Regulating Innovation with Uncertain Quality: Information, Risk and Access in Medical Devices

Matthew Grennan*
University of Pennsylvania, The Wharton School
grennan@wharton.upenn.edu

Robert J. Town
University of Pennsylvania, The Wharton School and NBER
rtown@wharton.upenn.edu

April 27, 2014

Abstract

This paper explores the role of regulation of new product entry when product quality is uncertain but market participants learn over time. We develop a model that captures the fundamental regulatory tradeoff between information generation, access, and risk: weak regulation inhibits learning and exposes consumers to the risk of unproven new products, but overzealous regulation increases entry costs and reduces access to a narrow choice set. Using new data and exogenous variation between EU and US medical device regulatory rules, we document patterns consistent with our model, and then take a structural approach to estimate the welfare implications of current and alternative regulatory policies. For the set of devices on which we have data, we estimate that both the US and EU are close to the optimal policy (though for the EU depends critically on free-riding off of US trials). We also estimate that embracing recent calls for more active "post-market surveillance" could further increase total surplus by as much as 19 percent. Relying on private incentives instead of regulator mandated trial lengths tends to lead to over-investment in information among the highest quality products.

*Data for this research was provided by Millennium Research Group (www.mrg.net). We owe thanks to Joe Harrington and seminar audiences at the FTC, Minnesota, Ohio State, Rice/Houston, UCLA, Wharton, and Yale for helpful comments. The Wharton Dean's Research Fund and Indian School of Business provided financial assistance. Ellie Prager, Mike Punzalan, Emi Terasawa, and Christine Yang provided excellent research assistance. Any errors are our own.

1 Introduction

Most innovative new products are brought to the market because their makers believe they provide new value. However, once in the hands of consumers, there is always some chance that the product will not operate as hoped and fail. The consequence of this failure ranges from the consumer losing her product expenditures to death. When product failure poses significant safety risks, products often must go through pre-market testing and become approved/certified by a formal body before entering the marketplace. The standard that this regulatory body uses to approve products has the potential to fundamentally alter market outcomes. In setting its product approval criteria, the regulator must weigh the benefits of reducing risk against the effect on product access (and potentially prices). We argue that a key decision the regulator makes is the information it requires the manufacturer to generate for the product to be approved. As first highlighted by Peltzman (1973) in the context of pharmaceuticals, higher informational standards increase product specific learning and lower consumption risk but also result in delayed access to fewer products and higher entry costs conditional on approval. Today such certification processes exist and are often a source of controversy in areas as diverse as electronics, airplanes, automobiles, finance, health care, and toys.[1]

This paper uses new data and exploits exogenous regulatory differences between the US and EU to quantify the tradeoff between access and risk for a large set medical devices introduced between 2004-2013. In the US medical devices are regulated by the Food and Drug Administration (FDA) while in the EU device approval is performed by organizations that contract with the EU called Notified Bodies. Importantly, the different regions apply different standards to medical device approval. Very roughly, the US applies a "safe and effective" standard while the EU only certifies safety of the product. This difference is material. Meeting the "effectiveness" standard often requires manufacturers to generate product performance information through randomized clinical trials. These trials are costly in both time and expense. Medical device manufactures (the vast majority of which are US-based) generally introduce their products in the EU well before they cross the Atlantic and seek FDA approval, if they decide to introduce the product in the US at all. According to the Boston Consulting Group, between 2005 and 2011, the average high risk and likely high value medical device was introduced in the US four years later than in the EU. The differences between the US and the EU

[1] See, for example in electronics "European Environmental Rules Propel Change in U.S.", The New York Times, July 06, 2004; in airplanes "Boeing Acknowledges Tests Underestimated 787 Battery Risks", The New York Times, April 23, 2013; in automobiles "U.S. Sues Chrysler After Auto Maker Refuses to Recall Cars", The New York Times, June 5, 1996; in finance "An FDA for Securities Could Help Avert Crises", Bloomberg, April 2, 2012; in toys "Toy Makers Fight for Exemption From Rules", The New York Times, September 28, 2010.

in the medical device approval process have led to calls for reform in both regions. In the US, the FDA has faced attacks from both sides, with some claiming that a slower, tougher approval process is crippling innovation; and others claiming that the approval process is too lax, allowing too many dangerous devices into the market.[2] Also, as rising incomes in the developing world lead to both greater incidence of "western" diseases and greater ability to afford the most advanced technologies, the debate on how to regulate medical devices has taken on global significance, drawing the interest of the UN and WHO.[3]

Despite the importance of the information and the access/risk tradeoff in markets where research and development leads to new products with uncertain quality, empirical research has been limited by two major difficulties: (1) assembling data that can quantify the returns from increased information relative to the cost of decreased access; (2) finding exogenous variation in regulatory regimes that can identify the tradeoff between these competing forces. In this study we address the second challenge by exploiting the fact that the EU approval process is both faster and less costly than the US process for any given device, and this difference is due largely to historical political processes. This allows us to measure the access/risk tradeoff using a newly constructed, detailed data set for a variety of medical devices available in the US and EU from 2004-2013.

The ideal way to address the first challenge would be to combine data on market outcomes (quantities and prices) with data on health and safety outcomes. Unfortunately, even in a highly regulated and documented industry such as medical devices, health outcome and safety data are not available at the product level. This forces us to ask what can be inferred with more commonly available data such as market prices and quantities, to which our answer is a substantial amount. We begin by constructing a theoretical model where products are invented with uncertain quality, market entry is regulated, and the market learns about product quality over time. The key feature of our model is that the rate of learning in premarket clinical trials can be greater than the rate of learning after market entry. This introduces a tradeoff where more regulation leads to more learning and less risk, but also delayed access and higher entry costs for innovative new products. The model clarifies patterns in the data that one should expect as a function of the distribution of product qualities invented, the rates of learning, consumer preferences, and regulatory rules.

Our data comes from Millennium Research Group (MRG), a medical device consult-

[2]For an example arguing the FDA is too lax "Report Criticized F.D.A. on Device Testing", The New York Times, January 15, 2009; and too tight "FDA Seeks to Toughen Defibrillator Regulations", The New York Times, March 22, 2013.

[3]"UN: Western Diseases a Growing Burden on Developing World," The Wall Street Journal, May 14, 2010. "Global Forum to Improve Developing Country Access to Medical Devices," press release, WHO, September 9, 2010.

ing firm. MRG collects detail, high frequency, hospital-product-level data for medical devices on prices, volumes and number of diagnostic procedures for hospitals in the US and the EU. Our analysis focuses on the market for coronary stents. We chose this segment as the coronary stent market is large an important with excellent market data and with constant innovations introduced over time. Coronary stents treat ischemic heart disease – the narrowing of the coronary artery cause by fatty deposits. Ischemic heart disease is the leading cause of global death accounting for 7 million deaths in 2010 (Lozano, 2012). In 2011 total, world-wide sales of coronary stents exceeded $7 billion with the vast majority of those sales occurring in the US and the EU.[4]

Our data analysis begins by documenting multiple patterns consistent with the model. Our analysis shows that the EU enjoys greater access to the best new medical technologies, while also bearing greater risk by allowing entry of a wider range of device qualities, earlier in each device's lifecycle. The greater access in the EU is evident in the fact that on average 47 percent of the stents used in the EU are unavailable in the US at that point in time. The greater risk in the EU is evident in the facts that on average products in the EU experience less usage overall and higher variance in usage patterns when first introduced, with this usage discount and variance decreasing and stabilizing over the first two years on the market (the US, by contrast, exhibits no such patterns). Differential learning rates between clinical trials and market use are identified by differential usage patterns over time for products with and without ongoing clinical trials.

To develop welfare measures and address policy questions regarding optimal regulation, we then proceed with a structural approach. We combine the data with our learning model of product choice to estimate the distribution of product qualities and risk between the EU and the US, as well as the speed of learning and preferences of consumers in the marketplace. With these parameters in hand, we estimate the impact of different regulatory rules on product introductions and consumer welfare.

We estimate that total surplus is maximized when the average premarket clinical trial is six months longer than the current EU requirements and four months shorter than current US requirements. Because total surplus as a function of time spent in premarket testing is relatively flat around the optimal, US policy is statistically equivalent to the optimal. By contrast, it at first appears that the EU could make welfare gains of up to 40 percent by increasing its standards—until one realizes that the EU is able to free-ride off of the information being generated in trials for US entry, which makes current EU policy a best-response to current US policy.

