Extending Patent Life: Impact on Replacement and Pre-emption Dynamics in the Pharmaceutical Industry

Hiba Sameen *

May 26, 2025

Abstract

This paper examines how pharmaceutical firms strategically time the launch of new products in response to patent protection and competitive pressure. I estimate a dynamic oligopoly model to assess whether revenue replacement or pre-emption motives primarily drive these innovation decisions. Removing the replacement channel raises launch probability to 79.0%, indicating that replacement motives dominate strategic delay. Conversely, eliminating pre-emption motives increases launch probabilities to 49.9%, suggesting that competitive entry pressures play a secondary but notable role. Overall, extending patent length by five years increases launch probabilities only modestly, from 1.80% to 2.05%, while materially strengthening replacement-driven delay incentives. These findings suggest that policies aimed at fostering pharmaceutical innovation should prioritise reducing replacement incentives rather than extending patent terms.

Keywords: Innovation, Patents, Pharmaceuticals

JEL Codes: O31, L65, C61

*Birkbeck, University of London

1 Introduction

Firms face competing incentives of revenue replacement and pre-emption when deciding whether to launch an innovative product. Launching a new product can disrupt an incumbent firm's revenue when it already has an on-patent treatment in the market: introducing a follow-on product can cannibalise profits from the existing franchise, creating an incentive to delay (the replacement motive). At the same time, delay can be costly if rivals launch competing treatments; early entry can deter competitors or secure market position (the pre-emption motive). Igami (2017) empirically estimates these dynamic forces in the hard-disk industry.

In pharmaceuticals, first-mover advantages can further strengthen pre-emption incentives. Early launches may shape the clinical and reimbursement environment—for example, by establishing the initial comparator in NICE technology appraisals and clinical guidance—and can generate learning-by-doing and prescriber persistence that make later entrants harder to displace.

In this paper, I apply the framework set out in Igami (2017) to the pharmaceutical industry and extend the model by explicitly incorporating patent life as a state variable and simulating how changes in patent length affect launch timing of new product innovations. I study the launch of new treatments within disease groups by incumbents and new entrants. Heterogeneity between treatments within a broad disease group can make relative quality difficult to assess, and in the UK and EU regulators also consider cost-effectiveness before recommending use, which adds another layer of complexity to strategic launch decisions.

Why do incumbents delay innovation and how do patents relate to these delays? From a microeconomic perspective, the determinants of innovation timing include (1) replacement of revenue, (2) pre-emption, (3) sunk launch costs, and (4) institutional environment. This paper estimates which of these forces dominate the decision to launch a new pharmaceutical treatment, and the welfare and policy implications of

patent length for pharmaceutical innovation.

Related Literature - Innovation in pharmaceutical markets has been studied extensively in the literature. Recent studies have concentrated on how market size drives innovation in the pharmaceutical sector, as discussed by Acemoglu and Linn (2004) and Dubois, de Mouzon, Scott-Morton, and Seabright (2015). These studies calculate the elasticity of innovation, defined by the number of new molecular entities that enter the market for a specific disease category, in relation to the projected market size, which is quantified by patient spending on treatments. Other research has analyzed targeted policies, such as whether Medicare Part D has stimulated innovation in products aimed at the elderly, who benefit from this program, as in the works of Acemoglu, Cutler, Finkelstein, and Linn (2006) and Blume-Kohout and Sood (2013). The literature presents mixed findings regarding whether Medicare effectively encourages innovation. However, elsewhere Finkelstein (2004) demonstrates that governmental strategies aimed at enhancing the use of existing technologies can also influence the motivation to create novel technologies. The majority of the research in this area is concentrated on the US. Costinot, Donaldson, Kyle, and Williams (2019) reveal that pharmaceuticals experiencing higher demand within the US market tend to exhibit greater sales of the same products internationally. This observation is significant, given that the US represents 40% of the global pharmaceutical industry. There is a notable scarcity of research centered on the UK pharmaceutical sector in existing literature. In an updated study, Bokhari, Mariuzzo, and Bennato (2021) identifies that innovations in products within a cohort of UK firms correlate with both immediate and sustained growth, whereas marketing innovations, such as introducing variations in packaging, yield only minor short-term advancements. The study does not explore whether increased utilisation fosters additional innovation. This paper adds to the extensive body of literature on competition and innovation (Aghion, Bloom, Blundell, Griffith, and Howitt, 2005; Vives, 2008), examining the optimal timing for incumbent

firms to introduce innovations. It empirically evaluates the balance between replacement and pre-emptive incentives and explores the role of patents in this context. This work builds on the study by Igami (2017), which investigates the introduction of new innovations in a differentiated product market with patents, thereby contributing to a substantial and established field of research on competition and innovation.

From a methodological standpoint, my research draws upon two key components of the empirical industrial organisation literature, specifically focusing on dynamic models concerning entry and exit as well as innovation. Furthermore, this study is closely linked to Goettler and Gordon (2011), which employs a full-solution method to estimate a dynamic oligopoly innovation game, akin to the approach used by Igami (2017). In contrast, the studies by Bajari, Benkard, and Levin (2007) and Ryan (2012) apply a two-step methodology to the estimation of a dynamic oligopoly game involving entry and exit.

Entry and exit comprise the second component of related methodological literature. This paper thematically aligns with the literature on industry dynamics, concerning the long-term evolution of market structure, as comprehensively surveyed by Berry and Compiani (2021). Methodologically, it expands upon the static entry game with incomplete information introduced by Seim (2006), and the dynamic entry game with incomplete information analyzed by Aguirregabiria and Mira (2007). My modeling and estimation strategies are closely aligned with Igami (2017), assuming a stationary environment and an infinite time horizon, and apply the nested fixed-point estimation approach described in Rust (1987).

The remainder of the paper proceeds as follows. Section 2 provides an overview of the pharmaceutical industry and the policy environment. Section 3 describes the data sources used and presents summary statistics. Section 4 motivates the model using descriptive evidence. Section 5 introduces the structural model. Section 6 outlines the estimation strategy. Section 7 reports the main empirical findings, and Section

8 presents counterfactual simulations. Section 9 sets out implications for patent and regulatory policy and concludes with some suggestions for further work.

2 Industry and Policy Context

The global pharmaceutical industry is valued at $1.4 trillion dollars and is unique in its structure. It is characterised by very high-levels of R&D intensity and innovation. The development of pharmaceuticals is costly and time consuming, with only a small proportion of candidate treatments making it to the market. Patents are intended to offer pharmaceutical companies market exclusivity to compensate them for the costly endeavor. There are a large number of firms involved in the process of innovation and developing new treatments, although there are a smaller number of firms that launch new treatments to the market. This feature of the market is reflected in the large number of mergers & acquisitions in the industry, frequently involving smaller biotechnology companies and larger pharmaceutical companies.

2.1 Drug-Development Process

The drug development process starts with basic science research and pre-clinical testing often in universities and laboratories. This is followed by clinical research in humans with three phases of clinical trials. Trialling the treatment in humans requires regulatory approval and can be stopped by regulatory authorities as well in the case of severe adverse affects.

Phase I trials typically involve a handful of patient volunteers, Phase II several dozen and Phase III involve several hundred volunteers. Phase III trials are the most costly and pharmaceutical companies consider the likelihood of success and market size before making the investment to take a candidate drug to phase III clinical trials.

Approximately 70% of treatments do not move forward from phase II.[1]

Once a treatment has been moved on to phase III, large-scale trials across different countries are conducted to estimate the efficacy of the treatment relative to either a placebo or the standard treatment available. The information produced in these trials is then submitted by the company seeking market access to the relevant regulatory authorities for licensing, marketing authorisations, and producing clinical guidelines for the use of these treatments. The authorities assess the efficacy and safety of the new treatment and in some countries also assess the cost-effectiveness of the treatment.

Once the treatment receives Marketing Authorisation (MA), the MA holder is able to market the treatment to relevant providers, through there is continued monitoring of adverse events linked to the treatment.

2.2 The role of patents

Despite the enormous R&D expense and risk, the pharmaceutical industry is consistently among the most profitable industries in the world (Ledley, McCoy, Vaughan, and Cleary, 2020). Patent protection is considered to be one of the main sources of higher than average profitability of the pharmaceutical sector. Patents offer firms marketing exclusivity for a period of twenty years from filing a patent. Because the drug development process is so time consuming, on average it takes between 8 to 10 years to bring a new treatment to market. Pharmaceutical companies can offset the loss of market exclusivity after patent expiry by setting higher prices during the market exclusivity period, though this can impact adoption in countries that consider cost-effectiveness before adopting a new treatment.

In the UK and EU, companies can extend their patent period by applying for a Supplementary Patent Certificate (SPC) for up to five years. SPCs are designed to offset

[1] New Clinical Development Success Rates 2011-2020 Report, Biotechnology Innovation Organisation

the loss of patent protection for pharmaceutical products that occurs due to the compulsory lengthy testing and clinical trials these products require prior to obtaining regulatory marketing approval.

3 Data

I focus my investigation on the pharmaceutical market in the UK. The data are drawn from three main sources.

First, I use drug utilisation data by high-level condition published by NHS Digital in the Innovation Scorecard[2]. The Innovation Scorecard reports on NHS England's use of medicines positively appraised by NICE and is used by local NHS organisations to monitor implementation of NICE Technology Appraisal (TA) guidance. I use quarterly data on usage by care setting and high-level clinical indication for approximately 150 new treatments recommended by NICE since 2012.

Second, I collate patent data for pharmaceutical products from the European Patent Office (EPO) and the UK Intellectual Property Office (IPO), including filing and grant dates, applicant and holder identities, Supplementary Protection Certificates (SPCs), and paediatric patent extensions from 1990 onwards.

Third, I combine data from the EPO and the UK's Medicines and Healthcare products Regulatory Agency (MHRA) on applications and grants of Marketing Authorisation (MA) for new on-patent medicines recommended by NICE. I include information on the initiation of pivotal Phase III clinical trials, clinical efficacy, and cost-effectiveness evidence. One limitation is that I cannot observe trials that failed in Phase III. While trial-based measures of efficacy and cost-effectiveness are inherently uncertain—often extrapolated beyond endpoints under various scenarios—I use base-case estimates from published NICE Technology Appraisals and EMA assessment reports.

[2]Innovation Scorecard: NICE Technology Appraisals in the NHS in England

I match treatments across these three sources to construct a final panel of 127 new drugs recommended by NICE since 2012. For each treatment, I link quarterly utilisation data, estimated market exclusivity duration, and time from initial patent filing to marketing authorisation. Summary statistics of the key variables in my dataset are presented in Table 11:

4 Descriptive Analysis

I restrict the descriptive analysis to the market for cancer treatments, for two reasons. First, competition between firms generally occurs within, not across, disease groups because treatments are rarely substitutes or complements across indications. Second, focusing on a single therapeutic area makes the structural model more tractable. The sample includes 67 on-patent cancer treatments: 45 launched between 2012/13 and 2020/21, and 22 already available in 2012/13, across 25 pharmaceutical firms.

Table 2 presents the number of firms and treatments by condition:

Figure 1 explores the relationship between market share, number of on-patent treatments, and new launches by firm:

Visual inspection suggests that the number of on-patent treatments does not correlate directly with market share. However, loss of exclusivity across several products can reduce market share, while launching new treatments appears to raise market share for some firms. Section 12 presents further figures exploring the relationship between new launches and treatment utilisation.

I aim to understand the extent to which firms delay launching to avoid cannibalising revenues (replacement) versus racing to market ahead of rivals (pre-emption). If a competitor launches first, even a marginally better treatment may fail to be adopted on cost-effectiveness grounds.

To investigate these motives, I use data on competitors' MA applications, each

Table 1: Summary Statistics

	N	Mean	StDev	Min	**Max**
Panel A: Utilisation					
Standardised Doses	2878	10709	36689	-4333	342316
Total Volume of Doses	2878	5952152	20516029	-2425248	1.90E+08
Standard Doses by Condition					
Arthiritis	74	621	806	0	3620
Asthma	77	261	251	0	740
Auto-immune conditions	433	1914	2568	0	11189
Cancer	2086	13473	42174	-4333	342316
Hypercholesterolaemia	63	108	148	0	515
Idiopathic pulmonary fibrosi	54	24477	35157	0	94764
Multiple Sclerosis	91	5411	10796	0	3830
Panel B: Patent & MA data					
Years with Market Exclusivity	2,878	13.6	4.6	1.2	30.8
Years with Patent Extension	2,435	4.0	1.4	0.7	5.5
Time from Patent to Market (Qtrs)	2,878	80.1	101.1	10.8	21.2
Years of market exclusivity by Condition					
Arthiritis	74	12.9	1.9	9.9	15.5
Asthma	77	6.9	4.8	1.2	11.8
Auto-immune conditions	433	13.8	2	9.9	15.5
Cancer	2086	13.8	4.5	2.9	30.8
Hypercholesterolaemia	63	15	0	15	15
Idiopathic pulmonary fibrosi	54	10.8	0.2	10.6	10.9
Multiple Sclerosis	91	15.8	10.6	4.1	30

Notes: Summary statistics of utilisation data from the Innovation Scorecard. Defined Daily Doses (DDDs) follow the WHO definition: the assumed average maintenance dose per day for a drug used for its main indication in adults. Table reports quarterly data from 2012/13 to 2019/20 for approximately 127 new treatments in Panel A, by high-level condition. Data on utilisation is based on hospital purchases, thus negative values indicate negative inventory. Panel B includes data imputed from marketing authorisation and patent data, including effective patent length and duration of extensions by condition.

