) \cdot \left(M_{jt}\right)^{1+\beta_D}$ and below by 0, we can conclude that the support of $\mathbf{\Pi}_{jt}\left(\mathcal{A}'_{jt}, Y_j, \psi_0\right)$ is bounded above by

$$\sum_{\tau=t}^{T_j} \sum_{j=1}^{N_j} p^{gov}_{jk\tau}(1) \cdot \left(M_{j\tau}\right)^{1+\beta_D}$$

and below by zero. Hence, it must have finite variance.

At this point we are ready to prove Theorem 1.

PROOF:

PROOF OF THEOREM 1:

For any given firm j, $\tilde{R}_t\left(S^o_j, S^o_{-j}, X_j\right)$ represents a draw from the distribution of $\mathbf{\Pi}_{jt}\left(\mathcal{A}^\star_{jt}\left(\mathcal{I}_{jt}\right), Y_j, \psi_0\right)$. By Corollary 6, the random variable $\mathbf{\Pi}_{jt}\left(\mathcal{A}^\star_{jt}\left(\mathcal{I}_{jt}\right), Y_j, \psi_0\right)$ has finite variance. Moreover, for all j, each random variable $\mathbf{\Pi}_{jt}\left(\mathcal{A}^\star_{jt}\left(\mathcal{I}_{jt}\right), Y_j, \psi_0\right)$ is independently distributed. Thus, our premise satisfies the Kolmogorov criterion, which implies that the strong law of large numbers applies to our sequence of random variables, and the sample average of the realized payoffs will converge to the average of their expected values.[66] Formally, for any $\epsilon > 0$, we can find M' such that

$$\frac{1}{M}\left|\sum_{j=1}^{M}\left(\tilde{R}_t\left(S^o_j, S^o_{-j}, X_j\right) - E\left[\tilde{V}_t\left(\mathcal{A}^\star_{jt}\left(\mathcal{I}_{jt}\right), Y_j, \psi_0\right)\middle|\mathcal{I}_{jt}\right]\right)\right| < \epsilon,$$

for all $M > M'$. This concludes the proof.

[66] The Kolmogorov criterion requires $\sum_{k=1}^{\infty} \frac{\sigma_k^2}{k^2} < \infty$, where σ_k^2 is the variance of the k^{th} random variable in the sequence.

C. *Proof of Theorem 2*

The proof of Theorem 2 is almost identical to the proof of Theorem 1. To see how, simply note that we can set $nsim = 1$, replace S_j with $S_j\left(\omega_j^1, \psi'\right)$, and repeat all the same steps from the proof above. Notice, however, that this is an extreme case, and the estimator will perform better if *nsim* is set high enough to make the simulation error for S_j small.

REFERENCES

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica* 75 (5): 1331–70.

Barseghyan, Levon, Maura Coughlin, Francesca Molinari, and Joshua C. Teitelbaum. 2020. "Heterogeneous Choice Sets and Preferences." arXiv:1907.02337.

Birg, Laura. 2016. "External Reference Pricing and the Choice of Country Baskets and Pricing Rules." Unpublished.

Borja, García Lorenzo. 2014. "Modeling Global Pricing and Launching of New Drugs." Doctoral thesis. Universidad de las Palmas de Gran Canaria.

Brekke, Kurt R., Chiara Canta, and Odd Rune Straume. 2016. "Reference Pricing with Endogenous Generic Entry." *Journal of Health Economics* 50: 312–29.

Brekke, Kurt R., Tor Helge Holmås, and Odd Rune Straume. 2015. "Price Regulation and Parallel Imports of Pharmaceuticals." *Journal of Public Economics* 129: 92–105.

Brekke, Kurt R., Ingrid Königbauer, and Odd Rune Straume. 2007. "Reference Pricing of Pharmaceuticals." *Journal of Health Economics* 26 (3): 613–42.

Cardell, N. Scott. 1997. "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity." *Econometric Theory* 13 (2): 185–213.

Carone, Giuseppe, Christoph Schwierz, and Ana Xavier. 2012. "Cost-Containment Policies in Public Pharmaceutical Spending in the EU." Unpublished.

Chaudhuri, Shubham, Pinelopi K. Goldberg, and Panle Jia. 2006. "Estimating the Effects of Global Patent Protection in Pharmaceuticals: A Case Study of Quinolones in India." *American Economic Review* 96 (5): 1477–1514.