Related to this issue of information generation after product approval, we also analyze a commonly suggested policy change that would relax premarket requirements

[4]Source: BCC Research.

but increase "post-market surveillance". We estimate that if it is possible to achieve post-approval learning rates close enough to those we observe from clinical trials at a comparable cost, the further benefits from such a policy change could be as high as 19 percent of total surplus. In the extreme case where post-approval learning is informative and not too costly, the optimal policy is to require no pre-approval trials at all.

Finally, we turn from optimal government policy to the question of manufacturers' private incentives to invest in information generation without regulatory mandates. This is an empirical question because theory suggests that firms may over invest because of business stealing incentives or underinvest because private returns are only a fraction of total social returns. We find evidence of both—firms with the highest quality products will tend to over invest while those with lower quality products will under invest.

Because our data collection, research design, and modeling efforts are focused on the issue of information generation, risk, and access, our analysis should be interpreted as holding other roles of the regulator—such as setting standards for what constitutes what constitutes acceptable evidence and verifying the information produced in trials—as fixed. Together, our results suggest that on the dimension of clinical trial requirements, the US and EU are both very close to optimal (in the EU case, conditional on US trials as a source of post-market information), at least for the product categories we have the data to analyze. We also discuss how our empirical results and theoretical model could inform extrapolation to a broader set of product categories.

Our work builds on recent empirical research on optimal regulation (Timmins 2002; Miravete, Seim, and Thurk 2013) and consumer learning (Roberts and Urban 1988; Erdem and Keane 1996; Ackerberg 2003; Crawford and Shum 2005; Ching 2010), and to our knowledge is the first to combine these two. This combination is essential in allowing us to address the policy question at the heart of this paper. Of course, Peltzman's (1973) pioneering work laid the foundation for our analysis. Using pre-/post-analysis, he argues that the 1962 FDA act which require clinical trials for pharmaceuticals prior to their introduction to the market harmed consumers as it reduced access to drugs without off-setting increases in product information. As our approach relies on established models and data that is relatively easily available, we hope provides an approach that future researchers might find useful in the area of entry regulation via product approval/certification processes. We also see this work as complementary to recent empirical research on the impact of patent length (another regulatory tool that impacts entry) on product introductions (Filson 2012) and innovative activity (Budish, Roin, and Williams 2013) as well as the literature on quality disclosure (Dranove and Jin, 2010).

Our analysis of the impact of different regulatory regimes not only speaks to the broad

questions of the economics of product quality regulation but to also informs policy with potentially large welfare consequences. The amount of economic activity regulated by the FDA and the Notified Bodies is significant. In the US the medical device market sales exceeded $150B in 2010 or 6% of total national health expenditures and approximately $130B in the EU or 7.5% of total health expenditures.[5]. Furthermore, the introduction of new medical technologies are responsible for much of the recent reductions in mortality and in so far as different regulatory regimes affect the availability of these technology, the welfare impact of these regulations likely extends beyond their direct impact on commerce.

The remainder of the paper is organized as follows: The next Section discusses the institutional background of medical device regulation in the US and EU. Section 3 develops a general model that captures the tradeoffs involved in regulating market entry of products with uncertain quality and derives testable predictions. Section 4 then tests these predictions in the data, finding evidence in support of the model. Section 5 takes a structural approach, explicitly estimating the parameters of the model and deriving welfare estimates for current as well as counterfactual regulatory regimes.

2 Medical Device Regulation in the US and the EU

Medical device regulation in the US began on May 28, 1976 with the passage of the Medical Device Amendments Act of 1976. This law established the regulator pathway for medical devices in the US, placing oversight authority within the Food and Drug Administration (FDA). The criteria the FDA is mandated to use is "safe and effective." Prior to the passage of the Act, medical devices were essentially unregulated. The Act established three classification of devices (I, II and III) which are assigned to each device based on the perceived risks associated with using the device. Class III devices are defined as those used in in supporting or sustaining human life, of substantial importance in preventing impairment of human health, or presents a potential unreasonable risk of illness or injury. Class I and Class II devices are lower risk devices for which there is a sufficient body of evidence demonstrating a performance standard for the design and manufacturing of the device.

There are two basic regulatory pathways within the FDA to bring a device to market: Pre-Market Approval (PMA) and the 510(k). The PMA process applies to Class III devices, while the 510(k) process generally applies to Class II and some Class I devices. Under the 510(k) process the manufacturer needs to demonstrate that the device is 'substantially equivalent' to a predicate device. Generally, bench testing data and perhaps a

[5]Donahoe and King, 2012; Medtech Europe, 2013

very small clinical study is all that is necessary for a device to demonstrate equivalency. While there is no standard timetable for 510(k) clearance, a straightforward clearance can typically be obtained within several months.

However, the approval process is much more complicated and costly for PMA devices. Approval of a PMA device requires the sponsor to provide data from a pivotal study. These are large, multi center, randomized clinical trials. These studies involve hundreds of patients and cost tens of millions of dollars to complete. PMA submission often reach thousands of pages in length. According to the Boston Consulting Group, the average cost of a PMA application approaches $100 million and often takes several years for the FDA to make a decision after the initial submission has been made. In 2012, only 37 PMAs were approved by the FDA.

In the EU the device approval process for Class III devices is very different than in the US.[6] Medical devices are regulated by three EU Directives. The main Directive is the Medical Devices Directive which has been adopted by each EU member state and passed in June, 1993. A medical device is approved for marketing in the EU once it receives a 'CE mark' of conformity. The CE mark system relies heavily on third parties know as "notified bodies" to implement regulatory control over devices. Notified Bodies are independent commercial organizations that are designated, monitored and audited by the relevant member states via "competent authorities." Currently, there are more than 70 active notified bodies within the EU.[7] A firm is free to choose any notified body designated to cover the particular type of device under review.[8] To obtain an CE mark a Class III medical device needs to only demonstrate safety and performance. Compliance with this standard usually can be demonstrated with much simpler and cheaper clinical trials than required by the FDA.

The differences in the two regulatory regimes is largely a consequence of different histories that lead up to the passing of the primary medical device legislation in the two regions. The Medical Device Directive, the centerpiece of the EU medical device regulatory framework, was passed in 1993 when there was keen interest in a new approach to harmonization regulatory frameworks across the member states. The EU had just undertaken a long and frustrating harmonization process for food and drugs. This new approach sought to avoid detailed and bureaucratic government approval processes, particularly duplicative approvals and was applied to other products including toys, pressure vessels and personal protective equipment. In contrast, the US medical device regula-

[6]Actually, there are four different classes of medical devices in the EU(Classes I, IIa, IIb and III). Class III devices in the EU closely map into Class III devices in the US.

[7]Recent regulatory reform of the Medical Device Directive now limits the ability of Notified Bodies to outsource their device reviews.

[8]See *Guidelines Relating to Medical Devices Directives, http://ec.europa.eu/health/medical-devices/documents/guidelines/.*

tory framework was established after the Dalkon Shield injured several thousand women. The FDA already had oversight on some aspects of medical devices and expanding that role was the most viable political option. At that time, a non-governmental approach to device regulation was never seriously considered by the Congress.

The differences between the two systems is the focus of a number of consulting, lobbying organizing and government reports. For example, a series of Boston Consulting Group reports shows that there is no difference in recalls between devices that are marketed in both the US and the EU. Of course, as we show below, the mix of devices that are introduced into the US is different and thus it is unclear what this study says about the impact of counterfactual regulations on device safety. In fact, the FDA countered the BCG study with their own case study of 10 devices that were approved in the EU that were not approved by the FDA and these devices lead to significant adverse events in patients. Of course, the FDA study only focused on the negative consequences of the EU's relatively lax regulatory standards and does not acknowledge the benefits of greater access to devices in the EU.

While the consequences of the different regulatory regimes has generated significant policy debate, what is less controversial is the there are significant lags in the introduction of devices between the US and the EU. Conditional on entry into both the US and the EU, BCG documents that medical devices are introduced into the US approximately four years after their EU introduction.[9] In the next section we develop a theoretical framework for assessing the trade-offs inherent in the different regulatory approaches. A notable advantage of our model is that the key parameters can be directly estimated from commonly available data, and thus the welfare of counterfactual policies can be assessed.

