Innovation through Recombination

Songyuan Teng

Job Market Paper

November 6, 2025

Please click here for the latest version

Abstract

New ideas often recombine existing ones; this insight is emphasized in recent economic growth theories, but evidence on its empirical relevance is scarce. This paper takes combinatorial growth to measurement by studying the pharmaceutical industry, where the distinction between novelty (discovering new building blocks) and recombination (assembling building blocks into products) is transparent. I document the rising importance of recombination, the firm life-cycle from knowledge accumulation to recombination, and the value premia for novelty. Motivated by these facts, I develop a theory of firm dynamics that distinguishes firm knowledge stocks from product portfolios. Innovation operates along two distinct yet intertwined margins: novel innovation expands knowledge, while combinatorial innovation deploys that knowledge to create new products. The calibrated model captures salient empirical patterns, implies sustained growth through rising recombination, and highlights sharp policy trade-offs: subsidizing novelty boosts short-run growth, while subsidizing recombination raises long-run growth with heterogeneous effects across firms.

JEL codes: O31, O47, O33, D22, L65

Keywords: Innovation, Combination, Growth, Firm Dynamics, Knowledge, Pharmaceutical

*I am grateful to my advisors Michael Peters, Fabrizio Zilibotti, and Samuel Kortum for their generous guidance and support. I thank Ana Cecilia Fieler, Joel Flynn, Teresa Fort, Charles Hodgson, Zhen Huo, Giuseppe Moscarini, Pascual Restrepo, Peter Schott, Winnie Van Dijk, and Yale macro lunch participants for their comments. I thank my peers Tuyetanh L. Tran for continued helpful discussions and Jack Liang for many insightful comments. I thank Scott Miller, Duc M. Chu, Renjing Huang, Qiyuan Li, Hoang Anh T. Phan, Xi Li, and Daojing Zhai for assistance in understanding chemistry backgrounds. Songyuan Teng: Yale, songyuan.teng@yale.edu.

1 Introduction

New ideas often arise by recombining existing ones. Modern software development, for example, combines reusable libraries to build new applications; drug discovery likewise proceeds by systematically varying and recombining known chemical structures to develop new therapies. In *The Theory of Economic Development*, Schumpeter emphasized that development is propelled by new combinations, repurposing and recombining ingredients of existing ones (Schumpeter, 1934).[1] This view raises a central question for economic growth: does recombination crowd out fundamentally novel ideas and slow progress, or does an expanding universe of recombinations sustain it? Recent theory formalizes the dynamics of combinatorial innovation and its implications for growth (e.g., Weitzman, 1998; Jones, 2023), but empirical evidence is thin: the literature largely treats "combinatorial growth" as a metaphor rather than a measured phenomenon. The policy implications are immediate: if innovation is primarily combinatorial, priorities should tilt toward knowledge diffusion and tools that scale search over recombinations (e.g., artificial intelligence, AI); if novelty is essential, targeted support for basic research should take precedence.

This paper takes the idea of combinatorial growth to measurement and quantifies its empirical relevance for aggregate growth. I proceed in three steps. First, I study the roles of novelty (discovering new ingredients) versus recombination (assembling ingredients into products) in the context of the pharmaceutical industry, where their distinction is transparent and directly observable. I uncover the rising importance of recombination, the firm life-cycle from knowledge accumulation to recombination, and the value premia for novelty. Second, motivated by these facts, I develop a new growth theory that formalizes combinatorial growth. I build on Klette and Kortum (2004) which, in its original form, conflates a firm's knowledge stock and product portfolio. This distinction takes center stage in my model: firms must both acquire knowledge, i.e., innovation capacity, through "novel innovation" and apply that capacity to create new products through "combinatorial innovation." Innovation and growth thus operate along two distinct but intertwined margins: knowledge acquisition and knowledge application. Third, I discipline and quantify the model using micro-level pharmaceutical data. The calibrated model identifies combinatorial innovation as the engine that sustains growth and reveals a salient policy trade-off: subsidizing novelty delivers a powerful but temporary boost in growth, whereas subsidizing recombination raises growth persistently.

To begin with, I focus on the pharmaceutical industry, an empirical setting where innovations

[1] The combinatorial nature of innovation has long been recognized across disciplines. Mathematician Henri Poincaré remarked that "to create consists precisely in not making useless combinations and in making those which are useful and which are only a small minority." Biologist François Jacob likened evolution to the work of "a tinkerer," constructing new functions by recombining existing parts. Computer scientist and economist Herbert Simon argued that complex systems are "nearly decomposable" and built from a "restricted alphabet of elementary terms," a modularity that enables powerful recombination.

are constructed from standard building blocks and the abstract notion of "recombination" becomes concrete and directly observable. I exploit a standard chemistry measure that decomposes each pharmaceutical innovation, i.e., a drug, into its constituent functional groups, discrete chemical "ingredients" that govern its molecular properties. This measure allows me to classify a drug as novel if it introduces a previously unseen functional group, and as a recombination if it creates a new co-occurrence of previously used functional groups. This approach applies to all small-molecule drugs, which account for 75% of drugs approved by the FDA's Center for Drug Evaluation and Research since 2000. The advantages of this setting are threefold: (1) functional groups are standard chemical concepts; their presence in a drug is discrete; off-the-shelf algorithms identify them from drugs (Ertl, 2017); (2) the pharmaceutical industry is highly innovation-intensive and offers rich, longitudinal data; (3) the drug discovery process is inherently combinatorial, making my measurement both intuitive and consistent with how chemists conceptualize drug development.

I apply this decomposition to the set of patented small-molecule anti-allergic drugs[2], a therapeutic area with relatively simple drug molecular structure and a steady R&D presence since the 1930s. Three patterns emerge. First, recombination is far more prevalent than novelty and has become more so over time. From 1990 – 2010, 79% of new drugs are recombinations of existing functional groups, while 13% introduce new functional groups; over time, novelty declines as recombination rises. Among recombinations, 35% are within-firm (combining only functional groups the firm has previously used) and 44% are across-firm (incorporating at least one functional group pioneered by another firm that is new to the focal firm). Second, a firm's accumulated knowledge stock, measured as the number of distinct functional groups previously used in its patented drugs, predicts both the level and the composition of its innovation: knowledge-rich firms develop more drugs and specialize in within-firm combination, whereas less knowledgeable firms are more likely to adopt functional groups new to the firm – a life cycle from early knowledge accumulation to later internal recombination. Third, conditional on entering clinical trials, drugs with novel functional groups command larger market and scientific premia, respectively measured as the stock market response to patenting and patent forward citations, than recombinations.

Motivated by these facts, I then propose a new Schumpeterian theory that formalizes combinatorial growth to quantify its importance and to study firms' trade-off between novelty and recombination. Building on Klette and Kortum (2004), my main departure is to distinguish a firm's knowledge stock from its product portfolio. I model firms as being characterized by two state variables: the number of products they sell (i.e., product portfolio) and the number of

[2]Molecules in my dataset span various development stages: preclinical research, clinical trials, and, in some cases, FDA-approval. In regulatory/industry usage, these molecules are referred to as "compounds", "drug candidates", "investigational drugs", or "approved drugs", depending on their development progress. For brevity, I refer to them as "drugs" when the context is unambiguous.

ingredients they own (i.e., knowledge stock). Firms innovate in two ways: (a) combinatorial innovation, which recombines existing ingredients to create a new product, and (b) novel innovation, which obtains an ingredient new to the firm and uses it to create a new product. A larger ingredient stock expands firm capacity to obtain new ingredients and create new products. The framework captures the conceptual trade-offs between basic and applied research and splits innovation and growth into these two distinct but connected margins. The model matches salient empirical features of the pharmaceutical industry: (1) firms with larger ingredient stocks tilt toward combinatorial innovation, implying a life cycle of initial ingredient accumulation through novel innovation to subsequent internal recombination; (2) firm value is additively separable into an ingredient component and a product component, implying that conditional on success, novel innovation is more valuable than recombination because the new ingredient increases the option value of future discovery and recombination.

I calibrate the model to microdata from the anti-allergic market and quantify the importance of combinatorial growth. The calibration centers on two elasticity parameters that govern to what extent richer knowledge improves a firm's capacity for novel and combinatorial innovation. I identify these elasticities using two moments: (1) the concentration of knowledge across firms, higher elasticities implying higher concentration, and (2) the value gap between successful novel versus combinatorial innovations, higher elasticities implying greater value of knowledge and thus a larger gap. Despite its parsimony, the calibrated model matches key untargeted micro-level patterns quantitatively, including the joint distribution of knowledge stocks and drug portfolios across firms and the elasticity of innovation outcomes with respect to knowledge stocks. The model implies that within-firm combination accounts for 30% of aggregate growth.

To examine the implications of declining novelty and rising recombination, I introduce "idea coincidence" into the baseline model: when a firm succeeds in novel innovation, the obtained ingredient may already exist elsewhere in the economy, pioneered earlier by another firm. As the aggregate ingredient pool expands over time, the probability of such coincidence rises, making genuinely new-to-market ingredients increasingly rare while expanding the scope for recombination. This full model yields three aggregate implications: declining novelty, rising recombination, and sustained growth powered by abundant combinatorial opportunities. Conversely, shutting down recombination, both across- and within-firm, reduces growth by 80% in the short run and to zero in the long run, as genuinely new-to-market ingredients become progressively harder to find.

I then use this full model to evaluate three policy instruments that respectively target (a) genuinely novel discoveries, (b) across-firm combination, and (c) within-firm combination. Rewards for novelty generate the largest short-run boost to entry and growth but leave long-run outcomes unchanged. Rewards for across- and within-firm combination raise the long-run growth rate but have opposite distributional effects: rewards for across-firm combination benefit less-

knowledgeable firms and reduce concentration, while rewards for within-firm combination benefit more-knowledgeable firms and increase concentration.

This paper makes three contributions. First, I bring the abstract discussion of combinatorial innovation to measurement by identifying the pharmaceutical industry as a natural empirical setting. Using a rich dataset, I document a series of market- and firm-level facts that deepen the understanding of how innovations through recombination and novelty unfold in practice. Second, I develop a theory that goes beyond the traditional view of firms as portfolios of products. I separate innovation investment and growth into two distinct but intertwined margins: generating new ideas and creating new products by applying existing ideas. The theory is thus an advance in its own right, beyond providing a unified explanation for the empirical patterns observed in the pharmaceutical industry. Third, the theory identifies combinatorial innovation as the primary engine of sustained growth and reveals sharp policy trade-offs between subsidizing novelty and subsidizing recombination. These implications speak directly to real-world policy instruments that differentially support basic versus applied research and suggest that the optimal choice depends on a policymaker's prioritized time horizon. While my empirical setting is the pharmaceutical industry, innovation by recombining granular building blocks is widely observed, appearing in settings such as materials science, software engineering, modular product design and, more broadly, citation networks of scientific publications and patents. The model mechanisms are general and therefore apply broadly.

Related literature. This paper builds on the classic view that innovation is fundamentally combinatorial (e.g., Schumpeter, 1934; Romer, 1992; Arthur, 2009). The closest antecedents are theoretical contributions that treat ideas as combinations of existing "ingredients." Weitzman (1998) develops a recombinant-growth model in which the space of ideas expands combinatorially, so that long-run growth is constrained only by the capacity to search and process this space. Jones (2023) models ideas as combinations of "ingredients" and shows that combinatorial draws from thin-tailed distributions can sustain exponential growth. These theories largely treat "combinations" as a metaphor rather than a measured phenomenon. I contribute by mapping "ingredients" to observable chemical structures in pharmaceuticals and using these empirics to discipline both theory and quantification.

Existing empirical measures of recombination include Uzzi et al. (2013), who infer combination of knowledge from citation networks; Akcigit, Kerr and Nicholas (2013) and Youn et al. (2015), who operationalize it via patent technology classes; and Dix and Lensman (2024), who study "combination therapy" at the drug-regimen level. My measure is closest to Dix and Lensman (2024) but goes deeper: whereas they treat the concurrent use of multiple drugs in a regimen as a combination, I decompose each drug into its constituent functional groups, thereby identifying each product as a combination of smaller "ingredients." The research focus also differs: Dix and

Lensman (2024) emphasize a particular inefficiency in combinatorial innovation, whereas I analyze the trade-off between novelty and recombination and, by embedding the mechanism in a macro framework, study the aggregate consequences of combinatorial innovation.

My work is related to the literature on incremental and radical innovation, which distinguishes modest, small-scale improvements to existing technologies from breakthrough, market-creating innovations (e.g., König, Lorenz and Zilibotti, 2016; Kelly et al., 2021; Acemoglu, Akcigit and Celik, 2022; Krieger, Li and Papanikolaou, 2022; König et al., 2022; Ribeiro, 2025). Unlike much of this literature, which often infers "incremental" vs. "radical" ex post from innovation outcomes, my measure observes directly whether a new product employs a novel ingredient. Using this measure, I show that firms with smaller knowledge stocks disproportionately adopt ingredients pioneered elsewhere, whereas knowledge-rich firms primarily recombine what they already know. This pattern echoes König et al. (2022), who document, in a different setting, that as firms accumulate capabilities, they transition from imitation and technology adoption toward internal innovation.

My finding of declining novelty over time speaks to the long-standing literature on "ideas getting harder to find" (e.g., Evenson, 1982; Kortum, 1993, 1997; Jones, 2009). More recently, Bloom et al. (2020) document sharp declines in research productivity across semiconductors, agriculture, and medical technologies; Fort et al. (2025) show that researchers' patenting productivity has risen, while output growth exhibits a secular decline conditional on patenting activities, indicating "growth is getting harder to find." My drug-level evidence aligns with the "harder to find" view, in the sense that genuinely novel ingredients become rarer, while aggregate progress can be sustained since the expanding stock of known ingredients supports more recombination.

I also build on the Schumpeterian endogenous-growth literature pioneered by Grossman and Helpman (1991) and Aghion and Howitt (1992) as embedded in the framework of Klette and Kortum (2004). This framework has proved empirically successful in matching firm-level facts (Lentz and Mortensen, 2008) and underpins applications ranging from heterogeneous innovation types (Akcigit and Kerr, 2018) to heterogeneous markups and misallocation (Peters, 2020). My departure is to separate firm knowledge stocks from product portfolios and to model an explicit choice between novel and combinatorial innovation. This separation implies that innovation outcomes are shaped more by a firm's knowledge stock than product portfolio and predicts a firm lifecycle from early-stage knowledge accumulation to later-stage recombination, a salient pattern borne out in the data.

Lastly, my analysis connects to biomedically grounded studies of idea production. Tranchero (2024) exploits genome-wide association studies (GWAS) as quasi-experimental, data-driven signals and shows that firms with deeper domain knowledge better screen noisy predictions from big data. This complementarity is consistent with my mechanism: while the feasible set of

functional-group combinations is broadly common across firms, domain knowledge, experience with particular functional groups, improves screening of the vast combinatorial space, increasing the likelihood of finding useful combinations. Relatedly, Krieger, Li and Papanikolaou (2022) document underinvestment in radical pharmaceutical innovation due to risk aversion, and Frankel et al. (2023) estimate the value of dynamic spillovers from drug discovery.

Roadmap. The rest of the paper proceeds as follows. Section 2 provides background on the pharmaceutical industry, introduces ingredient-based measurement, and details the data. In Section 3, I use my measurement of combinatorial growth to provide empirical evidence of its importance. Section 4 develops the baseline model. Section 5 applies it to the anti-allergic market and quantifies the contribution of within-firm combinatorial innovation to aggregate innovation dynamics. Section 6 introduces idea coincidence into the baseline framework and uses the full model to quantify the contribution of total combinatorial innovation (across-firm plus within-firm) and evaluate innovation policies. Section 7 concludes. Appendix contains measurement details, model extensions, proofs, robustness checks, and additional results.

2 Combinatorial Growth in the Pharmaceutical Industry

To precisely define ingredients and combinations, I focus on the pharmaceutical industry where the medicinal chemistry literature motivates a chemistry-based method to decompose a pharmaceutical innovation, i.e., a drug, into its constituent chemical "ingredients". This pharmaceutical setting offers three advantages: 1) the chemical "ingredients" to be introduced are standard concepts in medicinal chemistry, their presence/absence in a drug is discrete and identifiable through off-the-shelf algorithms, enabling clean measurement of recombination; 2) the pharmaceutical industry is among the most research-intensive industries and provides a rich, longitudinal dataset; 3) pharmaceutical innovation, especially drug discovery, often features an inherently combinatorial flavor with subfields that study how drug properties change when basic chemical structures are substituted or rearranged.

In this section, I describe the background of the pharmaceutical industry and the process of drug discovery. I explain how this context provides a chemistry-based method to decompose a drug into its constituent ingredients. I introduce my dataset and methodology to classify pharmaceutical innovations into different categories according to their novelty.

2.1 Industry Background

The pharmaceutical industry discovers, develops, and markets medicines to prevent, alleviate, or cure diseases. It is both an economic heavyweight and among the most research-intensive industries with 27 % – 34% of revenues ploughed back into research annually (Chandra et al.,

2024). Pharmaceutical companies are divided into two broad types: brand-name companies, those that are research-based and develop new medicines termed "brand-name drugs", and generic companies, those that typically spend little on research and manufacture "generic drugs", bioequivalent versions of brand-name drugs once their patents expire (Taylor, 2015). Because brand-name firms are the primary engine of pharmaceutical innovation, I focus on brand-name drugs in my analysis.[3]

Successful innovation in the pharmaceutical industry ultimately materializes as a new drug, what regulators call a new molecular entity (NME). These NMEs fall into two broad technological classes: Small-molecule drugs are compact organic compounds (typically $\leq$ 900 Da) such as acetaminophen, the active ingredient in Tylenol (see Figure 1). Biologics are much larger protein-based therapeutics such as vaccines and antibodies. Small-molecule drugs have accounted for 75% of all drugs approved by the FDA's Center for Drug Evaluation and Research since 2000, and 58% of pharmaceutical-sector revenue in 2023 (Arnum, 2024). I focus my study on small-molecule drugs because their relatively simple and well-defined chemical structures make it feasible to identify each drug's underlying "ingredients".

The development of a small-molecule drug begins with drug discovery, the first stage of the pharmaceutical R&D pipeline. In this stage, scientists search the vast chemical space for new "hits" that can modulate a biological target implicated in disease. It typically involves three core tasks: (i) identifying a therapeutic target (e.g., a protein); (ii) generating candidate molecules; and (iii) refining those candidates to maximize potency and minimize toxicity. To generate candidates, researchers often rely on combinatorial chemistry, using high-throughput screening technologies to assess up to millions of molecules. Only a few survive the discovery stage; those survivors are usually patented prior to clinical testing to secure intellectual property and deter rivals. The low survival rate highlights the extraordinary uncertainty inherent in drug development (Taylor, 2015). Following discovery, development proceeds to pre-clinical development, clinical trials, regulatory review, manufacturing, and post-marketing surveillance. My analysis focuses on patented drugs because they reflect idea generation at the discovery stage, where the underlying "ingredients" are created, combined, and refined.

2.2 Functional groups: the chemical "ingredients"

To study how firms innovate – by inventing new ingredients or recombining existing ones – I first represent chemical structure in a consistent, analyzable way. I now turn to functional groups, the key structural units that underpin drug behavior and allow me to classify innovations empirically.

[3]When a brand-name drug loses its patent protection, typically, generic companies enter with bioequivalent versions, which significantly lower drug price. In 2024, 90% of prescriptions are filled by generic drugs while they only account for 13% of prescription drug spending (Association for Accessible Medicines, 2024).

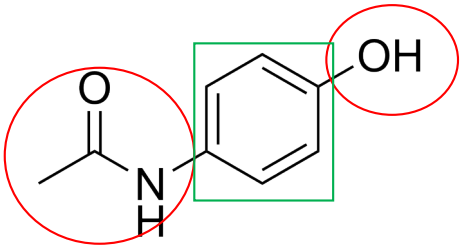

Note: Figure presents the chemical structure of an acetaminophen molecule. The carbon backbone, a benzene ring, is identified and boxed in green. Two functional groups, a secondary amide group (−NH−C(=O)−) and a phenolic hydroxyl group (−OH), are identified and circled in red.

Figure 1: Acetaminophen molecule, active ingredient of Tylenol.

In this application, "innovations" refer to small-molecule drugs, whose chemical structures, i.e., the arrangements of atoms, determine their behavior in the human body. A typical small-molecule drug consists of *functional groups* and a *carbon backbone*. In general, functional groups play the decisive role: these specific clusters of atoms determine chemical reactivity and drive the key interactions with biological targets (enzymes, receptors), thereby shaping drug potency, selectivity, and how human body absorbs and distributes the drug (Ertl, 2017). By contrast, the carbon backbone serves primarily as a stable framework that positions these functional groups in three-dimensional space and maintains drug overall molecular shape. Figure 1 illustrates the chemical structure of acetaminophen: the carbon backbone (a benzene ring) is boxed in green, and two functional groups, a secondary amide group (−NH−C(=O)−) and a phenolic hydroxyl group (−OH), are circled in red.