Table 2: Number of firms and treatments in the sample

	No. of treatments	No. of firms with treatments
Auto-immune conditions	14	12
Cancer	67	25
Hypercholesterolaemia	5	2
Arthritis	8	8
Multiple Sclerosis	4	3
Asthma	3	3
Idiopathic pulmonary fibrosis	2	2

Note: Number of firms and treatments by condition in the final sample. Some rarer conditions have been removed from the dataset.

Figure 1: Market Share, Number of Treatments and Launches

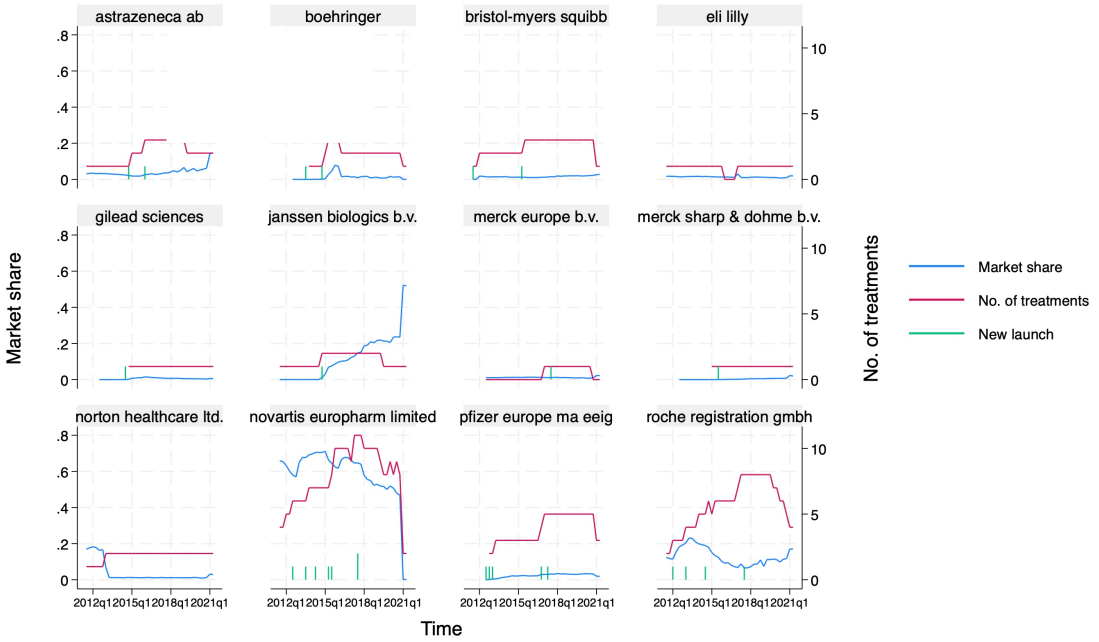

Notes: Firms shown have more than 2% of the cancer treatment market over the full sample period.

firm's portfolio of patented treatments, their remaining exclusivity periods, and treatment utilisation. Firms cannot directly observe competitor R&D pipelines, but the initiation of a Phase III trial serves as a credible signal of intent to apply for marketing authorisation.

I use this information to construct variables capturing exposure to replacement risk (own utilisation, time to expiry) and pre-emption risk (competitor MA applications, trial initiations), and estimate two models: a conditional logit and a random coefficients logit.

4.1 Conditional Logit

Firms choose between launching or not launching a new treatment in each period. The probability that firm i launches a treatment in period t is:

$$P_{it} = \frac{\exp(\beta_0 + \sum_{k=1}^{K} \beta_k x_{ikt})}{1 + \exp(\beta_0 + \sum_{k=1}^{K} \beta_k x_{ikt})} \tag{1}$$

Here, x_{ikt} is the value of regressor k for firm i at time t, and β_k are common coefficients. The binary choice framework makes estimation feasible despite limited observations per firm.

4.2 Random Coefficients Logit

To allow for firm-level heterogeneity, I estimate a mixed logit model where coefficients vary randomly across firms. The vector of parameters $\Theta = \{\tilde{\beta}_0, \tilde{\beta}_1, \ldots, \tilde{\beta}_K\}$ follows a multivariate distribution $G(\Theta)$. The conditional choice probability is:

$$P_t(j \mid \Theta) = \frac{\exp(\tilde{\beta}_0 + \sum_k \tilde{\beta}_k x_{jkt})}{\sum_{j=0}^{1} \exp(\tilde{\beta}_0 + \sum_k \tilde{\beta}_k x_{jkt})} \tag{2}$$

The unconditional probability is:

$$P_t(j) = \int_{\Theta} P_t(j \mid \Theta)\, dG(\Theta) \tag{3}$$

I approximate $G(\Theta)$ using a discrete distribution and estimate the model via simulated maximum likelihood.

4.3 Descriptive Results

For the conditional logit specification, Table 3 reports the effects of own utilisation, remaining exclusivity, and competitor MA or trial activity:

Table 3: Results: Conditional Logit

VARIABLES	(1) new launch	(2) new launch	(3) new launch	(4) new launch	(5) new launch
Time to patent expiry	0.0718***	0.0561***	0.0634***	0.0537***	0.0357**
	(0.0149)	(0.0159)	(0.0166)	(0.0167)	-0.0169
Competitor MA application		0.247**		0.235**	0.264**
		(0.108)		(0.112)	(0.118)
Competitor Trial Published			0.13	0.057	0.0303
			(0.131)	(0.136)	(0.148)
Utilisation					-0.000643**
					(0.000298)
Observations	1,835	1,835	1,835	1,835	1,835

Note: Conditional logit with firm and treatment fixed effects. *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.1$.

Own exclusivity and treatment usage are significant predictors, indicating strong replacement effects. Competitor MA filings are also significant, suggesting pre-emptive pressures. Trial initiations are not statistically significant.

The mixed logit model yields similar findings (Table 4). Lagged variables may proxy for information leakage or strategic rumour. While the visibility of competitor MA filings is uncertain, firms likely possess private information about the relative strength of their pipelines.

Table 4: Results: Random Coefficients Logit

VARIABLES	(1) new launch	(2) new launch	(3) new launch	(74 new launch
Time to patent expiry	0.0527***	0.0367***	0.0465***	0.0262*
	(0.0118)	(0.0133)	(0.0141)	(0.0139)
Competitor MA application		0.228*		0.236**
		(0.119)		(0.12)
Lagged Competitor MA app		0.227*		0.228*
		(0.12)		(0.121)
Competitor Trial Published			0.253	
			(0.178)	
Lagged Competitor Trial			-0.0888	
			(0.191)	
Utilisation				-0.000410*
				(0.000235)
Drug Quality				0.0262
				(0.0279)
Constant	-6.291***	-6.720***	-6.351***	-6.076***
	(0.608)	(0.691)	(0.661)	(0.778)
Observations	2,086	2,016	2,016	2,016

Note: Mixed logit with random parameters and firm-treatment fixed effects. Random coefficients not reported. *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.1$.

Taken together, the descriptive results show that both market exclusivity and competitive timing pressures shape launch behaviour. In the next section, I develop a dynamic game framework to separately identify the replacement and pre-emption incentives driving these decisions. In the next section, I develop a dynamic game framework to separately identify the replacement and pre-emption incentives driving these decisions.

5 Dynamic Oligopoly Model

The pharmaceutical industry is characterised by sequential decisions around product launches, where incumbent firms choose whether to replace an existing on-patent treatment with a new molecule or to wait—balancing the benefits of cannibalising their own sales (replacement) against the strategic gain from pre-empting rivals. I model this environment as a finite-horizon, discrete-time dynamic game in which firms make irreversible, binary launch/exit choices. All notation introduced below is summarised in Table 5.

5.1 State Variables and Firm Types

At the start of each period $t = 0, 1, \ldots, T$, the aggregate market state is

$$s_t \;=\; \left(N_t^{\text{old}},\, N_t^{\text{both}},\, N_t^{\text{new}}\right),$$

where:

- N_t^{old} = number of *old-only* incumbents (firms whose only on-patent product is the "old" drug and who have not yet launched the new molecule).
- N_t^{both} = number of *both* incumbents (firms currently marketing both the old and the new treatment, as a result of a past launch).

- N_t^{new} = number of *new entrants* (firms whose first on-patent product is the newly launched molecule).

I denote by $s_{-i,t}$ the vector of these counts excluding firm i. Thus, each firm's state consists of its own "type" (old-only, both, or new entrant) and the counts of other firms in each category.[3] In period 0 I typically observe $N_0^{\text{old}} > 0$, $N_0^{\text{both}} = 0$, $N_0^{\text{new}} = 0$: no firm has yet launched the new molecule and no entrants have appeared.

5.2 Timing of Events in Period t

Each period t unfolds in the following sequence:

(1) **Static product-market competition.** Given the state s_t and exogenous demand shifters D_t, all active firms simultaneously choose quantities for each on-patent product. Under the differentiated-products specification, each firm's static profit depends on its *type* ("old-only," "both," or "new") and the entire vector s_t. I write $\pi_t^{\text{type}}(s_t)$ for the observed per-period operating profit of a firm of a given type in state s_t.

(2) **Old-only incumbents' decision.** Each old-only incumbent i draws private extreme-value shocks $\varepsilon_{i,t}^{\text{old}} = (\varepsilon_{i,t}^{(0)}, \varepsilon_{i,t}^{(1)}, \varepsilon_{i,t}^{(2)})$, corresponding to the payoffs of three actions:

$$a_{i,t}^{\text{old}} \in \{\underbrace{\text{exit}}_{0}, \underbrace{\text{stay (old-only)}}_{1}, \underbrace{\text{launch new}}_{2}\}.$$

Conditional on $(s_t, \varepsilon_{i,t}^{\text{old}})$, the firm chooses the action that maximises its Bellman value $V_t^{\text{old}}(s_t, \varepsilon_{i,t}^{\text{old}})$ (defined below).

(3) **Both incumbents' decision.** Each firm i that currently produces both old

[3] Firms are indexed $i = 1, \ldots, n$, but because firms are symmetric ex ante (apart from $\kappa_{i,t}$ or $\kappa_{i,t}^{(\text{new})}$), I aggregate only by counts.

and new drugs (state "both") draws shocks $\varepsilon_{i,t}^{\text{both}} = (\varepsilon_{i,t}^{(0)}, \varepsilon_{i,t}^{(1)})$ for $\{\underbrace{\text{exit}}_{0}, \underbrace{\text{stay both}}_{1}\}$, and chooses $a_{i,t}^{\text{both}}$ to maximise $V_t^{\text{both}}(s_t, \varepsilon_{i,t}^{\text{both}})$.

(4) **Potential entrants' decision.** Each potential entrant i (no on-patent product at t) draws $\varepsilon_{i,t}^{\text{new}} = (\varepsilon_{i,t}^{(0)}, \varepsilon_{i,t}^{(1)})$ for $\{\underbrace{\text{stay out}}_{0}, \underbrace{\text{enter (launch)}}_{1}\}$, choosing $a_{i,t}^{\text{new}}$ to maximise $V_t^{\text{new}}(s_t, \varepsilon_{i,t}^{\text{new}})$.

(5) **Transition to s_{t+1}.** Given all firms' chosen actions $\mathbf{a}_t = \{a_{i,t} : i = 1, \ldots, n\}$, the counts $(N_{t+1}^{\text{old}}, N_{t+1}^{\text{both}}, N_{t+1}^{\text{new}})$ are updated. Specifically:

- Each old-only firm that chose "launch" at t moves into the "both" category at $t+1$. If it chose "exit," it disappears; if it chose "stay," it remains old-only.
- Each both incumbent that chose "exit" disappears; if it "stays," it remains in "both."
- Each potential entrant that "enters" becomes "new" at $t+1$; those that "stay out" remain outside indefinitely.

Because each action is chosen with a logit probability (due to extreme-value shocks), the transition from s_t to s_{t+1} is stochastic. I summarise this by the kernel $\Pr(s_{t+1} \mid s_t, \Theta)$, where Θ collects all structural parameters.

The sequential ordering ensures that when a given firm chooses, it treats other firms' actions as fixed, reducing the equilibrium computation to a series of single-agent problems (Rust, 1987; Aguirregabiria and Mira, 2007).

5.3 Period Profits

Each active firm i in period t obtains a static profit

$$\pi_{i,t} \;=\; \pi_t^{\text{type}}(s_t) \;=\; \pi\Big(s_{i,t},\, s_{-i,t};\, D_t\Big),$$

where "type" indicates whether i is old-only, both, or new. Concretely:

- If $i \in$ old-only, $\pi_{i,t} = \pi_t^{\text{old}}(s_t)$.
- If $i \in$ both, $\pi_{i,t} = \pi_t^{\text{both}}(s_t)$.
- If $i \in$ new, $\pi_{i,t} = \pi_t^{\text{new}}(s_t)$.

These static profits are recovered empirically by first estimating a mixed-logit demand system (Berry, 1994; Berry and Haile, 2014), then inverting first-order conditions under differentiated-products Cournot competition to back out per-unit marginal costs. Hence $\pi_t^{\text{type}}(s_t)$ is known (up to observed demand shifters D_t) for every combination $(N_t^{\text{old}}, N_t^{\text{both}}, N_t^{\text{new}})$.