3 A Model of Quality Uncertainty, Learning, Entry Regulation, and Consumer Choice

In this Section, we develop a model that captures the tradeoff between risk and access involved in regulating market entry of products with uncertain quality. In our model, products are developed with uncertain quality; this uncertainty is resolved over time via exogenous signals (e.g. from clinical trials or other research); a regulator restricts entry by requiring costly premarket clinical trials to accelerate learning about product quality; and risk-averse consumers choose from the available products in the market at any point in time.

[9]BCG (2012) *Regulation and Access to Innovative Medical Technologies.*

Our model captures many of the salient features of medical device markets and the role of the regulator, however the medical device sector is complicated and there are notable institutional features that we purposefully ignore in order to keep the model tractable and parsimonious. In particular, we do not model the possibility that the regulator will reject a device. We do, however, model manufacturers' optimal entry decisions in the face of clinical trial and entry costs. This amounts to an implicit assumption that no firm would want to enter with a product the regulator would want to reject. We believe this is reasonable because these costs are non-negligible for the majority of products.

We have also considered and decided not to study here other roles for medical device regulation. As we have modeled, medical device quality is uncertain and if manufacturers are differentially informed about their devices quality, device regulation could solve a lemons problem (Leland 1979). At the extensive margin of whether to have any regulation at all, the lemons problem is surely relevant. However, our focus is on the appropriate standards of that regulation not on whether the regulation should exist. The variation that we exploit aligns with this focus. The EU is more lax in their standard relative to the US yet we are unaware of any significant evidence that the device market in the EU 'unravels' more than in the US. In fact, the presences of many more device offerings in the EU strong suggests that the variation in regulations between the US and EU is not a margin that would induce a lemons type market failure.

The next several subsections lay out the model. Section 3.1 describes how market participants learn about product quality over time, Section 3.2 describes consumer behavior and how it is affected by uncertainty about product quality, Section 3.3 turns to supplier pricing and entry, and finally Section 3.4 lays out the role for a regulator to affect total surplus via information requirements and their effect on risk and access.

3.1 Modeling Learning

The key element of the model is the uncertainty over product quality and structure of learning over time. Our specification explicitly models the impact of clinical trials on the information set that physicians use to assess which product they should implant into their patients. Assume innovative new devices j are each developed with quality Q_j according to a distribution $F_t(Q)$[10]:

$$Q_j \sim F_t(Q) := N(Q_t, \sigma_Q^2). \tag{1}$$

[10]For simplicity, we assumed the prior and signal process to be normally distributed. In principle, these processes can be specified and empirically identified non parametrically. In practice, however, data limitations make a more parametric specification desirable, and we find that the simple normal model fits the data quite well with a small number of parameters.

where the subscript t allows for technological advancement over time.

Over time, unbiased but noisy signals A arrive regarding the product's quality as new data (from ongoing clinical trials and real world usage) are released and this information diffuses into the market (where here age a refers to the time since product j was introduced to the market, not the calendar month):

$$A_{ja} = Q_j + \nu_{ja} \qquad \text{where} \quad \nu_{ja} \sim \begin{cases} N(0, \sigma^2_{A^c}) & \text{if in clinical trials} \\ N(0, \sigma^2_A) & \text{if not} \end{cases} \tag{2}$$

Given these signals, beliefs about product quality are updated via Bayes' rule, and due to the normally distributed prior and signal, posterior beliefs are also distributed normal with mean:

$$Q_{ja+1} = \frac{\sigma^2_{ja}}{\sigma^2_{ja} + \sigma^2_{A^{ja+1}}} A_{ja+1} + \frac{\sigma^2_{A^{ja+1}}}{\sigma^2_{ja} + \sigma^2_{A^{ja+1}}} Q_{ja} \tag{3}$$

and variance:

$$\sigma^2_{ja+1} = \frac{\sigma^2_{A^{ja+1}}}{\sigma^2_{ja} + \sigma^2_{A^{ja+1}}} \sigma^2_{ja}. \tag{4}$$

With this uncertainty and learning as a backdrop, the regulator must make a decision regarding the required length of clinical trials, trading off the costs of later access versus the benefits of reduced risk. Once a product has been subjected to the required clinical trials, it is released to the market, and consumers (doctors and patients) make decisions about which product to use, given the current available choice set and information. Because the regulator weighs the implications for total surplus in it's decision, we begin with the consumers' problem and work backwards.

3.2 Modeling Consumer Choice

In the market, each device's perceived quality and uncertainty can be mapped into choice probabilities and welfare via a utility function that specifies the utility to patient/doctor combination i from using device j at time t (where here subscript t refers to the calendar month, which will be associated with different product age a for different products). We assume that the ex-ante expected (indirect) utility function takes the form

$$u_{ijt} = Q_{jt} - \frac{\rho}{2}\sigma^2_{jt} + \epsilon_{ijt}, \tag{5}$$

where ρ is the coefficient of risk aversion, and ϵ_{ijt} is an i.i.d. error term capturing the deviation of doctor preferences and/or patient appropriateness for device j relative to the population average. In our empirical exercise, we do not find price to be a statistically

significant determinant of demand. Because of this and the fact that we do not explicitly model supply, we leave price out of the specification here.

Assuming consumers choose the product j that maximizes expected utility from the set of products $\mathcal{J}_t$ available, the set of patients for whom a doctor chooses product j (in month t) is then $\mathcal{A}_{jt} := \{i | j = \arg\max_{k \in \mathcal{J}_t} u_{ikt}\}$. Then expected market shares are given by the choice probabilities:

$$s_{jt} = Pr[j = \arg\max_{k \in \mathcal{J}_t} u_{ikt}] = \int_{\mathcal{A}_{jt}} f_t(\epsilon) d\epsilon = \frac{e^{Q_{jt} - \frac{\rho}{2}\sigma_{jt}^2}}{\sum_{k \in \mathcal{J}_t} e^{Q_{kt} - \frac{\rho}{2}\sigma_{kt}^2}}, \tag{6}$$

where the last equality obtains from the standard "logit" assumption that ϵ is distributed i.i.d. extreme value type I with unit variance. The choice set always includes an outside option $j = 0$, with utility normalized to zero.

In a world where doctors and patients make choices with full information, the gains to greater access are unambiguous. However, the fact that product quality is uncertain at the time of regulatory approval and only revealed over time introduces distortions in choices and realized welfare due to lack of information and potentially risk aversion. The realized total surplus (not including fixed costs; in logit utils) over time periods $t = 1, \cdots, T$ will be given by

$$\widetilde{TS} = \sum_{t=1}^{T} \int_{\mathcal{A}_{jt}} u_{ijt} f_t(\epsilon) d\epsilon = \sum_{t=1}^{T} \ln \left(\sum_{j \in \mathcal{J}_t} e^{Q_{jt} - \frac{\rho}{2}\sigma_{jt}^2} \right) , \tag{7}$$

where the final equality obtains from the logit distributional assumption on ϵ.

3.3 Modeling Supplier Pricing and Entry

Fixed costs of market entry (in particular clinical trial costs) are substantial relative to the lifetime profits of the average medical device, making the decision to proceed with testing and launching a new product an important one. Once in the market, quantities are determined via consumer preferences as in the previous Section, and prices for the device are typically negotiated at the hospital or regional level.

We follow work by Grennan (2013) in medical devices and other recent work in negotiated price markets (Crawford and Yurukoglu 2012; Gowrisankaran, Nevo, and Town 2013; Ho and Lee 2013) in modeling prices as the outcome of a Nash Equilibrium of bilateral Nash Bargaining processes (Horn and Wolinsky 1988):

$$p_{jrt} = c_{jrt} + \frac{b_{jt}(r)}{b_{jt}(r) + b_{rt}(j)} \left[TS(\mathcal{J}_{rt}) - TS(\mathcal{J}_{rt} \setminus \{j\}) \right] , \tag{8}$$

where p_{jrt} is the price of product j in region r in month t, c is marginal cost, and b denotes Nash Bargaining weights. The pricing equation says that each product will capture a fraction $\frac{b_{jt}(r)}{b_{jt}(r)+b_{rt}(j)}$ of its marginal contribution $TS(\mathcal{J}_{rt}) - TS(\mathcal{J}_{rt} \setminus \{j\})$ to overall surplus. Because price does not enter demand (an assumption here that is born out in the empirical section), total surplus does not depend on how the surplus is split, making this a transferable utility game.