Cockburn, Iain M., Jean O. Lanjouw, and Mark Schankerman. 2016. "Patents and the Global Diffusion of New Drugs." *American Economic Review* 106 (1): 136–64.

Costinot, Arnaud, Dave Donaldson, Margaret Kyle, and Heidi Williams. 2019. "The More We Die, The More We Sell? A Simple Test of the Home-Market Effect." *Quarterly Journal of Economics* 134 (2): 843–94.

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

Crawford, Gregory S., and Ali Yurukoglu. 2012. "The Welfare Effects of Bundling in Multichannel Television Markets." *American Economic Review* 102 (2): 643–85.

Danzon, Patricia M., and Andrew J. Epstein. 2012. "Effects of Regulation on Drug Launch and Pricing in Interdependent Markets." *Advances in Health Economics and Health Services Research* 23: 35–71.

Danzon, Patricia M., Y. Richard Wang, and Liang Wang. 2005. "The Impact of Price Regulation on the Launch Delay of New Drugs—Evidence from Twenty-Five Major Markets in the 1990s." *Health Economics* 14 (3): 269–92.

Decarolis, Francesco. 2015. "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?" *American Economic Review* 105 (4): 1547–80.

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

Drummond, M. F. 2003. "Will There Ever Be a European Drug Pricing and Reimbursement Agency?" *European Journal of Health Economics* 4 (2): 67–69.

Dubois, Pierre, and Laura Lasio. 2018. "Identifying Industry Margins with Price Constraints: Structural Estimation on Pharmaceuticals." *American Economic Review* 108 (12): 3685–3724.

Dubois, Pierre, Ashvin Gandhi, and Shoshana Vasserman. 2018. "Bargaining and International Reference Pricing in the Pharmaceutical Industry." NBER Working Paper No. 30053.

Duggan, Mark, and Fiona M. Scott Morton. 2006. "The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing." *Quarterly Journal of Economics* 121 (1): 1–30.

Duso, Tomaso, Annika Herr, and Moritz Suppliet. 2014. "The Welfare Impact of Parallel Imports: A Structural Approach Applied to the German Market for Oral Anti-Diabetics." *Health Economics* 23 (9): 1036–57.

Eizenberg, Alon. 2014. "Upstream Innovation and Product Variety in the US Home PC Market." *Review of Economic Studies* 81 (3): 1003–45.

European Central Bank. 2002-2012. "Euro foreign exchange reference rates", https://www.ecb.europa.eu/stats/policy_and_exchange_rates/euro_reference_exchange_rates/html/index.en.html (accessed February 2023).

European Commission. 2012. "Impact Assessment Accompanying the Document Proposal for a Directive of the European Parliament and of the Council Relating to the Transparency of Measures Regulating the Prices of Medicinal Products for Human Use and Their Inclusion in the Scope of the Public Health Insurance Systems." Technical Report.

European Medicines Agency. 1995-2012. "European public assessment reports", https://web.archive.org/web/20151030181947/http://www.ema.europa.eu/ema/index.jsp?curl=pages/medicines/landing/epar_search.jsp (accessed February 2023).

Feng, Josh, Thomas Hwang, and Luca Maini. 2020. "Profiting from Most-Favored Customer Procurement Rules: Evidence from Medicaid." Unpublished.

Fong, Kyna, and Robin S. Lee. 2013. "Markov-Perfect Network Formation: An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks." Unpublished.

Grennan, Matthew. 2013. "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices." *American Economic Review* 103 (1): 145–77.

Grennan, Matthew, and Ashley Swanson. 2016. "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining." NBER Working Paper No. 22039.

Heads of Medicines Agency, 1995-2012. "MRI Product Index", https://mri.cts-mrp.eu/portal/home?domain=h (accessed February 2023).

Eurostat. 2002-2012. "GDP and main components (output, expenditure and income)", and "Population on January 1 by broad age group and sex", https://ec.europa.eu/eurostat/databrowser/view/namq_10_gdp/default/table?lang=en and https://ec.europa.eu/eurostat/databrowser/view/demo_pjanbroad/default/table?lang=en (accessed February 2023).

Ho, Kate, and Ariel Pakes. 2014. "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians." *American Economic Review* 104 (12): 3841–84.

Holmes, Thomas J. 2011. "The Diffusion of Wal-Mart and Economies of Density." *Econometrica* 79 (1): 253–302.

Houy, Nicolas, and Izabela Jelovac. 2015. "Drug Launch Timing and International Reference Pricing." *Health Economics* 24 (8): 978–89.