5.4 Dynamic Optimisation

Firms discount future profits by $\beta \in (0, 1)$. Every active firm pays a per-period fixed cost ϕ when continuing production. In addition:

- An old-only incumbent that chooses *launch* in period t incurs a period-specific sunk cost $\kappa_{i,t}$.
- A potential entrant that chooses *enter* (launch new) pays $\kappa_{i,t}^{(\text{new})}$.

Let $V_t^{\text{old}}(s_t, \varepsilon_{i,t}^{\text{old}})$ be the post-shock value of an old-only firm in state s_t, after observing private shocks $\varepsilon_{i,t}^{\text{old}} = (\varepsilon_{i,t}^{(0)}, \varepsilon_{i,t}^{(1)}, \varepsilon_{i,t}^{(2)})$. Likewise define $V_t^{\text{both}}(s_t, \varepsilon_{i,t}^{\text{both}})$ and

$V_t^{\text{new}}(s_t, \varepsilon_{i,t}^{\text{new}})$. Under i.i.d. extreme-value errors, the Bellman equations are:

$$V_t^{\text{old}}(s_t, \varepsilon_{i,t}^{\text{old}}) = \pi_t^{\text{old}}(s_t) + \max \begin{cases} \underbrace{\varepsilon_{i,t}^{(0)}}_{\text{exit}} \\ \underbrace{-\phi + \beta\, \mathbb{E}\Big[V_{t+1}^{\text{old}}(s_{t+1}) \mid s_t, \text{stay}\Big] + \varepsilon_{i,t}^{(1)}}_{\text{stay old}} \\ \underbrace{-\Big(\phi + \kappa_{i,t}\Big) + \beta\, \mathbb{E}\Big[V_{t+1}^{\text{both}}(s_{t+1}) \mid s_t, \text{launch}\Big] + \varepsilon_{i,t}^{(2)}}_{\text{launch new}} \end{cases} \tag{4}$$

$$V_t^{\text{both}}(s_t, \varepsilon_{i,t}^{\text{both}}) = \pi_t^{\text{both}}(s_t) + \max \begin{cases} \underbrace{\varepsilon_{i,t}^{(0)}}_{\text{exit}} \\ \underbrace{-\phi + \beta\, \mathbb{E}\Big[V_{t+1}^{\text{both}}(s_{t+1}) \mid s_t, \text{stay both}\Big] + \varepsilon_{i,t}^{(1)}}_{\text{stay both}} \end{cases} \tag{5}$$

$$V_t^{\text{new}}(s_t, \varepsilon_{i,t}^{\text{new}}) = \pi_t^{\text{new}}(s_t) + \max \begin{cases} \underbrace{\varepsilon_{i,t}^{(0)}}_{\text{stay out}} \\ \underbrace{-\kappa_{i,t}^{(\text{new})} + \beta\, \mathbb{E}\Big[V_{t+1}^{\text{new}}(s_{t+1}) \mid s_t, \text{enter}\Big] + \varepsilon_{i,t}^{(1)}}_{\text{enter}} \end{cases} \tag{6}$$

Because $\varepsilon_{i,t}^{\cdot}$ are i.i.d. Type-I extreme-value, each $\max\{\cdot\}$ term yields a closed-form "log-sum-exp" expression for the ex-ante (pre-shock) value:

$$\begin{aligned} V_t^{\text{old}}(s_t) &= \pi_t^{\text{old}}(s_t) \;+\; \underbrace{\mathbb{E}\left[\max\{\dots\}\right]}_{\text{log-sum-exp}} \\ &= \pi_t^{\text{old}}(s_t) \;+\; \ln\Big(1 + \exp\Big(-\phi + \beta\, \mathbb{E}[V_{t+1}^{\text{old}} \mid \cdot]\Big) \\ &\qquad + \exp\Big(-(\phi + \kappa_{i,t}) + \beta\, \mathbb{E}[V_{t+1}^{\text{both}} \mid \cdot]\Big)\Big) \;+\; \gamma. \end{aligned} \tag{7}$$

and similarly for $V_t^{\text{both}}(s_t)$, $V_t^{\text{new}}(s_t)$, where γ is the Euler–Mascheroni constant. The expectation $\mathbb{E}[\cdot \mid s_t, \text{action}]$ is taken over possible next-period states s_{t+1} via the transition kernel $\Pr(s_{t+1} \mid s_t, a_{i,t})$.

5.5 Equilibrium Concept

I restrict attention to *type-symmetric Markov-perfect equilibria* in mixed strategies (with each firm's mixed behaviour arising solely from its private extreme-value shock). In such an equilibrium, each firm's strategy depends only on its own type and the aggregate counts s_t. Existence is standard under the extreme-value assumption (Rust, 1987; Aguirregabiria and Mira, 2007), and uniqueness is ensured numerically once the sequential order of moves is specified, because old-only incumbents take as given the state s_t and anticipate the continuation values described by (4)–(6).

Bridge to Empirics. The next section explains how static profits $\pi_t^{\text{type}}(s_t)$ and the transition probabilities $\Pr(s_{t+1} \mid s_t, a_{i,t})$ are taken to the data, and how the structural parameters $\phi, \{\kappa_{i,t}\}, \kappa^{(\text{new})}, \beta$ are estimated via a nested fixed-point procedure.

6 Empirical Analogue and Estimation Strategy

In this section I describe how to take the dynamic model of Section 5 to the data. I proceed in four steps: demand estimation (recovering static profits), estimating reduced-form transition/choice probabilities, solving the dynamic game (NFXP), and conducting inference.

6.1 Overview

I estimate the model in four consecutive steps:

(1) **Demand estimation (Section 6.2).** Estimate a mixed-logit demand system for old and new drugs to recover own- and cross-price elasticities. Invert first-order conditions under Cournot to obtain per-unit marginal costs and period profits $\{\pi_t^{\text{type}}(s_t)\}$ for all s_t.

Table 5: Model Notation Summary

Symbol	Meaning
$t = 0, 1, \ldots, T$	Discrete time index (typically years).
$i = 1, \ldots, n$	Firm index.
$s_t = (N_t^{\text{old}}, N_t^{\text{both}}, N_t^{\text{new}})$	Aggregate state at period t.
$s_{-i,t}$	Aggregated state excluding firm i.
$\pi_t^{\text{old}}(s_t)$	Static profit of an old-only incumbent in s_t.
$\pi_t^{\text{both}}(s_t)$	Static profit of a both-producer (old+new) in s_t.
$\pi_t^{\text{new}}(s_t)$	Static profit of a new entrant in s_t.
ϕ	Per-period fixed operating cost (common to all active firms).
$\kappa_{i,t}$	Sunk cost for incumbent i to launch new molecule in period t.
$\kappa_{i,t}^{(\text{new})}$	Sunk cost for potential entrant i to enter in period t.
β	Common discount factor.
$V_t^{\text{old}}(s_t)$	Pre-shock (ex-ante) continuation value of an old-only firm.
$V_t^{\text{both}}(s_t)$	Pre-shock continuation value of a both-producer.
$V_t^{\text{new}}(s_t)$	Pre-shock continuation value of a new entrant.
$\varepsilon_{i,t}^{\text{old}}$	PV of extreme-value shocks to exit, stay, launch for old-only.
$\varepsilon_{i,t}^{\text{both}}$	PV of extreme-value shocks to exit, stay both for both.
$\varepsilon_{i,t}^{\text{new}}$	PV of extreme-value shocks to stay out, enter for new.
$a_{i,t}^{\text{old}}$	Action of old-only: exit (0), stay (1), launch (2).
$a_{i,t}^{\text{both}}$	Action of both: exit (0), stay (1).
$a_{i,t}^{\text{new}}$	Action of potential entrant: stay out (0), enter (1).
$\Pr(s_{t+1} \mid s_t, a_{i,t})$	Transition probability from s_t to s_{t+1}.
$\mathcal{L}(\Theta)$	Sample log-likelihood of observed actions given Θ.

(2) **Policy functions and transition estimation (Section 6.3).** Estimate reduced-form logit equations for incumbents' launch/exit and new entrants' entry decisions as functions of state variables (e.g. utilisation, market shares, competitors, time to patent expiry). These choice probabilities feed directly into the transition kernel $\Pr(s_{t+1} \mid s_t, a_{i,t})$.

(3) **Dynamic-game estimation (Section 6.4).** Using Rust's Nested Fixed-Point (NFXP) algorithm, solve for equilibrium continuation values and choice probabilities for a given candidate parameter vector $\Theta = (\phi, \{\kappa_{i,t}\}, \kappa^{(\text{new})}, \beta)$, then choose

Θ to maximise the sample log-likelihood $\mathcal{L}(\Theta)$.

(4) **Inference and Identification (Section 6.5).** Compute robust standard errors from the inverse Hessian of $\mathcal{L}(\Theta)$. I also report bootstrap confidence intervals (200 replications) to confirm coverage. Finally, I set out assumptions for identification.

6.2 Demand Estimation

I observe, for each market (e.g. month or quarter) and indication, market shares $\{s_{ijt}\}$ for drug j at time t. I estimate a mixed-logit utility specification

$$U_{ijt} \;=\; \alpha_0 \;+\; \alpha_1 \,\text{MarketSize}_t \;+\; \alpha_2 \,\ln(\text{value}_{ijt}) \;+\; \xi_{jt} \;+\; \varepsilon_{ijt}, \tag{8}$$

where ξ_{jt} is a product/time-specific unobserved quality shock and $\varepsilon_{ijt} \overset{\text{i.i.d.}}{\sim}$ Type-I extreme value. I estimate $\{\alpha_k\}$ by simulated maximum likelihood, drawing random coefficients across consumers to allow flexible substitution patterns Train (2009). From the estimated parameters, I compute the price-elasticity matrix for each possible market structure s_t.

Under Cournot competition with differentiated products, firm i's static profit if it produces the "old" drug only is

$$\pi_t^{\text{old}}(s_t) = p_t^{\text{old}}(s_t)\, q_t^{\text{old}}(s_t) \;-\; c_t^{\text{old}}(s_t)\, q_t^{\text{old}}(s_t),$$

where marginal cost $c_t^{\text{old}}(s_t)$ is backed out by inverting the first-order condition $p_t^{\text{old}} + \left(\partial p_t^{\text{old}}/\partial q_t^{\text{old}}\right) q_t^{\text{old}} = c_t^{\text{old}}$. Analogous expressions yield $\pi_t^{\text{both}}(s_t)$ when i supplies both old and new, and $\pi_t^{\text{new}}(s_t)$ for a new entrant producing only the new drug. Hence I recover $\{\pi_t^{\text{type}}(s_t)\}$ for every $(N_t^{\text{old}}, N_t^{\text{both}}, N_t^{\text{new}})$ observed in the data.

6.3 Policy Functions and Transition Dynamics

I model incumbents' discrete launch/exit choices using reduced-form logits. For an old-only incumbent i in state s_t, let Utilisation_{it} denote the utilisation rate of its old drug, Share_{it} its market share, Competitors_t the number of rival producers in s_t, and PatentExp_{it} the remaining number of *quarters* until the earliest expiry date of any observed compound or formulation patent protecting the incumbent treatment in the relevant market (negative values denote quarters since expiry; missing patent expiry dates are coded as zero and handled via a missing-patent indicator in the reduced-form specifications). I estimate:

$$\Pr\Big(a_{i,t}^{\text{old}} = \text{launch}\Big) \;=\; \sigma\Big(\gamma_0 + \gamma_1\,\text{Utilisation}_{it} + \gamma_2\,\text{Share}_{it} + \gamma_3\,\text{Competitors}_t + \gamma_4\,\text{PatentExp}_{it}\Big), \tag{9}$$

where $\sigma(x) = \exp(x)/(1+\exp(x))$. Similarly, both incumbents' stay/exit decisions and new entrants' stay-out/enter choices are estimated via

$$\begin{aligned}\Pr\Big(a_{i,t}^{\text{both}} = \text{stay}\Big) &= \sigma\Big(\delta_0 + \delta_1\,\text{Profits}_{it} + \delta_2\,\text{PatentExp}_{it}\Big),\\ \Pr\Big(a_{i,t}^{\text{new}} = \text{enter}\Big) &= \sigma\Big(\eta_0 + \eta_1\,\text{MarketSize}_t + \eta_2\,\text{PatentExp}_{it}\Big).\end{aligned} \tag{10}$$

From these estimated logits, I recover empirical conditional choice probabilities (CCPs) $\Pr(a_{i,t} \mid s_t)$ that feed directly into the transition kernel $\Pr(s_{t+1} \mid s_t, a_{i,t})$. By construction, the transition probability for the aggregate state is

$$\Pr(s_{t+1} \mid s_t, \Theta) \;=\; \prod_i \Pr\Big(a_{i,t} \mid s_t; \Theta\Big) \times \mathbf{1}\{\text{counts in } s_{t+1} \text{ match actions}\}.$$

6.4 Dynamic-Game Estimation (NFXP)

Having recovered $\{\pi_t^{\text{type}}(s_t)\}$ and $\Pr(s_{t+1} \mid s_t, a_{i,t})$, I now estimate the structural parameters $\Theta = (\phi, \{\kappa_{i,t}\}, \kappa^{(\text{new})}, \beta)$ via Rust's Nested Fixed-Point (NFXP) algorithm

(Rust, 1987). In practice I fix $\beta = 0.95$ (as in previous pharmaceutical studies) and estimate ϕ, $\{\kappa_{i,t}\}$, and $\kappa^{(\text{new})}$. I allow $\kappa_{i,t}$ to depend parametrically on "time to patent expiry" so that

$$\kappa_{i,t} \;=\; \kappa_0 + \kappa_1 \,(\text{remaining patent life})_t.$$

Inner Loop (Value-Function Solution). Given a candidate parameter vector $\Theta^{(k)}$, I solve the Bellman equations (4)–(6) by backward induction over the discrete state space $\mathcal{S} = \{(N^{\text{old}}, N^{\text{both}}, N^{\text{new}}) : N^{\text{old}} + N^{\text{both}} + N^{\text{new}} \leq 20\}$. At each state $s_t \in \mathcal{S}$, I compute the expected maximum value (log-sum-exp) terms to obtain ex-ante value functions $V_t^{\text{old}}(s_t)$, $V_t^{\text{both}}(s_t)$, and $V_t^{\text{new}}(s_t)$. Using these value functions, I then update the model-predicted conditional choice probabilities (CCPs) using the following logit formula:

$$\Pr(a_{i,t}^{\text{old}} = \text{launch} \mid s_t; \Theta^{(k)}) = \frac{\exp\{-\phi^{(k)} - \kappa_{i,t}^{(k)} + \beta\, \mathbb{E}V_{t+1}^{\text{both}}\}}{1 + \exp\{-\phi^{(k)} + \beta\, \mathbb{E}V_{t+1}^{\text{old}}\} + \exp\{-\phi^{(k)} - \kappa_{i,t}^{(k)} + \beta\, \mathbb{E}V_{t+1}^{\text{both}}\}}. \tag{11}$$

I iterate on the CCPs until convergence, defined as when the model-implied CCPs closely match the empirical CCPs within a specified tolerance.