Given expected prices and demand, each product's manufacturer will only decide to launch the product in a region if the expected profits of the product's expected lifetime A_{jr} will cover the fixed costs (including those associated with clinical trials) of entry $FC_r(T_r^c)$:

$$\text{enter iff} \qquad \sum_{a=1}^{A_{jr}} q_{jrt}(p_{jrt} - c_{jrt}) > FC_r(T_r^c) \ . \tag{9}$$

3.4 Modeling the Regulator's Tradeoffs

The total surplus equation 7 illustrates the main tradeoff between access and risk: the longer time T^c that products spend in premarket clinical trials, the lower the risk from uncertainty about product quality in the market σ_{jt}, but the less new technologies available in the consumer choice set $\mathcal{J}_t$ at any point in time and greater costs of entry. This tradeoff can be formalized mathematically by writing total surplus as a function of time spent in premarket clinical trials and considering the marginal return to increasing the amount of time spent in premarket testing to $T^c + 1$:

$$TS(T^c{+}1){-}TS(T^c) = \sum_{t=1}^{T} \ln\left(\frac{\sum_{j\in\mathcal{J}_t(T^c+1)} e^{Q_{jt}-\frac{\rho}{2}\sigma^2_{jt}(T^c+1)}}{\sum_{j\in\mathcal{J}_t(T^c)} e^{Q_{jt}-\frac{\rho}{2}\sigma^2_{jt}(T^c)}}\right) - FC(1)\left|\mathcal{J}_t(T^c+1)\setminus\mathcal{J}_t(T^c)\right| \ . \tag{10}$$

One way to very clearly see the tradeoff between access and risk as a function of learning is to consider the simplest scenario where there is no observational learning once a product enters the market and where there is no direct cost of premarket testing. In this case, the per-period marginal return to increasing premarket testing simplifies to

$$\frac{TS(T^c+1) - TS(T^c)}{T} = \frac{\rho}{2}(\sigma^2_{T^c} - \sigma^2_{T^c+1}) - \frac{1}{T}\ln\left(\frac{\sum_{j\in\mathcal{J}_0(T^c+1)} e^{Q_{j0}}}{\sum_{j\in\mathcal{J}_T(T^c)} e^{Q_{jT}}}\right) \ , \tag{11}$$

where the first term captures the per period utility gain from decreased risk, and the second term captures the total surplus generated by the rate of technological improvement in product quality over time.

4 Data and Preliminary Analysis of Access/Risk in US and EU

In this Section we introduce the data on product entry, usage, and pricing; and we document patterns in the data consistent with the model in Section 3 and suggesting the EU enjoys greater access to quality new devices, but also greater risk from lower quality products and the approval of products early on when quality is more uncertain.

The data used in this study consists of quantities and prices at the product-hospital-month level, collected by Millennium Research Group's (MRG) *MarketTrack* survey of hospitals across the US and EU from 2004-2013. This survey—covering approximately 10 percent of total market activity—is the main source of detailed market intelligence in the medical device sector, and its goal is to produce representative estimates of the distribution of market shares and prices by region. Though we use the hospital level data for some relevant summary statistics, for the majority of our analysis we aggregate the data to the region (US and EU) level in order to obtain accurate measures of market entry and overall usage of each device within a region.

4.1 The EU has Access to More, Newer Technologies

Table 1 and Figure 1 demonstrate the extent to which the EU enjoys greater access to medical devices than the US. During the time of our sample, the EU has over three times as many manufacturers and products in the market. For those products that eventually enter the US, the average lag time between EU and US introduction is 10 months (19 months for the more technologically advanced DES). Many of the products to which the EU has greater access are important, high-quality products. In the average month, 47 percent of the stents used in the EU are unavailable in the US at that point in time, and 32 percent will never be available in the US.

Table 1: Entry of manufacturers and products across US and EU

	US	EU
Mean manufacturers in market	4	14
Mean products in market	11	32
Total products in market (2004-13)	24	113
Mean months from EU to US entry	10	-
Mean months from EU to US entry (DES)	19	-

Figure 1: EU market share of products not available in US.

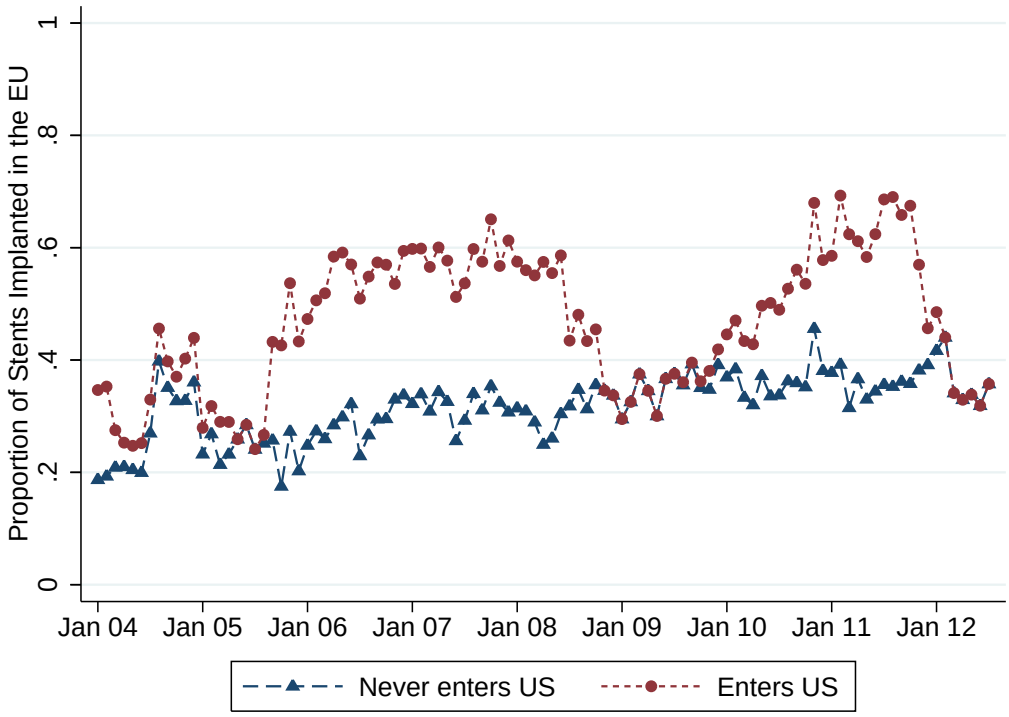

4.2 The EU Also Grants Access to More Technologies with Lower and More Uncertain Quality

Several statistics from the data suggest that the greater access enjoyed by the EU comes along with greater risk in the form of more low quality devices and more uncertainty regarding device quality at the time market access is granted.

The patterns from the data shown in Figure 2 suggest that the EU consumers bear more risk than those in the US by introducing a larger number of devices earlier in their life cycles with less information imparted about the quality of those devices. The left panel shows that in the EU the mean inclusive share, $(\frac{1}{J_a}\sum_j \ln(s_{ja}/s_{0a}))$, of a product is lower upon introduction, and gradually increasing with age until reaching a stable level after about two years in the market, whereas the mean inclusive share is constant with product age in the US. The right panel shows that the variance in inclusive shares across hospitals $(\frac{1}{H}\frac{1}{J_a}\sum_h\sum_j \left(\ln(s_{jha}/s_{0ha}) - \overline{\ln(s_{jha}/s_{0ha})}\right)^2)$ is larger early in a product's life, and gradually decreasing to a stable level over time. Again, this statistic is constant over the product lifetime in the US. Both of these patterns are consistent with greater uncertainty regarding product quality early in the product lifetime in the EU, which is gradually resolved over time via learning. The fact that the mean inclusive share

Figure 2: Evidence of Greater Risk and Learning Upon Market Entry in EU vs. US

(a) Mean Inclusive Share

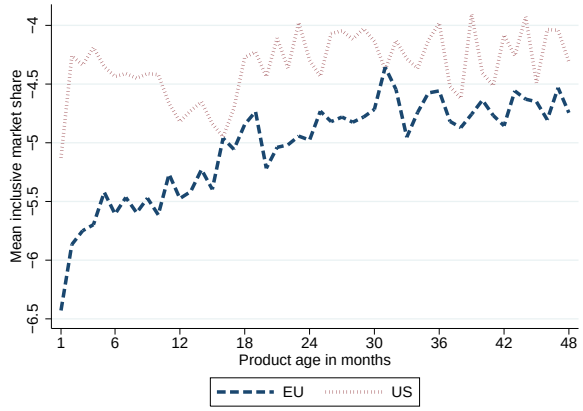

(b) Variance (Across Hospitals)

