Evergreening and Strategic Patenting: Evidence from the Pharmaceutical Industry

Shizhe Yu*

University of Wisconsin-Madison

shizhe.yu@wisc.edu

November 2, 2026

Preliminary draft. Please do not circulate without permission.

Abstract

This paper studies the welfare effects of patent evergreening, the increasingly prevalent practice of layering secondary patents around a brand drug to extend effective market exclusivity. Using the universe of small-molecule drugs approved from 1995 to 2019, I find that secondary patenting delays generic entry, in a pattern suggestive of strategic incentives, and is negatively associated with investment in primary R&D. A back-of-the-envelope welfare analysis built on a structural demand model suggests that for the patent types that block entry, the loss from deterred generic competition decisively outweighs the gain from the associated quality improvements, and amounts to an average of $90–$207 million per additional patent across the Medicare Part D market of the drug's therapeutic class, or $29–$68 per user.

*This paper benefited from the JPGI Summer Research Fellowship at the University of Wisconsin-Madison. The author has no conflicts of interest to disclose.

1 Introduction

The patent system has long been regarded as addressing innovators' "missing market" problem: by granting temporary exclusivity, it translates the future gains from an invention into upfront incentives to invest in it (Arrow, 1962). Yet its overall effect on innovation remains controversial. A growing empirical literature finds that, despite a sharp rise in patenting activity, stronger patent rights have not produced systematic improvements in research, productivity, or innovation quality (Lerner, 2009; Boldrin and Levine, 2013). While patents strengthen private incentives to invent, the resulting monopoly power can also dull competitive pressure from other innovators (Aghion et al., 2005).

Beyond its effect on the rate of innovation, the patent system can also create unintended consequences: by tying private returns to the scope of legal protection, it shapes firms' innovation efforts and can redirect effort away from socially valuable activities. Williams (2013) shows that intellectual-property restrictions on sequenced human genes reduced subsequent scientific research and product development on those genes by 20–30 percent; Galasso and Schankerman (2015) find that invalidating a patent raises follow-on citations to the focal invention by roughly 50 percent, with the effect concentrated among small firms blocked by large incumbents' portfolios; and Budish et al. (2025) document that weak property rights over new uses of existing drugs leave clinically promising research on those uses undone. Where the legal instrument and the underlying economics interact badly, the direction of innovative effort responds to the protection available rather than to social value (Budish et al., 2015).

This paper studies one such unintended consequence: patent evergreening. The term refers to a widely used strategy in the pharmaceutical industry, in which an incumbent accumulates secondary patents around the core compound patent of an existing brand drug, extending its effective market exclusivity beyond the expiration of that original, or primary, patent. Secondary patents do not claim the molecule itself but its other features: formulations, methods of use, salts and polymorphs, and delivery devices. Such patents need not be strong or directly extend the legal exclusivity of the underlying drug, but they raise the legal complexity and

entry barriers facing generic competitors, who must navigate or invalidate each listed patent to obtain approval. The practice therefore does more than reallocate innovative effort: it reshapes product-market competition by prolonging incumbents' market power and delaying the price competition that generic entry would otherwise bring (Hemphill and Sampat, 2012). Its scale is substantial: 78 percent of the drugs associated with newly issued patents are drugs already on the market rather than newly approved molecules, and among the roughly one hundred best-selling drugs, more than 70 percent had their protection extended at least once after launch, with almost half extended more than once (Feldman, 2018). As a result, it has drawn sustained legal and regulatory attention, from congressional action requiring the FDA to clarify listing standards to Federal Trade Commission challenges against improper listings and litigation ordering patents delisted (Federal Trade Commission, 2023).

Despite these concerns, the actual welfare consequences of the practice are theoretically ambiguous. On the one hand, by delaying generic entry, evergreening reduces consumer surplus through sustained high prices. On the other hand, the product refinements the patents protect may deliver genuine quality improvements that partially offset higher prices. The innovation side is equally ambiguous: the effort devoted to secondary patents may crowd out investment in fundamentally new molecules, but by raising the expected returns to maintaining an existing product, the practice could in principle encourage investment as well. Which of these forces dominates, and hence whether consumers are harmed on balance, is an empirical question that the existing literature has not resolved.

This paper answers that question with a combination of reduced-form and structural evidence from the U.S. pharmaceutical market. I assemble the universe of small-molecule new molecular entities approved between 1995 and 2019 by linking four datasets: patenting activity from the FDA Orange Book, drug characteristics and Paragraph IV challenge records from FDA approval data, firm-level clinical-trial activity from ClinicalTrials.gov, and prescription-drug utilization and cost-sharing data from the public Medicare Part D and Medicaid files. Four descriptive facts motivate the analysis. Secondary patenting has grown sharply across

approval cohorts while primary compound patenting shows no trend, so the growth is not driven by forces common to all patenting activity; secondary patents are filed progressively later in the product lifecycle, while primary patents bunch around approval—timing that gives firms a flexible instrument to deploy as competition develops; the distribution of secondary patents across drugs is heavily right-skewed, with a small set of drugs accumulating very large portfolios, suggesting that the practice is concentrated among a few heavy users rather than spread evenly across the industry; finally, generic entry bunches precisely at the expiration of a molecule's last secondary patent, indicating that the secondary portfolio is the binding legal barrier to competition. These patterns point to accumulation that is systematic and strategic rather than an incidental byproduct of ongoing innovation.

Motivated by these patterns, I conduct three sets of reduced-form analyses. First, exploiting cross-sectional variation in patent portfolios across molecules, I find that secondary patents do delay generic entry: compared to the compound patent, which is associated with roughly one additional year until the first generic entrant, each method-of-use patent—a patent covering an approved therapeutic indication rather than the drug itself—is associated with roughly 0.6 years, consistent with the widely documented "skinny label" problem, in which a generic that carves the patented indication out of its label to avoid infringement becomes a weaker substitute for the brand, and each patent covering the delivery device with over a year and a half. The associations are robust to confounders such as molecular complexity and firm capability. Second, the pattern of patenting activity is consistent with an entry-deterrence incentive: the intensity of new secondary listings has a non-monotonic, inverted-U relationship with the competitive state of a drug's therapeutic class, the signature of strategic investment identified by Ellison and Ellison (2011) and Dafny (2005). Third, within-firm increases in secondary-patent intensity predict significantly lower Phase 2 and Phase 3 clinical-trial activity, consistent with a reallocation of resources away from fundamental research. Together, these findings point to strategic patenting with real consequences for generic entry and for the allocation of innovation.

Building on these findings, and to answer the central question of this paper, the net welfare effect of evergreening on consumers, I conduct a back-of-the-envelope welfare analysis. Using public Medicare Part D data, I estimate a structural demand model for branded and generic drugs in which each type of secondary patent contributes separately to product quality, and use it to price one additional patent: the surplus cost of the generic competition it postpones against the surplus benefit of the quality it protects. For the patent types that delay entry, the cost dominates. One additional method-of-use patent reduces consumer welfare in 97 percent of the branded product-years on which it is evaluated, with a mean loss of $90 million (median $6 million) summed over the Medicare Part D market of the drug's therapeutic class, or $29 per user, with the postponed generic competition worth about $12.50 per user-year, against protected quality worth only about $0.43. The asymmetry is rooted in prices: generics enter at sufficiently low prices, while brand prices do not fall—and if anything rise—as secondary patents accumulate, so each year of postponed entry is a year of monopoly pricing with no offsetting relief. Because this calculation values only the products that exist, and assigns no cost to the new molecules displaced by the crowd-out documented above, these figures are a conservative bound on the total consumer cost.

Taken together, the evidence establishes both the phenomenon and its welfare costs. Secondary patenting around existing drugs is systematic and strategic—it intensifies precisely when deterrence pays—and it postpones the arrival of generic competition, the single largest source of price relief in pharmaceutical markets. Where the practice blocks entry, the cost to consumers decisively outweighs the quality the added patents protect; and the full burden is larger still, as the delayed competition comes on top of research resources diverted away from fundamentally new molecules. The consumer-welfare case against evergreening therefore turns not on fine judgments about the merit of individual patents but on its consequences: later competition, higher spending, and less fundamental innovation than the same resources would otherwise have produced.

This paper contributes to three strands of the literature. The first studies the interaction

of branded drugs with generics, including generic entry and the competition that follows. A body of work characterizes generic entry and post-entry pricing dynamics (Caves et al., 1991; Grabowski and Vernon, 1992; Reiffen and Ward, 2005), the determinants of firms' entry decisions (Morton, 1999), and the effective exclusivity that Hatch–Waxman leaves branded drugs (Grabowski and Kyle, 2007). A related literature studies how incumbents respond strategically to the threat of entry, including price adjustments, advertising, and product proliferation (Ellison and Ellison, 2011; Tenn and Wendling, 2014; Goolsbee and Syverson, 2008). Closest to this paper, Hemphill and Sampat (2011) and Hemphill and Sampat (2012) directly study the role of secondary patents in this setting, showing empirically that secondary patents and Paragraph IV challenges jointly shape the effective market life of branded drugs. I extend this work by further quantifying the evergreening problem and embedding the entry-delay estimates in a fuller welfare account.

A second strand examines the relationship between the patent system and innovation. General verdicts are mixed: stronger rights have not been shown to raise innovation systematically (Lerner, 2009; Boldrin and Levine, 2013), while a causal literature finds that intellectual-property rights can deter follow-on and new-use research (Williams, 2013; Galasso and Schankerman, 2015). On the market side, Duggan et al. (2016) show that extending pharmaceutical product-patent protection has measurable but modest price consequences—India's TRIPS-mandated adoption of product patents raised the prices of newly patented molecules by only a few percent—underscoring that the market power a patent confers depends on the competitive environment into which it is introduced, the margin this paper studies within the U.S. system. The mechanics by which layered rights raise entry barriers are studied in the patent-thicket literature (Shapiro, 2001; Farrell and Shapiro, 2008), and a largely descriptive literature documents pharmaceutical evergreening directly (Feldman, 2018; Righi and Simcoe, 2023). Closest in spirit, Gupta (2023) estimates that imperfect, uncertain-scope patents delay generic entry by roughly three years per drug, and Frakes and Wasserman (2025) document analogous thicketing in biologics. Adding to this literature, my within-firm evidence that secondary

patenting crowds out clinical-trial investment identifies a reallocation margin—incremental versus fundamental research within the incumbent firm—that this cross-firm, largely descriptive literature does not isolate.

A third strand studies how competition and market incentives shape the rate and direction of R&D. A market-size literature establishes that expected demand pulls innovation toward larger markets (Acemoglu and Linn, 2004; Dubois et al., 2015; Blume-Kohout and Sood, 2013; Myers and Pauly, 2019; Cockburn et al., 2016). A complementary literature shows that the direction of research is itself distorted by the incentives firms face: fixed patent terms bias effort toward quick-to-market projects (Budish et al., 2015), and financing frictions suppress novel drugs (Krieger et al., 2022). The dynamics of competition and innovation—and in particular incumbents' reluctance to cannibalize their own products—are studied by Aghion et al. (2005) and, structurally, by Goettler and Gordon (2011) and Igami (2017), and Lemus and Ozkul (2023) model how product hopping can dampen incentives for drastic innovation. To this strand of the literature, this paper contributes a quantitative welfare decomposition of pharmaceutical evergreening, together with evidence that the practice reallocates research effort away from fundamentally new molecules—a dynamic cost this paper documents but does not price.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on the Hatch–Waxman framework and the patents and exclusivities it governs. Section 3 describes the data and linkage. Section 4 presents descriptive evidence on the growth and timing of secondary patenting. Section 5 reports the reduced-form results on entry delay, the incentive to patent, and innovation crowd-out. Section 6 develops the back-of-the-envelope welfare decomposition. Section 7 concludes.

2 Institutional Background

This section describes the regulatory framework that governs generic entry and patent protection for small-molecule[1] drugs in the United States, and defines the institutional objects that the empirical analysis measures.

2.1 The Hatch–Waxman Framework

The U.S. market for small-molecule drugs is governed by the Drug Price Competition and Patent Term Restoration Act of 1984, known as the Hatch–Waxman Act, which balances innovation incentives against generic access. A brand-name firm brings a new drug to market through a New Drug Application (NDA) supported by its own clinical trials establishing safety and efficacy, and receives a period of exclusivity protected by its listed patents and by statutory exclusivity provisions.[2] In exchange, the Act also created an abbreviated pathway for generic entry: a generic firm may file an Abbreviated New Drug Application (ANDA) that need only establish bioequivalence to the original "reference listed drug" without repeating the clinical trials, which greatly shortens the path to generic entry. A generic rated therapeutically equivalent is automatically substitutable for the brand at the pharmacy under state substitution laws, and its entry typically collapses the price of the molecule almost immediately (Caves et al., 1991; Grabowski and Kyle, 2007).

The same Act also created the procedure by which a generic can enter before the listed patents expire. An ANDA applicant must certify, for each listed patent, either that it has expired, that entry will wait until it does, or that it is invalid or would not be infringed. The last of

[1] The analysis is confined to small-molecule drugs. Generic penetration is far weaker in biologics: substitution at the pharmacy is automatic only for therapeutically equivalent small-molecule generics, whereas a biosimilar requires a separate interchangeability designation. As a result, biosimilar entrants remain few and post-entry price erosion modest relative to the near-immediate collapse that follows small-molecule generic entry (Frakes and Wasserman, 2025).

[2] In addition to the patent system, there are also statutory exclusivities that block or delay generic approval independently of any patent. In general, there are five years of data exclusivity for a new chemical entity (NCE), seven years of market exclusivity for a drug designated for a rare "orphan" disease, and a six-month extension granted for conducting requested pediatric studies.

these is a Paragraph IV certification, named for the statutory subparagraph that defines it. A Paragraph IV filing counts as an act of infringement: if the brand firm sues within forty-five days, FDA approval of the generic is automatically stayed for up to thirty months while the litigation proceeds. To compensate challengers for bearing this litigation risk, the first Paragraph IV filer earns 180 days of generic marketing exclusivity if successful. Paragraph IV certifications are therefore the principal channel through which generic entry can occur before the patent envelope lapses.