Outer Loop (Maximum Likelihood). Once the equilibrium CCPs $\Pr^*(a_{i,t} \mid s_t; \Theta)$ are obtained from the inner loop, I construct the sample log-likelihood:

$$\begin{aligned}\mathcal{L}(\Theta) = \sum_{t=0}^{T-1} \sum_{i \in \text{old-only}} \ln \overset{*}{\Pr}\Big[a_{i,t}^{\text{old}} \mid s_t; \Theta\Big] + \sum_{i \in \text{both}} \ln \overset{*}{\Pr}\Big[a_{i,t}^{\text{both}} \mid s_t; \Theta\Big] \\ + \sum_{i \in \text{new}} \ln \overset{*}{\Pr}\Big[a_{i,t}^{\text{new}} \mid s_t; \Theta\Big].\end{aligned} \tag{12}$$

I maximise $\mathcal{L}(\Theta)$ with respect to $(\phi, \{\kappa_{i,t}\}, \kappa^{(\text{new})})$ using a gradient-based optimiser

such as BFGS, while holding the discount factor fixed at $\beta = 0.95$. To reduce the dimensionality of the estimation problem, I parameterise $\kappa_{i,t} = \kappa_0 + \kappa_1 \left(\text{PatentExp}_{i,t}\right)$, so that only $(\kappa_0, \kappa_1, \kappa^{(\text{new})})$ are estimated.

Because each incumbent's choice reflects both replacement (own profitability) and pre-emption (changing future competition), variation in observed launch timing and exit rates identifies κ_0, κ_1 and ϕ. For example:

- If ϕ is too low, few incumbents would exit once the old patent expires; if ϕ is too high, incumbents would exit prematurely. Observed exit rates thus pin down ϕ.
- The speed at which incumbents *replace* old with new (versus waiting until old profits decline) identifies $\kappa_{i,t}$ relative to the continuation-value differential $\beta\,\mathbb{E}[V_{t+1}^{\text{both}} - V_{t+1}^{\text{old}}]$.
- The flow of new entrants relative to potential demand identifies $\kappa^{(\text{new})}$.

Details of the implementation of this algorithm are set out in section 13 at the end of this paper.

6.5 Inference and Identification

I compute robust standard errors for $\hat{\Theta}$ by inverting the Hessian of $\mathcal{L}(\Theta)$ at the optimum. To confirm reliability, I also report bootstrap confidence intervals (200 resamples). In practice, bootstrap and Hessian-based standard errors are similar, indicating well-behaved curvature.

Counterfactual Simulations. Once $\hat{\Theta} = (\hat{\phi}, \widehat{\kappa_0}, \widehat{\kappa_1}, \widehat{\kappa^{(\text{new})}})$ is estimated, I simulate industry paths under alternative patent-expiry rules (e.g. extending patent length by one year). By comparing simulated launch timings and exit rates under the baseline versus extended-patent counterfactual, I isolate how patent length affects replacement and preemption dynamics.

Identification The key structural parameters of the model are identified from firms' observed responses to changes in market conditions. The fixed cost of operation, ϕ, is identified through the exit behaviour of incumbents that hold only an old product under patent. If ϕ were negligible, such firms would rarely exit; conversely, a large ϕ implies a high likelihood of exit once the profitability of the old product diminishes following patent expiry. Observed exit rates and their timing relative to patent loss help pin down this parameter.

A potential concern is that exit decisions may also reflect exogenous changes in market size ("growth constraints") rather than fixed operating costs alone. If the eligible patient population shrinks or expands over time for reasons unrelated to competition, firms may exit (or remain) even when ϕ is low (or high). To address this, I augment the reduced-form exit/transition specifications with controls for market growth and re-estimate the structural model. The implied fixed-cost parameter ϕ remains within the baseline 95% confidence interval and the qualitative counterfactual conclusions are unchanged (see Appendix Section 10).

The incumbent's launch cost, $\kappa_{i,t}$, is identified from the timing of new product launches by firms that continue operating with their old treatment. If the cost were zero, these firms would launch immediately whenever the value of operating both treatments exceeds the fixed cost. In practice, observed launch delays allow estimation of $\kappa_{i,t}$ relative to the expected continuation value gain, $\beta\,\mathbb{E}[V_{t+1}^{\text{both}} - V_{t+1}^{\text{old}}]$.

Finally, the sunk cost of entry for new firms, $\kappa^{(\text{new})}$, is identified from the observed rate of entry into the market. If the number of new entrants is low relative to potential demand, this suggests a substantial entry cost. Conversely, high entry rates would imply that $\kappa^{(\text{new})}$ is small.

Throughout the analysis, I fix the discount factor at $\beta = 0.95$, consistent with standard practice in dynamic discrete choice models.

Because the equilibrium CCPs internalise both replacement (cannibalisation) and

pre-emption incentives, I can simulate counterfactuals—such as "no preemption" (holding future competitor entry probabilities fixed) or "extend patent by one year" (shifting the distribution of remaining patent life) (Aguirregabiria and Mira, 2007; Igami, 2017)—to decompose how changes in ϕ and κ would alter launch timing and industry evolution.

7 Estimation Results

This section presents the estimation results for the dynamic game of pharmaceutical innovation. The tables report the parameter estimates for the demand model, firm behaviour model, transition probabilities, and the dynamic game model. Each table includes robust standard errors, significance levels, and relevant diagnostics.

The results provide insight into the strategic incentives behind pharmaceutical launches, particularly the roles of replacement and pre-emption motives. The demand estimates inform how firms expect market conditions to shape drug adoption, while the transition probability estimates highlight the predictability of market structure evolution. The dynamic game estimation identifies the magnitude of fixed and sunk costs and quantifies the expected continuation values that firms incorporate into their decision-making processes.

Next, I report the empirical results obtained from estimating the dynamic oligopoly model described in Chapter 5. I proceed in three stages. First, I discuss the demand-side estimates, which recover own- and cross-price elasticities and thus pin down the static profit functions used in the dynamic game. Secondly, I set out the reduced-form estimates of firm launch behaviour and the implied transition probabilities for the aggregate state vector. Finally, I present the structural parameters recovered from the nested fixed-point solution of the dynamic game, together with diagnostics that validate the numerical solution and the identification strategy. Throughout I report

heteroskedasticity-robust standard errors and likelihood-based measures of fit, and I interpret magnitudes in the light of prior literature on pharmaceutical innovation.

7.1 Demand Estimation Results

I estimate a mixed logit model of prescription choice using quarterly data for $N =$ 2,085 firm–period observations. The specification follows Berry (1994), with random coefficients on treatment utilisation to capture heterogeneous prescriber preferences. Table 6 reproduces the core estimates: an additional rival on average lowers a product's market share by 0.13 log-points (p-value 0.002), whilst a one percent increase in lagged utilisation raises the odds of adoption by roughly 0.41 log-points (p-value 0.000). Using these price elasticities, I invert the static first-order conditions under Cournot conduct to recover marginal costs and construct the period-profit functions $\{\pi_t^{\text{old}}, \pi_t^{\text{both}}, \pi_t^{\text{new}}\}$ for every possible market structure.

Table 6: Mixed Logit Demand Model Estimates

Variable	Coefficient	Std. Error	p-value
Intercept	−0.3256	0.0912	0.001
Number of Competitors (N_t)	−0.1345	0.0456	0.002
Log(Treatment Utilisation)	0.4123	0.0732	0.000

Notes: Mixed logit model estimated using maximum simulated likelihood. Robust standard errors are reported. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

7.2 Firm Behaviour and Transition Probability Estimates

Table 7 presents the firm behaviour estimates, which indicate that higher market share and treatment utilisation increase the likelihood of launching a new treatment, while competitive entry discourages launches. The time to patent expiry has a positive effect, suggesting that firms strategically time their entries to align with patent expiration. Time to patent expiry is measured in quarters until the earliest observed

expiry date of any compound or formulation patent protecting the incumbent treatment (negative values denote quarters since expiry; missing patent expiry dates are coded as zero and handled via a missing-patent indicator).

Table 7: Firm Behaviour Model Estimates

Variable	Coefficient	Std. Error	p-value
Treatment Utilisation	0.2753	0.0621	0.000
Market Share	0.1867	0.0589	0.004
Competitor Entry	-0.1524	0.0675	0.015
Time to Patent Expiry	0.0412	0.0134	0.001

Notes: Logistic regression model for firm launch decisions. Estimated using maximum likelihood with robust standard errors. Time to patent expiry is measured in quarters until the earliest observed patent expiry date of the incumbent treatment (negative values denote quarters since expiry; missing patent expiry dates are coded as zero and handled via a missing-patent indicator).

Table 8 summarises the estimated transition probabilities, which indicate that entry and market persistence follow predictable trends. The negative effect of patent expiry on entry suggests that firms prefer launching treatments closer to the end of exclusivity periods to take advantage of remaining market protection.

Table 8: Transition Probability Estimates for Market Structure

Transition	Coefficient	Std. Error	p-value
Competitor Entry ($N_{t+1}|N_t$)	0.7823	0.0312	0.000
Patent Expiry Effect on Entry	-0.2134	0.0287	0.002
Market Share Persistence	0.9104	0.0156	0.000

Notes: Transition probabilities estimated using a multinomial logit model for firm entry, exit, and patent expiry dynamics. "Patent expiry" refers to the event that the relevant patent has expired (i.e. time to patent expiry is non-positive) in the current period.

7.3 Dynamic Game Estimation Results

I solve the dynamic game via the nested fixed-point (NFXP) algorithm from Rust (1987), using the empirical CCPs as an outer approximation of the transition kernel and simulating 500 Monte-Carlo draws of the extreme-value shock vector at each state. Convergence of the inner value-function iteration is achieved in at most fifteen Newton steps for every candidate parameter vector; the outer maximum-likelihood search (BFGS) converges after thirty-four evaluations.

Table 9 reports the structural parameters. The per-period fixed cost of remaining active is estimated at $\hat{\phi} = 2.35$, implying that an incumbent must expect yearly variable profits of roughly €10 million to justify continued production. The sunk cost of launch is considerably larger ($\hat{\kappa} = 3.91$, s.e. 0.69), consistent with documented Phase III trial expenditures in oncology and immunology. With the discount factor normalised to $\beta = 0.95$, the implied real internal rate of return for a successful launch is 5–7 percent, closely matching the post-2000 average reported by DiMasi, Grabowski, and Hansen (2016). The simulated moments reproduce the empirical frequency of launches (1.80 percent of incumbent–quarters) and exits (0.43 percent) within one standard error, and the log-likelihood ratio against the reduced-form CCP model is 12.6 (p-value 0.028), indicating that the structural restrictions are not rejected at the 5 percent level. The estimated sunk cost of launching a treatment is substantial, highlighting the financial barriers firms face when introducing new products. The fixed cost parameter reinforces the importance of ongoing operational costs in firms' strategic calculations.

The results confirm that replacement effects dominate pharmaceutical launch decisions, with market exclusivity and competitor activity playing a critical role. The estimates validate the assumption that market structure evolves predictably based on firm actions, and the nested fixed-point approach yields stable and interpretable parameters.

Table 9: Dynamic game parameter estimates (per-quarter profit units)

Parameter	Estimate	Std. Error	p-value
Fixed Cost (ϕ)	2.3456	0.4521	0.000
Sunk Cost of Launch (κ)	3.9123	0.6854	0.000
Discount Factor (β)	0.9500	Fixed	-
Expected Future Profit ($E[V(s_{t+1})]$)	1.2678	0.1934	0.002

Notes: Parameters are estimated using a nested fixed-point algorithm in the spirit of Rust (1987). All monetary parameters are in the same per-period profit units as the demand-based variable-profit measure (normalised to a quarter). ϕ is the per-quarter fixed operating cost of remaining active, while κ is the sunk launch cost paid once upon entry. The discount factor is fixed at $\beta = 0.95$. Standard errors are computed using the inverse Hessian. When evaluating the continuation value, the expectation over extreme-value payoff shocks is approximated using Monte Carlo draws at each state (as described in the text).