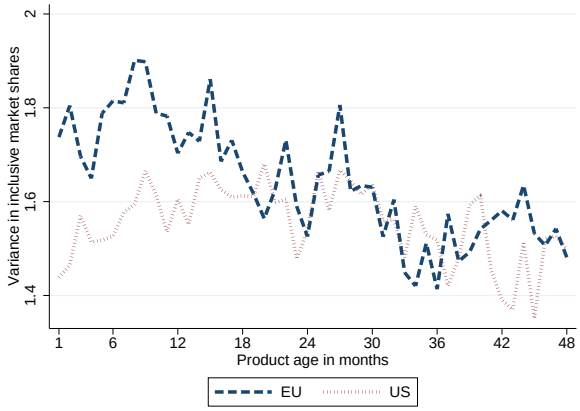

is lower early on further suggests that consumers are risk averse, discounting products whose quality is more uncertain.[11]

4.3 Regulatory Differences Don't Appear to be Driven by Differences in Disease or Treatment in EU vs. US

There is little evidence that these differences in usage patterns are being driven by other factors such as differences in disease incidence, preferences for angioplasty and stents, or price differences across the US and EU. If anything, the evidence and economic logic suggest that the differences documented in Figures 3 and 4 are results of the different market structures induced by the regulatory differences, rather than causes of regulatory differences. For example, the average ischemic heart disease mortality rate is very similar between the US and the EU suggesting that the disease incidences is also similar. Ischemic heart disease is caused by the accumulation of fatty deposits lining the inner wall of a coronary artery and coronary stents are used in treatment of this disease. The 2010 mortality rate in the US for ischemic heart disease was 126.5 deaths per 100,000. The corresponding figure for the EU is 130.0 per 100,000.[12] While there is significant variation across member countries in the EU and across states in the US in the prevalence of heart disease, there is little indication that there are meaningful differences between

[11] An alternative explanation of the patterns over the first few years after EU entry might be a ramping up of distribution and marketing over time. We find this explanation unlikely due to the fact that market entry and the subsequent product rollout is a highly anticipated event by manufacturers and consumers, and also due to the fact that we do not see a similar pattern upon US introduction.

[12] OECD *Health at a Glance, 2013.*

average disease prevalence across the two regions.

Figure 3: Comparison of diagnostic procedure patterns, EU vs. US.

(a) Diagnostic Angiographies

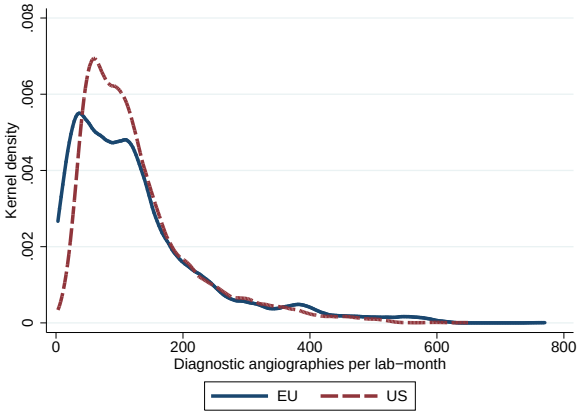

(b) Percent Resulting in Stenting

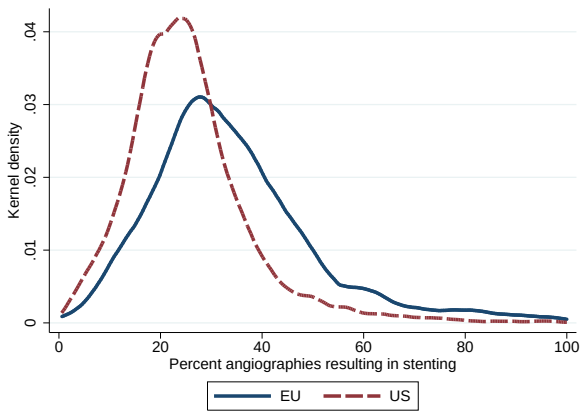

Figure 4: Comparison of usage and price patterns EU vs. US.

(a) DES Share of All Stents

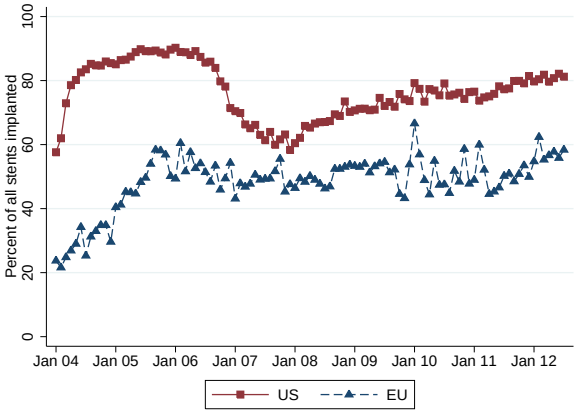

(b) Median Prices

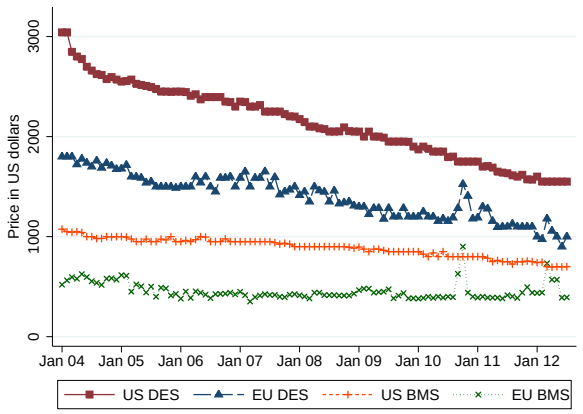

Prior to performing an angioplasty in which a stent may be inserted, the patient must undergo a diagnostic angiography. In this procedure, the blood flow through the coronary artery is visualized and this information is used to determine whether the patient should receive a stent or some other medical intervention. If the difference in the number of different stents available between the EU and the US was driven by higher demand for stents, then it should show up in the data with the EU performing a larger number of angiographies or having a higher rate of stenting conditional on the angiography rate. Figure 3 documents the distributions of the number of diagnostic

angiographies performed across the hospitals in our data and percent of those diagnostic procedures resulting in a stenting procedure across hospitals in the US and EU samples. The distributions are close to identical statistically, with the EU having a few more small volume hospitals and hospitals that are more likely to place a stent conditional upon a diagnostic procedure. In the EU, 32 percent of patients received a stent conditional on an angiography while in the US that figure was 28 percent. This modest differential seems unlikely to account for the stark differences of entry rates between the two regions.

Figure 4 documents that DES usage as a percentage of all stents used lower in the EU but follow similar patterns to the US over time. If the increased DES entry in the EU was driven by higher demand, we would expect the opposite pattern. Figure 4 also shows that the prices and hence profits per stent sold are lower in the EU. This is true for both BMS and DES and is true over our entire sample period. Both of these patterns are likely the result of lower reimbursement levels overall, lower DES reimbursement levels in particular, and more competing devices in the EU market. These results suggest that conditional upon FDA approval, average variable profit in the US is higher making it a more attractive entry target than the EU. This in turn suggests that the differential entry rates is driven by differences in regulation and not underlying demand.

5 Structural Identification, Estimation, and Results

The statistics presented in the previous Section are consistent with the model of regulation and learning developed in Section 3 and suggest that the EU is indeed less stringent than the US in regulating the entry of new medical devices. In this Section we estimate the parameters of our model in order to better understand the impact of this differential regulation. Using the quantity and price data across markets and over time, we estimate the distribution of product quality for innovations that could be introduced in the US and EU, the rates of learning over time, and risk aversion. We then use the model to quantify the welfare generated under different premarket clinical testing requirements (including those observed in the EU and US) and under a proposed alternative policy that would relax premarket requirements but increase the rate of observational learning through increased post-market approval data collection and reporting.