2.2 The Orange Book and Drug-Related Patents

The FDA's Approved Drug Products with Therapeutic Equivalence Evaluations, commonly referred to as the Orange Book, is the official registry of approved drug products and of the patents that protect them. Per regulation, a brand firm must list in the Orange Book every patent that covers the approved drug within thirty days of the patent's issuance. The listing is thus a public, dated record of the drug-related patents protecting each approved product.

Listed patents fall into two broad types. The primary, or compound, patent claims the active molecule itself and is generally filed early, before the drug reaches the market. Secondary patents claim features other than the underlying molecule, including formulations, crystalline forms, delivery devices, and methods of using the drug to treat a condition; by their nature, such patents are usually filed later in the drug's lifecycle.

2.3 Patent Evergreening

Because an ANDA applicant must clear the latest-expiring listed patent its product would infringe, secondary patents can extend the exclusivity envelope beyond the life of the compound patent when the molecule itself is unprotected. This layering of secondary patents around a single compound patent is the practice commonly referred to as patent evergreening. Crucially, the delay such patents produce requires neither that they be strong nor that they lengthen the envelope: every listed patent creates legal and time costs, as the entrant must clear the Para-

graph IV process for each one before legally entering the market, which effectively raises the entry barrier.

The practice is increasingly prevalent. Among the drugs associated with newly issued patents in the FDA's records, 78 percent are drugs already on the market rather than newly approved molecules, and among the roughly one hundred best-selling drugs, more than 70 percent had their protection extended at least once after launch, with almost half extended more than once (Feldman, 2018); continuation practice at the patent office supplies a steady flow of such filings long after the original application (Righi and Simcoe, 2023).

This issue has accordingly drawn sustained legal and regulatory attention. Regulators and courts have taken a series of actions against exploitation of the listing system: Congress required the FDA to clarify what may be listed in the Orange Book Transparency Act of 2020; the Federal Trade Commission has moved to police listing behavior directly, disputing large numbers of inhaler and auto-injector device listings as improper; and litigation has been brought against individual heavy users of the practice, most consequentially in Teva v. Amneal (Fed. Cir. 2024), where the Federal Circuit ordered device patents delisted on the ground that a listed patent must claim the active ingredient.

3 Data and Linkage

This section describes the data sources, how they are linked, and the resulting estimation samples.

3.1 FDA Orange Book (1995–2019)

The FDA Orange Book lists all patents and regulatory exclusivity periods associated with each approved drug product and is the primary source for tracking patenting behavior over a drug's lifecycle. I use the panel covering 1995–2019. Following Hemphill and Sampat (2012), patents are manually classified into five types: (i) primary patents cover the active ingredient;

(ii) formulation patents cover controlled-release and pharmaceutical-composition claims; (iii) method-of-use (MOU) patents cover therapeutic indications; (iv) ingredient-level patents cover salts, polymorphs, and enantiomers; and (v) device patents cover delivery devices and related new combinations. Categories (ii)–(v) together constitute secondary patents.

3.2 Drugs@FDA, ANDA, and Paragraph IV Records

Drugs@FDA provides approval dates and drug characteristics including new-molecular-entity (NME) status, application type, and exclusivity designations. Generic entry timing is measured from the FDA's first-time generic approval records, combining the agency's first-time generic approvals compilation (2001–2015) with its annual first-generic approval lists from 2016 onward, and recording the first year each active ingredient faced an approved generic entrant; because patents are observed through the 2019 Orange Book edition, entries after 2019 are treated as censored throughout. Related FDA records also report Paragraph IV certifications—generic applicants' assertions that a brand drug's listed patents are invalid or not infringed—including the date of the first challenge against each drug, the number of ANDAs filed, and whether a challenger qualified for the 180-day first-filer exclusivity.

3.3 ClinicalTrials.gov (AACT, 2008–2019)

ClinicalTrials.gov is the U.S. National Library of Medicine's public registry of clinical studies, on which industry sponsors are required to register applicable drug trials and report basic results. The Aggregate Analysis of ClinicalTrials.gov (AACT) provides firm-level Phase 1–3 trial counts from 2008 onwards, when FDAAA mandatory registration took effect for industry-sponsored trials. The dataset covers 40,794 trials across 79 large pharmaceutical firms over 2008–2019, and is used to construct the primary R&D outcome variables in the crowding-out analysis. Two caveats on completeness apply. First, the FDAAA mandate covers Phase 2–4 trials of FDA-regulated drugs with a U.S. site or conducted under an IND; Phase 1 trials are exempt, and audits of the trials submitted to the FDA for the drugs approved in 2012 find

that only a median of 57% per drug were registered, the gap consisting mostly of Phase 1 and ex-U.S. studies (Zarin et al., 2016; Miller et al., 2015). Second, against commercial trial databases the registry is not a superset: in a four-indication comparison, Citeline's Trialtrove held 31% more trials than ClinicalTrials.gov, and phase, status and site fields disagreed for a sizable share of matched records (Stergiopoulos et al., 2019). To gauge what this means for the firm-year counts, I compare AACT with a Trialtrove extract of U.S.-site trials (2010–2024) at the trial level through the NCT identifiers both carry. Among industry-sponsored Phase 1–3 trials starting in 2010–2019, 96% of Trialtrove's 17,829 trials appear in AACT's industry set, the two sources assign the same phase to 99% of shared Phase 2 and Phase 3 trials (95% of Phase 1), and the AACT-only trials are almost entirely trials with no U.S. site. Firm-year counts for the 79 panel firms correlate at 0.91–0.94 across the two sources. Re-estimating the crowding-out specification of Section 5.3 on the Trialtrove counts reproduces the Phase 3 pattern with a larger cumulative effect (a 6.9% reduction over lags one to six, against 4.1% with AACT on the same 2010–2019 window) and the same decay to zero at lag eight; Phase 1, the phase with the weakest registration coverage in both the literature and this comparison, remains the noisy column.

3.4 Medicare Part D and Medicaid Drug-Utilization Data

The structural demand analysis uses two public prescription-drug files. The Medicare Part D Prescribers by Geography and Drug file reports, for each drug and state, the number of claims and of 30-day-equivalent fills, gross drug cost, the number of beneficiaries, and the cost sharing of beneficiaries with and without the low-income subsidy, annually for 2013–2024; the out-of-pocket price is measured as beneficiary cost sharing per 30-day fill. The public Part D formulary and plan-benefit files (2019–2024) record each drug's tier placement and each plan's benefit design. The Medicaid State Drug Utilization Data report prescriptions and reimbursement by state program and National Drug Code over the same years. Drugs are mapped to products, defined as a molecule in a given formulation and brand/generic tier, and to ATC Level-3 classes.

The Medicare panel covers 1,532 molecules in 2,447 products across 163 classes, with about 2.3 billion fills a year; the Medicaid panel covers 1,743 molecules in 3,448 products across 171 classes. Appendix G describes how the two segments are built and stacked.

3.5 Sample Construction and Linkage

The unit of observation is the active ingredient: the unique chemical entity underlying a drug. This choice is natural on both regulatory and economic grounds. On the regulatory side, regulatory novelty and the single patent-term extension are conferred only once per active moiety. On the economic side and importantly for the mechanism studied here, generic substitution operates at the molecule level: an ANDA references a listed drug and, once approved and rated therapeutically equivalent, competes for the molecule's prescriptions. The two studies closest to this one adopt the same convention (Hemphill and Sampat, 2012; Gupta, 2023).

Two sample restrictions are imposed. First, I keep only NDA applications, which drops biologics licensed under the separate pathway. Second, I restrict to new molecular entities, defined as the first approved application per active ingredient, to focus on original branded drugs rather than reformulations. These restrictions give an analysis sample of 661 new molecular entities.

Each active ingredient in this sample is then linked to three records: its patent portfolio from the Orange Book, including the count of listed patents of each type and their expiration dates; its generic entry date from the first-time ANDA approval records; and its challenge history from the Paragraph IV records. Therapeutic classification comes from the World Health Organization's Anatomical Therapeutic Chemical (ATC) system, a hierarchy running from broad anatomical groups (Level 1) to chemical subgroups (Level 4). Because ATC codes attach to therapeutic uses rather than to molecules, an ingredient approved for several uses can carry more than one class; each ingredient is assigned its modal class, the one covering the largest number of its products. Appendix A describes how each of these links is built.

Table 1 provides the summary statistics for the resulting sample at the active-ingredient level. An Orange Book patent record exists for 88.2% of ingredients, and an ATC class is assigned to 93.0%. Secondary patents outnumber primary ones by roughly eighteen to one, with a mean secondary count of 5.1 against 0.3 primary: a molecule needs only one compound patent, filed before approval, while secondary patents accumulate over its commercial life. Within the secondary total no single type dominates, with means ranging from 0.1 for device patents to 2.3 for ingredient-level patents.

Panels B through D describe the competitive environment. Generic entry is observed for 199 ingredients (30.1%) by 2019, at a mean of 11.7 years after brand approval. Paragraph IV challenges reach 388 ingredients (58.7%), with a mean of 4.0 ANDAs filed against those challenged; challenged molecules are substantially more likely to see entry, with 41% experiencing generic entry against 15% of unchallenged molecules. Non-patent exclusivities are common: 92.9%[3] of ingredients carry new-chemical-entity status, one in four has an orphan designation, and one in four has received a pediatric extension.

Panel E turns to drug and firm characteristics. Approvals are spread evenly over the sample window, with both the mean and the median approval year at 2007. The sponsor experience measures display the familiar skew of the industry: the median sponsor brings ten prior NME approvals across all therapeutic classes but none in the focal drug's own ATC Level-2 class, while the most experienced firms bring over two hundred.

[3]A new molecular entity (NME) is a drug whose active moiety has not previously been approved by the FDA; new-chemical-entity (NCE) exclusivity is the five-year exclusivity that Hatch–Waxman grants such drugs. Because the sample's NME definition is operational—the first approved application per active ingredient—a small share of first-in-ingredient approvals, chiefly new salts, esters, and combinations of previously approved moieties, do not qualify for NCE exclusivity, which attaches to the active moiety.

	N	Mean	SD	Median	Min	Max
Panel A: Patent portfolio						
Has an Orange Book patent record	661	88.2%	—	—	—	—
# Primary (compound) patents	661	0.28	1.01	0.00	0.00	10.00
# Secondary patents	661	5.09	5.57	3.00	0.00	49.00
Formulation	661	1.66	3.27	0.00	0.00	34.00
Method-of-use	661	1.01	2.20	0.00	0.00	19.00
Ingredient-level	661	2.28	2.61	1.00	0.00	16.00
Device	661	0.13	1.09	0.00	0.00	17.00
Patent-term span (yrs, earliest to latest expiry)	661	7.00	6.82	6.00	0.00	36.00
Panel B: Generic entry						
Observed generic entry by 2019	661	30.1%	—	—	—	—
Years to first generic entry (if observed)	199	11.70	3.94	11.69	3.13	23.18
Panel C: Paragraph IV						
Received a Paragraph IV certification	661	58.7%	—	—	—	—
# ANDAs filed (if challenged)	388	4.04	5.29	2.00	0.00	51.00
Panel D: Non-patent exclusivity						
New chemical entity	661	92.9%	—	—	—	—
Orphan drug	661	24.8%	—	—	—	—
Pediatric extension	661	24.2%	—	—	—	—
Panel E: Drug and firm characteristics						
Brand approval year	661	2007	8	2007	1995	2019
Prior NME approvals by sponsor, same ATC Level-2 class	615	1.56	3.09	0.00	0.00	27.00
Prior NME approvals by sponsor, all classes	661	34.27	51.68	10.00	0.00	227.00

Notes: This table reports summary statistics for the 661 small-molecule new molecular entities approved in 1995–2019. The unit of observation is the active ingredient, and *N* is the number of active ingredients contributing to each row. Binary variables are reported as the share equal to one. Patent counts, the patent-term span, and the exclusivity indicators are set to zero for ingredients without an Orange Book patent or exclusivity record; same-class firm experience is defined for the 615 ingredients assigned an ATC class. Patent types follow the classification of Section 3; the four secondary types sum to the secondary total.

Table 1: Summary Statistics, Active-Ingredient Level

4 Descriptive Evidence

This section documents four facts that motivate the regression analysis. First, secondary patent counts per drug have risen steadily across approval cohorts while primary patenting shows no trend. Second, secondary patents are listed progressively later in the product life-cycle. Third, their distribution across drugs is heavily right-skewed. Fourth, generic entry bunches sharply at the expiration of a molecule's last-expiring secondary patent.

4.1 Secondary Patent Volume Has Increased, Primary Has Not

Figure 1 plots, by brand-approval-year cohort, the cross-sectional distribution of the number of patents listed per active ingredient, separately for primary patents and secondary patents.

Panels (a) and (b) count all patents ever listed for the molecule; because later cohorts have had less time to accumulate listings, panels (c) and (d) repeat the exercise counting only patents listed within five years of brand approval to account for the right-truncation of recent cohorts.