Given their central role in determining molecular function, I view a functional group as an ingredient and thus a small-molecule drug as a combination of its functional groups. This perspective aligns with how medicinal chemists design and interpret molecules[4], and allows me to decompose any small-molecule drug into a set of identifiable "ingredients". Using functional groups as the unit of analysis offers a key advantage: they are standard units in medicinal chemistry; their presence within a drug is discrete and observable, and the medicinal chemistry literature provides algorithms to identify them from drugs (e.g., Ertl, 2017). This makes the functional group–based representation both scientifically grounded and empirically tractable, allowing me to consistently measure innovation across time, firms, and therapeutic classes.

[4]In small-molecule drug discovery, recombination of chemical structures, in particular functional groups, is a central organizing principle: medicinal chemists commonly fix a core scaffold and swap or append functional groups to tune potency, selectivity, and pharmacokinetics. Combinatorial chemistry implements this directly by enumerating libraries of substituent combinations, while SAR (structure-activity relationships) and QSAR (quantitative SAR) model structure–activity patterns to predict which substitutions/combinations are most likely to yield active compounds (e.g., Hansch et al., 1962; Hansch and Fujita, 1964; Free and Wilson, 1964; Gallop et al., 1994; Macarron et al., 2011).

There are alternatives. In finance, recent work has used the "Tanimoto coefficient"[5] to measure chemical similarity between drugs (e.g., Krieger, Li and Papanikolaou, 2022). While effective for ranking overall similarity, Tanimoto coefficient does not reveal which chemical features are novel or reused, nor does it distinguish between innovation through entirely new components versus recombinations of known ones.

2.3 Data

I focus my analysis on small-molecule *anti-allergic* drugs. Narrowing the scope offers two advantages: (1) anti-allergic drugs generally have relatively simple chemical structures, reducing complexity in molecular comparisons; (2) innovation activity in this area is steady over my sample window (1990 – 2010), accounting for roughly 5% of FDA approvals, which ensures a rich sample of active development efforts.[6]

My primary data source is Clarivate Analytics' Cortellis Drug Discovery Intelligence database. This database is widely used in the pharmaceutical industry and the literature of economics and finance for empirical analysis of pharmaceutical innovation, investment decisions, and R&D spillovers (e.g., Krieger, Li and Papanikolaou, 2022; Frankel et al., 2023; Tranchero, 2024). Cortellis compiles information on drug candidates from publicly available sources such as patent filings, company reports, press releases, clinical trial registries, and FDA submissions. A drug typically enters the database either when it is first patented or when it is mentioned in a firm's pipeline disclosures. The database includes many later-stage preclinical drugs, which are often patented, though it may miss early-stage candidates that fail initial screening and are thus never publicly documented. For drugs that appear in the data, Cortellis backfills development timelines when possible to improve accuracy (Krieger, Li and Papanikolaou, 2022).

For each drug, I extract its molecular structure, development phase, therapeutic class, mechanism of action, originating and developing firms, and associated patents. I link each drug to its patent records and recover a drug's invention year as the earliest priority year among its associated patents. These variables allow me to trace the trajectory of a drug from early-stage research through clinical trials and to measure both the novelty of its functional groups and the organizational context of its development. In my sample, the number of drugs recorded increases sharply after 1985, likely reflecting improvements in digital patent records. I thus restrict my empirical analysis to the period 1990–2010.

An innovation of this paper is to decompose each drug into its constituent "ingredients" and

[5] Tanimoto coefficient measures similarity between two molecules based on their chemical features. It takes values in $[0, 1]$ with higher values indicating greater similarity (Nikolova and Jaworska, 2004).

[6] Over 1990 – 2010, allergic disease was a major locus of biomedical innovation. There was surging research output (Dwivedi, 2016), concentrated FDA-approval activity (Kinch and Merkel, 2015), and significant quality upgrades (Fein et al., 2019).

Table 1: Summary statistics of small-molecule anti-allergic drug market: 1990–2010

period	1990 - 1995	1996 - 2000	2001 - 2005	2006 - 2010	1990 - 2010
total # active firms	116	112	125	158	310
average # patented drugs per year	697	659	736	526	656
average # functional groups per drug	4.33	4.82	5.12	5.37	4.86
drug share, top 20% firms	79%	74%	70%	72%	78%

Notes: Values are either *totals* or *yearly averages* over the indicated period (inclusive of endpoints). "Active firms" are firms that patent ≥ 1 small-molecule anti-allergic drugs in that year. "Patented drugs" count the number of patented drugs. "Functional groups per drug" is the number of distinct functional groups per drug. "Drug share, top 20% firms" is the share of all patented drugs accounted for by the top quintile of firms within the period, ranking firms by total patented drugs.

thus observe exactly what it comprises. I do so by identifying functional groups from drugs using the Ertl (2017) algorithm, implemented in Python via the RDKit cheminformatics toolkit (Hall and Godin, 2017). The algorithm systematically extracts chemically meaningful substructures based on the presence of heteroatoms and reactive carbon environments (see details in Appendix A). The resulting dataset covers 13,786 drugs that were developed by 310 firms and spanned 1,044 distinct functional groups.

Table 1 reports summary statistics for my dataset. The number of active patenting firms rose from 116 (1990 – 95) to 158 (2006 – 10), while annual patented-drug counts were broadly stable at roughly 650 per year over 1990 – 2010. Molecular complexity increased over time: the average number of functional groups per drug rose from 4.33 to 5.37.[7] The market is concentrated: the top 20% of firms (ranked by number of patented-drug counts) account for 78% of all patented drugs over the full period.

Measure combinatorial innovation. I classify each drug in my sample into one of three mutually exclusive categories based on the functional groups it contains: novel, combination, or refinement. A novel drug introduces at least one functional group that has not appeared in any prior drug, representing a breakthrough at the ingredient level. A combinatorial drug uses only previously known functional groups, but they have not co-appeared in any prior drug patented by the developing firm.[8] Finally, a refinement drug reuses an existing combination of functional groups by its developing firm, reflecting incremental improvement of known drugs. This classification allows me to quantify the degree of novelty in drug discovery and to distinguish between foundational innovations, creative recombinations, and incremental improvements. I

[7]Drug molecules cannot be too sparsely or too heavily decorated with functional groups. Medicinal chemistry aims for a middle ground consistent with Lipinski's Rule of Five to balance potency, permeability, and safety.

[8]My definition of combinatorial drug is on the firm level. An alternative definition operates on the market level: a drug is combinatorial if it uses only previously known functional groups and they have not co-appeared in any prior drug in the market, regardless of the developer. Using the market-level definition leaves the empirical results qualitatively unchanged and only modestly affects magnitudes. For example, the average share of combinatorial drugs from 1990 – 2010 falls from 79% to 65%.

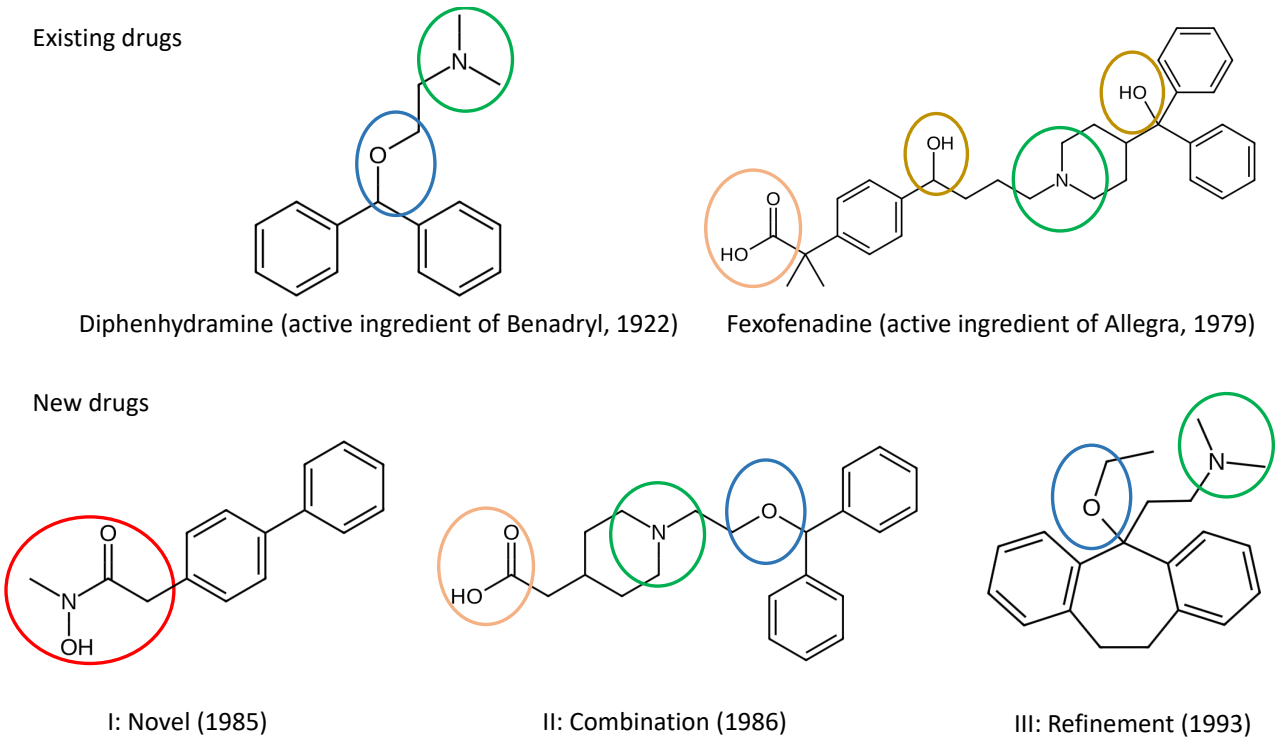

Note: Functional groups are highlighted with consistent color coding across all panels: green = tertiary amine, blue = ether group, orange = carboxyl group, gold = hydroxyl group, red = hydroxamic acid moiety. In the top row, diphenhydramine (1922) and fexofenadine (1979) are treated as the sole prior art and therefore define the initial, "known" functional-group pool. Assume the two drugs were discovered by different firms, A and B respectively. The bottom row shows three later drugs from the dataset and how they map onto my innovation taxonomy, assuming all three were developed by firm A: (i) novel introduces a functional group absent from the prior pool, (ii) combination unites only known groups, but they had not previously co-appeared in any patented drug by the developing firm A, and (iii) refinement reuses a co-appearance of known groups previously used by firm A in diphenhydramine.

Figure 2: Classification of innovation: an example.

further split combinatorial drugs into (i) within-firm combination: combinations formed exclusively from functional groups previously used by the focal firm, but in a new co-occurrence, and (ii) across-firm combination: combinations formed by at least one functional group previously used only by other firms (new to the focal firm) together with the firm's previously used functional groups. The within/across split matters because within-firm combinations leverage the firm's existing knowledge base, while across-firm combinations draw on external knowledge.

Figure 2 visualizes this functional-group-based taxonomy of pharmaceutical innovation. The top row shows two existing drugs, diphenhydramine (Benadryl®, 1922) and fexofenadine (Allegra®, 1979). Each colored circle marks one of the functional groups that anchor my analysis: tertiary amine (green), ether group (blue), carboxyl group (orange), and hydroxyl group (gold). Treating (hypothetically) these two drugs as the entire prior art, the highlighted functional groups define the initial ingredient pool available for subsequent recombination. Assume further that they were

discovered by two different firms, A and B respectively.

The bottom row selects three drugs from my dataset that exemplify each of the three innovation modes. Assume all three were developed by firm A.

I. Novel (1985). The first drug introduces a hydroxamic acid moiety (red) that is absent from prior functional-group pool, making it an example of novel innovation.

II. Combination (1986). Here the drug reunites the tertiary amine and ether group from diphenhydramine with the carboxyl group from fexofenadine (green + blue + orange). By recombining existing functional groups in a new constellation, it exemplifies combinatorial innovation. Moreover, it is classified as an across-firm combination, because the carboxyl group (orange) was used by firm B but new to firm A; if both prior drugs had been developed by firm A, it would instead be a within-firm combination.

III. Refinement (1993). This compound reuses the tertiary amine (green) and ether group (blue) from diphenhydramine but embeds them in a different carbon backbone. Because both functional groups trace to diphenhydramine, it is classified as a refinement.

3 Empirical Importance of Combinatorial Growth

In this section, I apply the measures developed in Section 2 to show that combinatorial innovation plays a central role in pharmaceutical innovation.

3.1 Fact 1: Contribution of combinatorial innovation

Figure 3 documents four time-series facts on aggregate innovation and its decomposition into novel versus combinatorial innovation. Panel A shows that the annual flow of newly patented drugs is fairly stable at roughly 650. Panel B shows a sharp, steady decrease in market-level introductions of new functional groups, indicating declining novelty. Consistent with this observation, panel C shows a pronounced shift from novelty toward recombination: the fraction of combinatorial drugs increases from 68% in 1990 to 82% in 2010, while the novel share decreases from 20% to 8% (the remainder are refinements). Panel D further splits recombinations into within-firm and across-firm cases. Both categories exhibit upward trends over time; on average, within-firm combination accounts for 35% of all new drugs and across-firm combination for 44%.

Overall, the decline in novel discoveries is offset by an increase in recombinations, leaving aggregate innovation roughly stable.

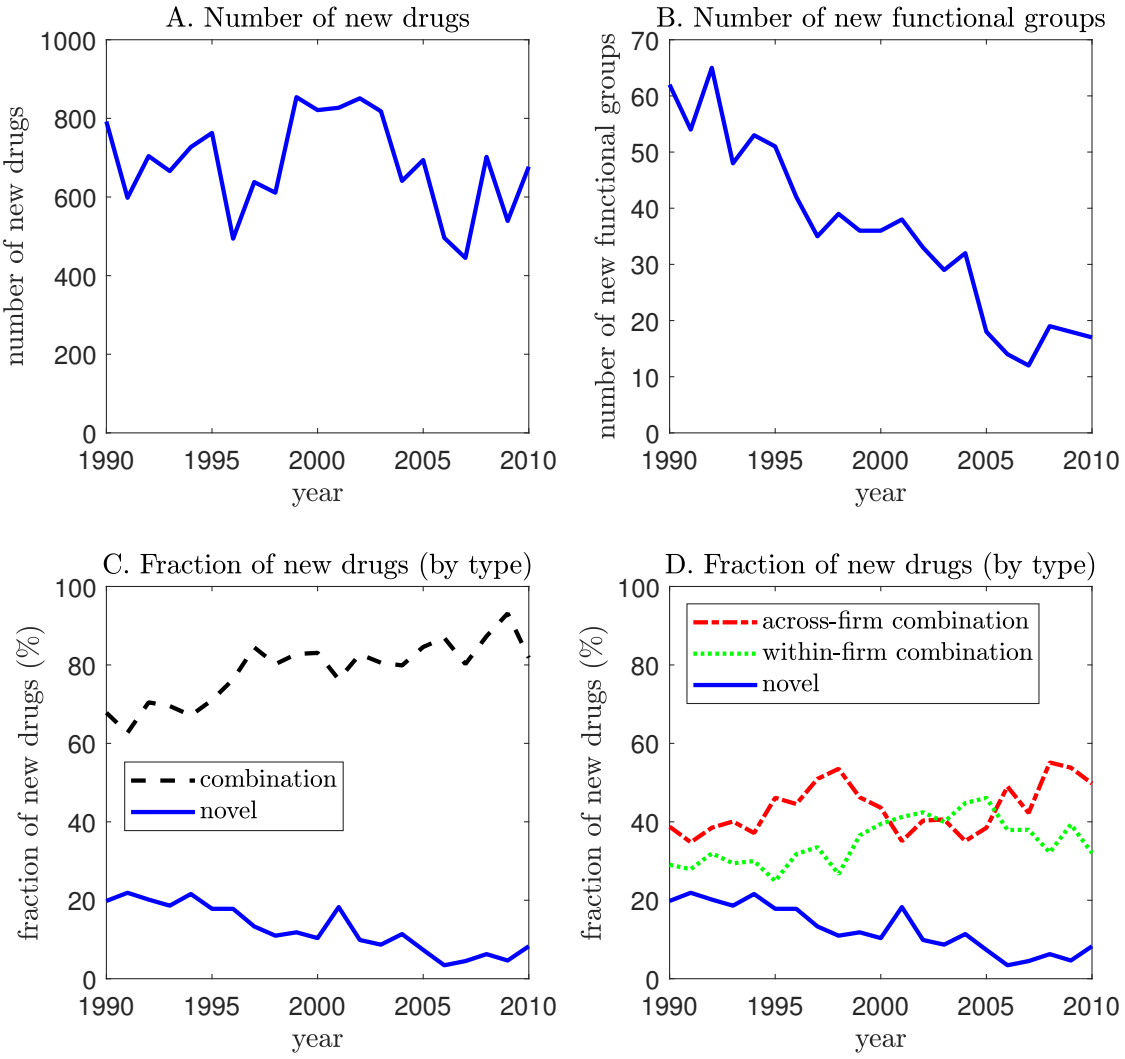

Note: This figure documents time-series patterns of anti-allergic market innovation using the Clarivate Cortellis database. I plot 1990 – 2010 to mitigate mis-classification arising from limited coverage before 1985. Panel A counts the number of newly patented drugs in year t across all firms. Panel B counts the number of functional groups that appear for the first time in patented anti-allergic drugs across all firms. Panel C reports the share of new drugs that are novel or combinatorial as defined in Section 2.3: a drug is novel if it contains at least one functional group not previously identified in any patented anti-allergic drug; a drug is combinatorial if it uses only previously known functional groups, but they have not co-appeared in any prior drug patented by the developing firm. The remaining share (refinement) uses only known functional groups in a co-appearance previously used by the firm. Panel D further decomposes combinatorial drugs into within-firm and across-firm cases. Across-firm combinations include at least one functional group previously used only by other firms (new to the focal firm). Within-firm combinations recombine only functional groups previously used by the focal firm, but in a new co-occurrence.

Figure 3: Contribution of novel versus combinatorial innovation: 1990 – 2010.

3.2 Fact 2: Recombination and knowledge across firms

I now examine the importance of combinatorial growth at the firm level. Specifically, I study how firm innovation intensity and specialization co-vary with its accumulated knowledge. If combinatorial forces are salient, firms with larger knowledge stocks should benefit from a richer set of feasible recombinations and, correspondingly, exhibit higher innovation intensity and tilt from obtaining external knowledge to internal recombinations.

Innovation intensity. I partition 1990 – 2010 into four five-year periods and construct firm-level variables for each period t as follows. I measure firm f's knowledge as the number of distinct functional groups used in its patented drugs by the start of period t, $i_{f,t}^{total}$.[9] Innovation outcomes over period t are: (1) the number of newly patented drugs, $n_{f,t}^{new}$, and (2) the number of functional groups appearing in those drugs that were absent from f's previously patented drugs, $i_{f,t}^{new}$.

I estimate the relationship between firm knowledge stock and innovation activity using a Poisson pseudo maximum likelihood (PPML) specification:

$$E[y_{f,t}|X_{f,t}] = \exp(\alpha + \beta_i \ln(i_{f,t}^{total}) + \gamma_t),$$

where $y_{f,t} \in \{n_{f,t}^{new}, i_{f,t}^{new}\}$ and $X_{f,t}$ denote regressors. The coefficient of interest β_i (subscript i for "ingredients") is the elasticity of innovation outcome with respect to knowledge $i_{f,t}^{total}$. I control for period fixed effects across all specifications to absorb aggregate shifts.[10]

The results are reported in Table 2. Column (1) shows that firms with larger knowledge stocks are more innovative: a 1% increase in knowledge (prior functional-group holdings) is associated with 1.07% more newly patented drugs in the subsequent five-year period. Column (2) shows a parallel pattern for new-to-firm ingredients: a 1% increase in knowledge is associated with 0.33% more newly adopted functional groups.

A potential omitted variable in this estimation is firm size: mechanically, larger firms tend to be more knowledgeable and more innovative. Columns (3) – (4) of Table 2 add a size control to distinguish the role of firm knowledge from size. I proxy size with the cumulative number of previously patented drugs up to period t, $n_{f,t}^{total}$.[11] The knowledge elasticities remain positive and

[9] Functional groups are basic chemical structures and thus not patentable per se. In practice, when a promising functional group is identified, firms often deter competition by filing families of composition-of-matter patents with broad Markush claims to ring-fence classes of compounds that embody that functional group (e.g., Cohen, Nelson and Walsh, 2000; Sternitzke, 2013; Wagner, Sternitzke and Walter, 2022). This can limit spillovers of functional groups across firms and leave the cumulative know-how around functional groups effectively proprietary.

[10] I do not control for firm fixed effects because $i_{f,t}^{total}$ is cumulative and highly persistent, leaving limited within-firm variation over four five-year periods.

[11] A more conventional size measure would be firm sales or employment, but consistent data are unavailable for many non-US and/or privately held firms in the sample. The cumulative number of prior patented drugs provides a comparable, consistently observed proxy.