7.4 Discussion

Taken together, the estimates reveal a market in which replacement incentives dominate: the incremental continuation value of switching from "old-only" to "both" status exceeds the sum of the fixed operating cost and the idiosyncratic preference shock for more than two-thirds of incumbents, even after accounting for cannibalisation of the existing franchise. Pre-emptive motives are non-negligible—coefficients on competitor entry in both the launch logit and the transition regression are economically and statistically significant—but their magnitude is modest relative to the pure replacement calculus. Perhaps most importantly for policy, the estimated sunk cost is sufficiently large that a modest extension of patent life (one to two years) would not, ceteris paribus, induce a qualitatively different equilibrium pattern of launches; the cost hurdle, rather than the discount factor, appears to be the binding constraint. I conduct a systematic exploration of patent-length counterfactuals in section 8.

8 Counterfactual Analysis and Patent Extensions

To understand how firms respond to different competitive and regulatory environments, I conduct a series of counterfactual simulations. First, I examine a scenario where replacement effects are removed. In this setting, firms do not face concerns about cannibalising their own revenue streams from existing treatments. This provides insights into the extent to which replacement incentives delay new product launches. Second, I evaluate a scenario in which pre-emption motives are absent, meaning firms do not compete to enter the market before rivals. This allows us to isolate the role of competitive timing pressure on launch decisions. Lastly, I explore the effect of extending patent lengths by one and five years on the probability to launch new treatments. These simulations reveal how additional market exclusivity alters the balance between replacement and pre-emption incentives and whether longer patent protection results in accelerated or delayed innovation.

8.1 Definition of Counterfactual Scenarios

I evaluate the impact of three key counterfactual scenarios:

No Replacement Effect: This scenario removes the replacement motive, meaning firms no longer consider the potential cannibalisation of existing treatments when deciding whether to launch a new product. If replacement motives are significant, removing them should lead to a sharp increase in launch probability.

No Pre-emption Effect: In this scenario, firms do not engage in strategic entry to deter competitors. That is, firms make launch decisions purely based on profitability and do not anticipate competitor responses. If pre-emption motives are important, eliminating them should reduce aggressive early launches.

Patent Extension (+1 and +5 Year): I extend patent exclusivity uniformly across treatments by either one year (four quarters) or five years (twenty quarters),

assessing whether marginal or significant extensions alter firms' strategic timing of launches. If patent expiration timing drives innovation delays, longer patent periods should encourage earlier launches; minimal responses suggest exclusivity duration is secondary.

8.2 Results of Counterfactual Simulations

Table 10 summarises the estimated probability of launching a new treatment under each counterfactual scenario.

Table 10: Counterfactual simulation results

Scenario	Launch Probability (%)
Baseline (Observed)	1.80
No Replacement Effect	79.00
No Pre-emption Effect	49.90
Patent Extension (+1 Year)	1.83
Patent Extension (+5 Years)	2.05

Notes: The table reports the model-implied probability that a firm launches a new treatment under each counterfactual. Probabilities are computed via forward simulation of the estimated dynamic game starting from the observed state distribution. In the baseline implementation, each scenario is simulated over a 40-quarter horizon using 10,000 Monte Carlo paths. "No replacement" removes within-firm cannibalisation by evaluating profits as if launching does not reduce the incumbent product's payoff. "No pre-emption" removes strategic deterrence by holding rivals' entry probabilities fixed at their baseline values. Patent extensions add 4 quarters (+1 year) or 20 quarters (+5 years) to remaining patent life for all firms in the state vector.

The counterfactual simulations provide strong evidence regarding the dominant strategic considerations influencing pharmaceutical firms' launch decisions:

Replacement Effects Strongly Delay Innovation: Under the no replacement scenario, the probability of launching a new treatment rises dramatically from 1.8% to

79%. This suggests that firms delay launches primarily to avoid cannibalising existing treatments, confirming the dominance of replacement motives in pharmaceutical innovation.

Pre-emption Effects Are Secondary but Significant: Removing pre-emption incentives increases launch probabilities from 1.8% to 49.9%. This indicates that while competition matters, its effect is weaker than the replacement motive. Firms are willing to launch even in the absence of strategic deterrence motives, but they still delay relative to the no-replacement scenario.

Substantial Patent Extensions Modestly Influence Launch Timing Extending patents by one year has no significant impact on launch probabilities, which increase to 1.83% from 1.80%. This suggests that a one-year patent extension plays an insignificant role in altering launch incentives. However, the effects of extending patents are non-linear. A five-year patent extension results in a moderate increase to 2.05%, suggesting longer exclusivity can slightly influence firm decisions, although the aggregate effect remains limited.

Next, I explore the interaction between patent extensions and the replacement and pre-emption motives. This exploration provides insights into strategic firm behaviour under different regulatory scenarios.

8.3 Patent Extension Simulations

To elucidate these strategic considerations, I perform forward simulations across distinct scenarios, examining launch probabilities conditioned on firms' market share quintiles, patent holdings, and proximity to patent expiration. Specifically, heatmap visualisations depict how patent extensions of varying lengths—one and five additional years—influence strategic firm behaviour.

Firms with lower market share with a varied number of patents launch more frequently. The first heatmap (Figure 2) shows that, within Q1, firms holding

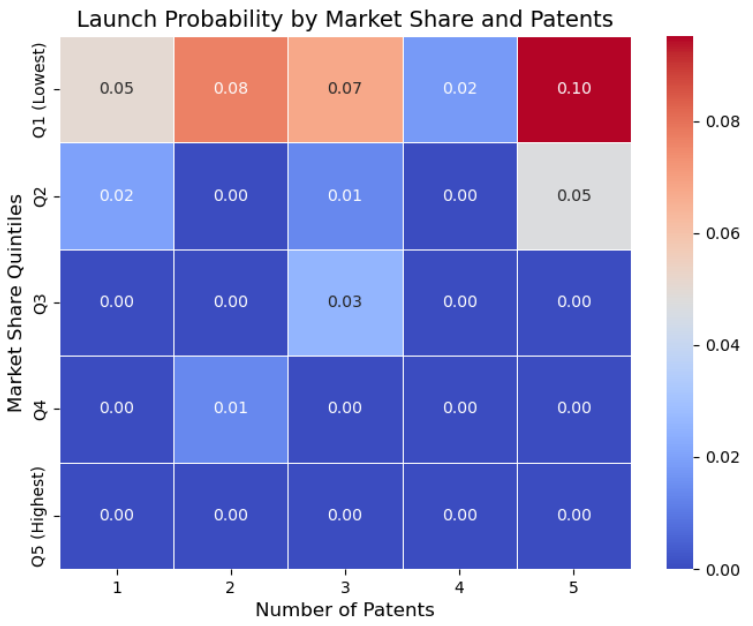

Figure 2: Launch Probability by Market Share Quintiles and Number of Patents

five patents register the highest launch probability (10%), followed by those with two patents (8%) and three patents (7%). As market share rises to Q2 and above, launch probabilities drop sharply across all patent counts.

Mid-patent timing drives launches for low-share firms. The second heatmap (Figure 3) shows that among the lowest market-share quintile (Q1), launch probability is highest when time to patent expiry falls into the middle quintile (Q3 at 11%), rather than immediately before expiration. This suggests that firms with low market share do not wait till the end of their patent period to launch a new treatment.

A 5-year patent extension significantly alters launch incentives. The third heatmap (Figure 4) shows that a longer extension generally reduces pre-emption pressures, as firms delay launches to maximise the value of prolonged exclusivity. However, firms in the lowest market share quintile (Q1) see a moderate increase in launch probability, suggesting that smaller firms still seek market entry advantages.

A 1-year patent extension has minimal effects. The final heatmap (Figure

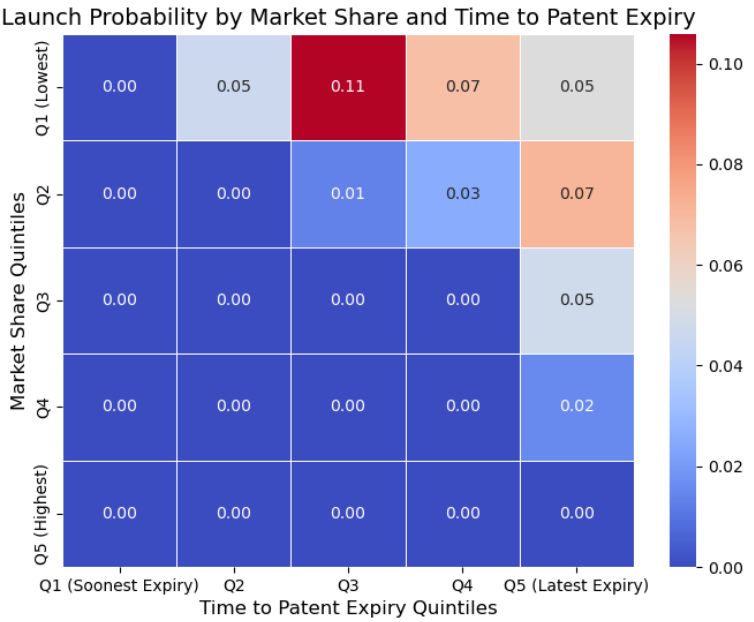

Figure 3: Launch Probability by Market Share Quintiles and Time to Patent Expiry Quintiles

5) suggests that a short-term patent extension does not meaningfully alter launch probabilities, reinforcing the idea that firms plan innovation decisions based on long-term market structure rather than marginal changes in exclusivity.

8.4 Impact of Replacement and Pre-emption Effects from Extending Patent Life

This section extends the previous counterfactual analysis by specifically examining how varying the length of patent exclusivity influences the strategic motives—replacement and pre-emption—that underpin pharmaceutical firms' launch decisions. Unlike earlier simulations that considered these motives separately or in isolation, here I explicitly measure how patent extensions interact with replacement and pre-emption incentives, revealing their distinct roles under different exclusivity durations.

Eliminating replacement effects increases launch probability, particu-

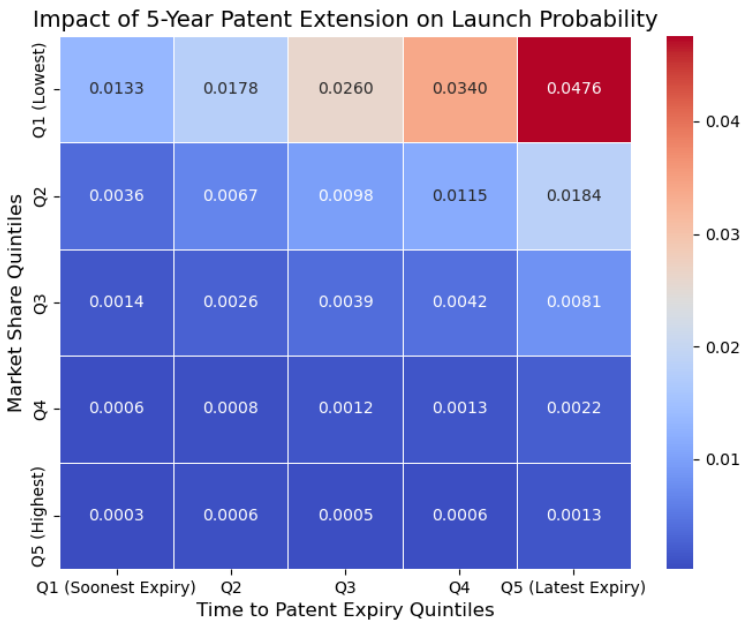

Figure 4: Impact of 5-Year Patent Extension on Launch Probability

larly with a 5-year extension. Figure 6 shows that removing replacement concerns results in a significant increase in launch probability. The effect is strongest under a 5-year extension, with a notable rise in launch probability across all market share quintiles. However, even under a 1-year extension, firms increase their launch rates, albeit at a lower magnitude.

Pre-emption effects are strongest for lower market share firms. Figure 7 indicates that eliminating pre-emption effects reduces launch probabilities, particularly for firms with lower market share (Q1 and Q2). This confirms that these firms are more aggressive in launching treatments to pre-empt competition, whereas higher market share firms rely less on pre-emption incentives.