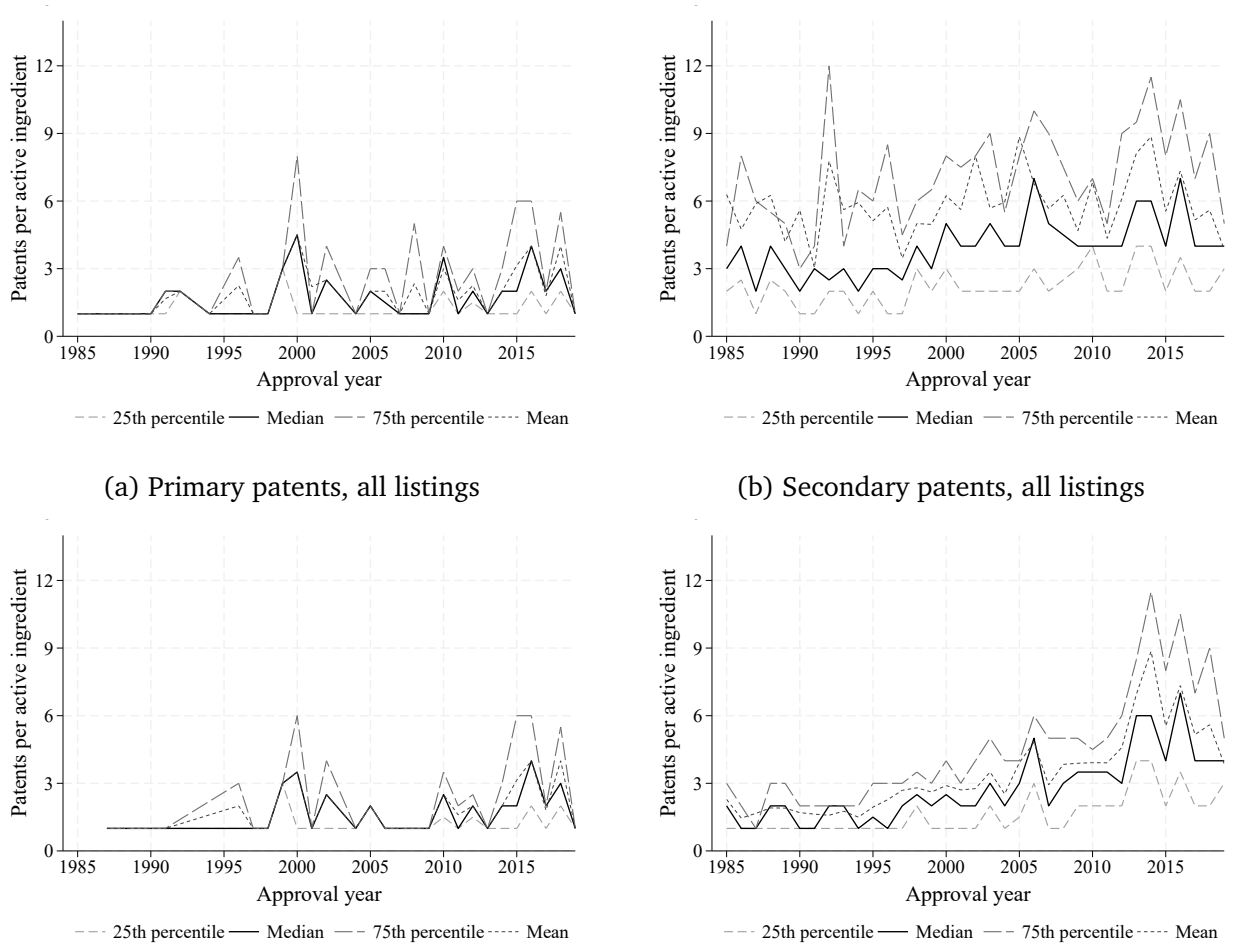
The contrast between the primary and secondary panels is striking. Primary patent counts per drug have remained essentially flat across cohorts. Secondary patent volume, by contrast, has increased substantially over time: both the median and the upper quartile of the secondary distribution rise steadily from the mid-1990s onward. The same pattern holds in the truncation-robust five-year window of panels (c) and (d). This growth is difficult to reconcile with a pure story of general patenting trends, such as rising molecular complexity, which would also affect primary patents; occurring independently of any trend in primary patenting, it is more suggestive of strategic accumulation.

This secular growth in secondary patenting is also consistent with the institutional evolution of life-cycle management documented in the literature. The Hatch–Waxman Act (1984) created the Orange Book listing system, generating structural incentives for post-approval patent accumulation that intensified with subsequent regulatory developments (Thomas, 2009).

4.2 Secondary Patents Are Listed Later in the Drug Lifecycle

A second descriptive fact concerns the timing of patent listings relative to brand approval. Figure 2 plots the distribution of listing timing separately for primary and secondary patents, with each bar being the share of all listed patents of that type entering the Orange Book in the indicated year since brand approval. Secondary patents are systematically listed later in a drug's commercial life than primary patents: primary compound patents are heavily concentrated at or near the brand approval date, consistent with patents filed during R&D that are already issued at the time of approval. Secondary patents, by contrast, are spread broadly across the product lifecycle, with a substantial fraction listed several years, and in some cases more than a decade, after initial approval.

This later timing is important for interpreting the evergreening mechanism. Primary patent



(a) Primary patents, all listings

(b) Secondary patents, all listings

(c) Primary patents, within 5 years of approval

(d) Secondary patents, within 5 years of approval

Figure 1: Number of Patents per Active Ingredient by Approval Cohort

Notes: Each panel plots, by brand-approval cohort, the 25th percentile, median, 75th percentile, and mean of the number of Orange Book patents listed per active ingredient. The sample is new molecular entities approved under New Drug Applications in 1985–2019; the regression analyses use the 1995–2019 cohorts. Panels (a) and (c) count primary (compound) patents; panels (b) and (d) count secondary patents (formulation, method-of-use, ingredient-level, and device), classified as in Section 3. Panels (a) and (b) count all patents ever listed; panels (c) and (d) count patents listed within five years of brand approval.

protection is established before or at the time of approval; secondary patents are added progressively over the product's commercial life, providing a flexible instrument for extending effective exclusivity in response to competitive developments.

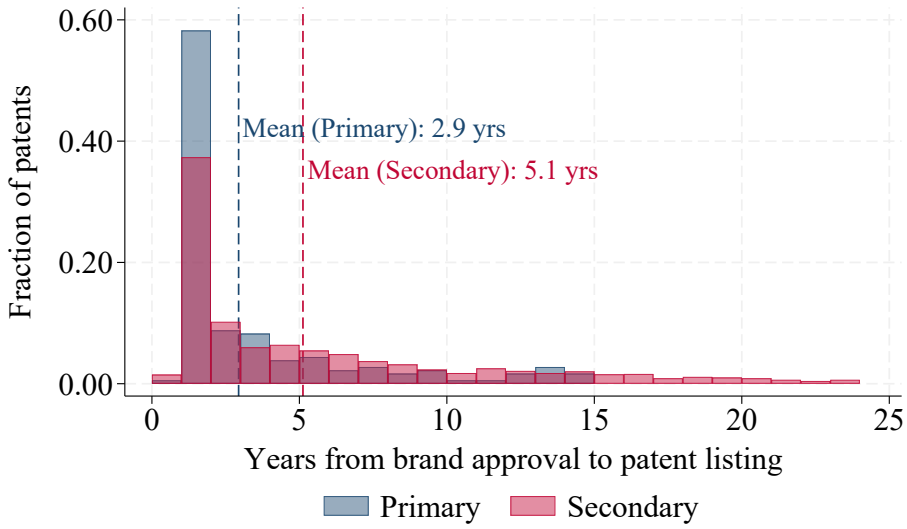


Figure 2: Patent Listing Timing Relative to Brand Approval

Notes: The figure plots the share of Orange Book patents of each type listed in each year-since-approval bin. The sample is patents listed for new molecular entities approved under New Drug Applications in 1995–2019. Patent types are classified as in Section 3.

4.3 Secondary Patent Counts Have a Heavy Right Tail

The distribution of secondary patent counts across active ingredients is heavily right-skewed, as shown in Figure 3. The unit of observation is the active ingredient, and each bar is the share of ingredients whose total count of listed patents of that type equals the indicated value over the drug's observed life; the primary and secondary distributions are overlaid for comparison. While many drugs carry few or no secondary patents, a meaningful share accumulates five, ten, or more, with the maximum in the sample reaching nearly fifty secondary patents for a single active ingredient. Primary patent counts, by contrast, are tightly concentrated, with most drugs holding one or two compound patents and a thin right tail. This skewness implies that the estimated relationship between secondary patent accumulation and market outcomes is driven substantially by the upper portion of the distribution, where the strategic incentive is most likely present.

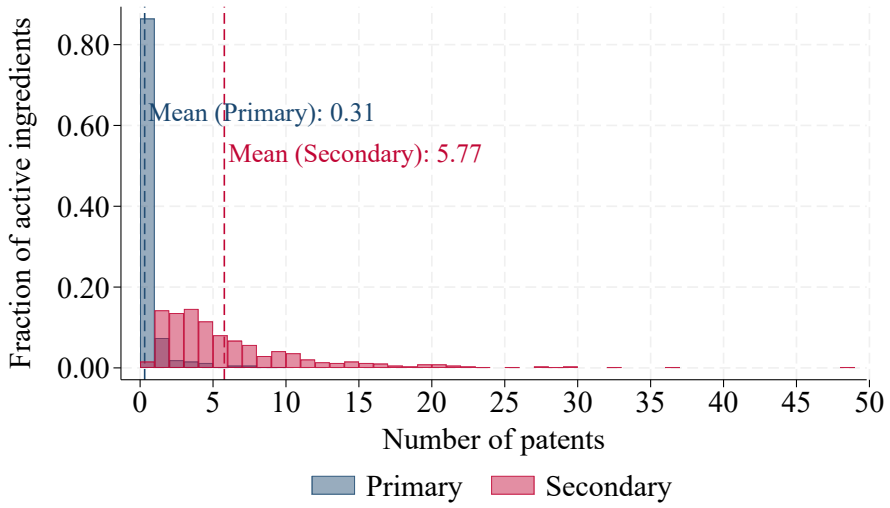


Figure 3: Distribution of Patent Counts per Active Ingredient

Notes: The figure plots the distribution of the number of Orange Book patents listed per active ingredient over the drug's observed life, for primary (compound) patents (blue) and secondary patents (red). The sample is new molecular entities approved under New Drug Applications in 1995–2019 with at least one Orange Book patent. Patent types are classified as in Section 3.

4.4 Generic Entry Clusters at Secondary Patent Expiration

A fourth descriptive fact directly links secondary patent coverage to the timing of generic entry. Figure 4 plots, for new molecular entities with both observed generic entry and at least one secondary patent, the distribution of the offset in years between an active ingredient's first generic entry and the expiration of its last secondary patent. If secondary patents were merely incidental to entry timing, this offset should be diffusely distributed. Instead, the distribution bunches at zero: entry in the same year the last secondary patent expires is the modal outcome, with a further mass just before zero reflecting successful Paragraph IV entry. This bunching is consistent with the secondary patent portfolio constituting the binding legal barrier to generic competition, with generic firms entering as the portfolio lapses or shortly before by Paragraph IV challenge.

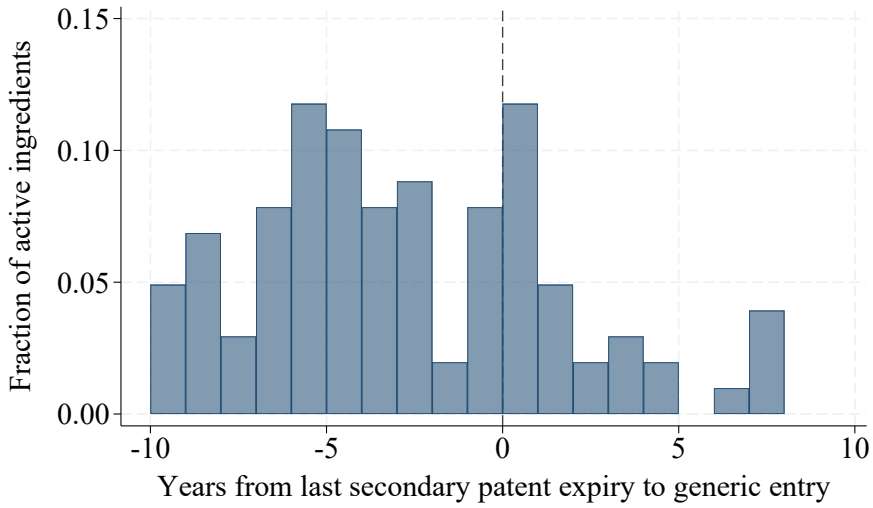


Figure 4: Bunching of Generic Entry Around Last Secondary Patent Expiration

Notes: The figure plots the distribution of the difference between an active ingredient's first generic entry year and the expiration year of its last-expiring secondary patent. The sample is new molecular entities approved under New Drug Applications in 1995–2019 with observed generic entry and at least one secondary Orange Book patent. First generic entry is from the FDA's list of first-time ANDA approvals; secondary patents and their expiration dates are from the Orange Book.

5 Reduced-Form Results

Motivated by the descriptive evidence, this section presents three sets of reduced-form results linking secondary patenting to generic entry and to the allocation of innovation investment. Section 5.1 estimates how a drug's secondary-patent portfolio relates to the timing of first generic entry. Section 5.2 turns to the incentive behind the accumulation, testing whether patenting intensity varies with the competitive state of a drug's therapeutic class in the inverted-U pattern that strategic entry deterrence predicts. Section 5.3 asks whether the accumulation displaces fundamental research, relating within-firm secondary-patent intensity to subsequent clinical-trial activity.

5.1 Secondary Patents and the Timing of Generic Entry

This section estimates the association between a drug's secondary-patent portfolio and the timing of first generic entry, exploiting cross-sectional variation in patent portfolios across

molecules.

The object of interest is a drug's effective exclusivity, defined as the span in years from brand approval to the first approved generic entrant. It is the economically relevant measure of protection and a central determinant of both innovation returns and consumer prices (Grabowski and Kyle, 2007). To assess whether secondary patents lengthen effective exclusivity while accounting for the censoring structure, the main specification is a Tobit regression that right-censors time to generic entry at 2019 — brand approval year for active ingredients with no observed entry as of 2019. The unit of observation is the active ingredient:

$$\text{Exclusivity}_i^* = \sum_k \theta_k N_{ik} + \mathbf{X}_i' \boldsymbol{\delta} + \eta_i, \tag{1}$$

where Exclusivity_i^* is the latent years from brand approval to first generic entry; N_{ik} is the count of patents of type k for drug i, with k varying across specifications. Tobit standard errors are computed from the observed information matrix.