Table 2: Innovation outcomes by firm knowledge: 1990–2010

	(1) $n_{f,t}^{new}$	(2) $i_{f,t}^{new}$	(3) $n_{f,t}^{new}$	(4) $i_{f,t}^{new}$
knowledge, $\ln(i_{f,t}^{total})$	1.07	0.33	0.88	0.47
	(0.08)	(0.05)	(0.25)	(0.17)
size, $\ln(n_{f,t}^{total})$			0.11	-0.09
			(0.16)	(0.10)
period fixed effect	✓	✓	✓	✓
R^2	0.65	0.18	0.65	0.18
observations	265	265	265	265

Note: The sample spans 1990 – 2010 and is split into four five-year periods. All columns use Poisson pseudo–maximum likelihood specification with period fixed effects: $E[y_{f,t}|X_{f,t}] = \exp(\alpha + \beta_i \ln(i_{f,t}^{total}) + \beta_n \ln(n_{f,t}^{total}) + \gamma_t)$. The dependent variables are the number of new drugs patented by firm f in period t in columns (1) and (3) and the number of newly adopted functional groups in those drugs in columns (2) and (4). Columns (1) - (2) include firm knowledge only ($\ln(i_{f,t}^{total})$), where $i_{f,t}^{total}$ is measured as the number of distinct functional groups used in f's patented drugs by the start of period t. Columns (3) - (4) include both firm knowledge and size ($\ln(i_{f,t}^{total}), \ln(n_{f,t}^{total})$), where $n_{f,t}^{total}$ is measured as its total number of patented drugs by the start of period t. Firms with no new patented drugs during the period are treated as exits and thus dropped. Standard errors are clustered at firm level and reported in parentheses.

statistically significant: a 1% increase in knowledge predicts 0.88% more newly patented drugs (column 3) and 0.47% more newly adopted functional groups (column 4) over the period. By contrast, the size coefficients are small and statistically insignificant in both outcomes (0.11 and -0.09, respectively).

Taken together, Table 2 indicates that knowledge – rather than size – primarily correlates with firm-level innovation intensity.[12]

Innovation specialization. I next examine how the composition of innovation varies with firms' knowledge stocks. For each firm-period observation (f, t), I compute four shares within $n_{f,t}^{new}$, corresponding to new drugs that are (1) novel, $\%\text{novel}_{f,t}$, (2) across-firm combination, $\%\text{comb}_{f,t}^{across}$, (3) within-firm combination, $\%\text{comb}_{f,t}^{within}$, and (4) refinements, $\%\text{ref}_{f,t}$, as defined in Section 2.3. For compactness, index these types by $k \in \{1, 2, 3, 4\}$, and write $y_{f,t}^{k}$ for the share of type k.

I estimate the relationship between firm knowledge stock and innovation specialization using

[12] Appendix Table E.16 estimates the same specification but disaggregates new drugs $n_{f,t}^{new}$ by innovation category (novel, combination, refinement); Appendix Table E.17 likewise disaggregates new functional groups $i_{f,t}^{new}$ into those new to the market versus new only to the firm. The same qualitative patterns hold.

Table 3: Innovation specialization by firm knowledge: 1990–2010

	(1)	(2)	(3)	(4)	(5)	(6)
	$\%\text{novel}_{f,t}$	$\%\text{comb}^{across}_{f,t}$	$\%\text{comb}^{within}_{f,t}$	$\%\text{novel}_{f,t}$	$\%\text{comb}^{across}_{f,t}$	$\%\text{comb}^{within}_{f,t}$
knowledge, $\ln(i^{total}_{f,t})$	-0.01	-0.14	0.14	0.07	−0.16	0.13
	(0.01)	(0.01)	(0.01)	(0.04)	(0.06)	(0.04)
size, $\ln(n^{total}_{f,t})$				−0.05	0.01	0.01
				(0.03)	(0.04)	(0.02)
period fixed effect	✓	✓	✓	✓	✓	✓
observations		265			265	

Note: The sample spans 1990 – 2010 and is split into four five-year periods. All columns use fractional multinomial logit specification with period fixed effects (see equations 3.1 – 3.2). The dependent variables are (i) share of drugs that includes a new-to-market functional group, $\%\text{novel}_{f,t}$; (ii) share of drugs that recombines across firms, using at least one functional group previously used in the market but not by firm f, $\%\text{comb}^{across}_{f,t}$; (iii) share of drugs that recombines only within the firm, using functional groups previously used by f but not previously co-appear in the same drug, $\%\text{comb}^{within}_{f,t}$. Refinement is the base category. Knowledge and size are defined as in Table 2. Firms with no new patented drugs during the period are treated as exits and thus dropped. Standard errors are clustered at firm level and reported in parentheses.

a fractional multinomial logit specification (Buis, 2017):

$$E[y^k_{f,t}|X_{f,t}] = \frac{\exp(\alpha^k + \beta^k_i \ln(i^{total}_{f,t}) + \gamma^k_t)}{1 + \sum_{j=1}^{3} \exp(\alpha^j + \beta^j_i \ln(i^{total}_{f,t}) + \gamma^j_t)}, \text{ for } k \in \{1,2,3\}, \tag{3.1}$$

$$E[y^4_{f,t}|X_{f,t}] = 1 - \sum_{j=1}^{3} E[y^j_{f,t}|X_{f,t}], \tag{3.2}$$

with refinement ($k = 4$) as the base category. The coefficients of interest β^k_i's capture how knowledge correlates with the composition of innovation. Period fixed effects are controlled for in all specifications to absorb aggregate shifts.

The results are reported in Table 3. Columns (1) – (3) show that larger knowledge stocks are associated with a lower share of across-firm combinations and a higher share of within-firm combinations: the coefficients on $\ln(i^{total}_{f,t})$ are -0.14 for $\%\text{comb}^{across}_{f,t}$ (column 2) and +0.14 for $\%\text{comb}^{within}_{f,t}$ (column 3), with small effects on $\%\text{novel}_{f,t}$ (column 1, -0.01). In terms of magnitude: moving from the median knowledge firm (23 functional groups) to one standard deviation above (77 functional groups), an increase of about 1.2 in $\ln(i^{total}_{f,t})$, implies roughly 17% fewer across-firm combinations (-0.14×1.2) and 17% more within-firm combinations (0.14×1.2).

Columns (4) – (6) add a size control, $\ln(n^{total}_{f,t})$. The knowledge pattern remains: the coefficients on $\ln(i^{total}_{f,t})$ are -0.16 for $\%\text{comb}^{across}_{f,t}$ (column 5) and +0.13 for $\%\text{comb}^{within}_{f,t}$ (column 6), implying that for the same $23 \to 77$ comparison, the increase in knowledge is associated with about 19% fewer across-firm combinations (column 5, -0.16×1.2) and 16% more within-firm combinations

(column 6, 0.13×1.2). By contrast, the size coefficients are small and statistically insignificant from zero across all three shares.[13]

Taken together, Table 3 indicates that knowledge stocks – rather than size – predict specialization; as knowledge accumulates, firms rely less on obtaining external knowledge and more on recombining what they already know. This pattern maps specialization along the knowledge distribution and suggests a firm life cycle: early-stage knowledge accumulation followed by later-stage internal recombination.[14]

3.3 Fact 3: Combinations are less valuable than novelty

Innovation is increasingly combinatorial at the market level. Whether this reallocation is benign or costly depends on the relative value of different innovation types. I therefore examine value measures and show that, conditional on entering clinical trials, drugs classified as novel command systematically higher value than combinatorial or refinement innovations.

Ideally, a drug's value is its expected profits net of development costs. Because most drugs in my sample never receive FDA approval, direct measurement is infeasible. Following the literature (e.g., Krieger, Li and Papanikolaou, 2022), I measure drug value using value of its associated patents. Firms typically file composition-of-matter patents on all promising drugs by the end of drug discovery stage and, if the project advances, layer on method-of-use, formulation, and process/manufacturing patents during development.[15] Consequently, the market value and scientific impact of these patents provide informative signals about a drug's appropriable economic returns and underlying knowledge value even when it does not reach FDA approval.

Clarivate Cortellis links each drug to its associated patents; in my sample, over 95% of drugs connect to at least one patent. I merge these links with (i) patent market value estimates and forward citations from Kogan et al. (2017) and (ii) patent 10-year breakthrough index from Kelly

[13]The coefficient on $\ln(i_{f,t}^{total})$ is 0.07 for $\%\text{novel}_{f,t}$ (column 4) and is marginally significant. This may reflect that one pathway to discovering new functional groups is by combining existing ones; accordingly, firms with larger knowledge stocks may also specialize in developing new functional groups. A fuller exploration of this mechanism is left for future work.

[14]A potential concern is a finite-alphabet effect: if the universe of functional groups is limited, firms that have already explored many functional groups will mechanically have fewer opportunities to adopt external ones and thus appear to tilt toward within-firm combination. Two pieces of evidence suggest this is not driving the results. First, over 1990 – 2010 the most knowledgeable firm had used about 32% of all functional groups observed in the market at the time, and the 95$^{\text{th}}$-percentile firm used about 11%, leaving ample scope for across-firm adoption. Second, a placebo experiment in which firms form drugs by randomly sampling functional groups reproduces the signs in Table 3 but with much smaller magnitudes; see Appendix Table E.18. Together, these facts bolster the view that abundant internal combinatorial opportunities – not a finite-alphabet constraint – shift firms from across-firm to within-firm combination.

[15]In my baseline specification, I include all patents linked to a drug, since each reflects aspects of the drug's expected payoff and costs. Results are robust to restricting the sample to patents issued within ten years of the earliest priority date among the drug's linked patents, which more closely approximates primary composition-of-matter patents.

et al. (2021).[16] The unit of observation is the patent-drug pair. I estimate the PPML specification in equation (3.3) with dependent variable $y_{p,d} \in \{\text{patent market value}, \text{patent forward citations}\}$ and the linear specification in equation (3.4) with $y_{p,d}$ equal to patent 10-year breakthrough index:

$$E[y_{p,d}|X_{p,d}] = \exp(\alpha + \beta_c \mathbb{1}_d^c + \sum_{k \in K} \beta^k \mathbb{1}_d^k + \sum_{k \in K} \beta_c^k \mathbb{1}_d^k \times \mathbb{1}_d^c + \gamma_{t(p)} + \mu_{f(p)}), \tag{3.3}$$

$$y_{p,d} = \alpha + \beta_c \mathbb{1}_d^c + \sum_{k \in K} \beta^k \mathbb{1}_d^k + \sum_{k \in K} \beta_c^k \mathbb{1}_d^k \times \mathbb{1}_d^c + \gamma_{t(p)} + \mu_{f(p)} + \epsilon_{p,d}, \tag{3.4}$$

where innovation types $K = \{\text{novel, across-firm combination, within-firm combination}\}$; refinement is the omitted type. $\mathbb{1}_d^k$ indicates innovation type k, and $\mathbb{1}_d^c = 1$ if drug d enters clinical trials. The specification rests on a simple intuition: novel innovations are potentially more valuable than other types of innovations when they embody useful knowledge. I use clinical-trial entry as an indicator of drug usefulness and thus $\mathbb{1}_d^c$ effectively partitions all drugs into two groups by their usefulness. The coefficients have a transparent interpretation: β^k compares type k to refinement among drugs that never entered clinical trials, while β_c^k captures the additional premia for type k versus refinement among drugs that did. All specifications include patent issue-year fixed effects $\gamma_{t(p)}$ and firm fixed effects $\mu_{f(p)}$.

Table 4 reports the estimation results. Column (1) documents the correlation between innovation types and the stock-market valuation of patents. Among drugs that never entered clinical trials, the type coefficients β^k are small and statistically insignificant, indicating no systematic differences across innovation types. Among drugs that did, the interaction coefficients β_c^k are positive and significant for all three types. Relative to refinements, the implied market-value premia are: 183% ($= \exp(1.02 + 0.02) - 1$) for novel innovation, 95% ($= \exp(0.74 - 0.07) - 1$) for across-firm combination, and 73% ($= \exp(0.56 - 0.01) - 1$) for within-firm combination. Thus, conditional on clinical-trial entry, novel innovation carries the largest market-value premia.

Column (2) focuses on forward citations. Similarly, among drugs that never entered clinical trials, there are no systematic differences in scientific value across innovation types. Among drugs that did, all three types earn positive scientific premia relative to refinements: novel 290% ($= \exp(1.31 + 0.05) - 1$), across-firm combination 197% ($= \exp(0.85 + 0.24) - 1$), and within-firm combination 123% ($= \exp(0.68 + 0.12) - 1$). Overall, novel innovation carries the highest scientific-value premia. Column (3) turns to 10-year breakthrough index. Among drugs that entered clinical trials, premia are again positive: relative to refinements, novel drugs score 10% ($= 0.11 - 0.01$) higher, across-firm combination 10% ($= 0.08 + 0.02$) higher, and within-firm combination 5% ($= 0.03 + 0.02$) higher, though the latter is not statistically significant.

[16] Kelly et al. (2021) construct a patent-level "importance" (breakthrough) score defined as the ratio of two similarities: (i) the similarity between the focal patent and future patents, and (ii) the similarity between that focal patent and past patents. A higher value indicates that the focal patent appeared novel at the time yet became a strong template for subsequent work, signifying greater importance.

Table 4: Innovation value by types: 1990 – 2010

	(1) market value$_{p,d}$	(2) citation$_{p,d}$	(3) ln(10-year breakthrough$_{p,d}$)
specification	PPML (equation 3.3)	PPML (equation 3.3)	linear regression (equation 3.4)
β^{novel}	0.02	0.05	-0.01
	(0.11)	(0.11)	(0.02)
$\beta^{comb,across}$	-0.07	0.24	0.02
	(0.09)	(0.13)	(0.02)
$\beta^{comb,within}$	-0.01	0.12	0.02
	(0.02)	(0.10)	(0.02)
β_c^{novel}	1.02	1.31	0.11
	(0.36)	(0.36)	(0.05)
$\beta_c^{comb,across}$	0.74	0.85	0.08
	(0.25)	(0.20)	(0.03)
$\beta_c^{comb,within}$	0.56	0.68	0.03
	(0.13)	(0.32)	(0.03)
R^2	0.77	0.38	0.48
observations	5,122	5,075	7,244

Note: Unit of observation: patent-drug. The sample covers anti-allergic drugs 1990 – 2010; patents are those linked to each drug in Clarivate Cortellis. Columns (1) – (2) use specification (3.3). Column (3) uses specification (3.4). $K = \{$novel, across-firm combination, within-firm combination$\}$ and the omitted category is refinement. $\mathbb{1}_d^k$ equals 1 if drug d falls in category k. $\mathbb{1}_d^c$ equals 1 if drug d goes into clinical trial, a usefulness indicator. The dependent variables are patent market value, forward citations, and log-transformation of breakthrough index from Kogan et al. (2017) and Kelly et al. (2021). β^k measures the type-k value difference relative to refinement among drugs that never entered clinical trials, and β_c^k is the additional premia for type k among drugs that did. All specifications include patent issue-year fixed effects $\gamma_{t(p)}$ and firm fixed effects $\mu_{f(p)}$. Standard errors are clustered at the firm level and reported in parentheses.

Conditional on clinical-trial entry, novel innovation earns the largest market-value and scientific premia; recombination also outperforms refinement, with across-firm combination surpassing within-firm combination. These findings align with the results of Krieger, Li and Papanikolaou (2022), who also document higher value for novelty (measured with Tanimoto distance); my contribution is to measure novelty at the functional-group level and to include the breakthrough index as an additional outcome.

Taking stock. Section 3 establishes three facts. First, aggregate innovation is stable but the composition shifts toward recombination. Second, firms with more accumulated knowledge innovate more and specialize in within-firm combination. Third, conditional on usefulness, recombinations are less valuable than novel innovation.

4 The Baseline Model

Motivated by the empirical evidence documented in Section 3, I build a theory to understand firms' tradeoff between novelty and recombination and quantify the importance of combinatorial growth. The framework builds on Klette and Kortum (2004) but makes a conceptual departure: I separate a firm's knowledge stock, here represented by a stock of ingredients, from its product portfolio. In the model, novel innovation expands firm knowledge, while combinatorial innovation deploys that knowledge to create new products without expanding the knowledge frontier. This distinction mirrors the conceptual trade-off between basic and applied research: the former builds up the capacity of innovation, while the latter uses this capacity to generate profit-earning products.

4.1 Firm innovation and creative destruction

Time is continuous. There is a unit mass of differentiated products; each product belongs to a firm and yields a constant flow profit $\bar{\pi}$, net of production cost. Firms are risk neutral, profit-maximizing, and discount future profits at an exogenous interest rate r.

Incumbent firms. An incumbent firm is characterized by two state variables: its knowledge stock, represented by the number of ingredients it owns i, and its product portfolio, represented by the number of products it currently produces (n).[17] I model innovation using the standard Poisson framework: a firm spends resources to generate innovations at a Poisson arrival rate, and higher spending raises this arrival rate. Specifically, firms expand through two types of innovations.

Combinatorial innovation. A combinatorial innovation draws on the firm's existing knowledge stock to create a new product. The key feature of the R&D technology is that innovation costs depend on the size of knowledge stock. Choosing a Poisson arrival rate λ_c (subscript c for "combinatorial") for combinatorial innovations requires flow R&D expenditure

$$R_c(i, \lambda_c) = \frac{A_c}{\phi_c} i^{-\zeta_c} \lambda_c^{\phi_c}, \quad \phi_c > 1, \quad \zeta_c > 0. \tag{4.1}$$

With the knowledge elasticity parameter $\zeta_c > 0$, a larger stock i *lowers* the cost of generating combinatorial arrivals, formalizing the idea emphasized in the theoretical literature that a richer "alphabet" of ingredients expands feasible combinatorial opportunities and makes successful combinations easier to find (e.g., Romer, 1992; Weitzman, 1998; Jones, 2023).[18][19] This specification

[17] Mapping to the empirical setting, an ingredient corresponds to a functional group, and a product corresponds to an anti-allergic drug. Drugs are "differentiated" in that each targets a specific allergy variant and patient population.

[18] Under this functional form, the marginal gains from combination diminish as knowledge grows, consistent with the notion that too many possible combinations can exhaust testing capacity (Weitzman, 1998).

[19] For tractability, I do not model an explicit choice over ingredient subsets, which would generate a 2^i-sized

also aligns with empirical evidence that firms with broader cumulative functional-group repertoires tend to subsequently patent more drugs, especially within-firm combinations (see Section 3.2). The parameter ϕ_c governs the convexity of costs in arrival rate: $\phi_c > 1$ implies increasing marginal cost of raising λ_c. The constant $A_c > 0$ is a scale parameter that captures technology difficulty specific to combinatorial R&D. A successful combinatorial innovation increases the firm's product count, $(i, n) \mapsto (i, n + 1)$.

Novel innovation. A novel innovation discovers a new ingredient (new to the focal firm) and combines it with the firm's existing stock of ingredients to create a product. I model this process symmetrically to combinatorial innovation. In particular, choosing a Poisson arrival rate λ_b (subscript b for "basic research") for novel innovations requires a flow R&D expenditure

$$R_b(i, \lambda_b) = \frac{A_b}{\phi_b} i^{-\zeta_b} \lambda_b^{\phi_b}, \quad \phi_b > 1, \quad \zeta_b > 0. \tag{4.2}$$

With knowledge elasticity $\zeta_b > 0$, the search for new ingredients becomes less costly as the existing knowledge stock grows, reflecting a "synergy"/absorptive-capacity effect in idea absorption (e.g., Arrow, 1962; Cohen and Levinthal, 1989; Griffith, Redding and Van Reenen, 2003). This formulation is consistent with empirical evidence that firms with larger knowledge stocks adopt more new functional groups in subsequent drugs (see Section 3.2). The parameter ϕ_b governs the convexity of costs in arrival rate: $\phi_b > 1$ implies increasing marginal cost of raising λ_b. Constant $A_b > 0$ is a scale parameter capturing novel-specific technology difficulty. A successful novel innovation increases both the ingredient stock and the product portfolio, $(i, n) \mapsto (i + 1, n + 1)$.[20]

Note that the effect of knowledge on innovation costs (equations 4.1-4.2) hinges on the knowledge elasticity parameters ζ_c and ζ_b. When ζ_c is high relative to ζ_b[21], increases in i reduce the unit cost of combinatorial search faster than that of novel search (R_c falls more steeply with i than R_b), so firms with larger knowledge stocks have a comparative advantage in combinatorial innovation. Note also that, at this stage, a novel innovation introduces an ingredient that is new to the focal firm, but not necessarily new to the market. Throughout this section and Section 5, I maintain this simple specification and show that the model matches a series of micro facts. The full model in Section 6 distinguishes new-to-market from new-to-firm ingredients and is able to

opportunity set as in Weitzman (1998) and Jones (2023), and instead model the arrival rate of combinatorial products. In Teng (2025), where I study learning in combinatorial innovation, I adopt the standard 2^i specification.

I do not let R&D expenditure $R_c(i, \lambda_c)$ depend on the number of products n because, empirically (see Section 3.2), a firm's innovation activity is not correlated with its prior drug count once knowledge is controlled for. Moreover, tying costs to n adds complexity and tightens the parameter space for well-behaved equilibria.

[20] For empirical alignment, I let a novel innovation generate a new product alongside the new ingredient: in the data I observe drugs, not standalone functional-group arrivals. This choice eases model-data mapping. Model tractability and qualitative implications are preserved under an alternative specification in which novel innovation adds only an ingredient.

[21] The parameter restriction also involves ϕ_b and ϕ_c, as later shown in Proposition 1.

reproduce the macro trends of declining knowledge expansion and rising recombination.