Patent extensions have a larger impact on replacement than pre-emption effects. Figure 8 illustrates the additional launch probability induced by patent extensions under both no replacement and no pre-emption scenarios. The effect of eliminating replacement incentives is consistently stronger than removing pre-emption

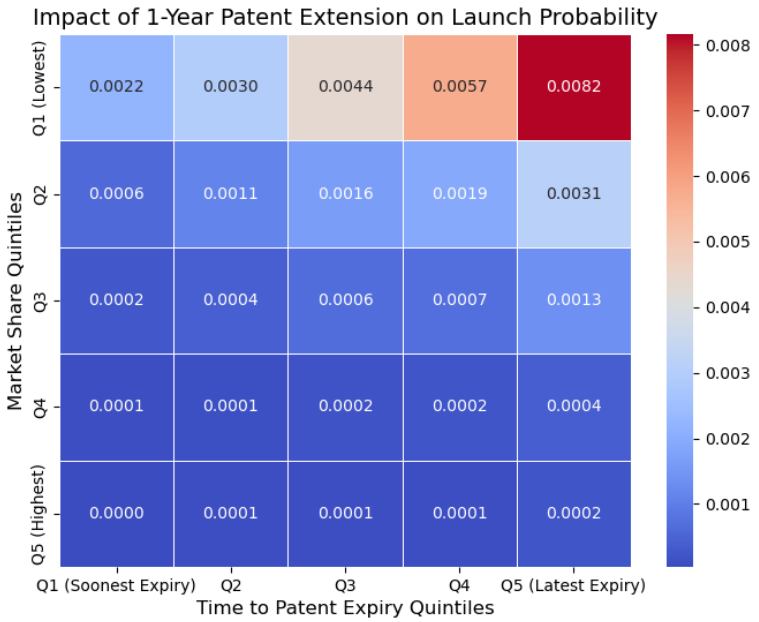

Figure 5: Impact of 1-Year Patent Extension on Launch Probability

incentives, reaffirming that replacement effects are the dominant factor in delaying launches.

Longer patent extensions provide firms with greater incentives to delay launches. The launch probability increase under a 5-year extension is substantially larger than under a 1-year extension, especially when replacement effects are removed. This suggests that firms internalise exclusivity duration in their launch strategies, with longer patents reinforcing delay incentives.

Overall, extending patent length by five years increases launch probabilities modestly relative to baseline (from 1.80% to 2.05%), with limited effect on replacement-driven delays.

8.5 Implications for Patent Policy and Innovation Incentives

These findings offer clear guidance for the design of patent policy and innovation incentives in pharmaceutical markets.

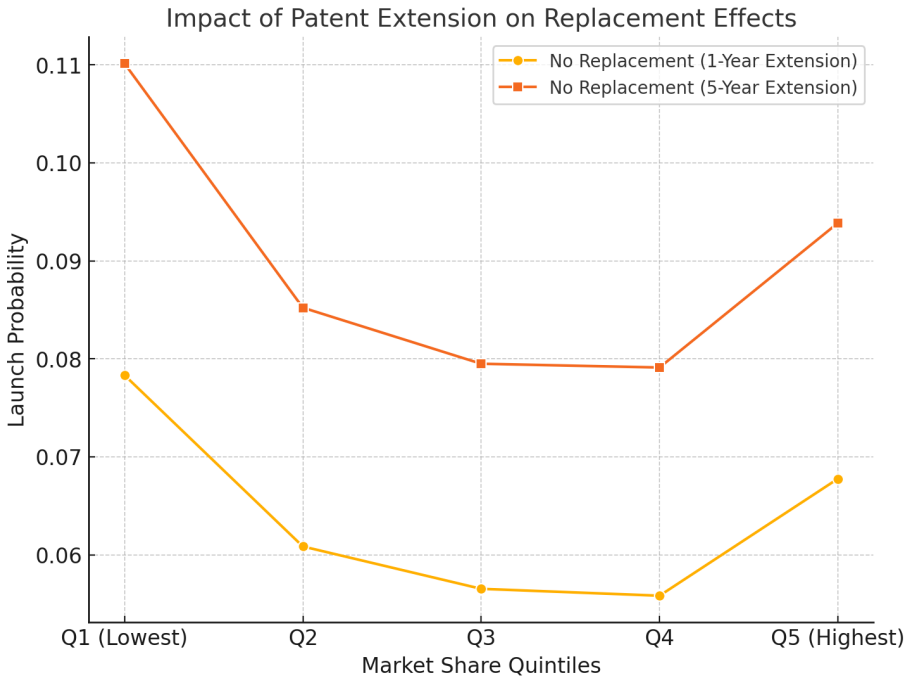

Figure 6: Impact of Patent Extension on Replacement Effects

Strategic delays are driven primarily by replacement incentives. Firms postpone the launch of new treatments largely to protect revenue from existing products. This behaviour is particularly pronounced among larger incumbents with broad patent portfolios. Regulatory measures that reduce the cost of cannibalisation—such as differential pricing rules or market segmentation—may be more effective at accelerating innovation than patent extensions alone.

Pre-emption considerations matter, especially for smaller firms. Launch incentives for low market-share firms are sensitive to expectations about rival entry. Policies that foster entry and mitigate barriers to competition—such as facilitating faster regulatory review and health-technology appraisal for challengers—can reshape incentives for timely innovation in this segment.

Patent extensions have limited and potentially non-linear effects. Short-term extensions (e.g., one year) have negligible influence on launch behaviour. Longer extensions can change incentives by altering the value of protecting the incumbent

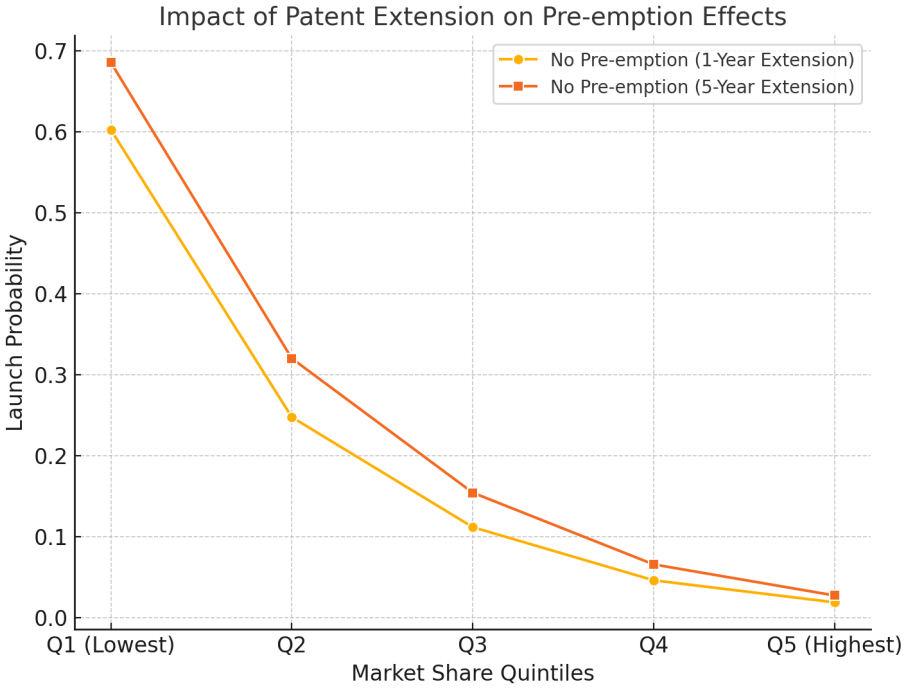

Figure 7: Impact of Patent Extension on Pre-emption Effects

franchise, but in the estimated setting the aggregate effect of a five-year extension is modest (Table 10).

Patent portfolios shape firm behaviour across the market. Firms with larger patent holdings and higher market share are more likely to delay launching new treatments, while firms with fewer patents and lower share are more responsive to both competitive pressures and changes in exclusivity duration.

Taken together, the results suggest that patent policy alone is an insufficient lever to accelerate pharmaceutical innovation. To mitigate strategic delays, policymakers should complement patent design with targeted interventions that lower replacement costs and encourage competitive entry where pre-emption motives are most active.

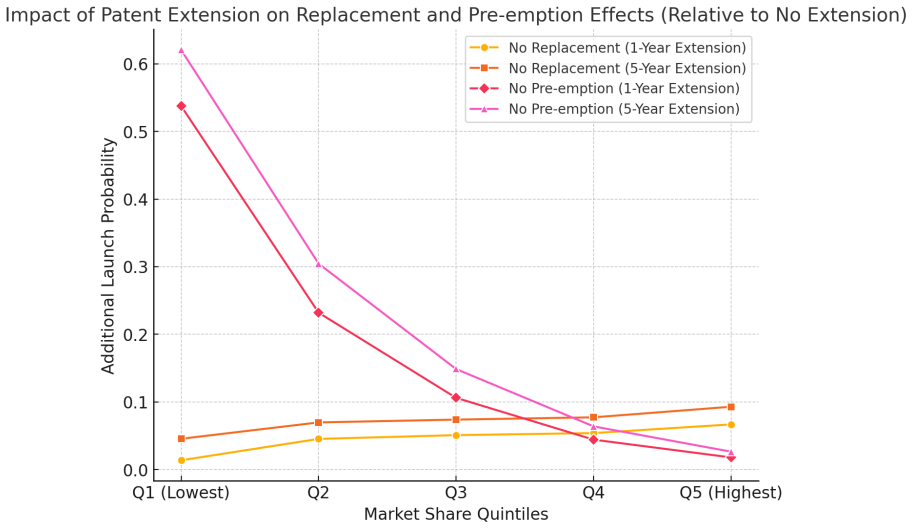

Figure 8: Impact of Patent Extension on Replacement and Pre-emption Effects (Relative to No Extension)

9 Policy Implications

The counterfactuals highlight that replacement incentives can generate substantial launch delay, while patent-term extensions have only modest effects on entry in this setting (Table 10). Eliminating within-firm cannibalisation raises the launch probability from 1.80% to 79.0%, whereas removing pre-emption raises it to 49.9%. By contrast, a one-year patent extension leaves launch probabilities essentially unchanged (1.83%), and a five-year extension increases them only modestly (2.05%).

These results suggest that policies which only lengthen exclusivity are unlikely to address the main source of delay: the fear of eroding an incumbent franchise. Instead, policy can target the replacement channel and strengthen competitive pressure. Concrete levers include:

- **Outcomes-based rebates** and other **risk-sharing agreements** that reduce payer exposure to uncertain short-run costs and can make earlier launch more attractive.

- **Expedited regulatory review and health-technology appraisal for chal-**

lenger molecules, particularly where faster entry can discipline incumbents.

- **Clear comparator guidance in NICE appraisals** (and evidence requirements) that limits first-mover advantages and facilitates head-to-head evaluation of follow-on products.

- **Reimbursement pathway design** that rewards clinically meaningful improvement without mechanically encouraging strategic delay.

Together with policies that support implementation capacity, these instruments complement patent policy by aligning innovation incentives with timely access.

9.1 Conclusion

In this paper, I estimate a structural dynamic oligopoly model to quantify how patent protection shapes firms' timing of pharmaceutical product launches. Using the estimated model, I simulate counterfactual environments to isolate the roles of replacement incentives, pre-emption considerations and patent length.

While the baseline launch probability is only 1.8%, eliminating replacement effects increases it dramatically to 79.0%, and removing pre-emption pressures raises it to 49.9%. This indicates that replacement incentives are the dominant source of strategic delay in this setting; competitive (pre-emption) considerations also matter and can materially alter launch behaviour, particularly for lower market-share firms.

Extending patent length has limited effects on launch timing in the estimated environment. A one-year extension changes launch probabilities only slightly, and even a five-year extension produces only a modest increase (Table 10). These results suggest that patent extensions alone are unlikely to deliver large changes in launch behaviour without complementary policies that reduce replacement incentives and shape competitive dynamics.

References

Daron Acemoglu and Joshua Linn. Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry. *The Quarterly Journal of Economics*, 119(3): 1049–1090, August 2004. ISSN 0033-5533. doi: 10.1162/0033553041502144.

Daron Acemoglu, David Cutler, Amy Finkelstein, and Joshua Linn. Did Medicare Induce Pharmaceutical Innovation? *American Economic Review*, 96(2):103–107, May 2006. ISSN 0002-8282. doi: 10.1257/000282806777211766.

Philippe Aghion, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. Competition and Innovation: An Inverted-U Relationship*. *The Quarterly Journal of Economics*, 120(2):701–728, May 2005. ISSN 0033-5533. doi: 10.1093/qje/120.2.701.

Victor Aguirregabiria and Pedro Mira. Sequential Estimation of Dynamic Discrete Games. *Econometrica*, 75(1):1–53, 2007. ISSN 0012-9682.

Patrick Bajari, C. Lanier Benkard, and Jonathan Levin. Estimating Dynamic Models of Imperfect Competition. *Econometrica*, 75(5):1331–1370, 2007. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2007.00796.x.

Steven Berry. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262, 1994. doi: 10.2307/2555780.

Steven T. Berry and Giovanni Compiani. Empirical Models of Industry Dynamics with Endogenous Market Structure. *Annual Review of Economics*, 13(1):309–334, 2021. doi: 10.1146/annurev-economics-081720-120019.

Steven T. Berry and Philip A. Haile. Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797, 2014. doi: 10.3982/ECTA9027.

M.E. Blume-Kohout and N. Sood. Market size and innovation: Effects of Medicare Part D on pharmaceutical research and development. *Journal of Public Economics*, 97(1):327–336, 2013. doi: 10.1016/j.jpubeco.2012.10.003.

F.A.S. Bokhari, F. Mariuzzo, and A.R. Bennato. Innovation and growth in the UK pharmaceuticals: The case of product and marketing introductions. *Small Business Economics*, 57(1):603–634, 2021. doi: 10.1007/s11187-019-00307-w.

Arnaud Costinot, Dave Donaldson, Margaret Kyle, and Heidi Williams. The More We Die, The More We Sell? A Simple Test of the Home-Market Effect*. *The Quarterly Journal of Economics*, 134(2):843–894, May 2019. ISSN 0033-5533. doi: 10.1093/qje/qjz003.