Beyond the patent counts, the control vector $\mathbf{X}_i$ collects: the span in calendar years from the earliest to the latest listed patent expiration, which captures the breadth of the patent portfolio independently of count; an indicator for whether the molecule ever received a Paragraph IV certification, which absorbs the selection of commercially attractive portfolios into challenges; the number of competing active ingredients, defined as the count of other approved active ingredients in the drug's ATC Level-2 class, which measures the intensity of therapeutic competition the drug faces; ATC Level-1 fixed effects, which absorb systematic differences in entry timing across broad therapeutic classes; firm fixed effects, which absorb time-invariant differences across parent firms; a brand-approval-year linear trend, which accounts for the secular change in exclusivity duration over the sample period; and the NCE, orphan, and pediatric exclusivity dummies, which control for the three main forms of non-patent regulatory protection under Hatch–Waxman and related statutes.

To complement the Tobit, I also estimate a Cox proportional-hazards model using the same covariate set except the firm fixed effects, with first generic approval as the failure event. A neg-

ative Cox coefficient implies a lower generic-entry hazard and hence longer effective exclusivity. The combination of a Tobit model and a Cox model provides a more robust characterization of the relationship between secondary patents and entry timing.

Table 2 reports the estimates. Columns 1 and 3 report the baseline result without the control set, estimated on the same sample, and columns 2 and 4 add the full controls. The results are intuitive. Even without controls (column 1), method-of-use patents predict later entry (0.629; $p < 0.01$) and primary patents enter positively (1.171; $p < 0.05$), while formulation and ingredient-level patents show negative raw associations (-0.285 and -0.224)—the signature of selection, as heavily patented molecules are commercially attractive and draw generic entry effort. Adding the controls (column 2) mitigates that selection and leaves the delay pattern in place: each additional method-of-use patent is associated with 0.569 additional years to first generic entry ($p < 0.01$), each primary compound patent with 1.015 years ($p < 0.05$), and each device patent with 1.610 years ($p < 0.01$), while the formulation coefficient is close to zero and the ingredient-level coefficient turns marginally positive (0.260; $p < 0.1$). The Cox estimates in columns 3–4 mirror the pattern in hazard terms. The magnitudes are economically meaningful: the compound patent extends exclusivity by roughly a year, each additional method-of-use patent roughly 0.6 years, and a device patent over a year and a half.

In particular, the size of the method-of-use coefficient is intuitive, as it is consistent with the prevalent "skinny label" problem. A method-of-use patent can be listed separately for each approved indication, so a brand firm accumulates a distinct legal barrier every time it adds a therapeutic use. A generic may in principle "carve out" the patented indication from its label—commonly referred to as a skinny label—rather than litigate it, but the carve-out is imperfect in practice, as a generic approved for a subset of indications is a weaker substitute for the brand. Each added indication can thus create a barrier that patent types cleared only once do not.

The estimates are robust to the main alternative explanations. Two confounders could generate the positive association without strategic patenting: secondary patents may proxy for the structural complexity of the underlying molecule, since complex drugs take longer to reverse-

	Tobit (years to generic)		Cox (log-hazard)	
Patent Type	(1)	(2)	(3)	(4)
# Primary (compound) patents	1.171**	1.015**	−0.206**	−0.250*
	(0.497)	(0.427)	(0.091)	(0.136)
# Formulation patents	−0.285***	−0.083	0.046***	0.017
	(0.104)	(0.101)	(0.018)	(0.017)
# Method-of-use patents	0.629***	0.569***	−0.109**	−0.102***
	(0.220)	(0.196)	(0.042)	(0.039)
# Ingredient-level patents	−0.224*	0.260*	0.050**	−0.015
	(0.134)	(0.138)	(0.021)	(0.027)
# Device patents	0.830	1.610***	−0.142	−0.233**
	(0.516)	(0.551)	(0.110)	(0.112)
Patent-term span	No	Yes	No	Yes
Paragraph IV certification	No	Yes	No	Yes
Class competition	No	Yes	No	Yes
Exclusivity indicators	No	Yes	No	Yes
Approval-year trend	No	Yes	No	Yes
ATC Level-1 FE	No	Yes	No	Yes
Firm FE	No	Yes	No	No
N	585		566	

Notes: This table reports estimates of equation (1). The unit of observation is the active ingredient; the sample is the small-molecule new molecular entities approved in 1995–2019 with a non-negative observation window and non-missing controls, and columns 1 and 3 use the samples of columns 2 and 4. Columns 1–2 report Tobit estimates; the dependent variable is years from brand approval to first generic entry, right-censored at 2019 minus the approval year for ingredients without observed entry. Columns 3–4 report Cox proportional-hazards log-hazard coefficients for the same event; a negative coefficient indicates a lower entry hazard. Patent counts are by type as classified in Section 3. Columns 2 and 4 add the patent-term span (years from the earliest to the latest listed patent expiration), an indicator for any Paragraph IV certification, the number of competing active ingredients in the ATC Level-2 class, indicators for new-chemical-entity, orphan-drug, and pediatric exclusivity, a linear approval-year trend, and ATC Level-1 fixed effects; column 2 also includes firm fixed effects. Tobit standard errors are from the observed information matrix; Cox standard errors are clustered by ATC Level-2 class. Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 2: Secondary Patents and Generic Entry: Tobit and Cox Estimates

engineer regardless of patenting, or for firm capability, since experienced firms file more patents of all types and also sustain stronger market positions. Appendix B addresses both: controlling for molecular complexity with the log Bertz complexity index—a graph-theoretic measure from computational chemistry that scores a molecule's structural complexity from the bond

structure and atom diversity of its molecular graph—only attenuates the method-of-use coefficient by roughly 17% (from 0.569 to 0.475), and adding firm-experience controls leaves it unchanged (0.569); in neither case does the estimate lose significance, and the Bertz index itself is insignificant in this cohort. In addition, the results are robust to the conditioning set—the method-of-use estimate remains significant under every combination of the two mediator controls, including with both dropped—and to replacing the approval-year trend with decade fixed effects (estimates not reported); and retaining molecules approved for indications spanning several therapeutic classes, if anything, attenuates the estimates, so the baseline is conservative on that margin (Appendix C). Controlling directly for the commercial size of the drug's market with a pre-entry Medicaid-reimbursement revenue proxy in the spirit of Newham et al. (2025), on the subsample where the proxy is measurable, attenuates the method-of-use estimate by roughly a tenth and leaves it significant at the one-percent level, while confirming that higher pre-entry revenue itself accelerates entry (Appendix B). Finally, because the counts include patents listed at any point in the drug's observed life, they are not predetermined with respect to entry; a landmark specification that measures each portfolio at a fixed drug age and restricts the sample to molecules still generic-free at that age leaves the method-of-use point estimate essentially unchanged (0.45–0.84 across landmark ages), though less precisely estimated at the shorter landmarks, since the predetermined windows retain little of the late-filed method-of-use variation (Appendix B).

5.2 The Incentive to Patent: Evidence from Class-Level Generic Competition

The results in Section 5.1 characterize the entry-deterrence effects of secondary patenting, but do not speak to why firms accumulate these patents in the first place. This section therefore examines the incentive to accumulate them, using variation in the competitive state of a drug's therapeutic class.

The intuition follows Dafny (2005) and Ellison and Ellison (2011): if evergreening is a

byproduct of ordinary product development rather than a strategic investment in exclusivity, its intensity should vary monotonically with the attractiveness of the market. Under strategic deterrence, by contrast, a firm has little incentive to erect barriers in a class no generic manufacturer is targeting, and equally little incentive once the class has been commoditized and the outcome is no longer in doubt, so the incentive is strongest at intermediate levels of competitive pressure. An inverted-U relationship between competitive pressure and defensive investment is therefore the signature of strategic behavior: monotonicity is the null, and an interior extremum is the object of interest.

Following this intuition, I test monotonicity directly in this setting. I measure competitive pressure by the number of other molecules in a drug's ATC Level-3 class that have lost exclusivity. Let $R_{c(i),t}$ denote the count of molecules in drug i's class, excluding i itself, for which a first generic entrant has been approved as of year t.

Two properties make this a useful source of variation. First, it is predetermined and exogenous to the focal firm: a rival's patents were filed earlier, and their expiration dates were fixed at grant. The timing of a rival molecule's loss of exclusivity is thus determined exogenously to the focal firm's patenting choices. Second, it is directly informative about the focal drug's own competitive position: each additional generic competitor in a drug's class raises the drug's own annual probability of facing a generic entrant by 0.028, comparable in size to the baseline annual hazard of 0.027 itself (estimated with molecule fixed effects on 9,395 drug-years with 250 entry events), so class-level entry is a leading indicator of own exposure rather than a substitute for it; Appendix D reports this first-stage relationship.

The estimating equation is a drug-year panel regression,

$$n_{it} = f\left(R_{c(i),t}\right) + \alpha_i + \mu_t + \lambda_{a(i,t)} + \varepsilon_{it}, \tag{2}$$

where n_{it} is the number of secondary patents newly listed for active ingredient i in year t; $f(\cdot)$ is left unrestricted; α_i is a molecule fixed effect; μ_t is a calendar-year fixed effect; and $\lambda_{a(i,t)}$

is a fixed effect for the drug's age in years since brand approval. The sample runs 2000–2019 and comprises 7,248 drug-year observations on the 585 molecules of the entry-delay sample.

Because the object of interest is the shape of f, I use the adaptive rank-based monotonicity test of Chetverikov (2019), which rejects when the data contain a statistically significant rising (or falling) stretch anywhere in the support of the regressor. The test requires no functional-form assumption and delivers precisely the object the Ellison–Ellison logic calls for: a direct test of monotonicity against the strategic alternative.

Table 3 reports the tests, applied to the drug-year panel of equation (2) with the count of newly listed secondary patents as the outcome. Non-monotonicity is established directly by the test: both one-sided nulls are rejected—the rising stretch with $p = 0.033$ and the falling stretch with $p = 0.036$—so monotonicity in either direction is rejected at $p = 0.036$. Moreover, the estimated profile displays exactly the inverted-U shape that the intuition of Ellison and Ellison (2011) and Dafny (2005) predicts. Figure 5 plots the partialled profile $\hat{f}(R)$: it lies above the uncontested-class level through the intermediate range, with statistically distinguishable positive deviations at three, four, and eleven generic competitors (+0.11, +0.13, and +0.37 new listings per drug-year, against a sample mean of 0.29), and falls sharply and significantly at the top of the support (−0.37 at fourteen competitors). The cluster-bootstrap 90 percent set for the peak location spans four to thirteen, suggesting an interior extremum.

Two robustness checks tighten the reading. First, the same tests applied to method-of-use listings—a secondary patent type with a robust entry-delay effect documented in Section 5.1—reject monotonicity in the baseline sample as well (rising stretch $p = 0.040$, falling stretch $p = 0.003$; joint $p = 0.040$). The margin firms adjust as their class commoditizes therefore includes a margin that postpones entry. Second, normalizing the state variable by class size using the share of the class that has gone generic, discretized to deciles of its positive support, produces robust results: the rising stretch remains strongly significant ($p = 0.004$), with a marginal falling stretch ($p = 0.079$) and a joint rejection at the ten-percent level. Since three generic competitors in a class of four is a different environment from three in a class of twenty,

Null hypothesis	T	p
f nonincreasing (rejection = rising stretch exists)	5.33	0.033
f nondecreasing (rejection = falling stretch exists)	5.06	0.036
Monotone in either direction (joint)	$p = \mathbf{0.036}$	
Molecule FE	Yes	
Year FE	Yes	
Drug-age FE	Yes	
Observations	7,248	
Molecules	585	

Notes: This table reports Chetverikov (2019) monotonicity tests of the profile in equation (2): the test statistic T and its p-value for each null hypothesis. The unit of observation is the active ingredient–year in 2000–2019. The dependent variable is the number of secondary Orange Book patents newly listed for the active ingredient in the year (mean 0.29), partialled out of molecule, calendar-year, and drug-age fixed effects jointly with unrestricted indicators for the state variable. The state variable is the number of other molecules in the same ATC Level-3 class with an approved first generic as of the year. The statistic is the studentized maximum of signed-pair sums with Epanechnikov weights at bandwidths $\{1.5, 3, 6, 14\}$; p-values are from a multiplier bootstrap with 999 draws and one standard-normal multiplier per molecule.

Table 3: Testing Monotonicity of Secondary Listings in Class-Level Generic Competition

the share is a defensible alternative measure.