As in the theoretical firm-dynamics literature, R&D efforts are undirected: when a firm succeeds in either novel or combinatorial innovation, it acquires one product by randomly displacing an existing product in the market, with each incumbent product equally likely to be displaced.[22] The displaced product exits its incumbent producer's portfolio. The Poisson hazard rate of such creative-destruction events, per product, is denoted by $\mu > 0$. I adopt this creative-destruction (rather than expanding-variety) formulation, because it aligns with my empirical setting of brand-name drugs, where new chemical entities frequently supplant earlier ones; for example, second-generation antihistamines replacing first-generation predecessors. Firms take μ as given; its value is determined in equilibrium.

Each ingredient depreciates at a Poisson hazard rate $\xi > 0$, capturing the intuition that the firm fully explores the potential use of an ingredient in finite time and the empirical reality that knowledge can become obsolete or be forgotten. Introducing ingredient depreciation ensures the existence of a stationary firm distribution; without it, the number of ingredients per firm would drift upward without bound.

Entry and exit. I model entry as typical in the firm-dynamics literature (e.g., Akcigit and Kerr, 2018). A unit continuum of potential entrants can enter through novel innovation. Choosing a Poisson arrival rate η for novel innovations requires a flow R&D expenditure

$$R_e(\eta) = \frac{A_e}{\phi_b}\eta^{\phi_b}, \tag{4.3}$$

where, for simplicity, the convexity parameter ϕ_b coincides with that for incumbents, but the scale parameter $A_e > 0$ captures entrant-specific difficulty in novel discovery. A successful entrant discovers a new ingredient, creates a new product built from that ingredient, and becomes an incumbent starting at $(i, n) = (1, 1)$ (one ingredient, one product). I assume entrants have little prior knowledge and therefore cannot enter through combinatorial innovation.

Firms that lose all their ingredients and products, $(i, n) = (0, 0)$, exit the market.[23]

4.2 Stationary equilibrium

I now characterize the Markov perfect equilibria of the economy in which firm strategies depend only on their individual states. Throughout this section and Section 5, I focus on a stationary equilibrium in which aggregate variables and cross-sectional firm distribution are time-invariant.

[22] Because each firm produces a finite number of products, the probability that a successful innovator displaces one of its own products is zero.

[23] This exit rule is natural: such firms have no profit-generating products, and their innovation cost is effectively infinite for both novel and combinatorial innovation (equations 4.1 – 4.2), implying a zero continuation value. This assumption also facilitates analytical tractability, as shown in Proposition 1.

The full model in Section 6 features a transitional setting.

Stationary equilibrium definition. A stationary equilibrium of the economy consists of incumbent innovation policy functions (λ_b, λ_c), entry rate η, creative destruction rate μ, ingredient depreciation rate ξ, interest rate r, and a joint distribution of firms $\Phi(i,n)$ over states (i,n), such that (i) incumbents choose (λ_b, λ_c) optimally (equation 4.4); (ii) entrants choose η optimally (equation 4.7); (iii) the total measure of products equals one $\sum_i \sum_n n\Phi(i,n) = 1$; (iv) the distribution $\Phi(i,n)$ evolves according to its law of motion and is time-invariant (equation 4.8).

Incumbent firm problem. Firm value functions determine optimal R&D choices. Consider a firm with i ingredients and n products. The firm takes (r, ξ, μ) as given and chooses R&D intensities λ_b (novel innovation) and λ_c (combinatorial innovation) to maximize the present value of profits.

Formally, the Bellman equation is:

$$
\begin{aligned}
r\,V(i,n) = \max_{\lambda_b,\lambda_c} \Bigg\{ & \underbrace{\bar{\pi} n}_{\text{flow profits}} - \underbrace{\frac{A_b}{\phi_b} i^{-\zeta_b} \lambda_b^{\phi_b} - \frac{A_c}{\phi_c} i^{-\zeta_c} \lambda_c^{\phi_c}}_{\text{innovation costs}} \\
& + \underbrace{\lambda_b \left[V(i+1,n+1) - V(i,n)\right]}_{\text{novel innovation}} + \underbrace{\lambda_c \left[V(i,n+1) - V(i,n)\right]}_{\text{combinatorial innovation}} \\
& - \underbrace{\xi i \left[V(i,n) - V(i-1,n)\right]}_{\text{ingredient depreciation}} - \underbrace{\mu n \left[V(i,n) - V(i,n-1)\right]}_{\text{creative destruction}} \Bigg\}.
\end{aligned} \tag{4.4}
$$

The first line on the right-hand side captures the firm's instantaneous flow payoff net of innovation costs. Each product generates the same profit flow $\bar{\pi}$, while the cost of innovation depends on the firm's knowledge stock i. The second line reflects the expected gains from successful innovation: novel innovation, at rate λ_b, expands both the ingredient set and the product portfolio, and combinatorial innovation, at rate λ_c, adds a product without expanding knowledge frontier. The third line captures expected losses in firm value. Ingredients depreciate at a constant Poisson rate ξ, shrinking the firm's knowledge stock and thus its future innovative capacity. Each product is displaced through creative destruction at Poisson rate μ, removing one product from the firm's portfolio. Creative destruction leaves the knowledge stock intact: the ingredient count i drops only via depreciation.

A key feature of this setup is that $V(i,n)$ is linear in n and admits a largely closed-form characterization. Because the R&D cost functions in equations (4.1) – (4.2) depend only on firm knowledge i but not on its product count n, acquiring an additional product does not change the marginal cost of future innovation. Once obtained, a product yields a constant profit flow $\bar{\pi}$ until it is displaced by creative destruction at Poisson rate μ. Consequently, the marginal value of a product is independent of the rest of the firm's state and equals the present value of its profit

stream net of expected loss from creative destruction. The following proposition formalizes this observation.

Proposition 1. *Consider the setup above. The value function is linear in the number of products:*

$$V(i,n) \;=\; a(i) \;+\; \frac{\bar{\pi}}{r+\mu}\, n \;=\; V(i,0) \;+\; V(0,n), \tag{4.5}$$

where $a(i)$ satisfies

$$\begin{aligned} r a(i) \;=\; & \frac{\phi_b - 1}{\phi_b} \left(A_b i^{-\zeta_b}\right)^{\frac{-1}{\phi_b - 1}} \left[a(i+1) - a(i) + \frac{\bar{\pi}}{r+\mu}\right]^{\frac{\phi_b}{\phi_b - 1}} \\ & + \frac{\phi_c - 1}{\phi_c} \left(A_c i^{-\zeta_c}\right)^{\frac{-1}{\phi_c - 1}} \left(\frac{\bar{\pi}}{r+\mu}\right)^{\frac{\phi_c}{\phi_c - 1}} \\ & - \xi i \Big[a(i) - a(i-1)\Big]. \end{aligned} \tag{4.6}$$

The corresponding optimal innovation intensities λ_b and λ_c are:

$$\lambda_b(i) = \left(\frac{a(i+1) - a(i) + \frac{\bar{\pi}}{r+\mu}}{A_b i^{-\zeta_b}}\right)^{\frac{1}{\phi_b - 1}} \quad , \quad \lambda_c(i) = \left(\frac{\frac{\bar{\pi}}{r+\mu}}{A_c i^{-\zeta_c}}\right)^{\frac{1}{\phi_c - 1}},$$

which depend solely on i and are independent of n.

Proof. See Appendix C. □

Proposition 1 shows that the value function is a separable sum of two terms: one term, $a(i)$, captures the option value of the current knowledge stock for generating future innovations; the other, proportional to n, reflects the annuity value of the current product portfolio. Note that, consistent with empirical evidence in Section 3.2, innovation policy functions (λ_b, λ_c) depend only on firm knowledge i, but not on product portfolio n once i is given.[24]

The characterization of $V(i,n)$ aligns with two salient empirical observations. First, conditional on usefulness, novelty is more valuable than combination (Section 3.3). A successful combination yields

$$v_c = V(i,n+1) - V(i,n) = \frac{\bar{\pi}}{r+\mu},$$

[24] Although $a(i)$ does not admit a closed form in general, the value function becomes explicit under the parameter restriction $\phi_b - \zeta_b = \phi_c - \zeta_c = 1$. In this case $a(i)$ is linear in i and the value function reduces to $V(i,n) = a\, i + \frac{\bar{\pi}}{r+\mu}\, n$, where a solves:

$$(r+\xi)a \;=\; \frac{\phi_b - 1}{\phi_b} A_b^{\frac{-1}{\phi_b - 1}} \left(a + \frac{\bar{\pi}}{r+\mu}\right)^{\frac{\phi_b}{\phi_b - 1}} + \frac{\phi_c - 1}{\phi_c} A_c^{\frac{-1}{\phi_c - 1}} \left(\frac{\bar{\pi}}{r+\mu}\right)^{\frac{\phi_c}{\phi_c - 1}}$$

In this knife-edge case, the marginal value of knowledge is constant, so $V(i,n)$ is linear in both state variables, and the innovation policy functions (λ_b, λ_c) are linear in i and independent of n. This result follows directly from Proposition 1.

while a successful novel innovation yields

$$v_b(i) = V(i+1, n+1) - V(i, n) = a(i+1) - a(i) + \frac{\bar{\pi}}{r+\mu}.$$

The gap

$$v_b(i) - v_c = a(i+1) - a(i) > 0$$

is the option value of the additional ingredient, i.e., the incremental value of future innovation unlocked by expanding the knowledge stock. Second, the separable form of $V(i, n)$ clarifies why startups can command positive market value even without revenue-generating products or patents: their knowledge, i.e., the stock of ingredients, embodies an option to create future products, a value captured precisely by $v_{startup} = V(i, 0) = a(i) > 0$.[25]

Entrant problem. Entrants choose novel R&D intensity η and solve

$$\max_{\eta} \quad \underbrace{-\frac{A_e}{\phi_b}\eta^{\phi_b}}_{\text{innovation costs}} + \underbrace{\eta\,[V(1,1) - V(0,0)]}_{\text{expected gains from entry}} \;\Rightarrow\; \eta = \left(\frac{V(1,1) - V(0,0)}{A_e}\right)^{\frac{1}{\phi_b - 1}}, \tag{4.7}$$

where the first term captures the cost of novel innovation and the second term reflects the expected gains from its success. Note that $V(0, 0) = 0$.

Equilibrium firm distribution. The cross-sectional distribution of firms over knowledge and products is endogenous and determined in equilibrium. Let $\Phi_t(i, n)$ denote the mass of firms at time t with knowledge i and product count n so that $\sum_{i=0}^{\infty}\sum_{n=0}^{\infty} n\Phi_t(i, n) = 1$ and that $\Phi_t(0, 0) = 0$. Given firms' innovation policies, creative destruction, and the exogenous processes for ingredient depreciation, $\Phi_t(i, n)$ evolves according to the Kolmogorov forward equation:

$$\begin{aligned}\dot{\Phi}_t(i,n) = &\underbrace{\lambda_b(i-1)\,\Phi_t(i-1,n-1)\,\mathbf{1}\{i \geq 1,\ n \geq 1\}}_{\text{novel inflow}} + \underbrace{\lambda_c(i)\,\Phi_t(i,n-1)\,\mathbf{1}\{n \geq 1\}}_{\text{combination inflow}} \\ &+ \underbrace{\xi\,(i+1)\,\Phi_t(i+1,n)}_{\text{ingredient depreciation inflow}} + \underbrace{\mu\,(n+1)\,\Phi_t(i,n+1)}_{\text{creative destruction inflow}} + \underbrace{\eta\,\mathbf{1}\{i=1,\ n=1\}}_{\text{entry}} \\ &- \underbrace{\left[\lambda_b(i) + \lambda_c(i) + \xi\, i + \mu\, n\right]\Phi_t(i,n)}_{\text{outflows}},\end{aligned} \tag{4.8}$$

where $\dot{\Phi}_t(i, n)$ is the time derivative. Intuitively, firms arrive at state (i, n) via: (i) a successful novel innovation from $(i-1, n-1)$ at rate $\lambda_b(i-1)$; (ii) a successful combinatorial innovation

[25] In the baseline model, successful entry delivers both a new ingredient and a product, so entrants arrive at $(i, n) = (1, 1)$; states of $(i, 0)$ arise only when products depreciate. In the full model (see Section 6), entry (and novel innovation) may deliver an ingredient without a product, $(i, n) = (1, 0)$.

from $(i, n-1)$ at rate $\lambda_c(i)$; (iii) ingredient depreciation from $(i+1, n)$ at rate $\xi(i+1)$; or (iv) creative destruction of a product from $(i, n+1)$ at rate $\mu(n+1)$. In addition, entry places mass at $(1, 1)$. Symmetrically, firms leave (i, n) due to: (i) their own novel or combinatorial success; (ii) ingredient depreciation; or (iii) creative destruction of one of their products. In a stationary equilibrium, $\dot{\Phi}_t(i, n) = 0$ for all (i, n).

Although a closed form for the joint stationary distribution $\Phi(i, n)$ is not available in general, the marginal distribution over knowledge is tractable. Define with a slight abuse of notation:

$$\Phi_t(i) \equiv \sum_{n \geq 0} \Phi_t(i, n), \text{ for } i \geq 1 \; ; \quad \Phi_t(0) \equiv \sum_{n \geq 1} \Phi_t(0, n), \tag{4.9}$$

where for $i = 0$, I sum only over $n \geq 1$ since firms with $(i, n) = (0, 0)$ exit. Summing equation (4.8) over n cancels terms that only move along the product dimension and yields a one-dimensional birth-death system in i:

$$\begin{aligned} \dot{\Phi}_t(i) = & \underbrace{\lambda_b(i-1)\,\Phi_t(i-1)\,\mathbf{1}\{i \geq 1\}}_{\text{novel inflow}} + \underbrace{\xi\,(i+1)\,\Phi_t(i+1)}_{\text{ingredient depreciation inflow}} + \underbrace{\eta\,\mathbf{1}\{i=1\}}_{\text{entry}} \\ & - \underbrace{\Big[\lambda_b(i) + \xi\, i\Big]\,\Phi_t(i)}_{\text{outflows}} - \underbrace{\mu\,\Phi_t(0,1)\mathbf{1}\{i=0\}}_{\text{outflows from creative destruction}}, \end{aligned} \tag{4.10}$$

where $\dot{\Phi}_t(i)$ denotes the time derivative. Intuitively, mass arrives at knowledge level i from $i-1$ via successful novel innovation and from $i+1$ via depreciation; entry adds mass at $i = 1$. Mass leaves i due to novel success to $i+1$ or depreciation to $i-1$. The term $\mu\Phi_t(0,1)\mathbf{1}\{i=0\}$ captures exit of zero-knowledge single-product firms when their lone product is creatively destructed. The next proposition characterizes the stationary marginal distribution $\Phi(i)$ for $i \geq 1$.

Proposition 2. *In a stationary equilibrium satisfying equation (4.8), the marginal distribution over knowledge defined in equation (4.9) satisfies*

$$\Phi(1) = \frac{\eta}{\xi}\,, \quad \text{and} \quad \lambda_b(i)\,\Phi(i) = \xi(i+1)\,\Phi(i+1) \quad \text{for all} \quad i \geq 1. \tag{4.11}$$

In particular, the entire interior $\{\Phi(i)\}_{i \geq 1}$ is pinned down by $(\lambda_b(i), \xi)$ alone; creative destruction μ and combination arrival rate $\lambda_c(i)$ affect $\Phi(i)$ only through boundary objects at $i = 0$.

Two implications follow. First, the marginal knowledge distribution can be characterized largely in closed-form: given the incumbent policy $\lambda_b(i)$ and the depreciation rate ξ, the entire interior $\{\Phi(i)\}_{i \geq 1}$ is determined by the recursion in Proposition 2; only the boundary mass $\Phi(0)$ depends on product-side dynamics. Second, the recursion links the tail of $\Phi(i)$ to the shape of $\lambda_b(i)$: if $\lambda_b(i)$ is convex (grows faster than linearly), the stationary knowledge distribution fails to

exist (mass drifts upward); if $\lambda_b(i)$ is linear, $\Phi(i)$ exhibits an exponential tail; if $\lambda_b(i)$ is concave, $\Phi(i)$ has a thinner tail.

Link creative destruction to growth. The baseline model features creative destruction but does not explicitly model aggregate growth. Appendix Section B presents a simple extension that embeds growth and shows that the growth rate equals the product of the innovation step size and the creative-destruction rate, as is standard in the firm-dynamics literature (Klette and Kortum, 2004; Akcigit and Kerr, 2018; Peters, 2020). Holding the step size fixed, there is a one-to-one mapping between creative destruction and growth. For parsimony, the baseline model therefore uses the creative-destruction rate as a proxy for the growth rate.

4.3 Model mechanism

I solve the model numerically and illustrate its mechanisms by reporting firm value and policy functions together with the stationary distribution.

Panel A of Figure 4 plots firm value function $V(i, n)$ against knowledge i. The value function is increasing in i: more ingredients raise the option value of future innovation. For large i, $V(i, n)$ turns convex in knowledge (two knowledge-positive firms are more valuable together than apart), implying a merger incentive often absent from standard firm-dynamics models. This is consistent with empirical observations that pharmaceutical companies often merge, acquire, and pool chemical libraries to co-develop drugs. The convexity arises because $\frac{\zeta_c}{\phi_c - 1} > 1$ introduces curvature in the value of knowledge $a(i)$ (see equation 4.6): i enters the second term on the right-hand side with exponent $\frac{\zeta_c}{\phi_c - 1}$. In numerical experiments, convexity disappears when $\frac{\zeta_c}{\phi_c - 1} \leq 1$. When $A_b = \infty$ (no novel innovation), one can prove that $\frac{\zeta_c}{\phi_c - 1} > 1$ implies $a(i)$ is supermodular.

Panels B and C of Figure 4 plot the optimal arrival rates for novel and combinatorial innovation, $\lambda_b(i)$ and $\lambda_c(i)$, respectively, both only dependent on i but not n. Quantitatively, $\lambda_b(i)$ and $\lambda_c(i)$ are both increasing in i, indicating that firms with a larger knowledge stock undertake more innovation of both types. The increase in novel innovation is concave and slower, whereas the increase in combinatorial innovation is convex and steeper, implying that, under chosen parameters, firms tilt progressively toward within-firm combination as knowledge accumulates. This pattern arises because $\frac{\zeta_b}{\phi_b - 1} < 1 < \frac{\zeta_c}{\phi_c - 1}$.

Figure 5 turns to the knowledge distribution in stationary equilibrium. Panel A plots the marginal knowledge distribution $\Phi(i)$. In line with Proposition 2, the shape of $\Phi(i)$ is pinned down by $\lambda_b(i)$ and ξ; for a concave $\lambda_b(i)$, the tail is thinner than exponential. Panel B reports the conditional knowledge distribution $\Phi(i|n)$ for $n \in \{1, 5, 9\}$. Firm knowledge and product counts are positively correlated in equilibrium: moving from single-product firms to firms with nine products, the distribution shifts to the right, reflecting higher average knowledge among larger product-portfolio firms.

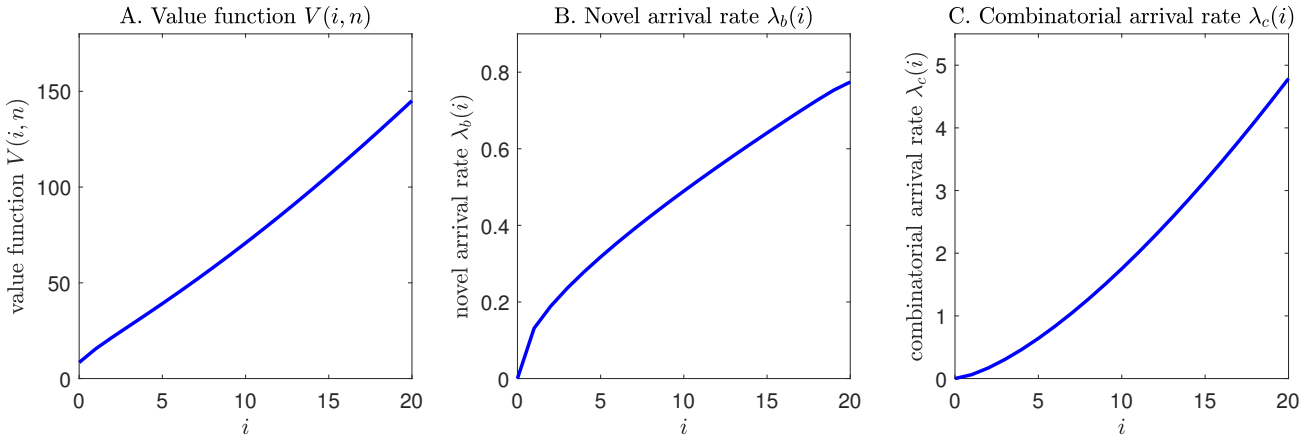

Note: This figure plots firm value function $V(i,n)$ at $n=1$ (panel A), and policy functions for novel innovation $\lambda_b(i)$ (panel B) and combinatorial innovation $\lambda_c(i)$ (panel C), respectively. The x-axis is firm number of ingredients i. Parameter values are set as in Table 5.

Figure 4: Firm value and policy functions.