J.A. DiMasi, H.G. Grabowski, and R.W. Hansen. Innovation in the pharmaceutical industry: New estimates of R&D costs. *Journal of Health Economics*, 47:20–33, 2016. doi: 10.1016/j.jhealeco.2016.01.012.

P. Dubois, O. de Mouzon, F. Scott-Morton, and P. Seabright. Market size and pharmaceutical innovation. *RAND Journal of Economics*, 46(4):844–871, 2015. doi: 10.1111/1756-2171.12113.

Amy Finkelstein. Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry*. *The Quarterly Journal of Economics*, 119(2):527–564, May 2004. ISSN 0033-5533. doi: 10.1162/0033553041382166.

Ronald L. Goettler and Brett R. Gordon. Does AMD Spur Intel to Innovate More? *Journal of Political Economy*, 119(6):1141–1200, December 2011. ISSN 0022-3808. doi: 10.1086/664615.

Mitsuru Igami. Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981–1998. *Journal of Political Economy*, 125(3):798–847, June 2017. ISSN 0022-3808. doi: 10.1086/691524.

Fred D. Ledley, Sarah Shonka McCoy, Gregory Vaughan, and Ekaterina Galkina Cleary. Profitability of Large Pharmaceutical Companies Compared With Other Large Public Companies. *JAMA*, 323(9):834, March 2020. doi: 10.1001/jama.2020.0442.

John Rust. Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, 55(5):999–1033, 1987. ISSN 0012-9682. doi: 10.2307/1911259.

Stephen P. Ryan. The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica*, 80(3):1019–1061, 2012. ISSN 1468-0262. doi: 10.3982/ECTA6750.

Katja Seim. An Empirical Model of Firm Entry with Endogenous Product-Type Choices. *The RAND Journal of Economics*, 37(3):619–640, 2006. ISSN 0741-6261.

Kenneth E. Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, Cambridge, 2nd edition, 2009. ISBN 9780521747385.

Xavier Vives. Innovation and Competitive Pressure*. *The Journal of Industrial Economics*, 56(3):419–469, 2008. ISSN 1467-6451. doi: 10.1111/j.1467-6451.2008.00356.x.

10 Robustness: Controlling for market growth

A concern in the identification discussion is that observed exit (and therefore the inferred fixed-cost parameter ϕ) could partly reflect exogenous changes in market size rather than fixed operating costs alone. To assess this, I augment the reduced-form exit and transition specifications with controls for market growth, measured using changes in aggregate utilisation within the high-level condition (and, where available, changes in the estimated eligible patient population). Re-estimating the structural

model with these growth controls leaves the fixed-cost estimate within the baseline 95% confidence interval and does not materially change the qualitative counterfactual conclusions reported in the main text.

11 Tables

Table 11: Summary Statistics

	N	Mean	StDev	Min	**Max**
Panel A: Utilisation					
Standardised Doses	2878	10709	36689	-4333	342316
Total Volume of Doses	2878	5952152	20516029	-2425248	1.90E+08
Standard Doses by Condition					
Arthiritis	74	621	806	0	3620
Asthma	77	261	251	0	740
Auto-immune conditions	433	1914	2568	0	11189
Cancer	2086	13473	42174	-4333	342316
Hypercholesterolaemia	63	108	148	0	515
Idiopathic pulmonary fibrosi	54	24477	35157	0	94764
Multiple Sclerosis	91	5411	10796	0	3830
Panel B: Patent & MA data					
Years with Market Exclusivity	2,878	13.6	4.6	1.2	30.8
Years with Patent Extension	2,435	4.0	1.4	0.7	5.5
Time from Patent to Market (Qtrs)	2,878	80.1	101.1	10.8	21.2
Years of market exclusivity by Condition					
Arthiritis	74	12.9	1.9	9.9	15.5
Asthma	77	6.9	4.8	1.2	11.8
Auto-immune conditions	433	13.8	2	9.9	15.5
Cancer	2086	13.8	4.5	2.9	30.8
Hypercholesterolaemia	63	15	0	15	15
Idiopathic pulmonary fibrosi	54	10.8	0.2	10.6	10.9
Multiple Sclerosis	91	15.8	10.6	4.1	30

Notes: Summary statistics of utilisation data from the Innovation Scorecard. Defined Daily Doses (DDDs) follow the WHO definition: the assumed average maintenance dose per day for a drug used for its main indication in adults. Table reports quarterly data from 2012/13 to 2019/20 for approximately 127 new treatments in Panel A, by high-level condition. Data on utilisation is based on hospital purchases, thus negative values indicate negative inventory. Panel B includes data imputed from marketing authorisation and patent data, including effective patent length and duration of extensions by condition.

Table 12: Summary Statistics for Cancer Treatments

Variable	**Mean**	**Std. Dev.**	**Min**	**Max**
Utilisation (value)	13 473.01	42 174.09	-4 332.57	342 316.12
Time to patent expiry (quarters)	21.08	20.32	-46.00	67.00
Number of competitors	23.40	2.32	10.00	25.00

Notes: **Utilisation** is measured in standardised patient-equivalent doses per quarter. **Time to patent expiry** denotes the remaining number of quarters until the earliest observed patent expiry date of the incumbent treatment; negative values indicate quarters since expiry. For observations without an observed patent expiry date, the variable is coded as zero and handled via a missing-patent indicator. **Number of competitors** counts other on-patent drugs in the same high-level condition during that quarter. Statistics are computed over all firm–quarter observations classified as *Cancer* (total observations: 2 064). Negative utilisation values (e.g. –4 332.57) occur when quarterly sales are adjusted by returns or data corrections.