5.3 Crowd-Out of Fundamental Innovation

Sections 5.1 and 5.2 establish that secondary patents delay entry and that their accumulation carries strategic incentives. Whether that accumulation comes at the expense of fundamental research is, however, a separate question. This section therefore goes a step further and asks whether secondary patenting crowds out investment in fundamental innovation, using within-firm variation in patenting intensity and subsequent clinical-trial activity. Because research budgets are set at the firm level, the unit of observation is a firm-year, and R&D commitment is measured by clinical trial counts:

$$\ln(\text{Trials}_{ft} + 1) = \phi_k \sum_{j=k-1}^{k+1} \text{SecPat}_{f,t-j} + \alpha_f + \mu_t + \varepsilon_{ft}, \tag{3}$$

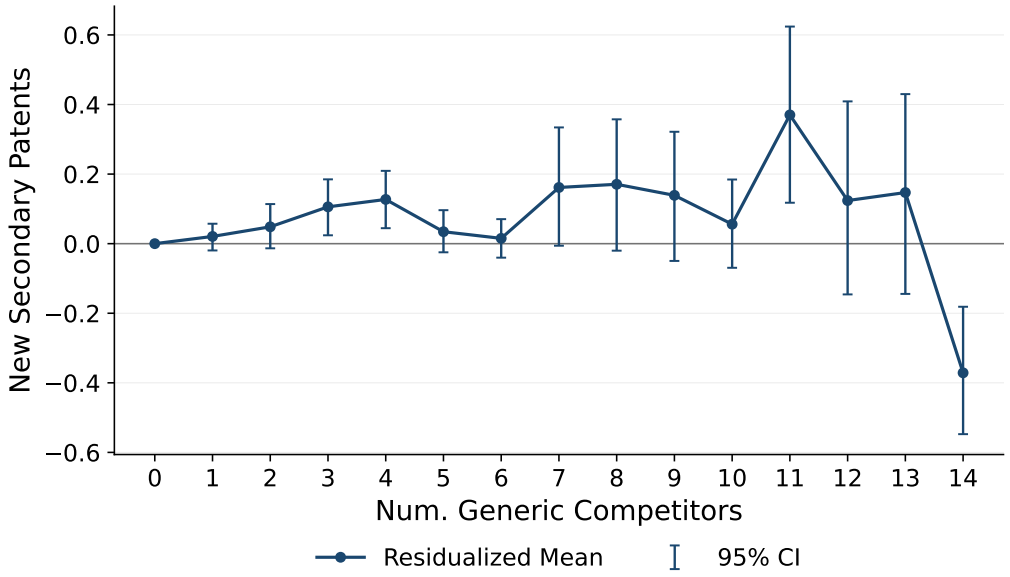


Figure 5: New Secondary Listings and Class-Level Generic Competition: The Inverted U

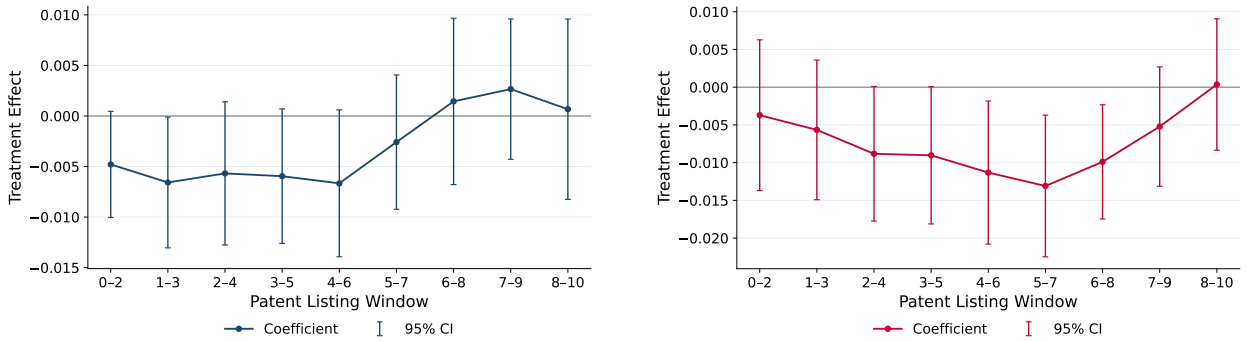
Notes: The figure plots the profile $\hat{f}(R)$ of equation (2): the mean number of secondary Orange Book patents newly listed per drug-year at each value of the class state, after partialling out molecule, calendar-year, and drug-age fixed effects jointly with unrestricted state indicators, normalized to zero at an uncontested class. The sample is molecules approved in 1995 or later. The state variable is the number of other molecules in the drug's ATC Level-3 class with an approved first generic as of the year. Vertical bars are 95 percent confidence intervals from a multiplier bootstrap with 1,999 draws and one standard-normal multiplier per molecule; zero rivals is the omitted category. Monotonicity tests are reported in Table 3.

where Trials_{ft} is the number of clinical trials of a given phase registered by firm f in year t, and the regressor is the firm's count of newly listed secondary patents summed over the three-year window centered k years earlier, winsorized at the 99th percentile. Because listings are lumpy at the firm-year level—a single product's patent cluster can dominate any one year—the three-year window smooths idiosyncratic listing timing; the equation is estimated separately for each window center $k = 1, \ldots, 9$, tracing out the horizon profile of the association. The firm fixed effects α_f absorb time-invariant differences in firm scale and research orientation, and the calendar-year fixed effects μ_t absorb aggregate trends in trial registration, so that ϕ_k is identified from within-firm changes in patenting intensity. Each regression is estimated separately for Phase 2 and Phase 3 trial counts on the post-2008 sample, where FDAAA mandatory registration ensures near-complete coverage; standard errors are clustered by firm. The profile brackets the horizon at which diverted development resources would surface in the trial

pipeline: progression from preclinical work to proof-of-concept and pivotal trials takes roughly four to seven years (DiMasi et al., 2016). In the standard development sequence, Phase 1 trials are small safety studies, Phase 2 trials provide the first evidence of efficacy in patients, and Phase 3 trials are the large pivotal studies that support approval; the analysis focuses on Phases 2 and 3 because they are the discretionary, capital-intensive commitments in which a reallocation of development resources would surface.

Figure 6 traces the profiles for Phase 2 and Phase 3 trials, the two most capital-intensive commitments in the primary innovation pipeline (about $59 million per compound in Phase 2 and $255 million in Phase 3, in 2013 dollars; DiMasi et al., 2016) and the margins most relevant for new-drug output. For Phase 2 (panel a), the association is negative from the earliest horizons: the coefficient sits between −0.005 and −0.007 for every window through 4–6 years, so a firm that lists ten additional secondary patents registers roughly five to seven percent fewer proof-of-concept trials over the following several years. From the 5–7 window onward the profile returns to zero. For Phase 3 (panel b), the profile is U-shaped: it deepens steadily from the 0–2 window to a trough of −0.0131 at 5–7 years ($p < 0.01$; the 2–4 through 6–8 windows are all significant at conventional levels), then decays back to zero by 8–10 years. At the trough, ten additional listings correspond to roughly thirteen percent fewer pivotal trials.

Two features of the profiles support the reallocation interpretation. First, the timing: the Phase 3 displacement is deepest five to seven years after the patenting activity, matching the four-to-seven-year horizon at which diverted resources would reach the pivotal-trial pipeline, while the Phase 2 displacement appears earlier, consistent with proof-of-concept commitments sitting nearer the start of the pipeline. Second, both profiles decay to zero at long horizons rather than settling at a permanently lower level: patenting effort displaces the trials it would otherwise have funded at the horizon where those trials would have entered the pipeline, rather than shrinking the firm's research capacity permanently. The underlying single-year lag estimates are reported in Appendix E and trace the same pattern.



(a) Phase 2 trials

(b) Phase 3 trials

Figure 6: Crowd-Out of Clinical Trial Activity: Windowed Lag Profiles

Notes: Each point reports the coefficient from a separate estimate of equation (3): ln(trial count+1) regressed on the firm's number of newly listed secondary patents summed over the indicated three-year window, winsorized at the 99th percentile. The unit of observation is the firm-year; the sample is 948 firm-years (79 firms) in 2008–2019, the period of mandatory ClinicalTrials.gov registration. All regressions include firm and year fixed effects. Standard errors are clustered by firm; vertical bars are 95 percent confidence intervals. Single-year lag estimates are reported in Appendix E.

6 Quantifying Welfare

The reduced-form evidence establishes that secondary patents delay generic entry and that firms adjust their accumulation to the competitive state of the class in the pattern strategic deterrence predicts, but the sign of the net effect on consumers is still uncertain, because the delay a patent causes and the quality it may protect are measured in different units and accrue on different time bases. This section quantifies both using a structural demand model.

6.1 Demand Model

The unit of analysis is a product d, defined as an active ingredient in a given formulation interacted with its brand/generic tier, with all generic manufacturers of a molecule pooled into a single composite generic product, observed in therapeutic market m, defined as an ATC Level-3 class in a state, and year t.

Each beneficiary–prescriber pair j chooses among the products available in $\mathscr{D}_{mt}$ plus an outside option of consuming no drug in the class. The indirect utility of pair j from product d,

and the decomposition of product quality, are

$$
\begin{aligned}
u_{jdmt} &= \underbrace{\alpha \ln p_{dmt} + \psi_{dt} + \tilde{\xi}_{dmt}}_{\delta_{dmt}} + \varepsilon_{jdmt}, \\
\psi_{dt} &= \sum_{k \in \mathcal{K}} \lambda_k \ln\left(1 + N^k_{dt}\right) + x'_{dt}\beta + \eta_d + \tau_t + e_{dt},
\end{aligned}
\tag{4}
$$

where δ_{dmt} is mean utility and ε_{jdmt} is an i.i.d. Type-I extreme-value shock. p_{dmt} is the beneficiary's out-of-pocket price per 30-day fill, and ψ_{dt} is the product's quality in year t, common to all states. Controls x_{dt} include indicators for generic version of the molecule, for holding a primary compound patent, and for event time since the drug's launch; η_d is a product fixed effect and τ_t a year effect. I additionally control for the log counts $\ln N^g$ and $\ln N^b$ of the sellers pooled into each tier of the molecule: N^g counts the generic manufacturers folded into the molecule's composite generic product, and N^b the branded products folded into its brand tier. These enter as aggregation controls, as discussed below.

The objects of interest are the coefficients λ_k on the active stocks N^k_{dt} of primary and secondary patents of each type $k \in \mathcal{K}$. On the demand side, types follow the Orange Book's own patent flags, which attach to every listing in every year of the demand sample: formulation patents are those flagged as claiming the drug product, method-of-use patents those listed with a use code, and the remaining secondary patents form a residual type. Device patents, which claim the drug product as delivered, fall under the drug-product flag, while ingredient-level patents on salts, polymorphs, and enantiomers have no flag of their own. The stocks enter in logs, so $\lambda_k \ln 2$ is the utility of doubling $1 + N^k_{dt}$, and the value of one more patent declines with the size of the stock. A positive λ_k implies that patents of type k are associated with quality improvement from the consumer's perspective, while $\lambda_k = 0$ means they function purely as exclusivity instruments.

The model is a plain logit. The public files report drug-level aggregates by state rather than person-level records, so there is no individual variation with which to discipline random coefficients or demographic interactions; heterogeneity across markets is instead absorbed by

the class-by-state-by-year fixed effects of the estimation below.

Pooling all generic manufacturers of a molecule into a single composite product follows Ching (2010): physicians and patients typically cannot choose which manufacturer's generic is dispensed, because the decision rests with the pharmacist under state substitution laws (Song and Barthold, 2018). Pooling has a mechanical consequence for measured utility. Under Type-I extreme-value errors, a composite of N symmetric alternatives that share mean utility δ has inclusive value $\delta + \ln N$, so the pooled tier's utility rises with the number of sellers behind it purely through aggregation. The log seller counts $\ln N^g$ and $\ln N^b$ absorb this variety term, which is equivalent to modeling the composite as N symmetric products rather than one (Ackerberg and Rysman, 2005, see also Gowrisankaran and Rysman, 2012). They thereby discipline the generic indicator, which would otherwise conflate the utility of the generic tier with the mechanical effect of the number of pooled sellers.

6.2 Estimation

The model is estimated in two steps on market shares and out-of-pocket prices constructed from the Part D file (Section 3). The first step regresses the logit inversion $\ln s_{dmt} - \ln s_{0mt}$ on $\ln p_{dmt}$ with product-by-year fixed effects, which recover ψ_{dt}, and class-by-state-by-year fixed effects, so that α is identified from differences in cost sharing for the same product in the same year across states. The second step regresses the recovered qualities $\hat{\psi}_{dt}$ of branded products on their log active patent stocks, the controls, and product and year fixed effects, so that λ_k is identified from within-product changes in the product's own active patents as patents are listed and expire.

Price is instrumented with a formulary-placement shifter: for each product, state, and year, the out-of-pocket cost per fill implied by the product's tier placement and the benefit design of every Part D plan offered in the state, weighted by plan enrollment. The first stage is strong (cluster-robust $F = 27.4$, clustering by state). Appendix F describes the instrument in detail.

The potential market M_{mt} is twelve times the number of Part D enrollees in the state and

year, so that shares are 30-day fills per enrollee-month. It is the same for every class, so with class-by-state-by-year fixed effects it does not affect the estimates and matters only for scaling welfare, which is reported per user of the class and in total over the class's market. Appendix G describes the construction of the Medicare segment and of the companion Medicaid segment.

6.3 Demand Estimates

Table 4 reports the estimates of equation (4). The coefficient on log out-of-pocket price is −0.836 and precisely estimated. Generic products carry higher net utility than their branded counterparts at an equal price (+1.562); conditional on the out-of-pocket price, this premium is best read as the institutional steering that out-of-pocket prices do not capture, as tiered formularies, state substitution laws, and default generic dispensing channel prescriptions toward the generic tier at any given copay. The generic aggregation term $\ln N^g$ enters with +1.346, so a molecule's generic tier gains utility as manufacturers enter, and utility rises slightly with years since the product's own market entry (+0.044).

The patent coefficients are the substantive result. Formulation patents enter with +0.393 and method-of-use patents with +0.327, so a larger active stock of either type raises a branded product's mean utility; the residual type of secondary patents enters negatively (−0.350). Primary compound patents enter with +0.422. Consumers therefore appear to value the improvements that formulation and method-of-use patents protect.

6.4 The Welfare Effect of One Additional Patent

To value one additional secondary patent of type k for a given molecule, I compare the discounted consumer-surplus stream the market delivers with and without the patent, evaluated at observed prices. The patent does two things: it protects a quality improvement worth $\Lambda_k = \lambda_k[\ln(1 + N^k) - \ln N^k]$, the increment in the brand's logged stock of type k that the patent completes, in the brand's mean utility, and it postpones generic entry by X_k years, the per-patent delay taken from the Tobit regression of Table 2. Discounting at $r = 3\%$ with

$\delta = 1/(1+r)$ over a fixed calendar horizon of $Z = 20$ years—the 3% rate is the standard discounting convention in health-economic evaluation, and the twenty-year horizon approximates a drug's patent-protected commercial life, matching the statutory patent term—the net consumer-welfare effect of the patent is

$$\Delta W_k = \sum_{t=0}^{X_k-1} \delta^t \underbrace{w_2}_{\substack{\text{additional patent,}\\ \text{no generic entry}}} + \sum_{t=X_k}^{Z-1} \delta^t \underbrace{w_3(t-X_k)}_{\substack{\text{additional patent,}\\ \text{generic competition}}} - \sum_{t=0}^{Z-1} \delta^t \underbrace{w_1(t)}_{\substack{\text{no additional patent,}\\ \text{generic competition}}}, \tag{5}$$

where each w is per-consumer surplus in the indicated state: w_1 under the baseline patent portfolio with generic competition present, w_2 with the additional patent's quality but no generic entry, and w_3 with the additional patent and generic competition. Intuitively, the first two terms are the surplus stream with the patent: for its first X_k years consumers enjoy the protected quality improvement but generic competition is postponed, and thereafter the generic is present alongside the improvement. The third term is the stream without the patent, in which generic competition runs over the full horizon. Post-entry flows are indexed by time since entry, so w_1 and w_3 are not constants, and the fixed calendar window means the with-patent stream forfeits the last X_k years of post-entry surplus off the end of the horizon—the sense in which delay is costly even though the patent eventually expires.