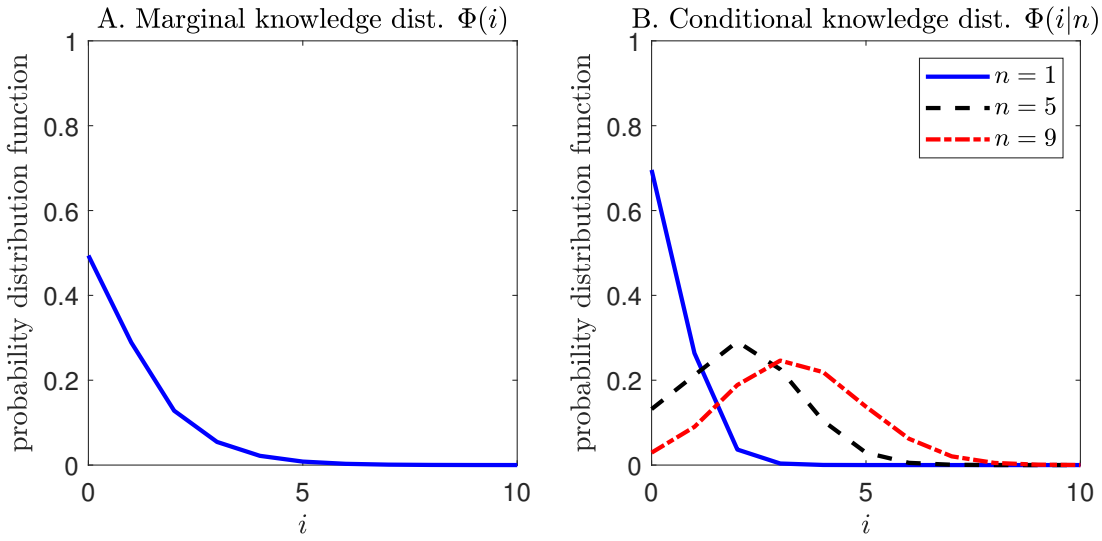

Note: Panel A plots the stationary marginal distribution of knowledge, $\Phi(i)$. Panel B plots the conditional knowledge distribution, $\Phi(i|n)$, for product counts $n \in \{1,5,9\}$. The x-axis is the number of ingredients i; the y-axis is the probability mass. Distributions are computed from the stationary solution to the law of motion in equation (4.8) (using the policy functions) with parameters as in Table 5.

Figure 5: Firm distribution in stationary equilibrium.

While my empirical setting is pharmaceuticals, the model is kept general and its mechanisms are thus broadly portable to other domains where recombination is central, such as materials science, software engineering, modular product design, and the citation networks of scientific publications and patents.

Table 5: Calibration and parameter choice.

Panel A: Parameters calibrated externally					
Parameter	Symbol	Value	Parameter	Symbol	Value
interest rate	r	2%	R&D curvature	(ϕ_b, ϕ_c)	2
Panel B: Parameters calibrated through moment matching					
Parameter	Symbol	Value	Moments	Data	Model
novel R&D, scale	A_b	110	pharma R&D / net income	0.8	0.8
comb. R&D, scale	A_c	135	share of within-firm comb. drugs	38%	37%
entry novel R&D, scale	A_e	750	entrants / active firms	26%	21%
novel R&D knowledge elas.	ζ_b	0.55	ing. share, top 20% firms	56%	56%
comb. R&D knowledge elas.	ζ_c	1.45	value: novel / combinatorial	1.6	1.7
depreciation rate	ξ	0.15	ingredient hazard rate	0.15	0.15

5 Quantitative Analysis

I now apply the theory to microdata from the anti-allergic drug market. Despite its parsimony, with only six internally calibrated parameters, the model is able to fit the micro-level patterns. I then use the model to quantify the contribution of within-firm combination to aggregate innovation.

5.1 Calibration

I begin by calibrating the model. After normalizing the product flow profits $\bar{\pi} = 1$, the framework has nine structural parameters. I pin down these parameters using a combination of external estimates from the literature and internal calibration via moment matching, as detailed in Table 5.

Externally calibrated parameters. I set the interest rate $r = 2\%$. The pharmaceutical industry's cost of capital and equity are estimated to be remarkably high at around 11%, reflecting substantial risk in drug development and marketing (Sertkaya et al., 2024; Damodaran, 2025). In the model, this feature is captured endogenously: as indicated in equation (4.5), a product's cash flow $\bar{\pi}$ is effectively discounted at $r + \mu$. The stationary equilibrium implies $\mu = 10\%$, so the effective discount rate is $r + \mu = 2\% + 10\% = 12\%$, in line with industry estimates.

Calibrating the curvature parameters (ϕ_b, ϕ_c) that map R&D outlays into Poisson arrival rates is challenging in the absence of detailed data on firms' R&D spending in anti-allergic drugs. Empirical studies estimating the elasticity of patents with respect to R&D expenditures typically find values around 0.5, which corresponds to quadratic cost curvature (see, e.g., Blundell, Griffith and Windmeijer, 2002; Acemoglu et al., 2018). Following the literature, I externally set

$\phi_b = \phi_c = 2$ and provide robustness in Appendix Section D.3. This calibration strategy is also adopted in other papers with similar model structures (e.g., Akcigit and Kerr, 2018; Peters, 2020; De Ridder, 2024).

The remaining parameters are calibrated by matching relevant moments. I first solve and simulate the model, then compute the model-implied moments, and adjust the parameters until these moments align with their empirical counterparts. The model matches the targeted moments well. Because each moment is influenced by all parameters, the parameters are determined jointly. Below, I provide a heuristic identification argument that links each parameter to the moment most directly informing its value.

Innovation cost function: scale parameters. The joint magnitude of (A_b, A_c) determines the overall level of innovation costs and hence influences the aggregate R&D-to-net-income-ratio. Their relative magnitude shapes the relative cost of novel versus combinatorial innovation, and thus the share of novel versus combinatorial products. I thus jointly calibrate (A_b, A_c) to two empirical moments: the pharmaceutical industry's R&D-to-net-income-ratio and the share of drugs being within-firm combination.[26] Drawing on published pharmaceutical industry estimates, R&D spending is of the same order as net income; I target a moderate R&D-to-net-income-ratio at 0.8 and show robustness in Appendix Section D.1.[27] "Combinatorial innovation" in the baseline model corresponds to within-firm combination in the empirical analysis (Section 3), whose share, among non-refinement drugs, is 38% $(= 35\%/(79\% + 13\%))$.[28]

For the entrant scale parameter A_e, a natural target would be the entry rate (entrants as a share of all firms). Because the data do not directly record firm entry/exit – only which firm patents which drug in which year – I instead target a related rate: the fraction of firms that patent their first drug in year t divided by all firms that patent at least one drug in year t. Averaged over 1990-2010, this rate is 26%.[29]

Innovation cost function: knowledge elasticity parameters. The knowledge elasticities (ζ_b, ζ_c) govern how a firm's knowledge stock scales its capacity for both novel and combinatorial

[26] I target R&D-to-net-income-ratio rather than R&D intensity because the model does not have a direct counterpart of revenue: flow profits $\bar{\pi}$ capture revenue net of production cost.

[27] U.S. pharmaceutical R&D intensity is estimated to be approximately 18% between 2000 – 2019 by the Congressional Budget Office (Office, 2021) and to be 18% for PhRMA members (mostly big pharmas) and 13% for firms covered by Business Enterprise Research and Development Survey of the National Science Foundation from 2008 – 2019 by Sertkaya et al. (2024). US pharmaceutical net-income margin is estimated to be approximately 18% for the largest 25 firms and 15% for a broader sample between 2006 – 2015 by the Government Accountability Office (U. S. Government Accountability Office, 2017) and to be 14% for the largest 35 firms between 2000 – 2018 by Ledley et al. (2020). Overall, R&D spending is at the same order as net income.

[28] In the baseline model, "novel" means new to the focal firm (not necessarily new to the market). Accordingly, novel innovation in the model maps to both "novel drug" (a drug including new-to-market functional groups) and "across-firm combination" (a drug recombining functional groups across firms) in the empirical analysis of Section 3, whereas combinatorial innovation in the model maps to "within-firm combination" (a drug recombining functional groups only within the firm).

[29] The model implies an entry rate of 4.3%, close to the estimate of 5.8% in Akcigit and Kerr (2018).

innovation. They pin down two aspects of the model. First, they determine the value of knowledge, which maps into the value difference between novel and combinatorial innovations. Second, they shape the process of knowledge and product accumulation and hence the equilibrium concentration of knowledge and products. I therefore discipline (ζ_b, ζ_c) with two moments aligned to these aspects.

First, I target the value difference between novel and combinatorial drugs. In Table 4, column (1) focuses on market value and implies that, conditional on clinical-trial entry, novel drugs command on average 63% higher market value than within-firm combinations ($\exp((1.02+0.02)-(0.56-0.01))-1$), and across-firm combinations (which the baseline model classifies as "novel") command 13% higher market value than within-firm combinations ($\exp((0.74-0.07)-(0.56-0.01))-1$). Column (2) focuses on scientific value and yields analogous premia of 75% and 34% respectively. Aggregating these pieces, I target a novel-to-combinatorial value ratio of 1.6 in the baseline calibration and show robustness in Appendix Section D.2. Second, I target the ingredient share held by the top 20% of firms (ranked by ingredient stock), which is 57% in the data. Section 5.2 details the construction of this moment in the model.

Ingredient depreciation rate. Knowledge in pharmaceuticals is subject to obsolescence, as technologies and scientific understanding evolve rapidly. In the data, I observe that firms frequently stop developing drugs with older functional groups and shift toward newer ones. I estimate the hazard rate at which a firm discontinues the use of a functional group using a maximum likelihood approach. Specifically, for each firm f, I collect all functional groups fi used between 1990 and 2010, and record the first and last years of use as (F_{fi}, L_{fi}). An ingredient is considered *depreciated* ($D_{fi}=1$) if $L_{fi} \leq 2010$. I define its observed duration as

$$S_{fi} = \min\{L_{fi}, 2010\} - F_{fi}.$$

The hazard rate $\hat{\xi}$ is then estimated by the maximum likelihood estimator for an exponential duration model,

$$\hat{\xi} = \frac{\sum_{fi} D_{fi}}{\sum_{fi} S_{fi}},$$

which equals the ratio of the total number of depreciation events to the total observed functional group-years in the sample. I estimate $\hat{\xi} = 0.15$ (standard error, 0.002) and calibrate the model's ingredient depreciation rate to this value.

5.2 Model validation

This section evaluates the model against a set of untargeted moments. I report two complementary exercises: one is cross-sectional, and the other traces innovation dynamics over time. First, I

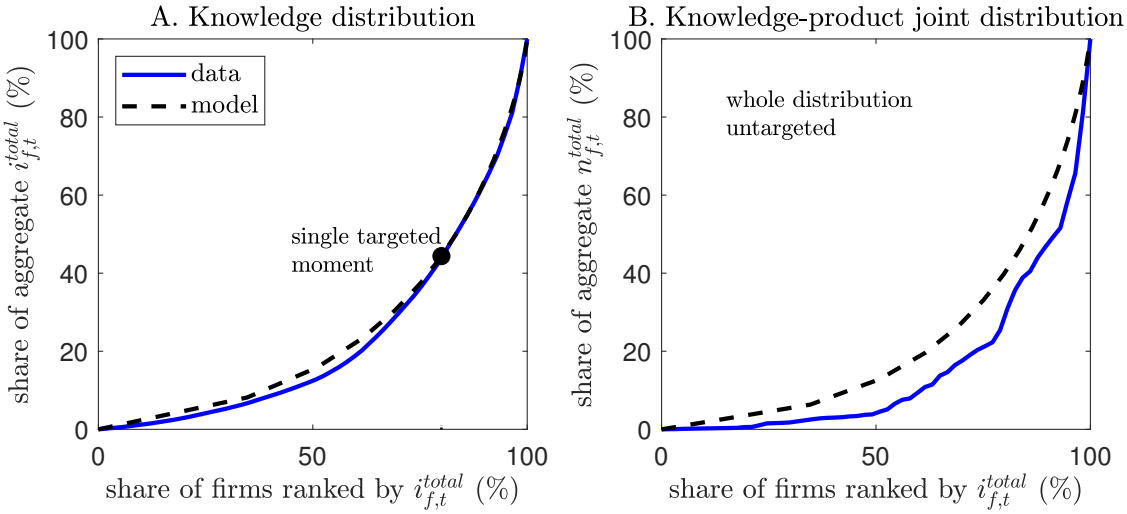

Note: This figure compares the model and 1990 data for the distribution of knowledge (i.e., functional-group holdings, panel A) and the joint distribution of knowledge and product portfolio (i.e., patented drugs, panel B). Firms on the x-axis are ranked by their cumulative knowledge stock $i^{total}_{f,t}$ as of 1990; the x-axis shows the cumulative share of firms from lowest to highest $i^{total}_{f,t}$. Panel A plots a Lorenz-style curve: the y-axis is the cumulative share of aggregate $i^{total}_{f,t}$ as of 1990, capturing the knowledge distribution. Panel B keeps the same x-axis ranking but changes the y-axis to the cumulative share of products as of 1990, capturing the knowledge-product joint distribution. Solid lines depict data; dashed lines depict the model. Because firm entry and exit are not directly observed in the data, I restrict the sample for data construction to firms that patented at least one drug both before 1990 and during 1990 – 1995.

Figure 6: Knowledge distribution and knowledge-product joint distribution: data versus model

compare the cross-sectional marginal distribution of knowledge and the cross-sectional joint distribution of knowledge and product (drug) portfolios in the data to those implied by the model. Second, I estimate model-implied elasticities of innovation outcomes with respect to firm knowledge and compare them to empirical evidence in Section 3.2.

Knowledge distribution and knowledge-product joint distribution. Figure 6 reports, from the data, the cross-sectional marginal distribution of knowledge (panel A, blue solid line) and the joint distribution of knowledge and products, i.e., drugs, (panel B, blue solid line). For each firm f that had patented at least one drug by 1990, I construct cumulative counts through 1990 of (i) the number of patented drugs ($n^{total}_{f,1990}$) and (ii) the number of distinct functional groups used in those drugs ($i^{total}_{f,1990}$). The notations are consistent with those in Section 3.2. I rank firms by $i^{total}_{f,1990}$; I then plot the cumulative share of aggregate $i^{total}_{f,1990}$ (panel A) and of aggregate $n^{total}_{f,1990}$ (panel B). The market is concentrated: the top 20% of firms ranked by knowledge account for 56% of the aggregate knowledge stocks and 71% of all patented drugs.

On the model side, natural proxies for $(i^{total}_{f,1990}, n^{total}_{f,1990})$ are firm state variables $(i_{f,t}, n_{f,t})$. However, the data use cumulative objects that do not subtract for depreciation of knowledge or creative destruction of products. To ensure comparability, I augment the model simulation to track cumulative measures of knowledge and products $(i^{total}_{f,t}, n^{total}_{f,t})$ alongside the states; these measures

accumulate new ingredients and products without netting out subsequent losses. Because every firm exits in finite time with probability one in the model[30], these cumulative measures remain bounded, and the joint distribution of $(i_{f,t}, n_{f,t}, i^{total}_{f,t}, n^{total}_{f,t})$ converges to a stationary distribution.

Panels A and B of Figure 6 overlay the model-implied distributions (black dashed lines). Although the calibration targets only a single moment – the top-20% share of knowledge (marked in panel A) – the model tracks the full cross-sectional knowledge distribution and knowledge-product joint distribution well. In the data, the top 20% firms account for 56% of knowledge and 71% of drugs; in the model, the corresponding shares are 56% and 60%. At the top 10% firms, the data report 36% of knowledge and 53% of drugs, versus 37% and 40% in the model. Overall, beyond the single targeted moment, the model reproduces the shape of the knowledge distribution and captures much of the knowledge-product concentration observed in the data.

Heterogeneous innovation across firms. I next examine, in the model, how firms at different points of the knowledge distribution innovate differently – both how much they innovate and which types of innovation they specialize in – and compare these implications to the evidence in Section 3.2. To mirror the empirical design, I simulate an economy of 50,000 firms for 20 years and divide the simulation into four five-year periods. Along the simulation, I record $(i_{f,t}, n_{f,t}, i^{total}_{f,t}, n^{total}_{f,t})$ as defined above. For each firm-period, I then construct the same firm-level outcomes as in the data: the number of newly patented drugs, $n^{new}_{f,t}$, and the number of new ingredients appearing in those drugs, $i^{new}_{f,t}$.

Innovation intensity. I first estimate PPML specifications that relate innovation to knowledge, paralleling the regressions in Section 3.2. Using the state variable,

$$E[y_{f,t}|X_{f,t}] = \exp(\alpha + \beta_i \ln(1 + i_{f,t}) + \gamma_t), \tag{5.1}$$

with $y_{f,t} \in \{n^{new}_{f,t}, i^{new}_{f,t}\}$ and γ_t denoting period fixed effects. The coefficient of interest β_i (subscript i for "ingredients") is the elasticity of innovation with respect to knowledge $i_{f,t}$.

The results are reported in Table 6. Columns (1) – (2) show that firms with larger knowledge stocks are more innovative: a 1% increase in $i_{f,t}$ is associated with 1.38% more new drugs (column 1) and 0.92% more new ingredients (column 2) over the next five years. These positive elasticities reflect $\zeta_b > 0$ and $\zeta_c > 0$, i.e., knowledgeable firms are more capable of innovation, and are qualitatively consistent with evidence in Section 3.2.

Innovation specialization. I next examine composition of innovation. In the data, I decompose new drugs into novel, across-firm combination, within-firm combination, and refinement. In the

[30] As shown in Figure 4, the novel R&D rate $\lambda_b(i)$ is concave in i, lying everywhere below the depreciation flow ξi, which is linear in i. Consequently, the knowledge stock follows a birth-death process with negative drift and, starting from any $i \geq 1$, hits $i = 0$ in finite time almost surely. Once at $i = 0$, the firm cannot innovate further ($\lambda_b(0) = \lambda_c(0) = 0$), and its product count evolves as a pure death process with per-product hazard μ; thus n also reaches zero in finite time.

model, I compute two shares within $n_{f,t}^{new}$: (i) the share that originates from novel innovation, $\%\text{novel}_{f,t}$ (corresponding to novel plus across-firm combination in the data); and (ii) the share from combinatorial innovation, $\%\text{comb}_{f,t}$ (corresponding to within-firm combination in the data). Because $\%\text{novel}_{f,t} + \%\text{comb}_{f,t} = 1$, I estimate the relationship between innovation specialization and firm knowledge using a fractional logit specification:

$$E[y_{f,t}|X_{f,t}] = \frac{\exp(\alpha + \beta_i \ln(1 + i_{f,t}) + \gamma_t)}{1 + \exp(\alpha + \beta_i \ln(1 + i_{f,t}) + \gamma_t)}, \tag{5.2}$$

with $y_{f,t} \in \{\%\text{novel}_{f,t}, \%\text{comb}_{f,t}\}$ and γ_t denoting period fixed effects. The coefficient of interest β_i captures how knowledge correlates with specialization. This formulation is equivalent to a multinomial logit specification with either $\%\text{novel}_{f,t}$ or $\%\text{comb}_{f,t}$ as the base category.

The results are reported in Table 6. Columns (3) – (4) show that richer knowledge stocks tilt firm toward combinatorial innovation: the coefficients on $\ln(1 + i_{f,t})$ are -0.25 for $\%\text{novel}_{f,t}$ (column 3) and 0.25 for $\%\text{comb}_{f,t}$ (column 4). Quantitatively, moving from the median-knowledge firm to one standard deviation above, an increase of 0.32 in $\ln(1 + i_{f,t})$, implies roughly 8% fewer novel innovations and 8% more combinatorial ones, consistent with a life-cycle pattern in which firms transition from early knowledge accumulation to later internal recombination. This pattern is consistent with empirical evidence (see Section 3.2). It arises because in my calibration, $\frac{\zeta_c}{\phi_c - 1} > \frac{\zeta_b}{\phi_b - 1}$: as knowledge accumulates, combinatorial costs decline faster than novel costs.

Control for portfolio size. Columns (5) – (8) of Table 6 add control for firm product portfolio, $\ln(1 + n_{f,t})$. The knowledge coefficients are essentially unchanged, while the product-portfolio coefficients are near zero and statistically insignificant across outcomes, echoing the theoretical result that optimal innovation arrival rates depend on knowledge i but not on the number of products n. The insignificance is also consistent with evidence reported in Tables 2 and 3. Table E.19 in the Appendix measures firm knowledge and product portfolio using $(\ln(i_{f,t}), \ln(n_{f,t}))$ and show that the results are robust.

State versus cumulative measures. Because the empirical results use cumulative measures of knowledge and products that do not net out depreciation or creative destruction, I repeat this estimation using the model's cumulative counterparts $(i_{f,t}^{total}, n_{f,t}^{total})$ as regressors.