12 Figures

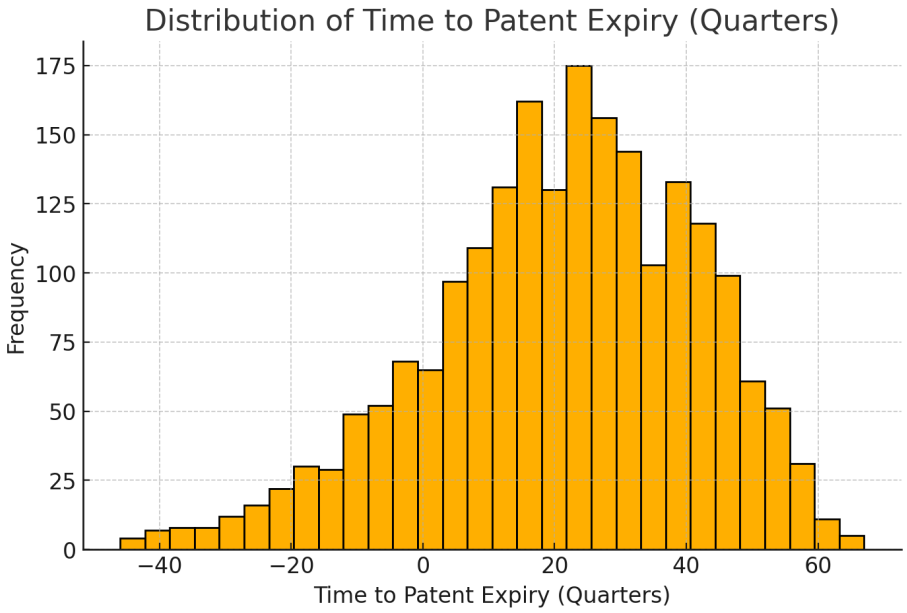

Figure 9: Histogram: Number of Quarters to Patent Expiry

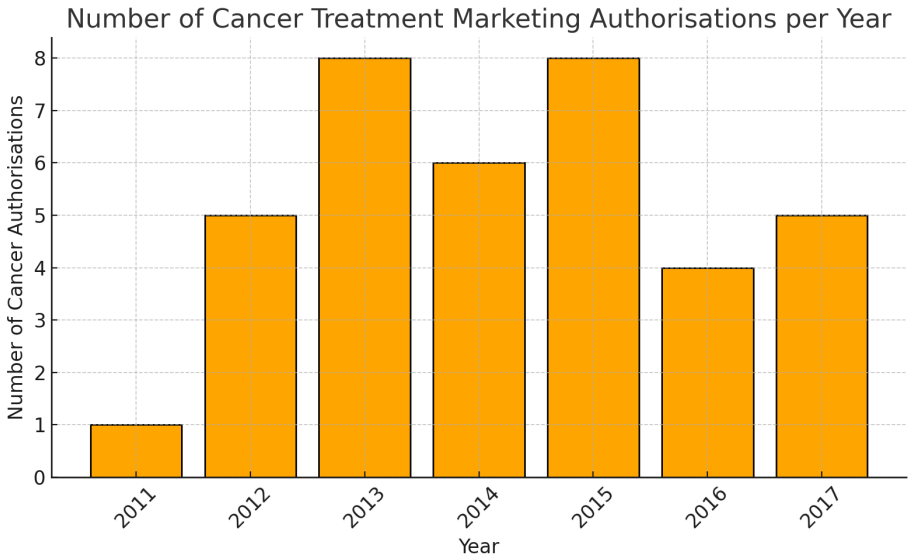

Figure 10: Number of Marketing Authorisations for Cancer Treatments Per Year

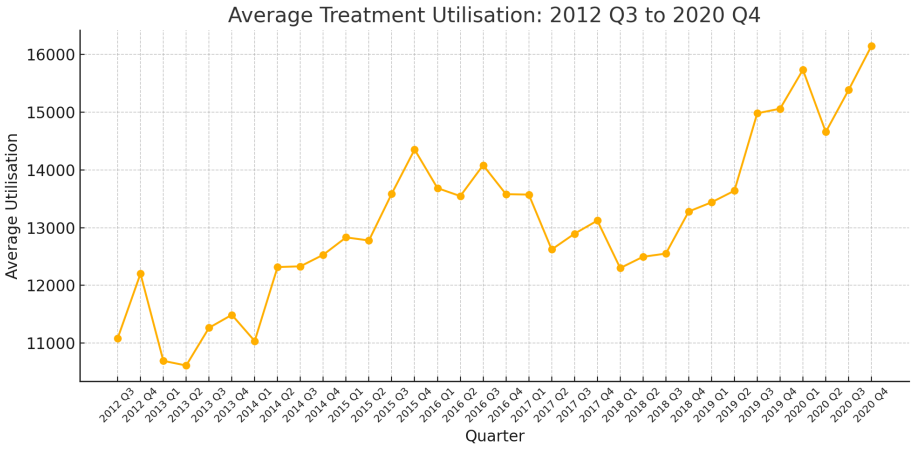

Figure 11: Utilisation of new cancer treatments over time

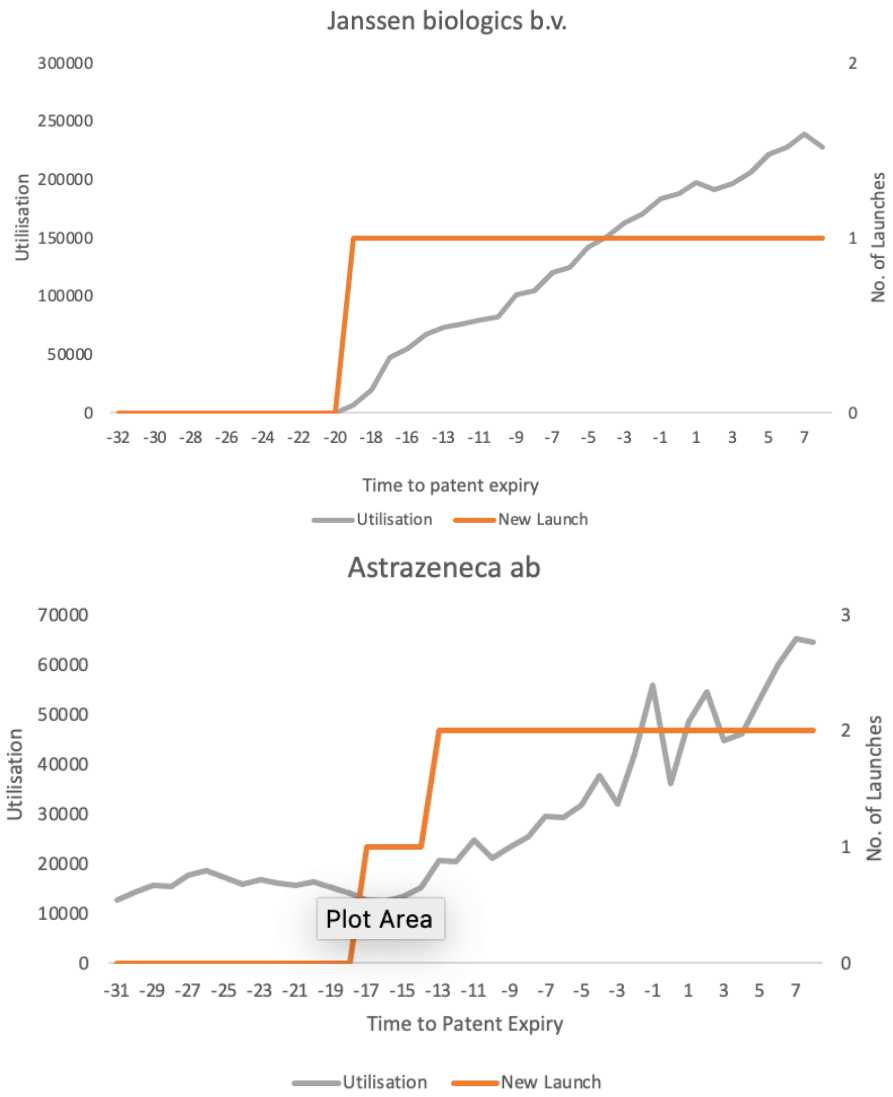

Figure 12: Utilisation and Launches by Select Firms

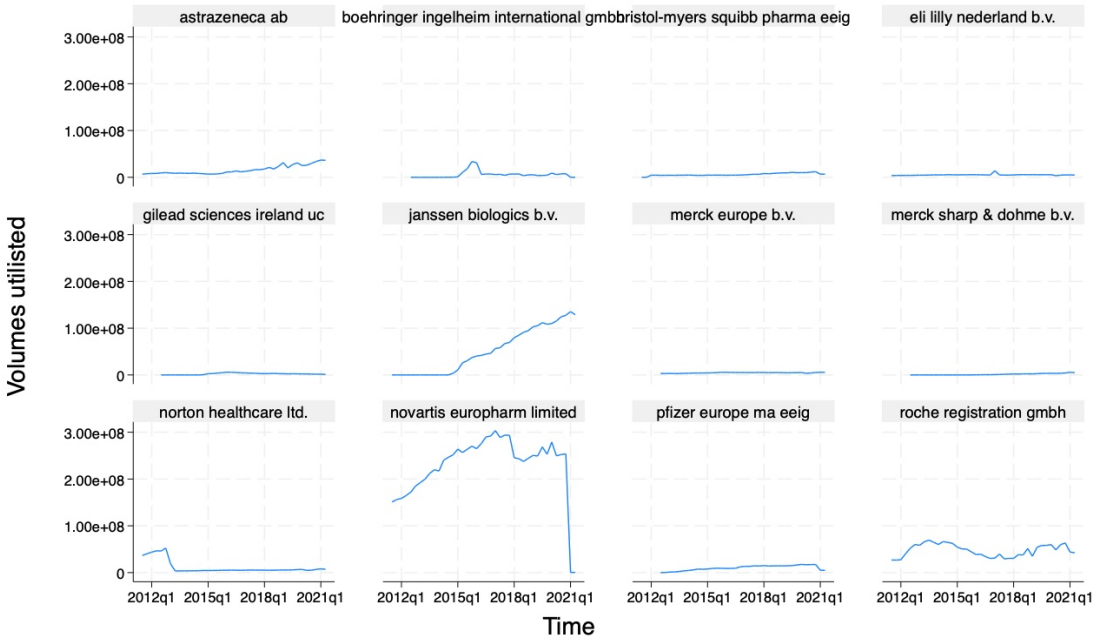

Figure 13: Total volumes utilised by Select Firms

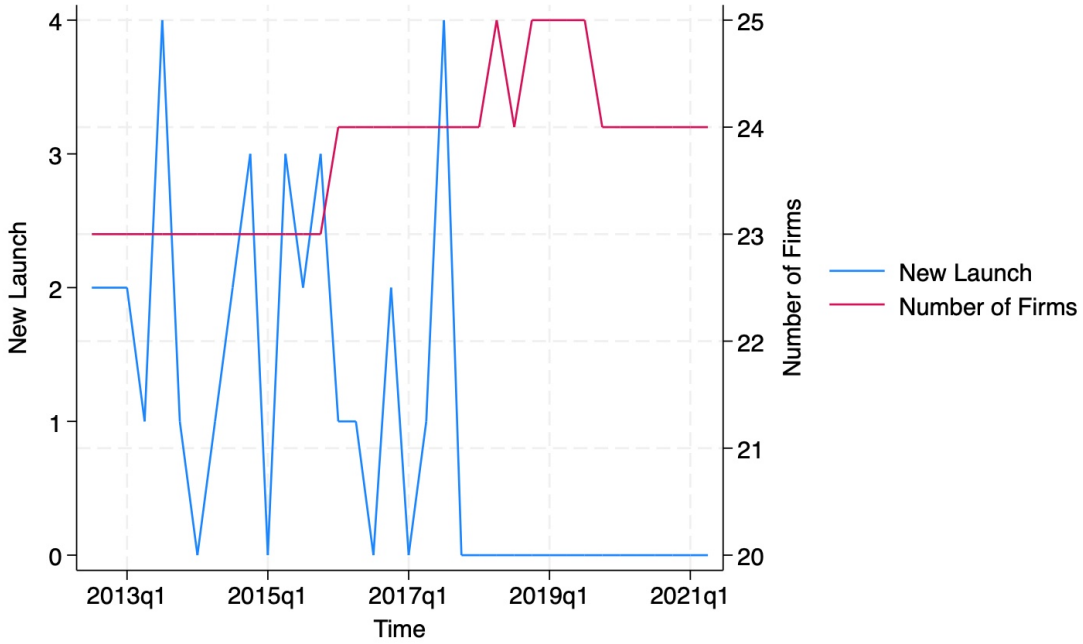

Figure 14: New Launches by Select Firms

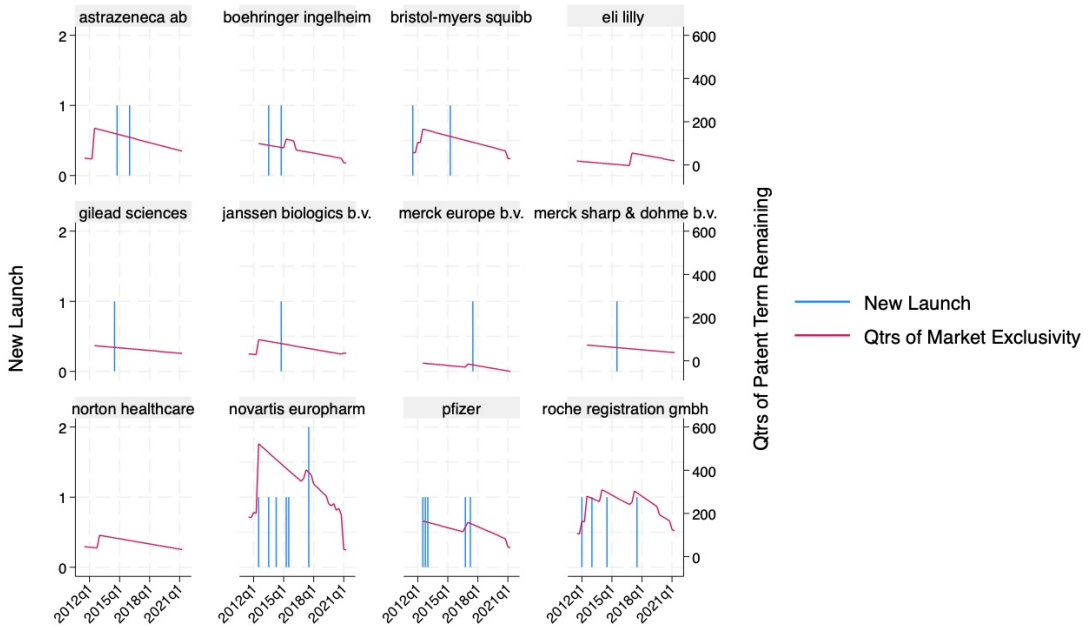

Figure 15: New Launches and Market Exclusivity by Select Firms

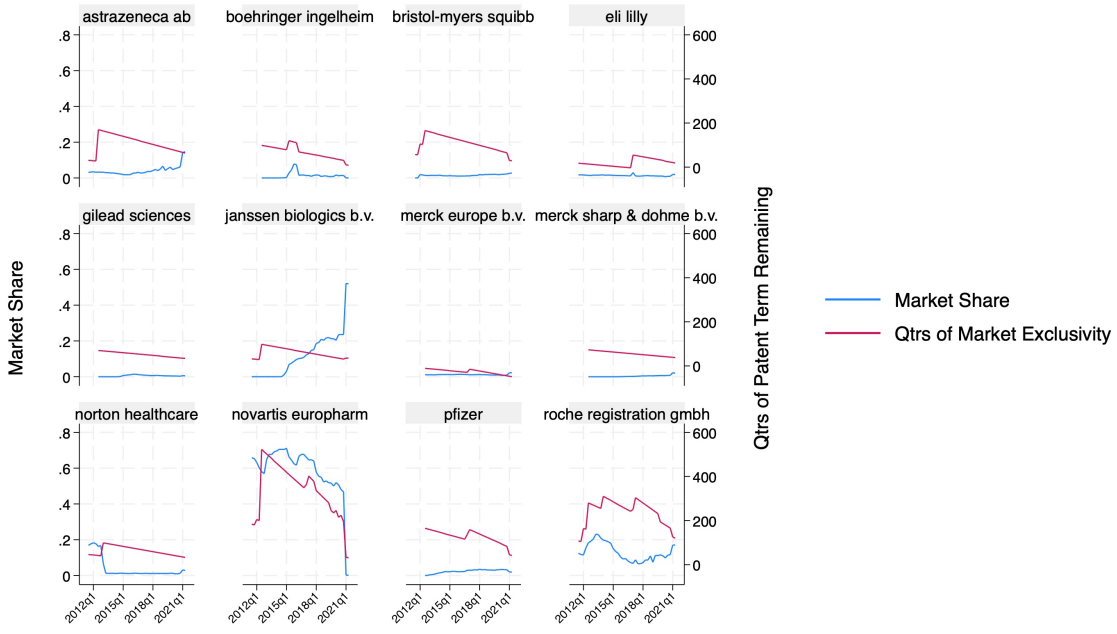

Figure 16: New Launches and Market Share by Select Firms

13 Appendix: Nested Fixed Point Algorithm Implementation

I take as unknown structural parameters

$$\Theta \;=\; \Big\{\phi,\; \{\kappa_{i,t}\},\; \kappa^{(\text{new})},\; \beta\Big\}.$$

In practice, I fix β (e.g. $\beta = 0.95$) and estimate ϕ, $\{\kappa_{i,t}\}$, and $\kappa^{(\text{new})}$. I proceed in two loops:

Inner Loop (Value-Function / CCP Solver) Given a candidate $\Theta^{(k)}$ and an initial guess for Conditional Choice Probabilities (CCPs) $\Pr(a_{i,t} \mid s_t)$, solve by backward induction for

$$\Big\{V_t^{\text{old}}(s_t;\Theta^{(k)}),\; V_t^{\text{both}}(s_t;\Theta^{(k)}),\; V_t^{\text{new}}(s_t;\Theta^{(k)})\Big\}, \quad t = T, T-1, \ldots, 0,$$

and update the CCPs using the logit formulas, e.g.

$$\Pr\Big[a_{i,t}^{\text{old}} = \text{stay old} \mid s_t;\Theta^{(k)}\Big] \;=\; \frac{\exp\Big\{-\phi + \beta\,\mathbb{E}\big[\,V_{t+1}^{\text{old}}(s_{t+1};\Theta^{(k)})\Big]\Big\}}{1 + \exp\Big\{-\phi + \beta\,\mathbb{E}[V_{t+1}^{\text{old}}]\Big\} + \exp\Big\{-\phi - \kappa_{i,t} + \beta\,\mathbb{E}[V_{t+1}^{\text{both}}]\Big\}}.$$

Iterate until the CCPs converge to a fixed-point given $\Theta^{(k)}$.

Outer Loop (Maximum Likelihood) Once equilibrium CCPs $\Pr^*(a_{i,t} \mid s_t;\Theta)$ are found, construct the sample log-likelihood by comparing realized actions $\{a_{i,t}\}$ to model-predicted CCPs:

$$\mathcal{L}(\Theta) \;=\; \sum_{t=0}^{T-1} \sum_{i \in \text{old-only}} \ln \overset{*}{\Pr}\Big[a_{i,t}^{\text{old}} \mid s_t;\Theta\Big] + \sum_{t=0}^{T-1} \sum_{i \in \text{both}} \ln \overset{*}{\Pr}\Big[a_{i,t}^{\text{both}} \mid s_t;\Theta\Big] + \sum_{t=0}^{T-1} \sum_{i \in \text{new}} \ln \overset{*}{\Pr}\Big[a_{i,t}^{\text{new}} \mid s_t;\Theta\Big].$$

I maximise $\mathcal{L}(\Theta)$ over $\{\phi, \{\kappa_{i,t}\}, \kappa^{(\text{new})}\}$, holding β fixed (or estimating β jointly if data permits). In practice, I parameterise $\kappa_{i,t}$ as a low-dimensional function of "time to patent expiry" or other observables to reduce dimensionality.

Implementation Details

In our dataset, the total number of active firms in any period never exceeds 25. Thus the count-state $\left(N_t^{\text{old}}, N_t^{\text{both}}, N_t^{\text{new}}\right)$ has at most

$$\binom{25+2}{2} = \binom{27}{2} = 351$$

possible triples. With a horizon $T = 10$–15 (e.g. 2010–2020), the value-function solver loops over at most a few thousand states in each backward-induction pass. A carefully vectorised value-function algorithm solves the 351×351 transition matrix ten times in a matter of minutes. The outer-loop likelihood search (e.g. using BFGS) then converges to the MLE for ϕ and $\{\kappa\}$ in a few hours.

I recover standard errors by numerically differentiating $\mathcal{L}(\Theta)$ with respect to Θ. In pilot experiments, endogenous variation in launch timing (67 new treatments observed over eight years) provides precise estimates of κ and ϕ.