Illanes, Gastón. 2016. "Switching Costs in Pension Plan Choice." Unpublished.

IMS Health. 2002–2012. "MIDAS Database."

Institute for Health Metrics and Evaluations. 2002-2012. "Global Burden of Disease Study", https://vizhub.healthdata.org/gbd-results/ (accessed February 2023).

Jaffe, Sonia, and Mark Shepard. 2017. "Price-Linked Subsidies and Health Insurance Markups." Unpublished.

Kanavos, Panos, Sotiris Vandoros, Rachel Irwin, Elena Nicod, and Margaret Casson. 2011. Differences in Costs of and Access to Pharmaceutical Products in the EU. Brussels: European Parliament.

Katz, Michael. 2007. "Essays in Industrial Organization and Technological Change." PhD diss. 2007.

Kyle, Margaret K. 2006. "The Role of Firm Characteristics in Pharmaceutical Product Launches." *RAND Journal of Economics* 37 (3): 602–18.

Kyle, Margaret K. 2007. "Pharmaceutical Price Controls and Entry Strategies." *Review of Economics and Statistics* 89 (1): 88–99.

Kyle, Margaret K. 2011. "Strategic Responses to Parallel Trade." *B.E. Journal of Economic Analysis and Policy* 11 (2).

Kyle, Margaret K., and Yi Qian. 2013. "Intellectual Property Rights and Access to Innovation: Evidence from TRIPS." Unpublished.

Maini, Luca, and Fabio Pammolli. 2023. "Replication of Reference Pricing as a Deterrent to Entry: Evidence from the European Pharmaceutical Market." *American Economic Association* [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.38886/E160781V1.

Matteucci, Giorgio, and Pierfrancesco Reverberi. 2017. "Drug Innovation, Price Controls, and Parallel Trade." *International Journal of Health Economics and Management* 17 (2): 159–79.

Morales, Eduardo, Gloria Sheu, and Andrés Zahler. 2019. "Extended Gravity." *Review of Economic Studies* 86 (6): 2668–2712.

OECD. 2008. OECD Health Policy Studies, Pharmaceutical Pricing Policies in a Global Market. Paris, France: OECD.

Pakes, A. 2010. "Alternative Models for Moment Inequalities." *Econometrica* 78 (6): 1783–1822.

Pakes, A., J. Porter, Kate Ho, and Joy Ishii. 2015. "Moment Inequalities and Their Application." *Econometrica* 83 (1): 315–34.

Pharmaceutical Industry Competitiveness Task Force. 2006. "Competitiveness and Performance Indicators 2005." Unpublished.

Ridley, David B., and Chung-Ying Lee. 2020. "Does Medicare Reimbursement Drive Up Drug Launch Prices?" *Review of Economics and Statistics* 102 (5): 980–93.

Stargardt, Tom, and Jonas Schreyögg. 2006. "Impact of Cross-Reference Pricing on Pharmaceutical Prices: Manufacturers' Pricing Strategies and Price Regulation." *Applied Health Economics and Health Policy* 5 (4): 235–47.

Toumi, Mondher, Cécile Rémuzat, Anne-Lise Vataire, and Duccio Urbinati. 2013. External Reference Pricing of Medicinal Products: Simulation-Based Considerations for Cross-Country Coordination. Brussels: European Commission.

Towse, Adrian, Michele Pistollato, Jorge Mestre-Ferrandiz, Zeba Khan, Satyin Kaura, and Louis Garrison. 2015. "European Union Pharmaceutical Markets: A Case for Differential Pricing?" *International Journal of the Economics of Business* 22 (2): 263–75.

Vogler, Sabine, Lena Lepuschütz, Peter Schneider, and Verena Stühlinger. 2015. Study on Enhanced Cross-Country Coordination in the Area of Pharmaceutical Product Pricing - Final Report. Brussels: European Commission.

Wenzl, Martin, and Suzannah Chapman. 2019. "Performance-Based Managed Entry Agreements for New Medicines in OECD Countries and EU Member States: How They Work and Possible Improvements Going Forward." OECD Health Working Paper No. 115.

Wollmann, Thomas G. 2018. "Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles." *American Economic Review* 108 (6): 1364–1406.

This article has been cited by:

1. Juan Pablo Atal, José Ignacio Cuesta, Felipe González, Cristóbal Otero. 2024. The Economics of the Public Option: Evidence from Local Pharmaceutical Markets. *American Economic Review* **114**:3, 615-644. [Abstract] [View PDF article] [PDF with links]