5.1 Demand and Learning Model Estimation

The parameters of the utility function—and by extension the parameters of the device quality distribution and learning process—can be estimated by a revealed preference assumption and data on device market shares in each month. Matching the choice

probabilities implied by utility maximization and the market share data, and inverting the system as in Berry (1994) to recover the mean utility parameters gives

$$\ln(s_{jt}/s_{0t}) = \delta_{jt} := Q_{jt} - \frac{\rho}{2}\sigma_{jt}^2 := Q_j - \frac{\rho}{2}\sigma_{jt}^2 + \xi_{jt} \ , \tag{12}$$

where the unobservable ξ_{jt} in the final equation includes any errors in the current expected quality estimate Q_{jt} as well as any idiosyncratic market preferences. The main challenge here is that none of the variables on the right hand side of this equation are directly observed in the data. Our strategy will be to use variation over time and across products in a region to estimate the product qualities Q_j, the mean Q and variance σ_Q^2 of the product quality distribution, the signal variances σ_A^2 and $\sigma_{A^c}^2$, and the risk aversion parameter ρ.

5.1.1 Identification and Estimation of Demand and Learning

We estimate the parameters via a generalized method of moments algorithm (detailed in Appendix A). A simple and semi-parametric way to estimate equation 12 would be to regress the inclusive shares $\ln(s_{jt}/s_{0t})$ on product and age fixed effects (age fixed effects interacted with whether a product is in clinical trials or not to allow for differential learning rates from trials and observation). The age fixed effects would then capture the combined effect of learning and risk aversion on utility. However, because we are interested in questions that involve altering the observed learning rates, we need to add structure via the learning model to disentangle these forces. Comparison to the fixed-effect model provides a useful benchmark for assessing the fit of the more parsimonious and parametric learning model.

Like all learning models, the identification of the signal variance depends on fitting the model to the *shape* of how choice behavior changes with the age of the product. The risk aversion parameter is then identified as the multiplicative shifter that best fits that shape to the observed choices. Despite the fact that data only exists post market approval, we are able to separately estimate the rates of learning in clinical trials σ_{A^c} and observationally σ_A because we observe all products post market approval in the EU, and a subset of these products are concurrently involved in clinical trials required for eventual FDA approval. The learning and risk parameters are estimated using the within-product variation, as they are all conditional on the product fixed effects whose parameters provide estimates of the product qualities Q_j.

We use the empirical distribution of the product fixed effects estimated from the EU data to estimate the mean and variance of the distribution of product qualities developed. This amounts to an assumption that all products that a firm might want to introduce

to the market are in fact introduced in the EU. This is plausible as the EU has some products with very low market shares and profits that are likely near the threshold at which fixed costs of product development and entry are just covered.

5.1.2 Results of Demand and Learning Estimation

Table 2: Structural parameter estimates of learning model

μ_Q	σ_Q^2	$1/\sigma_A^2$	$1/\sigma_{A^c}^2$	ρ
-5.72	1.37	0.00	0.12	0.58
(0.01)	(0.01)	(0.02)	(0.02)	(0.02)

$N = 3252$. Standard errors clustered at the month level ($N_T = 103$).

The parameters estimates from the model are presented in Table 2. The first observation is that the coefficient of variation on the distribution of product quality, $|\frac{\sigma_Q^2}{\mu_Q}|$ is relatively high at 0.24. There is meaningful underlying variation in product quality that exposes consumers to risk. The second observation is that the learning rates vary according to whether the product is under clinical trial or not. Interestingly, the parameter estimate indicate that there is virtually no experiential market learning occurring. Finally, the implied coefficient of risk aversion is quite sensible. The parameter estimate in Table 2 is not directly interpretable as it is in utility units. However, if we convert that estimate into a dollar equivalent by normalizing the total surplus per stenting procedure to \$50,000 (the estimated dollars in quality adjusted life years from the procedure), then the estimated risk aversion parameter is $\rho_{\$} = 1.4 \cdot 10^{-4}$. This is within the range of estimates of risk aversion in other studies such as Cohen and Einav (2007).

Figure 5: Comparison of estimates from fixed effect and learning models.

(a) Product Quality, Q_j **(b)** Uncertainty Discount, $-\frac{\rho}{2}\sigma_{jt}^2$

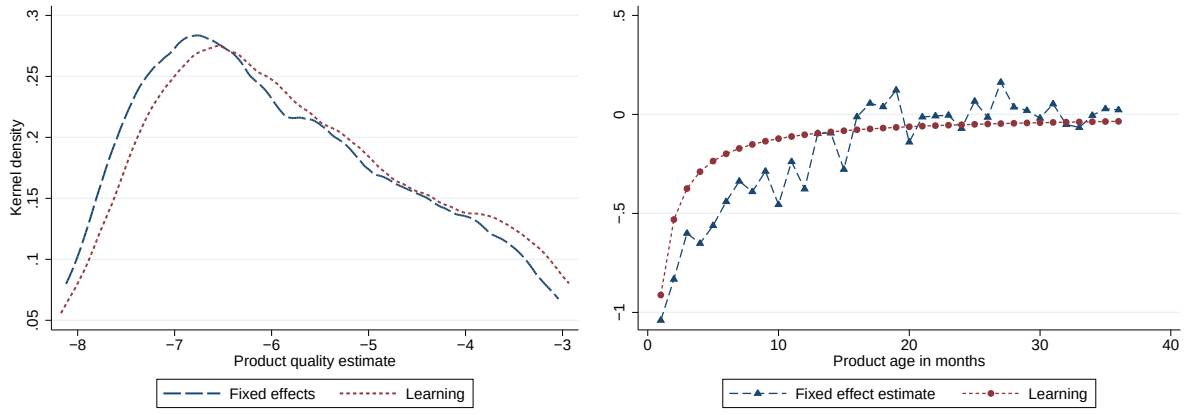

Figure 5 shows the estimated distribution of product qualities Q_j and uncertainty discounts $-\frac{\rho}{2}\sigma_{jt}^2$ for both the learning model and the more flexible model with product and age fixed effects. Despite its parsimony, the simple learning model fits the data nearly as well as the much more nonparametric fixed effects model (RMSE of 0.946 vs. 0.955). Said somewhat differently, our highly parameterized model does not seem to be imposing structure on the problem that is inconsistent with the patterns in the data.

5.2 Pricing and Entry Model Estimation

5.2.1 Results of Pricing ad Entry Model Estimation

Table 3: Structural parameter estimates of supply model

mean TS (\$)	mean AV (\$)	γ_{BMS} (\$)	γ_{DES} (\$)	μ_b	σ_b	FC_{US}
50,000	1303	100	325	0.47	0.30	30,221,774
(6,040)	(2)	-	-	(0.00)	(0.00)	(126,940)

$N = 3252$. Standard errors clustered at the month level ($N_T = 103$).

Table 3 summarizes several estimates from the demand model that are important inputs to supply estimation as well as the supply parameter estimates themselves. Because we find that price does not influence demand, we do not have the standard price coefficient available to scale demand estimates from logit utils to dollars. Instead we take advantage of the fact that like many medical technologies, the procedure of angioplasty with a stent has been subject to numerous studies attempting to value the average quality adjusted life years added by the procedure in dollar terms. We use \$50,000 (published estimates range from \$32,000 to \$80,000) to calibrate the mean total surplus generated per procedure into dollars. Then the marginal contribution (sometimes also called added value) $AV = TS(\mathcal{J}_{rt}) - TS(\mathcal{J}_{rt} \setminus \{j\})$ to be bargained over is a realistic \$1303 for the average stent in the market. Our bargaining parameter estimates indicate that on average the supplier obtains nearly half of this surplus, but there is a great deal of variation with standard deviation of 0.30.

As noted in Grennan (2013), the large added values and final prices make cost estimation difficult because their is very little data near the intercept of the pricing equation. Because we have the advantage of having price data for many more devices, we do obtain some situations with prices in the range of what industry insiders estimate marginal costs to be. Thus we calibrate marginal costs to be equal to the minimum BMS and DES prices observed in the data, respectively. All of our main results and policy implications are robust to marginal costs as low as zero.

Having data on products that enter the EU but not US, and further having EU price

and quantity data, provide us with an especially good setting for estimating fixed costs of US entry. The resulting value of 30 million dollars predicts entry in the data almost perfectly, and matches well with what industry publications estimate to be the costs of various launch phases (in particular of the Makower et. al. (2010) survey that reports the average pivotal trial required by the FDA to cost 1.6 million dollars per month, and our average EU-US entry lag of 10 months).