Consumer surplus in each state is the standard logit log-sum, converted to money-metric terms by the marginal utility of income:

$$\mathrm{CS}_{mt} = \frac{\bar{p}_{mt}}{|\alpha|} \ln\left(1 + \sum_{d \in \mathscr{D}_{mt}} e^{\delta_{dmt}}\right), \tag{6}$$

where, because price enters in logs, the marginal utility of a dollar is $|\alpha|/p$; it is evaluated at $\bar{p}_{mt}$, the share-weighted average out-of-pocket price per fill in the observed market, and held fixed across the three states. Dollar magnitudes are reported directly at the estimated price coefficient, with no external calibration.

Base observations are branded product-years in which the molecule's own generic has not

yet entered, so the observed choice set corresponds to the no-generic state. They are evaluated in national Medicare markets, an ATC Level-3 class in a year with fills, costs, and enrollees summed over states, for the molecules of the analysis sample over 2013–2024. Welfare is computed per potential consumer and reported in two measures: the total over the class's market, the value per potential consumer multiplied by the number of Part D enrollees in the market; and the value per user of the class, that total divided by the number of Part D beneficiaries filling each product of the class in the year, summed over products. For w_1 and w_3 the generic must be simulated. Its out-of-pocket price is set at 0.370 times the brand's, the median paired ratio among molecule-years where both tiers are observed, and its mean utility is constructed from the projected generic-tier, aggregation, age, and primary-patent terms of Table 4, carrying over the molecule's unobserved quality net of its secondary-patent quality; generic penetration comes in over three years through the $\ln N^g$ term.

6.5 Results

Table 5 reports the calculation for the method-of-use, device, and formulation types. The headline rows are method-of-use and device patents, the two types with a robust entry delay in Table 2. For each type, the table decomposes the per-patent effect into its components: the annual value of generic availability G, the annual value of the protected quality Q_k, and the delay X_k, and reports the median, mean, and 10th–90th percentile range of the discounted net effect ΔW_k across base observations, together with the share of observations for which the effect is positive.

Method-of-use patents reduce consumer surplus decisively. One additional such patent buys 0.569 years of postponed generic competition, worth \$12.51 per user of the class per year, and returns quality worth \$0.43 per year, or 3.4% of the generic-availability value. The net effect is negative for 96.62% of base observations, with a median loss of \$9.66 per user over the twenty-year horizon; the middle 80 percent of observations spans −\$1,007 to −\$0.02, and the mean loss per user, total loss divided by total users across observations, is \$29.34, well above

the median because the loss scales with the value of generic availability in the drug's market, which is heavily right-skewed—most markets lose small amounts, while a minority of large markets loses much more. Device patents are costlier still: the delay is 1.610 years against protected quality worth \$0.55 per user per year, and the effect is negative for 97.79% of base observations, with a median loss of \$23.27 per user, a middle-80-percent range of −\$2,374 to −\$0.08, and a mean loss of \$67.88 per user. Summed over the class's market (Panel B of Table 5), one additional method-of-use patent costs Medicare beneficiaries a median of \$5.9 million and a mean of \$89.7 million over the twenty-year horizon, and one additional device patent a median of \$14.1 million and a mean of \$207.4 million.

Formulation patents, with no detectable delay, are valued through their quality flow alone and are positive by construction (median \$3.15, mean \$7.44 per user). The ingredient-level type of the entry-delay analysis has no counterpart here: the Orange Book flags that classify the demand-side patent stocks do not separate salts, polymorphs, and enantiomers, which fall into a residual type whose coefficient is negative and cannot be read as protected quality.

The magnitudes are easiest to understand as a race between a large, brief cost and a small, long-lived benefit. Generic availability is worth roughly \$12.51 per user of the class per year, while the quality a method-of-use or device patent protects is worth about \$0.43 or \$0.55 per year—more than twenty times smaller. The delay cost is borne immediately, while the quality accrues only slowly over the product's remaining life, which is why even fractions of a year of postponed competition dominate the comparison. Two features of the calculation make these figures conservative. First, the counterfactuals hold prices fixed, although brand prices, if anything, rise rather than fall as secondary patents accumulate, so fixing prices understates the cost of delay. Second, the calculation values only the products that exist and assigns no cost to the new molecules never developed because of the crowd-out documented in Section 5.3; any plausible valuation of a displaced drug adds substantial further loss through that channel. The figures are therefore best read as a lower bound on the consumer cost per patent.

	Estimate	(SE)
Price and product characteristics		
Log out-of-pocket price (α)	−0.836	(0.114)
Generic indicator	+1.562	(0.146)
Years since own market entry	+0.044	(0.004)
$\ln N^g$ (generic composite)	+1.346	(0.071)
$\ln N^b$ (brand composite)	+0.409	(0.176)
Patent stocks (active)		
# Primary (compound)	+0.422	(0.053)
$\ln(1 + \#\,\text{formulation})$ (λ_{form})	+0.393	(0.069)
$\ln(1 + \#\,\text{method-of-use})$ (λ_{mou})	+0.327	(0.070)
$\ln(1 + \#\,\text{other secondary})$	−0.350	(0.134)
Observations	412,044	

Notes: This table reports estimates of equation (4) on public Medicare Part D data. A product is an active ingredient in a given formulation and brand/generic tier, with all generic manufacturers of a molecule pooled; a market is an ATC Level-3 class in a state and year. The first step regresses the logit inversion on log out-of-pocket cost per 30-day fill with product-by-year and class-by-state-by-year fixed effects, instrumenting price with the formulary-placement cost of Appendix F; standard errors are clustered by state. The second step regresses the recovered product-year quality of branded products on the log active stocks of secondary patents, classified by Orange Book flags, indicators for generic availability, for holding a primary patent, and for event time relative to the molecule's first generic entry, and product and year fixed effects; standard errors are clustered by product. The generic indicator, years since entry, $\ln N^g$, $\ln N^b$, and the primary-patent count are from a projection of every product's price-free quality on these characteristics with molecule-formulation and class-by-year fixed effects; standard errors are clustered by molecule-formulation. Observations are the product–state–year cells of the first step. Standard errors in parentheses.

Table 4: Demand Estimates

Patent type	Quality λ_k	Entry delay X_k (yrs)	Generic value G (per yr)	Quality Q_k (per yr)	ΔW_k median	ΔW_k mean	ΔW_k p10–p90 range	Share with $\Delta W_k > 0$
Panel A. Per user of the class ($)								
Method-of-use	0.327	0.569	12.51	0.43	−9.66	−29.34	[−1,007.40, −0.02]	3.38%
Device	0.393	1.610	12.39	0.55	−23.27	−67.88	[−2,374.10, −0.08]	2.21%
Formulation	0.393	—	12.39	0.55	3.15	7.44	[0.01, 297.71]	100%
Panel B. Total over the class's market: value per potential consumer × Part D enrollees ($ million)								
Method-of-use	0.327	0.569	38.24	1.31	−5.94	−89.68	[−150.68, −0.04]	3.37%
Device	0.393	1.610	37.87	1.68	−14.07	−207.45	[−341.90, −0.09]	2.21%
Formulation	0.393	—	37.87	1.68	1.84	22.75	[0.02, 41.92]	100%

Notes: This table reports the consumer-welfare effect of one additional secondary patent, equation (5). The sample is the 2,760 branded product-years (366 molecules, 881 national Medicare class-year markets, 2013–2024) in which the molecule's own generic has not yet entered. Panel A reports dollars per user of the drug's ATC Level-3 class, where users are the Part D beneficiaries filling each product of the class in the year, summed over products; it omits six observations without a user count, and its means are weighted by users. Panel B reports millions of dollars summed over the class's market: the value per potential consumer times the number of Part D enrollees in the class-year market (median 43.6 million). λ_k is the coefficient on the log stock of the corresponding Orange Book type in Table 4; device patents carry the drug-product (formulation) flag. X_k is the entry delay per patent from Table 2. $G = w_1 - w_2$ is the annual value of generic availability net of the patent's quality and $Q_k = w_3 - w_1$ the annual value of the patent's quality, averaged over observations. ΔW_k is the net effect discounted at $r = 3\%$ over $Z = 20$ years, with generic penetration ramping in over three years and the quality improvement attributed entirely to the patent; the median and 10th–90th percentile range are across observations.

Table 5: Consumer Welfare of One Additional Secondary Patent

7 Conclusion

This paper asks how the accumulation of secondary patents affects generic entry and innovation, what incentives drive the accumulation, and what it costs consumers. Four findings emerge.

First, patents delay generic entry, and the delay tracks the legal blocking power of the patent type. Each compound patent is associated with roughly one additional year of effective exclusivity, each method-of-use patent with roughly 0.6 years, and each device patent with over a year and a half. The method-of-use asymmetry has an institutional explanation: such patents can be listed separately for each approved indication, and the section viii carve-out that would in principle let a generic avoid them has proved imperfect in practice.

Second, the intensity of secondary patenting varies with the competitive state of a drug's therapeutic class in the inverted-U pattern that strategic deterrence predicts and ordinary product development does not. Monotonicity is rejected outright ($p = 0.036$), with an interior peak in the range Ellison and Ellison (2011) identify as the signature of strategic investment.

Third, there is suggestive evidence that the accumulation crowds out fundamental research: within firms, heavier secondary patenting predicts fewer Phase 2 and especially fewer Phase 3 trials over the following three to seven years, with the Phase 3 displacement deepest five to seven years out—exactly the horizon at which diverted development resources would reach the pivotal-trial pipeline—and both profiles fading to zero at longer horizons, consistent with a temporary reallocation of the research budget away from new molecules.

Fourth, for the patents that display deterrence effects, the cost to consumers outweighs the quality gain. The per-patent welfare effect is negative for at least 96 percent of the observations on which it is evaluated, with mean losses of roughly $90 million to $207 million per additional patent, summed over the Medicare Part D market of the drug's therapeutic class over a twenty-year horizon, or $29 to $68 per user (medians of $6 million to $14 million, the means pulled down by a tail of large markets), while the patent types with no detectable delay protect quality that consumers appear to value.

Two limitations bound these conclusions. The welfare counterfactuals hold prices fixed; because brand prices do not fall as secondary patents accumulate, this if anything understates the consumer harm. And the static calculation assigns no cost to molecules never developed, a channel that is plausibly substantial but not quantified here.

References

Acemoglu, D. and Linn, J. (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *Quarterly Journal of Economics*, 119(3):1049–1090.

Ackerberg, D. A. and Rysman, M. (2005). Unobserved product differentiation in discrete-choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics*, 36(4):771–788.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., and Howitt, P. (2005). Competition and innovation: An inverted-u relationship. *The quarterly journal of economics*, 120(2):701–728.

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pages 609–626. Princeton University Press, Princeton, NJ.

Bertz, S. H. (1981). The first general index of molecular complexity. *Journal of the American Chemical Society*, 103(12):3599–3601.

Blume-Kohout, M. E. and Sood, N. (2013). Market size and innovation: Effects of medicare part d on pharmaceutical research and development. *Journal of Public Economics*, 97:327–336.

Boldrin, M. and Levine, D. K. (2013). The case against patents. *Journal of Economic Perspectives*, 27(1):3–22.

Budish, E., Durvasula, M. M., Roin, B. N., and Williams, H. L. (2025). Missing markets for innovation: Evidence from new uses of existing drugs. Technical report, National Bureau of Economic Research.

Budish, E., Roin, B. N., and Williams, H. (2015). Do firms underinvest in long-term research? evidence from cancer clinical trials. *American Economic Review*, 105(7):2044–2085.

Caves, R. E., Whinston, M. D., and Hurwitz, M. A. (1991). Patent expiration, entry, and competition in the u.s. pharmaceutical industry. *Brookings Papers on Economic Activity: Microeconomics*, pages 1–66.

Chetverikov, D. (2019). Testing regression monotonicity in econometric models. *Econometric Theory*, 35(4):729–776.

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

Cockburn, I. M., Lanjouw, J. O., and Schankerman, M. (2016). Patents and the global diffusion of new drugs. *American Economic Review*, 106(1):136–164.

Dafny, L. S. (2005). Games hospitals play: Entry deterrence in hospital procedure markets. *Journal of Economics & Management Strategy*, 14(3):513–542.

DiMasi, J. A., Grabowski, H. G., and Hansen, R. W. (2016). Innovation in the pharmaceutical industry: New estimates of r&d costs. *Journal of Health Economics*, 47:20–33.

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

Duggan, M., Garthwaite, C., and Goyal, A. (2016). The market impacts of pharmaceutical product patents in developing countries: Evidence from india. *American Economic Review*, 106(1):99–135.

Ellison, G. and Ellison, S. F. (2011). Strategic entry deterrence and the behavior of pharmaceutical incumbents prior to patent expiration. *American Economic Journal: Microeconomics*, 3(1):1–36.

Farrell, J. and Shapiro, C. (2008). How strong are weak patents? *American Economic Review*, 98(4):1347–1369.

Federal Trade Commission (2023). Federal trade commission statement concerning brand drug manufacturers' improper listing of patents in the orange book. Policy statement, September 2023.

Feldman, R. (2018). May your drug price be evergreen. *Journal of Law and the Biosciences*, 5(3):590–647.

Frakes, M. D. and Wasserman, M. F. (2025). Strategic patenting: Evidence from the biopharmaceutical industry. NBER Working Paper 34024, National Bureau of Economic Research.