The results are reported in Table 7. Columns (1) – (2) show that a 1% increase in $i_{f,t}^{total}$ is associated with 0.36% more new drugs and 0.23% more new ingredients; columns (3) – (4) show a 7% shift from novel to combinatorial innovation. Controlling for the cumulative number of products $n_{f,t}^{total}$ in columns (5) – (8) leaves the patterns intact; the portfolio coefficients remain small. The attenuation of coefficients on knowledge from Table 6 to Table 7 is expected for two reasons: (1) mechanically, $i_{f,t} \leq i_{f,t}^{total}$ (strict inequality in most cases), so a given percentage change in the smaller state variable translates into a larger proportional effect on outcomes; (2)

$i_{f,t}$ directly governs innovation choices in the model, whereas $i_{f,t}^{total}$ are correlated with innovation outcomes only through its (imperfect) correlation with $i_{f,t}$.

Comparison to the data. The estimation results using cumulative measures $i_{f,t}^{total}$ and $n_{f,t}^{total}$ (see Table 7) are directly comparable to empirical results from Section 3.2. Columns (5) – (6) of Table 7 imply that a 1% increase in cumulative knowledge predicts 0.53% more new drugs and 0.35% more new ingredients, versus 0.88% and 0.47% in the data (see Table 2). Columns (7) – (8) imply that moving from the median-knowledge firm to one standard deviation above, an increase of 0.78 in $\ln(i_{f,t}^{total})$, yields roughly a 7% decrease in novel innovation and a 7% increase in combinatorial innovation, compared with 11% and 15% empirically[31]; the differences are statistically indistinguishable.

In sum, estimation results from the model align well with empirical evidence both qualitatively and quantitatively: innovation intensity and specialization are driven by knowledge rather than size, and knowledge accumulation systematically rediects innovation effort toward internal recombination.

5.3 Contribution of combinatorial innovation

In the model, novel innovation expands a firm's knowledge stock; combinatorial innovation then converts that stock into additional products. Without novelty, there is no new knowledge and, hence, little scope for recombination. In this sense, novelty is the primitive margin, while recombination scales with the accumulated knowledge stock. I now use the model to quantify the importance of combinatorial growth. Specifically, I consider a counterfactual economy where combinatorial innovation becomes prohibitively costly ($A_c \to \infty$), forcing firms to rely exclusively on novel innovation.

It is tempting to conclude that shutting down recombination would mechanically reduce aggregate growth. However, in my model, two additional forces operate: a reinforcing force, without the capability to recombine existing knowledge, the payoff to expanding knowledge falls, depressing investment in novel innovation; and a countervailing general-equilibrium force, lower innovation activities reduce equilibrium creative destruction, and longer product lives raise expected per-product profits, partly restoring incentives for novel innovation and encouraging

[31] The corresponding empirical results are reported in Appendix Table E.20. Mapping model to data, "novel" in the baseline model corresponds to either a novel drug or an across-firm combination in the data, while "combinatorial" in the baseline model corresponds to within-firm combination. Accordingly, I construct three mutually exclusive shares in the data: (i) novel plus across-firm combination, (ii) within-firm combination, and (iii) refinement. I then estimate a fractional multinomial logit specification analogous to equations (3.1 – 3.2) using refinement as the base category. In terms of magnitudes from an analogous experiment in the data: moving from the median knowledge firm to one standard deviation above, an increase of 1.2 in $\ln(i_{f,t}^{total})$, implies roughly 11% fewer patented drugs that are classified as novel or across-firm combination (-0.09×1.2) and 15% more drugs classified as within-firm combination (0.13×1.2).

Table 6: Model-implied regression coefficients for innovation outcomes: $(i_{f,t}, n_{f,t})$.

	(1) $n^{new}_{f,t}$	(2) $i^{new}_{f,t}$	(3) $\%\text{novel}_{f,t}$	(4) $\%\text{comb}_{f,t}$	(5) $n^{new}_{f,t}$	(6) $i^{new}_{f,t}$	(7) $\%\text{novel}_{f,t}$	(8) $\%\text{comb}_{f,t}$
knowledge, $\ln(1+i_{f,t})$	1.38	0.92	-0.25	0.25	1.39	0.93	-0.25	0.25
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)
product, $\ln(1+n_{f,t})$					-0.01	-0.01	-0.00	0.00
					(0.01)	(0.01)	(0.00)	(0.00)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.08	0.03	–		0.08	0.03	–	
observations	86,818	86,818	50,644		86,818	86,818	50,644	

Note: The simulated dataset is constructed to match the empirical dataset in time frame, variable definitions, and sample structure, but with a larger simulated population of 50,000 firms for improved precision. Columns (1) – (2) use PPML specification (5.1). Columns (3) – (4) use fractional logit specification (see equation 5.2). Columns (5) – (6) use PPML specification (5.1) but include the product portfolio size control $\ln(1+n_{f,t})$. Columns (7) – (8) use fractional logit specification (5.2) but include the product portfolio size control $\ln(1+n_{f,t})$. The dependent variables are the number of new drugs patented over the subsequent five-year period for columns (1) and (5); the number of new ingredients used in those drugs for columns (2) and (6); the fraction of those drugs being novel for columns (3) and (7); and the fraction of those drugs being combinatorial for columns (4) and (8). Knowledge $\ln(1+i_{f,t})$ and product portfolio $\ln(1+n_{f,t})$ are measured as the log-transformations of firm state variables (i, n).

Table 7: Model-implied regression coefficients for innovation outcomes: $(i^{total}_{f,t}, n^{total}_{f,t})$.

	(1) $n^{new}_{f,t}$	(2) $i^{new}_{f,t}$	(3) $\%\text{novel}_{f,t}$	(4) $\%\text{comb}_{f,t}$	(5) $n^{new}_{f,t}$	(6) $i^{new}_{f,t}$	(7) $\%\text{novel}_{f,t}$	(8) $\%\text{comb}_{f,t}$
knowledge, $\ln(i^{total}_{f,t})$	0.36	0.23	-0.07	0.07	0.53	0.35	-0.09	0.09
	(0.01)	(0.01)	(0.00)	(0.00)	(0.02)	(0.02)	(0.01)	(0.01)
product, $\ln(n^{total}_{f,t})$					-0.15	-0.10	0.02	-0.02
					(0.02)	(0.02)	(0.00)	(0.00)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.03	0.01	–		0.03	0.01	–	
observations	86,818	86,818	50,644		86,818	86,818	50,644	

Note: The simulated dataset, specifications, and dependent variable definitions are the same as Table 6. Knowledge $\ln(i^{total}_{f,t})$ and product portfolio $\ln(n^{total}_{f,t})$ are cumulative measures of knowledge and products without netting out subsequent losses from depreciation and creative destruction.

entry.

I report the results in Table 8. In this experiment, the aggregate creative-destruction rate, and thus growth rate, falls by 30% (from 0.099 to 0.069). The decline is sharper for incumbent-driven displacement – 41% (from $0.078 = 0.041 + 0.037$ to $0.046 = 0.046 + 0$) – while entry-driven displacement rises by 10% (from 0.021 to 0.023), reflecting the improved returns from longer-lived products. Because, in the baseline, more knowledgeable firms specialize in combinatorial

Table 8: Counterfactual without within-firm combination ($A_c \to \infty$): equilibrium outcomes

	Baseline	No recombination	Δ (level)	Δ (%)
Product displacement dynamics				
aggregate creative-destruction rate	0.099	0.069	−0.030	−30%
incumbent-driven novel displacement	0.041	0.046	+0.005	+12%
incumbent-driven recomb. displacement	0.037	0	−0.037	−100%
entry-driven displacement	0.021	0.023	+0.002	+10%
Concentration				
top 20% (by knowledge) share of products	60%	55%	−5%	−8%

Note: "No recombination" sets within-firm combination costs prohibitively high ($A_c \to \infty$); all other parameters are held fixed. I recompute the stationary equilibrium. Displacement dynamics are rates per unit time. Concentration is the product share held by the top 20% of firms ranked by knowledge. Reported percentage changes reflect the discussion in the text.

innovation to expand product counts, shutting down recombination hurts them disproportionately. Market concentration therefore falls: the top 20% of firms by knowledge hold 55% of products in this experiment, compared with 60% in the baseline.

Combinatorial innovation, i.e. within-firm combination, amplifies the returns to knowledge: it is a key mechanism linking knowledge accumulation to product proliferation and, in equilibrium, to market concentration. Removing it weakens this link, reduces overall innovation and displacement, and shifts activity toward entry rather than incumbent product turnover. In this section, I quantify the importance of within-firm combination. In the next section, I introduce the full model that endogenizes across-firm combination and allows me to quantify the importance of full combinatorial innovation (within- plus across-firm combination).

6 Full Model: Idea Duplication by Coincidence

In Sections 4 – 5, I keep the model parsimonious, focus on the stationary equilibrium, and show its capability in matching a series of micro-level facts. At this stage, however, the model is silent on two salient empirical patterns documented in Section 3.1. First, the rising importance of recombination. Does recombination crowd out novel discoveries and thereby is detrimental to long-run growth? Or does it instead operate as a powerful engine of growth when other types of innovation slow down? What is the total contribution of combinatorial growth (across-firm plus within-firm)? Second, the decline in novelty. The data indicate a slowing expansion of the knowledge frontier, yet the current model does not specify whether a "new" ingredient is new to the market or merely new to the focal firm, and thus lacks a well-defined notion of knowledge frontier. Will the knowledge frontier eventually cease to expand, with consequences for long-run

growth? Can policy instruments sustain novelty when it becomes scarce?

In this section, I address these aspects by augmenting the baseline model with a single, pervasive force in innovation: idea duplication by coincidence. When a firm succeeds in novel innovation, the obtained ingredient may already exist elsewhere in the economy, pioneered earlier by another firm. The probability of such coincidence rises with the size of the aggregate knowledge pool. This mechanism is natural and prominent both in pharmaceutical R&D and in research broadly. The sociology literature coins this phenomenon "multiples" or "multiple discoveries" and documents its wide presence (e.g. Ogburn and Thomas, 1922; Merton, 1961, 1963; Lamb, 1984).[32] Throughout this section, I refer to new ingredients with no prior market appearance as new-to-market ingredients; they define the knowledge frontier. New ingredients that were previously pioneered elsewhere are merely new to the firm; I call them new-to-firm ingredients.

This duplication force makes new-to-market ingredients progressively harder to find. Yet because combinatorial innovation can capitalize on the expanding stock of known ingredients, the economy still exhibits sustained growth, with equilibrium innovation staying positive and bounded away from zero. Overall, the model delivers an "unbalanced" growth path with rising importance of combinatorial growth. The full model provides two further advantages. First, it introduces theoretical counterparts to the three empirical innovation types – novel, across-firm combination, and within-firm combination – enabling a clearer analysis of their trade-offs. It also allows me to quantify the full contribution of combinatorial innovation (across-firm plus within-firm). Second, the full model reveals a sharp policy trade-off: rewards for novelty generate a sharp by temporary boost in growth, whereas rewards for combination raise growth persistently with heterogeneous effects across firms.

6.1 Idea coincidence

The baseline model in Section 4 assumes that novel innovation generates an ingredient that is new to the focal firm but does not specify whether this ingredient is new to the market. I now introduce this distinction in a way that largely preserves the analysis in Section 4. The full model delivers a growth path that is unbalanced – declining novelty alongside rising recombination – so in this section I index variables by time t.

Idea coincidence. Let $I_t \geq 0$ denote the aggregate measure of distinct ingredients ever discovered, a continuous aggregate state variable, not a count. When a firm succeeds in novel

[32] In pharmaceutical research, firms scan common literature and attend common conferences, converging to similar ideas of drug design. Beyond pharmaceuticals, this coincidence pattern is widespread across innovative activities. Classic work by Ogburn and Thomas (1922) notes that such coincidence goes beyond famous examples such as the invention of calculus (Newton versus Leibniz) to "many important scientific discoveries not as well known." Similarly, Merton (1961, 1963) emphasize that many discoveries arrive as "multiples," wherein independent researchers reach similar results nearly simultaneously or rediscover earlier findings.

innovation, the accompanying new ingredient coincides with any known ingredient in I_t with intensity γ, so the probability that this ingredient is new-to-market (appearing for the first time in the market) follows

$$\Pr(\text{new-to-market}|\gamma, I_t) = \exp(-\gamma I_t).$$

For all successful novel innovation across firms and over time, I assume γ follows a Gamma distribution with shape parameter θ and unit scale parameter, $\gamma \sim \text{Gamma}(\theta, 1)$. Integrating over γ yields

$$\pi(I_t) \equiv \Pr(\text{new-to-market}|I_t) = (1 + I_t)^{-\theta}. \tag{6.1}$$

Equation (6.1) has an intuitive interpretation: as the aggregate knowledge pool expands, a "new" idea becomes likelier to coincide with existing knowledge, making truly new-to-market discoveries increasingly rare while increasing the chance of rediscovering an ingredient pioneered by other firms.

To fix language, I retain the terms "novel innovation" (arrival rate λ_b) and "combinatorial innovation" (arrival rate λ_c) from Section 4. A success from novel innovation is a novel success. The resulting ingredient is either new-to-market, in which case the product is a novel product, or new-to-firm (previously used by other firms but not by the focal firm), in which case the product is an across-firm combination. By contrast, a success from combinatorial innovation is always a within-firm combination that uses only ingredients already in the firm's own knowledge stock.

Moreover, I assume that an across-firm combination earns flow profits $\bar{\pi}$ only with probability p; with probability $1 - p$, the new product is blocked by the pioneering firm of the new ingredient and earns zero profits. Thus, although a new-to-firm success always adds an ingredient to the firm's knowledge stock, it translates into a new product only with probability p. This assumption is consistent with industry facts: when a promising functional group is identified, pioneering firms often file families of composition-of-matter patents with broad claims to deter follow-on entry and raise imitation costs (e.g., Cohen, Nelson and Walsh, 2000; Sternitzke, 2013; Wagner, Sternitzke and Walter, 2022). The possible innovation outcomes and the associated state transitions are summarized in Figure 7.

Note that when $p = 1$, i.e., when novel products and across-firm combinations yield the same profits, I_t ceases to matter at the firm level: firms are indifferent to whether a new ingredient is new-to-market or merely new-to-firm. In this special case, the firm problem collapses to the analysis in Section 4 and the model is stationary. The economy still exhibits declining novelty and rising recombination through $\pi(I_t)$, but there is no meaningful transition in the rate of entry, creative destruction or growth. Below, I focus on the empirically relevant case $p < 1$.

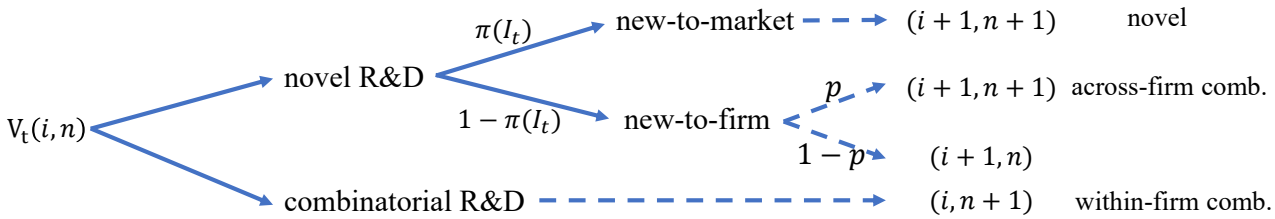

Figure 7: Innovation type illustration

Equilibrium. I now characterize the Markov perfect equilibria of the economy where firm strategies depend on the time-series of aggregate knowledge pool $(I_t)_t$ and creative destruction rate $(\mu_t)_t$.

Equilibrium definition. An equilibrium of the economy consists of incumbent innovation policy functions $(\lambda_{t,b}, \lambda_{t,c})$, entry rate η_t, creative destruction rate μ_t, ingredient depreciation rate ξ, interest rate r, aggregate distinct-ingredient stock I_t, and a joint distribution of firms $\Phi_t(i,n)$ over states (i,n), such that (i) incumbents choose $(\lambda_{t,b}, \lambda_{t,c})$ optimally (equation 6.2); (ii) entrants choose η_t optimally (equation 6.5); (iii) the total measure of products equals one $\sum_i \sum_n n\Phi_t(i,n) = 1$; (iv) I_t evolves according to $\dot{I}_t = \pi(I_t)\left(\eta_t + \int_{(i,n)} \lambda_{t,b}(i)\mathrm{d}\Phi_t(i,n)\right)$, since new-to-market ingredients arise from novel innovation by either entrants or incumbents; (v) the distribution $\Phi_t(i,n)$ evolves according to its law of motion (equation 6.6).

Incumbent firm problem. Consider a firm with i ingredients and n products. The firm takes $[r, \xi, (I_t)_t, (\mu_t)_t]$ as given and chooses R&D intensities for novel and combinatorial innovation (λ_b, λ_c) to maximize the present value of profits:

$$
\begin{aligned}
r\,V_t(i,n) = \max_{\lambda_b,\lambda_c}\Bigg\{ & \underbrace{\bar{\pi}\,n}_{\text{flow profits}} + \underbrace{\dot{V}_t(i,n)}_{\text{capital gain}} - \underbrace{\frac{A_b}{\phi_b}\,i^{-\zeta_b}\lambda_b^{\phi_b} - \frac{A_c}{\phi_c}\,i^{-\zeta_c}\lambda_c^{\phi_c}}_{\text{R\&D costs}} \\
& + \underbrace{\lambda_b\Big[k_t\,V_t(i{+}1,n{+}1) + (1-k_t)\,V_t(i{+}1,n) - V_t(i,n)\Big]}_{\text{novel R\&D} \to \text{novel + across-firm comb. + new product blocked}} + \underbrace{\lambda_c\Big[V_t(i,n{+}1) - V_t(i,n)\Big]}_{\text{comb. R\&D} \to \text{within-firm comb.}} \\
& + \underbrace{\xi\,i\Big[V_t(i{-}1,n) - V_t(i,n)\Big]}_{\text{ingredient depreciation}} + \underbrace{\mu_t\,n\Big[V_t(i,n{-}1) - V_t(i,n)\Big]}_{\text{creative destruction}} \Bigg\}. \qquad (6.2)
\end{aligned}
$$

The first line on the right-hand side collects the firm's instantaneous flow payoff net of innovation costs: each product yields flow profit $\bar{\pi}$; the term $\dot{V}_t(i,n)$ is the capital gain from moving through time; and the innovation costs depend on the firm's knowledge i and the chosen arrival rates (λ_b, λ_c). The second and third lines capture the expected value changes from the firm's Poisson events.

A novel success arrives at rate λ_b. With probability $\pi(I_t)$, the new ingredient is new to market, expanding both the ingredient set and the product portfolio $(i,n) \to (i+1, n+1)$. The product is novel. With probability $1-\pi(I_t)$, the new ingredient is new to firm; conditional on this, the firm obtains a product, an across-firm combination, with probability p (so $(i,n) \to (i+1, n+1)$) and the product is blocked with probability $1-p$ (so $(i,n) \to (i+1, n)$). Thus, the probability of obtaining a profit-generating product from a novel success is $k_t = \pi(I_t) + (1-\pi(I_t))p$. Combinatorial innovation arrives at rate λ_c and adds a product, a within-firm combination, using existing ingredients, $(i,n) \to (i, n+1)$. The remaining terms reflect losses in value: ingredients depreciate at rate ξ, reducing the knowledge stock $(i,n) \to (i-1, n)$; and creative destruction occurs at rate μ_t per product, displacing one product from the portfolio $(i,n) \to (i, n-1)$. Note that creative destruction does not destroy the firm's knowledge stock: the number of ingredients i is unaffected.

Value function $V_t(i,n)$ preserves the feature of being linear in n and its clear interpretation. Because R&D costs depend only on knowledge i but not on the product count n, the marginal value of a product is independent of (i,n) and equals the present value of its profit stream net of expected loss from creative destruction. The following proposition formalizes this observation

Proposition 3. *Consider the setup above. The value function is linear in the number of products:*

$$V_t(i,n) \;=\; a_t(i) \;+\; b_t\, n, \quad \textit{where} \quad r b_t = \bar{\pi} + \dot{b}_t - \mu_t b_t \tag{6.3}$$

and $a_t(i)$ satisfies

$$\begin{aligned} r a_t(i) - \dot{a}_t(i) \;=\; & \frac{\phi_b - 1}{\phi_b} (A_b i^{-\zeta_b})^{\frac{-1}{\phi_b - 1}} (a_t(i+1) - a_t(i) + k_t b_t)^{\frac{\phi_b}{\phi_b - 1}} \\ & + \frac{\phi_c - 1}{\phi_c} (A_c i^{-\zeta_c})^{\frac{-1}{\phi_c - 1}} b_t^{\frac{\phi_c}{\phi_c - 1}} \\ & + \xi i (a_t(i-1) - a_t(i)) \end{aligned} \tag{6.4}$$

The corresponding optimal innovation intensities $\lambda_{t,b}$ and $\lambda_{t,c}$ are:

$$\lambda_{t,b}(i) = \left(\frac{a_t(i+1) - a_t(i) + k_t b_t}{A_b i^{-\zeta_b}} \right)^{\frac{1}{\phi_b - 1}} \quad , \quad \lambda_{t,c}(i) = \left(\frac{b_t}{A_c i^{-\zeta_c}} \right)^{\frac{1}{\phi_c - 1}} ,$$

which depend solely on i and are independent of n.