5.3 Welfare Implications of Regulatory Policy

With the model and estimated structural parameters, we can examine the impact of different regulatory regimes on welfare. The first exercise we perform is examining the optimal regulatory standard for clinical trail length. Figure 6 plots expected total surplus per patient treated in the market, $\sum_{t=1}^{T} \ln\left(\sum_{j\in\mathcal{J}_t(T_c)} e^{Q_{jt}-\frac{\rho}{2}\sigma_{jt}^2(T_c)}\right)$, versus the required length of time spent in clinical testing (relative to the current EU required clinical testing). The results suggest that the optimal tradeoff of access vs. risk is reached at approximately $T_c^* = 5$ months of premarket clinical testing. An interesting feature of the estimated total surplus as a function of time in premarket clinical testing is that it is relatively flat for a wide range of trial lengths near the optimum—for trial lengths between four and twelve months, total surplus remains within five percent of the maximum obtained at five months. Thus the US average of $T_c^{US} = 10$ extra months spent in clinical testing after EU introduction seems too burdensome, although because of the flatness in the total surplus as a function of trial time in this range, the US policy is not statistically different from the optimal.

Figure 6: Estimated Total Surplus as a Function of Time in Premarket Clinical Testing

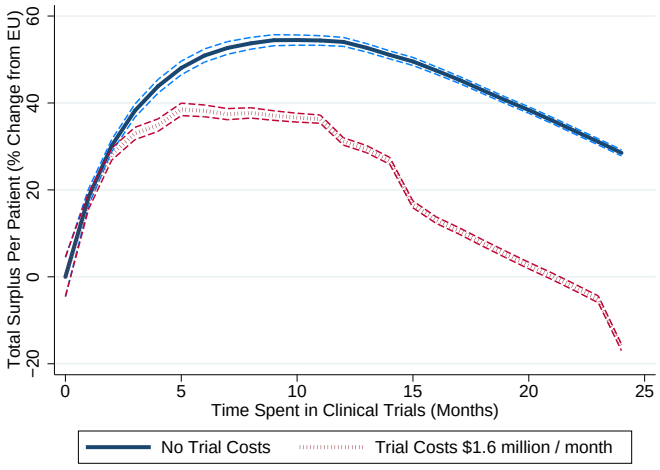

Outside of the flat range, however, surplus drops rapidly with zero month trials and twenty month trials both resulting in an almost 40 percent drop in surplus relative to the optimal. At first this seems to suggest that the EU could make welfare gains of up to 40 percent by increasing its standards—until one realizes that the EU is able to free-ride off of the information being generated in trials for US entry, which makes current EU policy a best-response to current US policy. In effect, the EU is getting free post-approval learning. This issue of "post-market surveillance" and the learning it could induce post-approval has actually been on the policy table in the US, and so in the next section we use our model to conduct a more rigorous examination of its merits.

5.4 Alternative Policy: Shorter Trials with Increased Post-market Learning

We estimated the post market approval observational learning rate is zero for the set of products in our data. There are several potential reasons for the lack of post-market approval learning. For some products, observational learning from real world use (not having the randomization into treatment and control as in a clinical trial) may make it difficult to infer product quality. For other products, though—and likely for those in our sample—the problem is simply a lack of systematic data collection and sharing of information.

One frequently suggested regulatory policy is to relax requirements on premarket clinical trials but increase requirements on post-market surveillance, including data collection, analysis, and reporting. This policy has a direct connection to our model in the sense that it's intention is to increase the rate of post-market approval observational learning—in the language of our model, this means decreasing the variance σ_A^2 of the signals that arrive outside of clinical trials. We analyze this policy by taking the estimated model, varying σ_A^2, and calculating the corresponding optimal trial length $T_c^*(\sigma_A)$ and total surplus generated $TS(\sigma_A, T_c^*(\sigma_A))$. Figure 7 displays the results.

When observational learning is as fast as clinical trial learning, there is no reason to run clinical trials at all, and total surplus is highest–19 percent higher than with no observational learning–because there is no tradeoff to be made between access and learning. As the noise of observational learning increases (relative to clinical trial learning), it becomes optimal to require longer clinical trial periods prior to market access in order to take advantage of the faster learning rate of clinical trials. This transition happens relatively rapidly. Once the noise of observational learning has a standard deviation of three times that of clinical trial learning ($\sigma_A = 3\sigma_{A^c}$), the optimal clinical trial length is seven months and total surplus is only four percent higher than the no observational

Figure 7: Optimal Trial Length and Total Surplus as a Function of Observational Learning

(a) Optimal Trial Length, $T_c^*(\sigma_A)$

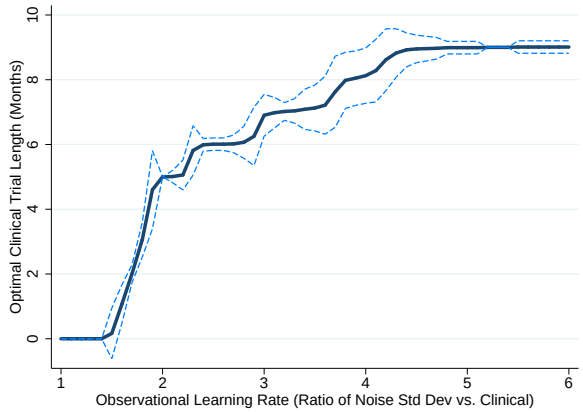

(b) Total Surplus, $TS(\sigma_A, T_c^*(\sigma_A))$

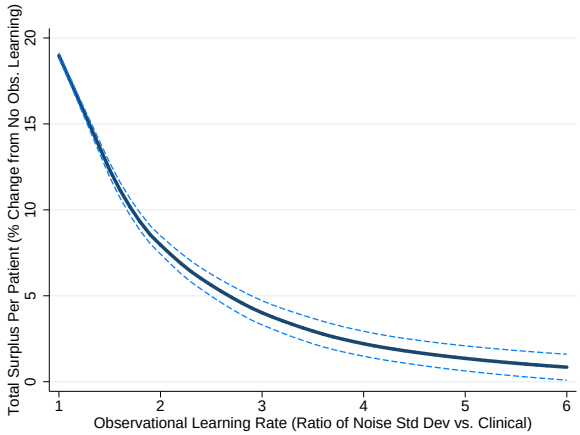

learning case. After this point the gains flatten out, and for $\sigma_A > 4.5\sigma_{A^c}$ the gains from observational learning are close to zero with an optimal clinical trial length of 10 months, the same as in the case with no observational learning.

5.5 Alternative Policy: Allow Manufacturers to Determine Trial Lengths

Why is a regulatory body required to regulate medical device entry? As mentioned previously, the FDA and similar institutions serve functions beyond simply mandating the amount of information (size and length of clinical trials) that must be generated before a product is allowed on the market. Neither our data nor model provide the tools to answer the full question of what the market might look like with no regulation at all. However, we can consider a market where the amount of information generated is a choice variable for each product, rather than for the regulator.

The fully specified game this induces among products is challenging to solve because it involves a continuous choice variable over 113 products. To minimize the computational burden, we begin to look at this problem by fixing the choice of all firms but one focal firm, and looking at the best-response function of that focal firm. What we find is that for higher quality products the business stealing incentive tends to lead to over-investment in learning; whereas for lower quality products the fact that profits are only a fraction of the social surplus created leads to under-investment. Figure **??** shows the results for two representative products.