Galasso, A. and Schankerman, M. (2015). Patents and cumulative innovation: Causal evidence from the courts. *Quarterly Journal of Economics*, 130(1):317–369.

Goettler, R. L. and Gordon, B. R. (2011). Does amd spur intel to innovate more? *Journal of Political Economy*, 119(6):1141–1200.

Goolsbee, A. and Syverson, C. (2008). How do incumbents respond to the threat of entry? evidence from the major airlines. *Quarterly Journal of Economics*, 123(4):1611–1633.

Gowrisankaran, G. and Rysman, M. (2012). Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, 120(6):1173–1219.

Grabowski, H. and Kyle, M. (2007). Generic competition and market exclusivity periods in pharmaceuticals. *Managerial and Decision Economics*, 28(4-5):491–502.

Grabowski, H. G. and Vernon, J. M. (1992). Brand loyalty, entry, and price competition in pharmaceuticals after the 1984 drug act. *The journal of law and economics*, 35(2):331–350.

Gupta, C. (2023). One product, many patents: Imperfect intellectual property rights in the pharmaceutical industry. Working paper, UCLA Anderson School of Management. SSRN Working Paper.

Hemphill, C. S. and Sampat, B. N. (2011). When do generics challenge drug patents? *Journal of Empirical Legal Studies*, 8(4):613–649.

Hemphill, C. S. and Sampat, B. N. (2012). Evergreening, patent challenges, and effective market life in pharmaceuticals. *Journal of health economics*, 31(2):327–339.

Igami, M. (2017). 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.

Krieger, J., Li, D., and Papanikolaou, D. (2022). Missing novelty in drug development. *Review of Financial Studies*, 35(2):636–679.

Lemus, J. and Ozkul, E. (2023). Product hopping and innovation incentives. *Journal of Competition Law and Economics*, 19(4):580–609.

Lerner, J. (2009). The empirical impact of intellectual property rights on innovation: Puzzles and clues. *American Economic Review*, 99(2):343–348.

Miller, J. E., Korn, D., and Ross, J. S. (2015). Clinical trial registration, reporting, publication and FDAAA compliance: A cross-sectional analysis and ranking of new drugs approved by the FDA in 2012. *BMJ Open*, 5(11):e009758.

Morton, F. M. S. (1999). Entry decisions in the generic pharmaceutical industry. *The Rand journal of economics*, pages 421–440.

Myers, K. and Pauly, M. V. (2019). Endogenous productivity of demand-induced r&d: evidence from pharmaceuticals. *RAND Journal of Economics*, 50(3):591–614.

Newham, M., Seldeslachts, J., and Banal-Estañol, A. (2025). Common ownership and market entry: Evidence from the pharmaceutical industry. *American Economic Journal: Microeconomics*, 17(4).

Reiffen, D. and Ward, M. R. (2005). Generic drug industry dynamics. *Review of Economics and Statistics*, 87(1):37–49.

Righi, C. and Simcoe, T. (2023). Patenting inventions or inventing patents? continuation practice at the uspto. *RAND Journal of Economics*, 54(3):480–508.

Shapiro, C. (2001). Navigating the patent thicket: Cross licenses, patent pools, and standard setting. In *Innovation Policy and the Economy*, volume 1, pages 119–150. MIT Press.

Song, Y. and Barthold, D. (2018). The effects of state-level pharmacist regulations on generic substitution of prescription drugs. *Health Economics*, 27(11):1717–1737.

Stergiopoulos, S., Getz, K. A., and Blazynski, C. (2019). Evaluating the completeness of ClinicalTrials.gov. *Therapeutic Innovation & Regulatory Science*, 53(3):307–317.

Tenn, S. and Wendling, B. W. (2014). Entry threats and pricing in the generic drug industry. *Review of Economics and Statistics*, 96(2):214–228.

Thomas, J. R. (2009). Patent "evergreening": Issues in innovation and competition. *Congressional Research Service*, 1(1):10–11.

Williams, H. L. (2013). Intellectual property rights and innovation: Evidence from the human genome. *Journal of Political Economy*, 121(1):1–27.

Zarin, D. A., Tse, T., Williams, R. J., and Carr, S. (2016). Trial reporting in ClinicalTrials.gov—the final rule. *New England Journal of Medicine*, 375(20):1998–2004.

A Data Linkage

This appendix describes how the data sources of Section 3 are joined to a common spine of active ingredients. The linkage proceeds in seven steps.

1. *Spine.* Drugs@FDA product records are collapsed to one row per active ingredient, keeping the earliest approval date as the brand approval date. Throughout the FDA sources, products are identified by the pair of application and product numbers, and this pair is the merge key wherever it is available.

2. *Orange Book.* The Orange Book patent and exclusivity files are matched to the spine on the application–product pair. Patent records are then de-duplicated to one row per ingredient–patent pair, keeping the latest listed expiration date, because a single patent may be listed against several strength or route variants of the same molecule. Each patent's type is assigned by the manual classification of Section 3, which reads the patent's Orange Book use code, its title and abstract, and its claim text, and falls back on the Orange Book drug-substance and drug-product flags when the text is uninformative. Exclusivity records are collapsed to ingredient-level indicators for new-chemical-entity, orphan, and pediatric exclusivity.

3. *First generic entry.* The FDA's list of first-time generic approvals is matched to the spine on the application number, giving each ingredient the date of its first approved generic; years to entry is the difference from the brand approval year.

4. *Paragraph IV challenges.* The FDA's Paragraph IV list is reported at the drug level. It is lifted to the ingredient level by taking the earliest challenge date across the molecule's applications, the total number of certified ANDAs, and whether any challenger qualified for the 180-day first-filer exclusivity.

5. *Therapeutic class.* Each product's National Drug Code is mapped to its RxNorm concept, then to WHO ATC classes through the National Library of Medicine's RxClass relations,

aggregated to the ingredient, and parsed into the Level-1 through Level-4 hierarchy. Because the ATC system classifies therapeutic uses rather than molecules, an ingredient approved for several uses can map to more than one class. Each ingredient is assigned a single modal class: the ATC Level-2 code, together with its Level-3 and Level-4 subdivisions, that appears across the largest number of the ingredient's products, with ties broken alphabetically. All class-based measures below—the class-competition control, the clustering of standard errors, and the class competitive state—use this modal class. The number of distinct Level-4 codes an ingredient maps to is retained as its indication count, and Appendix C examines the sensitivity of the entry-delay estimates to the molecules spanning several classes. Coverage is 93.0% of the 1995–2019 sample.

6. *Parent firms and clinical trials.* Sponsor names are mapped to parent firms through a manual crosswalk, and the parents define the firm fixed effects. ClinicalTrials.gov records attach at the sponsor level only—no molecule-level link exists—which is why the crowd-out analysis of Section 5.3 is conducted at the firm-year level.

7. *Medicare Part D and Medicaid.* The Part D file identifies drugs by generic and brand name only. Each name pair is mapped to a product—a molecule, its formulation, and its brand/generic tier—through the National Drug Codes that carry those names, and the product takes the modal ATC Level-3 class of its codes; the Medicaid file reports National Drug Codes directly and is mapped through the same code-to-product table. The maps cover 95.2% of Part D fills and 89.5% of Medicaid prescriptions. Branded products attach Orange Book patents through their application numbers, a patent counting toward year t's active stock if it was listed by t, had not expired by January 1 of t, and had not been removed from the book before t.

B Robustness of the Entry-Delay Estimates: Molecular Complexity, Firm Capability, Patent-Count Timing, and Market Size

This appendix reports the four robustness exercises referenced in Section 5.1, which address the possibility that the positive MOU–exclusivity association reflects the structural complexity of the underlying molecule, the capability of the sponsoring firm, the mechanical accumulation of listings over the drug's observed life, or the commercial size of the drug's market rather than strategic patenting. The first two exercises augment the baseline specification of equation (1) on the baseline sample; the third re-measures the patent counts themselves; and the fourth adds a pre-entry market-size control on the subsample where it is measurable.

Molecular complexity. To control for molecular complexity, I augment the baseline Tobit of equation (1) with the log Bertz complexity index (log(Bertz CT)), a graph-theoretic measure of structural complexity, available for 531 of the 585 baseline molecules. Adding this control attenuates the MOU coefficient by roughly 17% (from 0.569^{***} to 0.475^{**}); the coefficient remains significant at the five-percent level. The Bertz index itself is positive but statistically insignificant in this cohort, so molecular complexity is not an important independent determinant of exclusivity duration here, and it does not account for the method-of-use delay. Table A1 reports the full results.

Firm capability. To control for firm capability, I add firm-level experience measures: the number of prior NME approvals by the same parent firm in the same ATC Level-2 class and total prior NME approvals across all classes. Patent coefficients are virtually unchanged after adding these controls: the MOU coefficient is 0.569^{***} with and without them. Firm ATC experience enters positively but insignificantly (0.197–0.224 across specifications), and total firm experience is small and insignificant, so firm capability does not explain the delay pattern

	Spec (A) Total		Spec (B) Pri/Sec		Spec (D) 5-way	
	Baseline	+Bertz	Baseline	+Bertz	Baseline	+Bertz
Total patent count	0.164**	0.135**	—	—	—	—
	(0.069)	(0.069)				
# Primary (compound) patents	—	—	0.775*	0.627	—	—
			(0.423)	(0.444)		
# Secondary patents	—	—	0.135*	0.111	—	—
			(0.074)	(0.072)		
# Method-of-use patents	—	—	—	—	0.569***	0.475**
					(0.196)	(0.196)
log(Bertz CT)	—	0.561	—	0.524	—	0.676
		(0.496)		(0.496)		(0.496)
Patent-term span	Yes	Yes	Yes	Yes	Yes	Yes
Paragraph IV certification	Yes	Yes	Yes	Yes	Yes	Yes
Class competition	Yes	Yes	Yes	Yes	Yes	Yes
Exclusivity indicators	Yes	Yes	Yes	Yes	Yes	Yes
Approval-year trend	Yes	Yes	Yes	Yes	Yes	Yes
ATC Level-1 FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
N	585	531	585	531	585	531

Notes: This table reports Tobit estimates of equation (1) without (Baseline) and with (+Bertz) the log Bertz index of molecular complexity (Bertz, 1981). The unit of observation is the active ingredient; the sample is the new molecular entities approved in 1995–2019, of which 531 have a Bertz index. The dependent variable is years from brand approval to first generic entry, right-censored as in Table 2. Specification (A) uses the total patent count, (B) separate primary and secondary counts, and (D) counts by type, of which only the method-of-use coefficient is shown. All specifications include the controls of Table 2. Standard errors from the observed information matrix in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A1: Robustness: Tobit with Bertz Molecular Complexity Control

beyond what the firm fixed effects already absorb. Table A2 presents the results.

Patent-count timing. The baseline counts include every patent listed over the drug's observed life, so they are not predetermined with respect to the entry outcome: a drug that remains generic-free longer has more years in which to accumulate listings. To address this, I re-estimate equation (1) in a landmark design. For each landmark age $K \in \{3, 5, 7\}$, the five patent-type counts are recomputed using only patents listed within K years of brand approval—

	Spec (A) Total		Spec (B) Pri/Sec		Spec (D) 5-way	
	Baseline	+FirmExp	Baseline	+FirmExp	Baseline	+FirmExp
Total patent count	0.164**	0.170**	—	—	—	—
	(0.069)	(0.069)				
# Primary (compound) patents	—	—	0.775*	0.781*	—	—
			(0.423)	(0.422)		
# Secondary patents	—	—	0.135*	0.142*	—	—
			(0.074)	(0.074)		
# Method-of-use patents	—	—	—	—	0.569***	0.569***
					(0.196)	(0.195)
Firm ATC experience	—	0.224	—	0.216	—	0.197
		(0.162)		(0.161)		(0.158)
Firm total experience	—	0.065	—	0.069	—	0.067
		(0.066)		(0.065)		(0.065)
Patent-term span	Yes	Yes	Yes	Yes	Yes	Yes
Paragraph IV certification	Yes	Yes	Yes	Yes	Yes	Yes
Class competition	Yes	Yes	Yes	Yes	Yes	Yes
Exclusivity indicators	Yes	Yes	Yes	Yes	Yes	Yes
Approval-year trend	Yes	Yes	Yes	Yes	Yes	Yes
ATC Level-1 FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
N	585					

Notes: This table reports Tobit estimates of equation (1) without and with two firm-experience controls: the sponsoring parent firm's number of prior new-molecular-entity approvals in the same ATC Level-2 class (firm ATC experience) and across all classes (firm total experience). The unit of observation is the active ingredient; the sample is the new molecular entities approved in 1995–2019. The dependent variable is years from brand approval to first generic entry, right-censored as in Table 2. All specifications include the controls of Table 2. Standard errors from the observed information matrix in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A2: Robustness: Tobit with Firm Experience Controls

and are therefore predetermined at the landmark—and the sample is restricted to molecules still generic-free at age K, the risk set to which a portfolio measured at that age can legitimately be compared. Table A3 reports the estimates alongside the baseline. Two patterns emerge. First, the primary-patent coefficient strengthens to roughly 1.4–1.9 years and is significant at every landmark, as expected for the one patent type that is fully predetermined even in the baseline, since compound patents are filed before approval. Second, the method-of-use point estimate is stable across landmarks (0.445–0.836, around the baseline 0.569) but

loses precision at the shorter landmarks: method-of-use patents are filed late in the lifecycle (Figure 2), so an early predetermined window retains little of their variation and the standard error roughly doubles; at $K = 7$, where the window captures more of the lifecycle, the estimate is significant at the five-percent level. Re-estimating the baseline lifetime counts on the same landmark samples reproduces the baseline coefficients almost exactly (0.572^{***}, 0.568^{***}, and 0.670^{***} at $K = 3$, 5, and 7), confirming that the loss of precision reflects the measurement window rather than the sample restriction. The landmark estimates thus indicate that the method-of-use association is not an artifact of outcome-dependent accumulation, while also underscoring that a fully predetermined cross-sectional count cannot capture the late-life listings that constitute the phenomenon itself.