Proof. See Appendix C. □

Proposition 3 shows that the value function is separable into two components: (i) $a_t(i)$, the option value of the current knowledge stock for generating future innovations; and (ii) $b_t\, n$, the annuity value of the current product portfolio. This mirrors Proposition 1, except that, because the environment is nonstationary, both a_t and b_t vary over t. Three implications from the baseline model carry over: (1) the innovation policy functions $(\lambda_{t,b}, \lambda_{t,b})$ depend only on firm knowledge i,

but not on product portfolio n once i is given; (2) combinatorial successes (from $\lambda_{t,c}$) are less valuable than novel successes (from $\lambda_{t,b}$); and (3) startups can command positive market value even without revenue-generating products or patents, since $V_t(i,0) = a_t(i) > 0$.

Entrant problem. At time t, entrants choose novel R&D intensity η and solve

$$\max_{\eta} \quad \underbrace{-\frac{A_e}{\phi_b}\eta^{\phi_b}}_{\text{innovation costs}} \quad + \underbrace{\eta\,[k_t V_t(1,1) + (1-k_t)V_t(1,0) - V_t(0,0)]}_{\text{expected gains from entry}}, \tag{6.5}$$

where the first term is the flow cost of novel innovation and the second term is the expected surplus from a novel success: with probability $k_t = \pi(I_t) + (1-\pi(I_t))p$ the entrant obtains both an ingredient and a product, otherwise only an ingredient.

Firm distribution. The cross-sectional distribution of firms over knowledge and products is endogenous and determined in equilibrium. Denote by $\Phi_t(i,n)$ the mass of firms at time t with knowledge i and product count n so that $\sum_{i=0}^{\infty}\sum_{n=0}^{\infty} n\Phi_t(i,n) = 1$ and that $\Phi_t(0,0) = 0$. Given firms' innovation policies, creative destruction, and the exogenous processes for ingredient depreciation, $\Phi_t(i,n)$ evolves according to the Kolmogorov forward equation:

$$\begin{aligned}
\dot{\Phi}_t(i,n) = {} & \underbrace{k_t\,\lambda_{t,b}(i-1)\,\Phi_t(i-1,n-1)\,\mathbf{1}\{i\geq 1,\, n\geq 1\}}_{\text{novel R\&D} \to \text{novel + across-firm comb.}} + \underbrace{(1-k_t)\,\lambda_{t,b}(i-1)\,\Phi_t(i-1,n)\,\mathbf{1}\{i\geq 1\}}_{\text{novel R\&D} \to \text{new product blocked}} \\
& + \underbrace{\lambda_{t,c}(i)\,\Phi_t(i,n-1)\,\mathbf{1}\{n\geq 1\}}_{\text{comb. R\&D} \to \text{within-firm comb.}} + \underbrace{\xi\,(i+1)\,\Phi_t(i+1,n)}_{\text{ingredient depreciation}} + \underbrace{\mu_t\,(n+1)\,\Phi_t(i,n+1)}_{\text{creative destruction}} \\
& + \underbrace{k_t\,\eta_t\,\mathbf{1}\{i=1,n=1\} + (1-k_t)\,\eta_t\,\mathbf{1}\{i=1,n=0\}}_{\text{entry}} - \underbrace{\Big[\lambda_{t,b}(i) + \lambda_{t,c}(i) + \xi\,i + \mu_t\,n\Big]\Phi_t(i,n)}_{\text{outflows}}.
\end{aligned} \tag{6.6}$$

Mass arrives at state (i,n) through five channels: (i) novel R&D with a product (either new-to-market with prob. $\pi(I_t)$ or new-to-firm and unblocked with prob. $(1-\pi(I_t))p$) moves firms from $(i-1,n-1)$ to (i,n) at rate $k_t\,\lambda_{t,b}(i-1)$; (ii) novel R&D, new-to-firm but blocked (prob. $(1-\pi(I_t))(1-p) = 1-k_t$) adds an ingredient but no product, moving $(i-1,n)\to(i,n)$ at rate $(1-k_t)\,\lambda_{t,b}(i-1)$. (iii) within-firm combination adds a product using existing ingredients, moving $(i,n-1)\to(i,n)$ at rate $\lambda_{t,c}(i)$; (iv) ingredient depreciation reduces knowledge, so firms flow from $(i+1,n)$ to (i,n) at rate $\xi(i+1)$; and (v) creative destruction of a product moves firms from $(i,n+1)$ to (i,n) at rate $\mu_t(n+1)$. Entry places mass at $(1,1)$ with probability k_t and at $(1,0)$ with probability $1-k_t$ at rate η_t. The final bracket subtracts outflows from (i,n) due to the firm's own novel and combinatorial successes, ingredient depreciation, and product destruction. Indicator functions ensure feasibility of the source states.

Numerical algorithm. The economy features an unbalanced transition in which the composition of innovation, entry, and aggregate creative destruction evolves over time. I compute the transition

with a standard backward-forward (HJB-KFE) routine, following the spirit of Achdou et al. (2022) and Buera and Shin (2013). There are four steps: (1) solve the initial and the terminal stationary equilibrium, which anchor the transition; (2) guess aggregate paths by posting a time path for (I_t, μ_t) along the transition; (3) given the guess, solve the value problem backward to obtain policies; then, starting from the initial steady state, simulate the distribution forward to the terminal date; (4) construct new $(I_t, \mu_t)_t$ from the simulated path, damp, and iterate until convergence. This procedure delivers a time path for policies, distributions, and aggregates consistent with the non-stationary environment.

Unlike standard transition exercises that start from a baseline stationary distribution and add a shock, my environment has no such "pre-shock" steady state: the aggregate knowledge stock I_t is ever increasing and $\pi(I_t)$ is ever decreasing. I therefore take a stand on the initial condition. Specifically, I choose a level I_0 (calibrated to an empirical moment) and compute the stationary equilibrium that would obtain if I were held fixed at I_0; I take the resulting stationary distribution as the initial cross-section. Along the transition, however, I_t is allowed to rise from I_0 and accumulate over time.

On the terminal side, the model has a well-defined limit: as $t \to \infty$, $I_t \to \infty$. Hence the new steady state corresponds to $\pi(I_t) \to 0$: every novel success is new-to-firm (never new-to-market), and the probability of obtaining a profit-generating product from a novel success converges to $k_t = p$.[33]

6.2 Calibration and quantitative analysis

I now apply the full model to microdata from the anti-allergic drug market. The model reproduces the macro patterns of declining novelty and rising recombination and highlights them as a feature of growth: as the knowledge pool expands, new-to-market ingredients become harder to find, so innovation dynamics drift down. Simultaneously, abundant combinatorial opportunities sustains progress and keeps the long-run growth rate bounded away from zero. I document these transitional dynamics and use the model to quantify the contribution of combinatorial growth. Crucially, the complete framework allows me to measure the full impact of recombination – not only within-firm combination but also across-firm combination.

Calibration. To make results consistent and comparable with the baseline model, I fix the parameters shared with the baseline model and calibrate only the new ones. Three parameters are new: (i) the stock of distinct ingredients in the initial stationary equilibrium, I_0; (ii) the shape parameter of Gamma distribution θ governing $\pi(I_t) = (1 + I_t)^{-\theta}$; and (iii) the probability p of a

[33] Proof sketch: if I_t converged to some finite $\bar{I}$, then $\pi(I_t) \to \pi(\bar{I}) > 0$. With positive novel arrival rate $\lambda_{t,b}(i) > 0$, aggregate novel successes would continue to add new-to-market ingredients at a rate non-diminishing to zero, implying further non-diminishing growth in I_t – a contradiction.

new-to-firm ingredient yielding a profit-generating (unblocked) product. These parameters pin down the model-implied dynamics of innovation composition between "novel" and across-firm combinations. In the model, the share of novel within novel + across-firm combination at time t is

$$\frac{\%\text{novel}_t}{\%\text{novel}_t + \%\text{comb}_t^{across}} = \frac{\pi(I_t)}{\pi(I_t) + (1 - \pi(I_t))p},$$

where $\pi(I_t)$ is the probability that a novel draw is new-to-market, and $(1-\pi(I_t))p$ is the probability it is new-to-firm and unblocked (thus generating an across-firm combination). The parameter θ also governs the growth dynamics of I_t by determining how quickly genuinely new-to-market ingredients become harder to find. Accordingly, I calibrate (I_0, θ, p) to match three moments: (1) the above novel/(novel+comb^{across}) ratio in 1990 (initial cross-section); (2) the same ratio in 2010 as implied by the 20-year model transition; and (3) the overall decline in the number of new ingredients over time.

Figure 8 shows that the full model reproduces the empirical targets closely. Panel A reports the share of novel products as of novel products and across-firm combinations. Year 0 on the x-axis represents 1990 in the data. The data (blue solid) drifts downward from 1990 to 2010, indicating waning novelty; the model path (black dashed) closely tracks this 20-year transition. Panel B plots the market-level introduction of new (new-to-market) ingredients (normalized to 1990 = 1); the model captures the secular decline in discoveries of new-to-market ingredients. Together, these panels confirm that calibrating (I_0, θ, p) is sufficient to match the shift from novelty to across-firm combination and the decline in novelty over time.

Transitional dynamics. Figure 9 displays the 200-year transition implied by the calibrated model. Panel A shows that the cumulative stock of ever-discovered ingredients I_t rises monotonically. As the knowledge base expands, the probability that a novel success yields a new-to-market ingredient falls (panel B), so the flow of novel discoveries declines over time (panel C).

Panel D shows that the creative destruction rate μ_t declines over time and converges to its new stationary level. Two forces are at work. First, as $\pi(I_t)$ decreases, genuinely novel products become less likely, which reduces incentives for novel innovation and lowers μ_t. Second, a lower μ_t lengthens product lifetimes and raises product profits; this general-equilibrium effect induces more R&D effort on the margin and partially offsets the fall in μ_t. In the long-run equilibrium, as $I_t \to \infty$ and $\pi(I_t) \to 0$, the creative destruction rate converges. Consistent with this mechanism, panel E shows a compositional shift: genuinely novel products become increasingly rare, the share of across-firm combinations rises, and the share of within-firm combinations increases modestly due to the general equilibrium effect. The entry rate η_t declines gradually along the transition and approaches its new steady state (panel F).

Overall, the model implies: growing knowledge makes genuine breakthroughs scarce and tilts

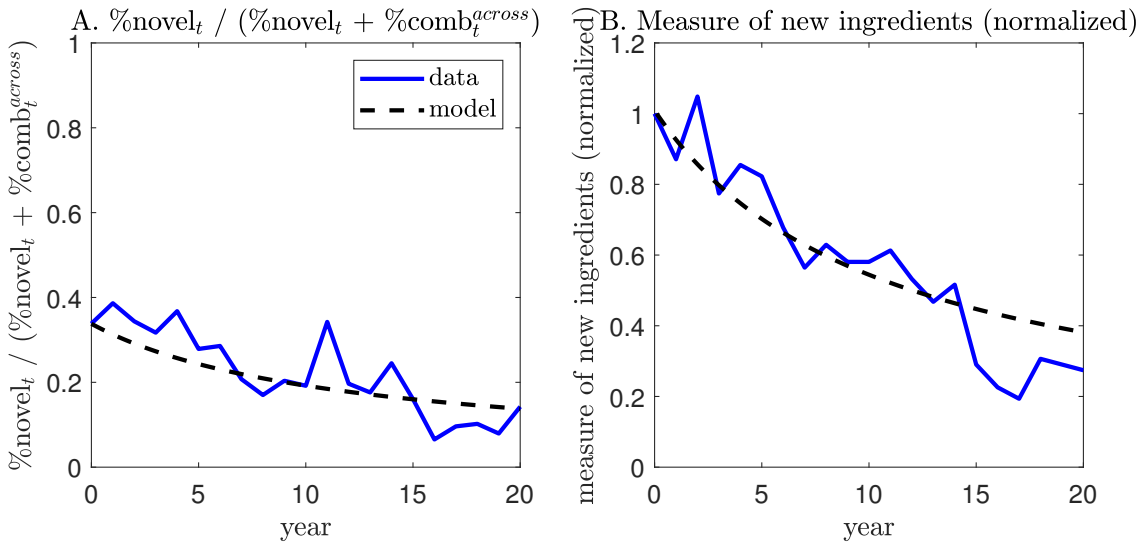

Note: Blue solid lines = data; black dashed lines = model. Panel A reports the fraction of novel drugs within "new-knowledge" drugs, defined as %novel / (%novel + %comb^{across}). Panel B reports the annual count of new-to-market ingredients, normalized to 1990 = 1. Model series depict a 20-year transition starting from the stationary distribution at aggregate ingredient stock I_0. Parameters shared with the baseline model are held fixed; (I_0, θ, p) are calibrated to match (i) the 1990 novel/(novel+across) ratio, (ii) the 2010 ratio implied by the transition, and (iii) the observed decline in new-ingredient introductions. Data: anti-allergic small-molecule drugs, 1990 – 2010."

Figure 8: Calibration fit: novelty share and decline in new ingredients (data versus model)

innovation toward recombination. In the long run, novelty fades, but the expanding opportunities of recombination sustain innovation and growth.

Contribution of combinatorial innovation. In Section 5.3, I use the baseline model to quantify the contribution of within-firm combination to aggregate innovation. The full model separates novel innovation from across-firm combination, allowing me to evaluate the total contribution of recombination (within-firm + across-firm) to aggregate dynamics. Specifically, I consider a counterfactual in which both types of combination become prohibitively costly ($p \to 0$ and $A_c \to \infty$), forcing firms to rely exclusively on genuine novelty.

Figure 10 compares this no-combination counterfactual (black dashed line) with the baseline calibration (blue solid line). In panel A, eliminating recombination triggers a sharp drop in aggregate creative destruction μ_t (e.g., from 0.093 to 0.017 at $t = 0$). This is immediate because, in the baseline, roughly 80% of new products at $t = 0$ comes from recombinations; removing them greatly dampens innovation dynamics. Over time, as the knowledge pool expands and $\pi(I_t)$ declines, genuinely novel products become extremely rare, so $\eta_t \to 0$ (panel B) and $\mu_t \to 0$ (panel A). Panel C examines market concentration, measured as the share of cumulative products $n_{f,t}^{total}$ held by the top 20% of firms ranked by knowledge $i_{f,t}^{total}$. The baseline series is fairly stable. Under the no-combination counterfactual, concentration is hump-shaped: it rises early as low-knowledge firms exit, while high-knowledge incumbents can coast on their larger existing portfolios; later,

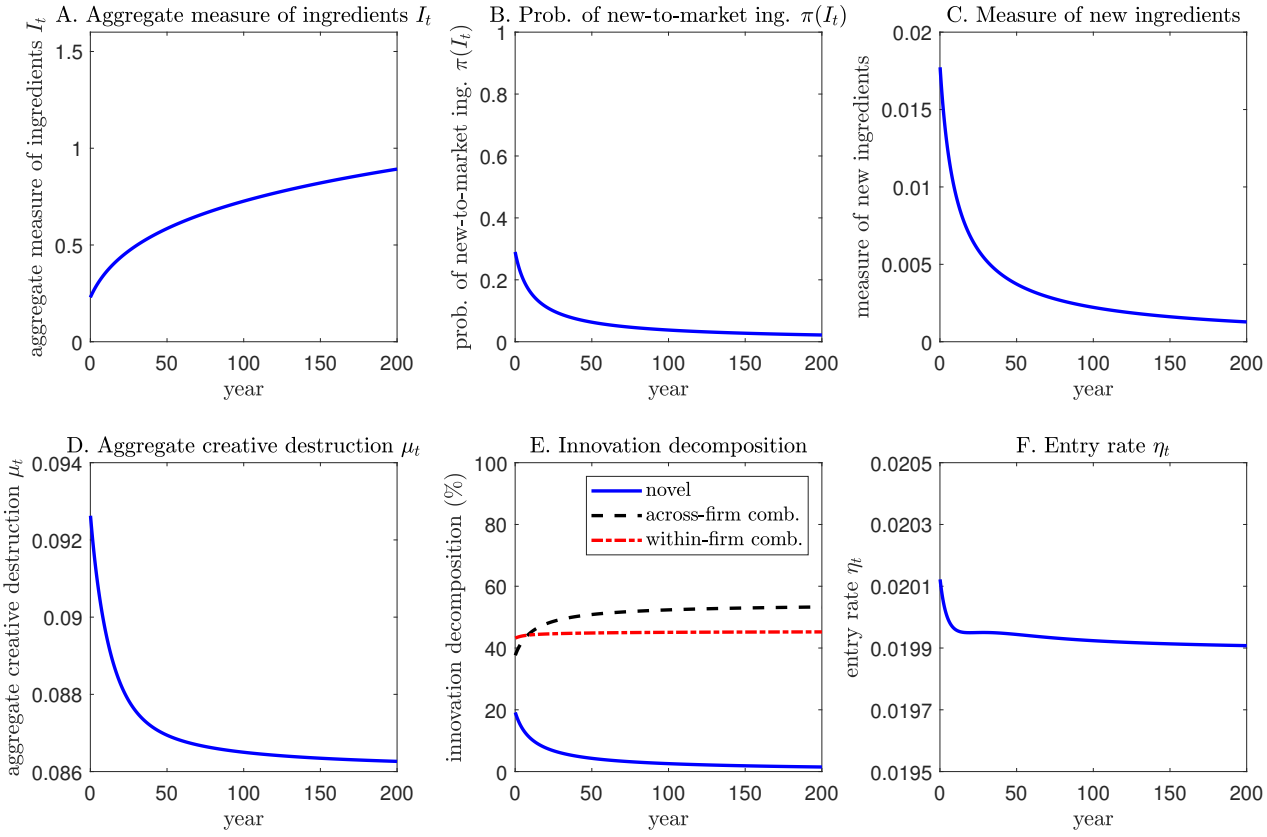

Note: Model-implied 200-year transition from the calibrated initial stationary equilibrium. Parameters shared with the baseline model are held fixed and reported in Table 5; (I_0, θ, p) and other transition-relevant parameters follow the calibration in Section 6.2. Axes show model years from the initial date $t = 0$.

Figure 9: Transitional dynamics.

as $\mu_t \to 0$ and entry dries up, even leaders cannot renew their product sets, portfolios gradually depreciate, and concentration falls toward a lower long-run level.

Taken together, the figure shows that combinatorial innovation is crucial for sustaining creative destruction, supporting entry, and – through continued turnover – maintaining the observed level of concentration.

6.3 Heterogeneous policy implications

Designing public policies to encourage innovation and promote technological and economic growth is a central policy concern. In pharmaceuticals, for example, the FDA encourages novelty by granting drugs that contain a new molecular entity (NME) a longer exclusivity period, thereby extending their market life.

In this section, I use the full model to evaluate how various innovation policies affect aggregate

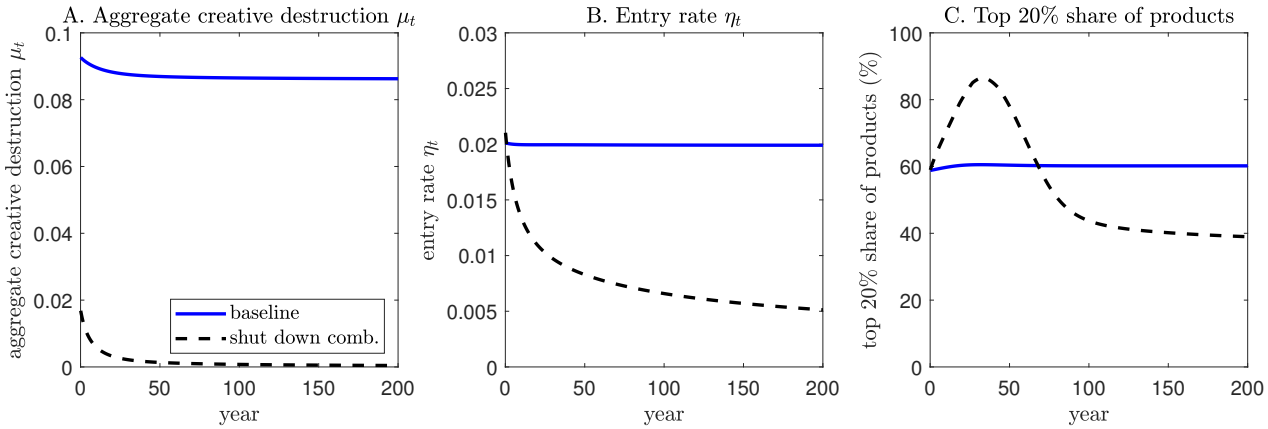

Note: Baseline (blue solid) versus a counterfactual that shuts down all recombination (black dashed) by setting $p \to 0$ and $A_c \to \infty$. Panel C plots share of cumulative products $n_{f,t}^{total}$ held by the top 20% of firms ranked by knowledge $i_{f,t}^{total}$, excluding firms with zero knowledge ($i_{f,t} = 0$). All series are generated by the full model.

Figure 10: Contribution of combinatorial innovation.

entry and growth, and to quantify their heterogeneous impacts across firms. Specifically, I evaluate three policies that aim to correct under-investment in R&D: fixed cash prizes awarded to firms that achieve (i) a novel invention, (ii) an across-firm combination, or (iii) a within-firm combination.[34] I label these policies \$novel, \$comb^{across}, and \$comb^{within}. Rather than solving the social planner's problem, I fix a common government budget across the three policies and compare their implications. For comparability with the baseline in Section 6.2, I keep the initial firm distribution unchanged and vary only the policy along the transition.