Figure 8: Optimal Trial Lengths Based on Private Incentives

(a) Higher Quality Product

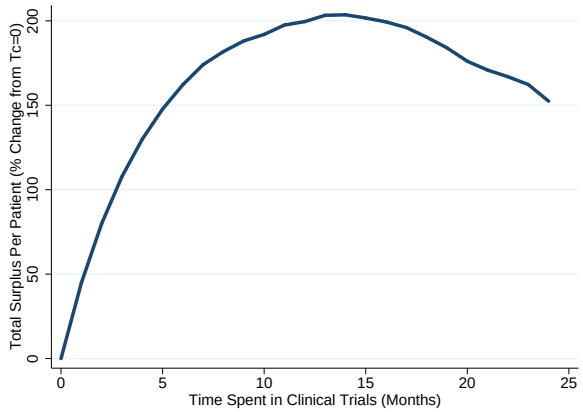

(b) Lower Quality Product

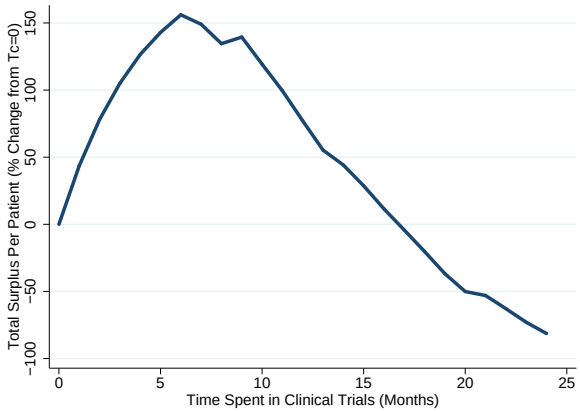

6 Conclusion

The tradeoff between access and risk in regulating the market entry of new products is important in a variety of industries, and in particular in medical devices, where it is an active topic of policy debate in almost every country in the world. In this paper we develop a model with products introduced when quality is still uncertain, learning over time, and regulator (and manufacturer) decisions regarding market entry. We show that the empirical predictions of the model are borne out in market share data in the US and EU medical device markets and are consistent with the beliefs that the US regulatory environment is more restrictive than the EU. We then estimate the structural parameters of the model for use in welfare and counterfactual analysis.

For the set of devices on which we have data, we estimate that both the US and EU are close to the optimal policy (though for the EU depends critically on free-riding off of US trials). We also estimate that if it is possible to achieve post-market learning rates close enough to those we observe from clinical trials at a comparable cost, then embracing recent calls for more active post- market surveillance could further increase total surplus by as much as 19 percent. Relying on private incentives instead of regulator mandated trial lengths tends to lead to over-investment in information among the highest quality products and under-investment among the lowest quality products.

Of course, our analysis is limited in the set of devices for which detailed market data is available, and extrapolating to policy for all devices should be done with care. The theoretical model we develop provides some guidance for how this extrapolation should depend on the uncertainty in quality of new product introductions, the rate of technological improvement, the learning rate in clinical trials, and the observational

learning rate for any type of device being considered.

While estimating the welfare effects of the access/risk tradeoff for an exogenously given set of innovations is an important step towards better understanding this phenomenon, a more complete understanding would allow for the regulatory regime to effect the types of innovations firms develop for the market. A more dynamic analysis of this type would require a significant extension to the theory, and would also require detailed data on innovative activities of the firms in a market.

A Estimation Algorithm Details

A.1 Demand/learning estimation algorithm

1. Compute $\delta_{jt} = \ln(s_{jt}/s_{0t})$ for all product-months.
2. Construct an initial estimator for σ_Q using the empirical equivalent from the distribution of δ_{jt}.
3. Guess an initial value for σ_A.
4. Compute the full vector of ${\sigma_{jt}}^2 = \frac{\sigma_A^2}{a_{jt}{\sigma_Q}^2+\sigma_A^2}{\sigma_Q}^2$.
5. Least squares then gives you an estimator for ρ and the product qualities Q_j as a function of the guess for σ_A, where $[Q_j;\rho](\sigma_A) = (X'X)^{-1}X'Y$ with $X = [1_j, -\frac{1}{2}{\sigma_{jt}}^2]$ and $Y = \ln(s_{jt}/s_{0t})$. (Here Q_j represents the vector of coefficients on product dummy variables, and 1_j the matrix of product dummy variables.)
6. We need to make sure that the distribution of Q_j is consistent with the prior σ_Q by recomputing σ_Q from the current Q_j from 5, and repeating 4-6 until σ_Q converges.
7. Compute the residuals $\xi_{jt} = \ln(s_{jt}/s_{0t}) - Q_j + \frac{\rho}{2}{\sigma_{jt}}^2$.
8. Evaluate GMM objective function based on $E[\xi' Z] = 0$ where $Z = \left[\frac{1}{a_{jt}} \;\; \frac{1}{a_{jt}^2}\right]$.
9. Repeat 4-8 until we find the value of σ_A that minimizes the GMM objective function.

B Direct and Indirect Costs of Clinical Trial Length

When clinical trials have a fixed cost that depends on trial length $FC(T_c)$, then this cost enters total surplus directly, and also indirectly through entry decisions and thus the set of products available in the market $\mathcal{J}_t(T_c, FC(T_c))$. Equation 13 incorporates both of these:

$$TS(T_c) = \sum_{t=1}^{T} \ln\left(\sum_{j\in\mathcal{J}_t(T_c,FC(T_c))} e^{Q_{jt}-\frac{\rho}{2}\sigma_{jt}^2(T_c)}\right) - FC(T_c)|\cup_{t=1}^{T}\mathcal{J}_t \setminus \mathcal{J}_1| \; . \tag{13}$$

The size of the effect of trial costs on total surplus will then depend on the size of the fixed cost of a trial $FC(T_c)$ relative to the distribution of product qualities (and thus profits and total surplus) and the rate at which new products are created.

Because profits of a product will typically depend on the set of other products available in the market, a full specification of pricing and entry models would

Figure 9: Distribution of Profits Over Time and Across Products.

(a) Average Profits Over the Product Lifecycle

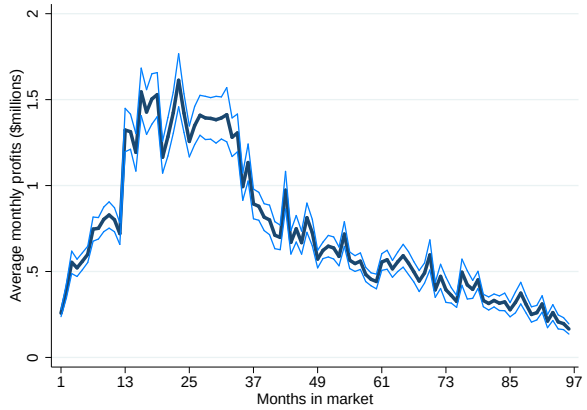

(b) Distribution of Lifetime Profits over Products

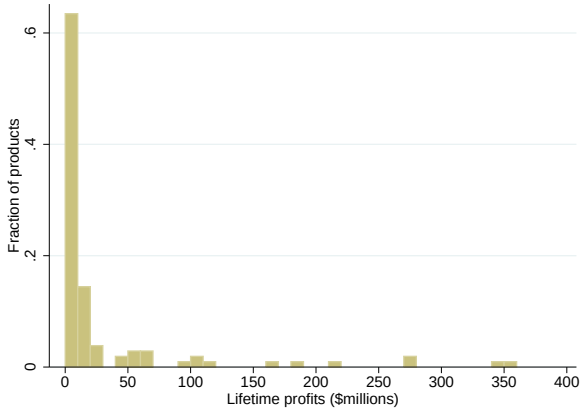

? ? ? ? ? ? ? ? ? ? ?

References

[Ackerberg, 2003] Ackerberg, D. A. (2003). Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination. *International Economic Review*, 44(3):1007–1040.

[Berry and Waldfogel, 1999] Berry, S. and Waldfogel, J. (1999). Free entry and social inefficiency in radio broadcasting. *RAND Journal of Economics*, 30(3):397–420.

[Berry, 1994] Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2).

[Ching, 2010] Ching, A. T. (2010). A dynamic oligopoly structural model for the prescription drug market after patent expiration. *International Economic Review*, 51(4):1175–1207.

[Cohen and Einav, 2007] Cohen, A. and Einav, L. (2007). Estimating risk preferences from deductible choice. *The American Economic Review*, pages 745–788.

[Crawford and Shum, 2005] Crawford, G. S. and Shum, M. (2005). Uncertainty and learning in pharmaceutical demand. *Econometrica*, 73(4):1137–1173.

[Erdem and Keane, 1996] Erdem, T. and Keane, M. P. (1996). Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets. *Marketing Science*, 15(1):1–20.

[Grennan, 2013] Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, 103(1):145–177.

[Miravete et al., 2013] Miravete, E. J., Seim, K., and Thurk, J. (2013). Complexity, Efficiency, and Fairness in Multi-Product Monopoly Pricing. Technical report, Working paper.

[Roberts and Urban, 1988] Roberts, J. H. and Urban, G. L. (1988). Modeling multi-attribute utility, risk, and belief dynamics for new consumer durable brand choice. *Management Science*, 34(2):167–185.

[Timmins, 2002] Timmins, C. (2002). Measuring the Dynamic Efficiency Costs of Regulators' Preferences: Municipal Water Utilities in the Arid West. *Econometrica*, 70(2):603–629.