Market size. To control for the commercial size of the drug's market directly, I construct a pre-entry revenue proxy in the spirit of Newham et al. (2025): total national Medicaid reimbursements for the drug's brand products, from the CMS State Drug Utilization Data (1991–2019), matched through National Drug Codes and averaged over drug ages four and five—a window that is predetermined with respect to entry. The proxy is well defined for 186 of the 585 baseline molecules; the remainder either lack a window closing by 2019, cannot be matched (chiefly older cohorts whose discontinued NDCs are absent from the current NDC directory), or record no Medicaid brand reimbursement. Table A4 reports the results. Revenue behaves exactly as the selection story predicts: conditional on the controls, higher pre-entry revenue significantly accelerates entry, with a Tobit coefficient of -0.732 on log revenue and a Cox log-hazard of 0.205. The method-of-use coefficient is nonetheless robust to it: within the revenue-defined subsample, adding the control attenuates the estimate by roughly a tenth (from 0.869 to 0.791), and it remains significant at the one-percent level. Because the subsample skews toward recent and commercially successful molecules—visible in the larger subsample coefficients—the informative comparison is within the subsample rather than against the full-sample baseline.

	Baseline	Landmark age (window counts)		
		$K = 3$	$K = 5$	$K = 7$
# Primary (compound) patents	1.015**	1.915**	1.837***	1.421**
	(0.427)	(0.754)	(0.668)	(0.635)
# Formulation patents	−0.083	0.126	0.251	0.418*
	(0.101)	(0.283)	(0.243)	(0.222)
# Method-of-use patents	0.569***	0.445	0.458	0.836**
	(0.196)	(0.426)	(0.401)	(0.374)
# Ingredient-level patents	0.260*	0.360	0.381*	0.259
	(0.138)	(0.248)	(0.216)	(0.195)
# Device patents	1.610***	3.324*	3.391*	3.056**
	(0.551)	(1.986)	(1.992)	(1.491)
Patent-term span	Yes	Yes	Yes	Yes
Paragraph IV certification	Yes	Yes	Yes	Yes
Class competition	Yes	Yes	Yes	Yes
Exclusivity indicators	Yes	Yes	Yes	Yes
Approval-year trend	Yes	Yes	Yes	Yes
ATC Level-1 FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	585	482	427	363

Notes: This table reports Tobit estimates of equation (1) with patent counts measured at fixed drug ages. The unit of observation is the active ingredient. The dependent variable is years from brand approval to first generic entry, right-censored as in Table 2. The Baseline column reproduces column 2 of Table 2, with counts of all patents listed over the drug's observed life. In the landmark columns, counts include only patents listed within K years of brand approval, and the sample is restricted to baseline molecules without generic entry by age K; patents without an observed listing date are excluded. On the landmark samples, the lifetime counts give method-of-use coefficients of 0.572***, 0.568***, and 0.670***. All specifications include the controls and firm fixed effects of Table 2. Standard errors from the observed information matrix in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A3: Robustness: Tobit with Predetermined (Landmark) Patent Counts

	Full baseline	Revenue-defined subsample		Cox
	Tobit	Tobit	+ln(Revenue)	+ln(Revenue)
# Primary (compound) patents	1.015**	1.440**	1.463**	−0.501*
	(0.427)	(0.594)	(0.582)	(0.278)
# Formulation patents	−0.083	−0.051	0.052	−0.026
	(0.101)	(0.146)	(0.145)	(0.030)
# Method-of-use patents	0.569***	0.869***	0.791***	−0.105**
	(0.196)	(0.231)	(0.215)	(0.050)
# Ingredient-level patents	0.260*	0.155	0.154	−0.036
	(0.138)	(0.204)	(0.195)	(0.061)
# Device patents	1.610***	3.949*	3.040	−0.210
	(0.551)	(2.150)	(2.013)	(0.178)
ln(Medicaid revenue)	—	—	−0.732***	0.205**
			(0.269)	(0.081)
Patent-term span	Yes	Yes	Yes	Yes
Paragraph IV certification	Yes	Yes	Yes	Yes
Class competition	Yes	Yes	Yes	Yes
Exclusivity indicators	Yes	Yes	Yes	Yes
Approval-year trend	Yes	Yes	Yes	Yes
ATC Level-1 FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	No
N	585	186	186	186

Notes: This table reports estimates of the entry models of Table 2 with a pre-entry market-size control. The unit of observation is the active ingredient. In the Tobit columns, the dependent variable is years from brand approval to first generic entry, right-censored as in Table 2; the final column reports Cox log-hazard coefficients. Following Newham et al. (2025), the control is the log of total national Medicaid reimbursement for the drug's brand products (CMS State Drug Utilization Data, pre-rebate), matched by National Drug Code and averaged over drug ages four and five; it is defined for the 186 baseline molecules with positive matched brand reimbursement in a window closing by 2019. Column 1 reproduces the baseline of Table 2; the remaining columns use the molecules with the control, without and with it. All specifications include the controls of Table 2. The Tobit includes firm fixed effects with standard errors from the observed information matrix; the Cox excludes firm fixed effects and clusters standard errors by ATC Level-2 class. Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A4: Robustness: Tobit with Pre-Entry Market-Size (Medicaid Revenue) Control

C Robustness: Multi-Indication Molecules

The reduced-form analysis treats the active ingredient as the unit even when a molecule is approved for indications spanning more than one therapeutic class, pooling those indications into a single molecule-level observation. This appendix reports the sensitivity of the entry-delay estimates to that choice by comparing the baseline sample—which retains multi-indication molecules—to the subsample restricted to single-indication molecules, holding the specification of equation (1) fixed.

Retaining the multi-indication molecules mildly attenuates the secondary-patent and method-of-use coefficients relative to the single-indication subsample, which contains 537 of the 585 baseline molecules: the pooled secondary count moves from 0.152 (0.080) in the single-indication subsample to 0.135 (0.074) in the baseline, and the method-of-use coefficient from 0.603 (0.223) to 0.569 (0.196); the ordering across patent types—method-of-use patents carrying a robust delay—is unchanged. The direction is expected: a molecule serving two indications pools two distinct entry clocks into one right-censored observation, so the pooled coefficient is an attenuated, and hence conservative, measure of the per-indication delay. The baseline consequently does not overstate the entry-delay input to the welfare calculation on this margin.

D Class-Level Entry and Own Entry Exposure

Section 5.2 uses the number of generic competitors in a drug's therapeutic class as a measure of the competitive pressure the drug faces. This appendix verifies the premise: that class-level entry is informative about the focal drug's own exposure.

I estimate a discrete-time hazard of the focal drug's first generic entry on the same class-competition measure, over 2000–2016. The risk set includes all cohorts, not only the 1995–2019 baseline sample, to maximize the number of entry events; the premise being verified concerns the informativeness of class-level entry, not the sample definition. It contains 9,395

drug-years on 839 molecules with 250 entry events, a baseline hazard of 0.0266 per year.

	(1) ATC Level-1 FE	(2) Molecule FE
Generic competitors in class	+0.0137***	+0.0284***
	(0.0039)	(0.0077)
Lead: competitors arriving in next 2 yrs	+0.0175***	+0.0162***
	(0.0048)	(0.0058)
Drug-age FE	Yes	Yes
Calendar-year FE	Yes	Yes
Observations	9,395	9,395

Notes: This table reports linear probability estimates of the focal drug's first generic entry in year t. The unit of observation is the drug-year; the sample runs from the later of approval and 2000 until own entry or 2016. The regressor is the number of other molecules in the drug's ATC Level-3 class with generic entry as of year t; the lead is the corresponding count for entries in the following two years. All columns include drug-age and calendar-year fixed effects; column 1 adds ATC Level-1 and column 2 molecule fixed effects. Standard errors clustered by active ingredient in parentheses. $^{***}p < 0.01$.

Table A5: Class-Level Generic Entry and the Focal Drug's Own Entry Hazard

Column (2) is the specification of interest: with molecule fixed effects, each additional generic competitor in the class raises the drug's own annual entry hazard by 0.028, comparable in size to the baseline rate itself. Class-level entry is therefore a leading indicator of own exposure rather than a substitute for it, which is what generic manufacturers building class-specific regulatory and manufacturing capability would produce.

The comparison also shows why the fixed-effect level matters. With only ATC Level-1 effects, cross-class differences in entry propensity remain in the error—the class-level entry rate alone predicts own entry at +0.106 ($t = 13.6$; not reported in the table)—and the lead is significant, the signature of a common class-level shock rather than a response. The within-molecule specification of column (2) mitigates that problem, and its lead coefficient, although significant, is about half the contemporaneous effect; the premise of Section 5.2 rests on that specification.

E Crowd-Out: Single-Year Lag Estimates

This appendix reports the single-year lag estimates underlying the windowed profiles of Figure 6. Each cell of Table A6 is a separate two-way fixed-effects regression of ln(trial count + 1) on the firm's count of newly listed secondary patents at the indicated single-year lag, winsorized at the 99th percentile, on the same 948 firm-year sample. The single-year estimates are noisier than the windowed ones but trace the same shape: the Phase 2 association bottoms out at the five-year lag and fades afterward, and the Phase 3 association is negative throughout the three-to-seven-year range. The isolated positive Phase 2 coefficient at the seven-year lag is straddled by statistically insignificant estimates at the adjacent six- and eight-year lags, and no three-year window containing it is distinguishable from zero in Figure 6.

	Phase 2	Phase 3
Sec. patents $(t-3)$	−0.00598	−0.00746
	(0.00382)	(0.00459)
Sec. patents $(t-5)$	−0.0105**	−0.0119*
	(0.00430)	(0.00671)
Sec. patents $(t-7)$	0.00873**	−0.0103**
	(0.00431)	(0.00486)
Sec. patents $(t-10)$	0.00336	0.00411
	(0.00681)	(0.00710)
Firm FE	Yes	Yes
Year FE	Yes	Yes
N	948 firm-year observations (79 firms, 2008–2019)	

Notes: This table reports single-year-lag versions of equation (3); each cell is a separate regression. The unit of observation is the firm-year; the sample is 2008–2019, the period of mandatory ClinicalTrials.gov registration. The dependent variable is ln(trial count + 1) for Phase 2 and Phase 3 trials (columns). The regressor is the firm's number of newly listed secondary patents at the indicated lag, winsorized at the 99th percentile. All regressions include firm and year fixed effects. Standard errors clustered by firm in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A6: Crowd-Out: Secondary Patents and Clinical Trial Pipeline, Single-Year Lags

F Demand Instrument

This appendix describes the price instrument used in the demand estimation of Section 6.

Out-of-pocket price is treated as endogenous, since the unobserved demand shock $\tilde{\xi}_{dmt}$ can enter both mean utility and the cost sharing that plans negotiate. The instrument is a formulary-placement shifter built from the public Part D formulary and plan-benefit files, available from 2019. For each plan, the product's tier placement and the plan's cost-sharing rule for that tier give the out-of-pocket cost per 30-day fill that an enrollee without the low-income subsidy would face in the initial coverage phase—the copayment, or the coinsurance rate times the product's national gross cost per fill, bounded by the tier's minimum and maximum—and products off the formulary receive the full gross cost. Averaging across the plans offered in a state, weighted by January plan enrollment, gives the instrument for the product in that state and year. It shifts the price a state's beneficiaries face through plan design and placement decisions taken at the time of the annual bid, rather than through the year's local demand for the product, and it varies across states for the same product and year because the regional plan menus differ. The first stage is strong, with a cluster-robust F-statistic of 27.4 (clustering by state).

G Public Drug-Utilization Data: The Medicare and Medicaid Segments

The demand data stack two public segments into one panel of prescription-drug markets. A market is a segment, an ATC Level-3 class, a state, and a year, and products carry the same identifiers in both segments, so a molecule in a given formulation and tier is the same product whether it is dispensed under Medicare or Medicaid. The two segments differ in four respects that govern how they are combined.

Coverage and geography. The Medicare segment is built from the Part D Prescribers by

Geography and Drug file and covers all drugs dispensed under Part D, with the state defined by the prescriber's location. The Medicaid segment is built from the State Drug Utilization Data, which report the prescriptions and reimbursement of each state Medicaid program by National Drug Code, with the state defined by the paying program. Because Part D pays for the drugs of beneficiaries dually eligible for Medicare and Medicaid, the Medicaid file covers non-dual enrollees, and the two segments do not double-count prescriptions. The mapping of drugs to products covers 95.2% of Part D fills (84.7% of gross drug cost) and 89.5% of Medicaid prescriptions (93.6% of reimbursement).

Potential markets. The potential market of a segment-state-year is twelve times its enrollees, so that shares are 30-day fills per enrollee-month and the outside option is a month without a fill in the class. Medicare enrollees are Part D enrollees by state (35.1 million in 2013 and 53.7 million in 2024 across the states and the District of Columbia); Medicaid enrollees are residents aged 0–64 with Medicaid or other means-tested coverage in the American Community Survey (49.7 million in 2013 and 60.0 million in 2024), the age restriction removing most dual eligibles. The same potential market serves every class of a segment; with class-by-state-by-year fixed effects it affects only the scaling of welfare.

Prices. Medicare beneficiaries face cost sharing that varies across plans, tiers, states, and years, which the Part D file reports by low-income-subsidy status and from which the out-of-pocket price per fill is computed. Medicaid cost sharing is nominal by statute and is not reported, so the Medicaid segment carries no usable price variation. The price coefficient is therefore identified from the Medicare segment alone, and the demand estimates and welfare figures reported in the paper are for the Medicare segment; in the stacked panel, Medicaid rows enter with the price term set to zero, so that they can inform the patent and product terms without affecting α.

Suppression. Both files suppress small cells: Part D withholds counts based on fewer than eleven claims or beneficiaries, and the Medicaid data withhold state-drug cells with fewer than eleven prescriptions. Suppressed cells are dropped, which removes thinly prescribed products

in small states. For the welfare calculation, national markets are formed by summing fills, costs, and enrollees over states within a segment, class, and year.