I keep the government budget at unity, interpreted as the total profits of all products net of production costs for a year. For each policy, I solve for the fixed cash prize s so that the total discounted prize payout equals unity. For example, for policy \$novel, I solve

$$\int_{t=0}^{\infty} \exp(-rt) \left[s \times \pi(I_t) \left(\eta_t + \int_{(i,n)} \lambda_{t,b}(i) \mathrm{d}\Phi_t(i,n) \right) \right] \mathrm{d}t = 1,$$

where the left-hand side is the present value of expected prize payments under policy \$novel. At time t, total novel innovation intensity is $\eta_t + \int_{(i,n)} \lambda_{t,b}(i) \mathrm{d}\Phi_t(i,n)$; each success yields a novel product with probability $\pi(I_t)$, and each novel product is rewarded s. Prize magnitudes reflect relative frequencies of successes: novel products are rare (thus $s = 5$), across-firm combinations are less frequent than within-firm combinations in early years (thus $s = 1$ versus $s = 0.6$).

[34] The extension in Appendix Section B clarifies why such policies are of interest. With the additional innovation step-size structure, the extended model features a standard "standing on giants' shoulders" force: the private value of an innovation does not internalize the option value it creates for future innovators, so equilibrium innovation intensity can be inefficiently low relative to the social optimum. Absent this force, the model unambiguously features inefficiently high innovation: the innovating firm does not internalize the business-stealing loss it imposes on others, so private value exceeds social value.

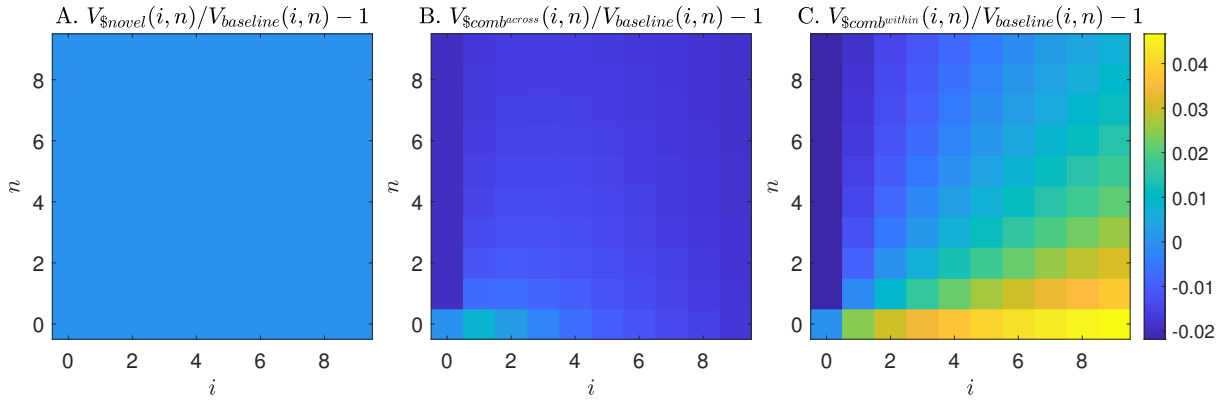

Note: Heatmaps report the long-run (new stationary) percentage change in firm value relative to baseline, $V_{policy}(i,n)/V_{baseline}(i,n)-1$, for each firm state (i,n). Values are computed in the limiting stationary economy reached under each policy (i.e., after the transition). Color indicates magnitude according to the shared color bar. Policies \$novel, \$comb^{across}, and \$comb^{within} respectively reward a fixed prize s to: (i) successful novel product containing a new-to-market ingredient; (ii) successful across-firm combinations (adopting an external ingredient); and (iii) successful within-firm combinations (combining ingredients already held). For comparability, the government budget is held fixed across policies and s is chosen so the present value of expected prize payments equals that budget.

Figure 11: Value changes across firms: before and after policy.

Figure 11 reports the long-run percentage change in firm value under each policy relative to baseline, $V_{policy}(i,n)/V_{baseline}(i,n)-1$, across firm state (i,n). Values are evaluated in the new stationary equilibrium reached in the long run, i.e., in the limiting economy where the aggregate ingredient stock is large ($I_t \to \infty$). Panel A focuses on policy \$novel. Value changes are identically zero. In the limiting stationary equilibrium $\pi(I_t)=0$, novel product never appears, so the prize is never paid and the policy is neutral. Panel B examines policy \$comb^{across}. The prize for across-firm combinations induces adoption of external ingredients through novel innovation. Gains are concentrated among low-knowledge firms (especially with small n), who can climb the knowledge ladder by importing others' ingredients, while the heightened creative destruction redistributes surplus away from high-knowledge incumbents, who experience value losses. Panel C shows that policy \$comb^{within}, by rewarding internal combinations, primarily benefits high-knowledge firms, especially those with relatively few products (large i, small n), because they hold many internal combinatorial opportunities to activate and can translate knowledge into products quickly. Firms with little knowledge have few internal ingredients to exploit and gain little (or may lose slightly via intensified competition).

Overall, the three policies have starkly different distributional effects: \$novel is long-run neutral, \$comb^{across} shifts value toward knowledge-poor firms by accelerating diffusion and catch-up, and \$comb^{within} shifts value toward knowledge-rich firms by enhancing returns to internal recombination.

Figure 12 plots the transition path under each policy against the baseline. Policy \$novel

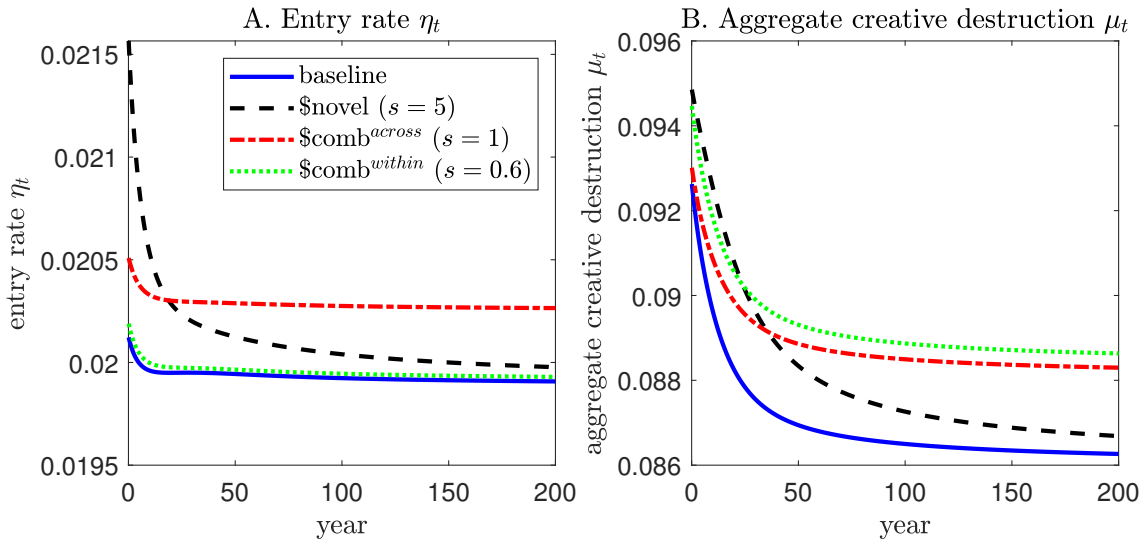

Note: Model-implied transition paths under four regimes: baseline (blue solid), \$novel (black dashed, $s = 5$), \$comb^{across} (red dash-dot, $s = 1$), and \$comb^{within} (green dotted, $s = 0.6$). Panel A plots the entrant flow rate η_t. Panel B plots the aggregate creative-destruction rate μ_t. For each policy, the prize s is chosen so that the present value of expected payouts equals a common government budget; other parameters and the initial firm distribution match Section 6.2. Time is in model years.

Figure 12: Entry and creative destruction implied by different policy scenarios.

(black dashed line) offers the highest reward per success. It triggers an immediate jump in entry rate (panel A) and, through heightened competition from entrants, a short-run rise in creative destruction (panel B). In the long run, however, both η_t and μ_t drift back to baseline because the probability of genuine novelty vanishes ($\pi(I_t) \to 0$), so the prize is rarely paid and does not alter steady-state dynamics. Policy \$comb^{across} (red dash-dotted line) provides a moderate reward and thus produces a moderate increase in entry and a smaller rise in creative destruction in the short run. Knowledge-poor firms respond by importing external ingredients, while knowledge-rich incumbents (specialized in internal recombination) react less. Both η_t and μ_t converge to levels above baseline in the long run. Policy \$comb^{within} (green dotted line) yields little change in entry – entrants arrive via novel innovation and are not directly rewarded – but it materially raises creative destruction from $t = 0$ because knowledge-rich incumbents monetize numerous internal recombinations. Creative destruction also settles above baseline in the new steady state.

Taken together, the policies trade off short-run stimulus against long-run impact and have heterogeneous effects across firms with differing knowledge stocks. Rewarding true novelty generates the largest short-run boost but cannot lift long-run growth. By contrast, \$comb^{across} raises long-run innovation mainly by helping knowledge-poor firms climb the knowledge ladder, whereas \$comb^{within} does so by amplifying the returns to recombination among knowledge-rich firms with abundant internal opportunities.

These implications map directly onto policy instruments that differentially support basic

versus applied research and indicate that the optimal mix depends on the policymaker's prioritized time horizon and tolerance for market concentration.

7 Conclusion

New ideas often arise by recombining existing ones. Do such recombinations slow or sustain economic growth? While recent theory formalizes the growth implications of recombination, direct evidence on its empirical relevance remains scarce.

This paper measures combinatorial growth and evaluates its importance empirically and in a quantitative model. I focus on the pharmaceutical industry, where the novelty-recombination distinction is transparent. I use a chemistry-based measure to decompose each drug into its constituent functional groups, discrete chemical building blocks that shape its properties. This approach applies to all small-molecule drugs, which account for 75% of drugs approved by the FDA's Center for Drug Evaluation and Research since 2000. This setting is suited to studying recombination: functional groups are standard chemical concepts, their presence/absence in a drug is discrete and observable, data are rich and longitudinal, and the drug discovery process is inherently combinatorial.

Focusing on anti-allergic drugs, I document three patterns. First, recombination is more prevalent than novelty. From 1990 – 2010, 79% of new drugs are recombinations of existing functional groups, while 13% contain new functional groups; novelty declines as recombination rises. Second, firms with less accumulated knowledge, measured as fewer distinct functional groups in their patented drugs, tend to develop drugs by adopting new functional groups, while more knowledgeable firms recombine their used functional groups. This points to a firm life cycle: early-stage knowledge accumulation to later-stage internal recombination. Third, conditional on clinical trial entry, recombinations are less valuable than novel innovation.

Motivated by the documented empirical patterns, I introduce the distinction between knowledge stock and product portfolio into Klette and Kortum (2004). Novel innovation builds knowledge, while combinatorial innovation deploys knowledge to create new products. These two margins of innovation are distinct yet connected, channeling growth through two processes: knowledge generation and knowledge application. I calibrate this model to the microdata from anti-allergic market. Despite its parsimony, the model reproduces untargeted empirical patterns in the joint distribution of knowledge and products, and in firms' innovation intensity and specialization across the knowledge distribution. The model attributes 30% of aggregate creative destruction and growth to within-firm combination.

Speaking to the macro patterns of declining novelty and rising recombination, I further introduce knowledge duplication into the model: as the aggregate knowledge stock expands, a "new"

idea increasingly overlaps with existing knowledge, making genuine novelty rarer while expanding the scope for recombination. The full model generates declining novelty, rising recombination and sustained growth driven by unlimited combinatorial opportunities. Recombination (within-firm and across-firm) can account for more than 80% of aggregate growth rate. Policy counterfactuals reveal heterogeneous effects. Rewards for novelty deliver the largest short-run boost to entry and competition but do not raise long-run growth because novelty becomes progressively rare. Rewards for combinations raise the long-run growth rate, with heterogeneous effects across firms with differing knowledge stocks: rewards for across-firm combination benefit less-knowledgeable firms and reduce concentration, while rewards for within-firm combination benefit more-knowledgeable firms and increase concentration.

References

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik. 2022. "Radical and Incremental Innovation: The Roles of Firms, Managers, and Innovators." *American Economic Journal: Macroeconomics*, 14(3): 199–249.

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, Reallocation, and Growth." *The American economic review*, 108(11): 3450–3491.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2022. "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach." *The Review of Economic Studies*, 89(1): 45–86.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth through Creative Destruction." *Econometrica : journal of the Econometric Society*, 60(2): 323–351.

Akcigit, Ufuk, and William R. Kerr. 2018. "Growth through Heterogeneous Innovations." *The Journal of political economy*, 126(4): 1374–1443.

Akcigit, Ufuk, William Kerr, and Tom Nicholas. 2013. "The Mechanics of Endogenous Innovation and Growth: Evidence from Historical U.S. Patents."

Arnum, Patricia Van. 2024. "Bio/Pharma Watchlist: Small-molecule Drugs."

Arrow, Kenneth J. 1962. "The Economic Implications of Learning by Doing." *The Review of Economic Studies*, 29(3): 155–173.

Arthur, Brian. 2009. *The Nature of Technology: What It Is and How It Evolves.* New York:Free Press.

Association for Accessible Medicines. 2024. "2024 U.S. Generic & Biosimilar Medicines Savings Report."

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review*, 110(4): 1104–44.

Blundell, Richard, Rachel Griffith, and Frank Windmeijer. 2002. "Individual Effects and Dynamics in Count Data Models." *Journal of econometrics*, 108(1): 113–131.

Buera, Francisco J., and Yongseok Shin. 2013. "Financial Frictions and the Persistence of History: A Quantitative Exploration." *Journal of Political Economy*, 121(2): 221–272.

Buis, Maarten L. 2017. "FMLOGIT: Stata Module Fitting a Fractional Multinomial Logit Model by Quasi Maximum Likelihood." *Statistical Software Components.*

Chandra, Amitabh, John Drum, Michael Daly, Henry Mirsberger, Samuel Spare, Ulrich Neumann, Silas Martin, and Noam Kirson. 2024. "Comprehensive Measurement of Biopharmaceutical R&D Investment." *Nature Reviews Drug Discovery*, 23(9): 652–653.

Cohen, Wesley M., and Daniel A. Levinthal. 1989. "Innovation and Learning: The Two Faces of R & D." *The Economic Journal*, 99(397): 569–596.

Cohen, Wesley, Richard Nelson, and John Walsh. 2000. "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)." National Bureau of Economic Research w7552, Cambridge, MA.

Damodaran, Aswath. 2025. "Cost of Equity and Capital (US)."

De Ridder, Maarten. 2024. "Market Power and Innovation in the Intangible Economy." *The American economic review*, 114(1): 199–251.

Dix, Rebekah, and Todd Lensman. 2024. "Combining Complements: Theory and Evidence from Cancer Treatment Innovation." Yale University, Department of Economics Job Market Paper.

Dwivedi, Sandhya. 2016. "Global Allergy Research during 1994-2013: A Scientometric Study."

Ertl, Peter. 2017. "An Algorithm to Identify Functional Groups in Organic Molecules." *Journal of cheminformatics*, 9(1): 36–7.

Evenson, Robert E. 1982. "International Invention: Implications for Technology Market Analysis."

Fein, Michael N., David A. Fischer, Andrew W. O'Keefe, and Gord L. Sussman. 2019. "CSACI Position Statement: Newer Generation H1-antihistamines Are Safer than First-Generation H1-antihistamines and Should Be the First-Line Antihistamines for the Treatment of Allergic Rhinitis and Urticaria." *Allergy, Asthma, and Clinical Immunology : Official Journal of the Canadian Society of Allergy and Clinical Immunology*, 15: 61.

Fort, Teresa C, Nathan Goldschlag, Jack Liang, Peter K Schott, and Nikolas Zolas. 2025. "Growth Is Getting Harder to Find, Not Ideas."

Frankel, Alexander P., Joshua L. Krieger, Danielle Li, and Dimitris Papanikolaou. 2023. "Evaluation and Learning in R&D Investment." National Bureau of Economic Research NBER Working Paper 31290, Cambridge, MA.

Free, Spencer M., and James W. Wilson. 1964. "A Mathematical Contribution to Structure-Activity Studies." *Journal of Medicinal Chemistry*, 7(4): 395–399.

Gallop, Mark A., Ronald W. Barrett, William J. Dower, Stephen P. A. Fodor, and Eric M. Gordon. 1994. "Applications of Combinatorial Technologies to Drug Discovery. 1. Background and Peptide Combinatorial Libraries." *Journal of Medicinal Chemistry*, 37(9): 1233–1251.

Griffith, Rachel, Stephen Redding, and John Van Reenen. 2003. "R&D and Absorptive Capacity: Theory and Empirical Evidence*." *The Scandinavian Journal of Economics*, 105(1): 99–118.

Grossman, Gene M., and Elhanan Helpman. 1991. "Quality Ladders in the Theory of Growth." *The Review of Economic Studies*, 58(1): 43–61.

Hall, Richard, and Guillaume Godin. 2017. "IFG: Identify Functional Groups Implementation in Rdkit."

Hansch, Corwin., and Toshio. Fujita. 1964. "P-σ-π Analysis. A Method for the Correlation of Biological Activity and Chemical Structure." *Journal of the American Chemical Society*, 86(8): 1616–1626.

Hansch, Corwin, Peyton P. Maloney, Toshio Fujita, and Robert M. Muir. 1962. "Correlation of Biological Activity of Phenoxyacetic Acids with Hammett Substituent Constants and Partition Coefficients." *Nature (London)*, 194(4824): 178–180.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?" *The Review of Economic Studies*, 76(1): 283–317.

Jones, Charles I. 2023. "Recipes and Economic Growth: A Combinatorial March down an Exponential Tail." *The Journal of political economy*, 131(8): 1994–2031.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring Technological Innovation over the Long Run." *The American economic review. Insights*, 3(3): 303–320.

Kinch, Michael S., and Janie Merkel. 2015. "An Analysis of FDA-approved Drugs for Inflammation and Autoimmune Diseases." *Drug Discovery Today*, 20(8): 920–923.

Klette, Tor Jakob, and Samuel Kortum. 2002. "Innovating Firms and Aggregate Innovation."

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *The Journal of political economy*, 112(5): 986–1018.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth*." *The Quarterly Journal of Economics*, 132(2): 665–712.

König, Michael, Jan Lorenz, and Fabrizio Zilibotti. 2016. "Innovation vs. Imitation and the Evolution of Productivity Distributions." *Theoretical economics*, 11(3): 1053–1102.

König, Michael, Kjetil Storesletten, Zheng Song, and Fabrizio Zilibotti. 2022. "From Imitation to Innovation: Where Is All That Chinese R&D Going?" *Econometrica : journal of the Econometric Society*, 90(4): 1615–1654.

Kortum, Samuel. 1993. "Equilibrium R&D and the Patent–R&D Ratio: U.S. Evidence." *The American Economic Review*, 83(2): 450–457.

Kortum, Samuel. 1997. "Research, Patenting, and Technological Change - The Econometric Society."

Krieger, Joshua, Danielle Li, and Dimitris Papanikolaou. 2022. "Missing Novelty in Drug Development." *The Review of financial studies*, 35(2): 636–679.

Lamb, David. 1984. *Multiple Discovery: The Pattern of Scientific Progress.* [Avebury, Buckinghamshire]:Avebury.

Ledley, Fred D., Sarah Shonka McCoy, Gregory Vaughan, and Ekaterina Galkina Cleary. 2020. "Profitability of Large Pharmaceutical Companies Compared With Other Large Public Companies." *JAMA : the journal of the American Medical Association*, 323(9): 834–843.

Lentz, Rasmus, and Dale T. Mortensen. 2008. "An Empirical Model of Growth through Product Innovation." *Econometrica : journal of the Econometric Society*, 76(6): 1317–1373.

Macarron, Ricardo, Martyn N. Banks, Dejan Bojanic, David J. Burns, Dragan A. Cirovic, Tina Garyantes, Darren V. S. Green, Robert P. Hertzberg, William P. Janzen, Jeff W. Paslay, Ulrich Schopfer, and G. Sitta Sittampalam. 2011. "Impact of High-Throughput Screening in Biomedical Research." *Nature Reviews Drug Discovery*, 10(3): 188–195.

Merton, Robert K. 1961. "Singletons and multiples in scientific discovery: A chapter in the sociology of science." *Proceedings of the American Philosophical Society*, 105(5): 470–486.

Merton, Robert K. 1963. "Resistance to the systematic study of multiple discoveries in science." *European Journal of Sociology/Archives Européennes de Sociologie*, 4(2): 237–282.

Nikolova, Nina, and Joanna Jaworska. 2004. "Approaches to Measure Chemical Similarity – a Review." *QSAR & Combinatorial Science*, 22(9-10): 1006–1026.

Office, Congressional Budget. 2021. "Research and Development in the Pharmaceutical Industry | Congressional Budget Office."

Ogburn, William F, and Dorothy Thomas. 1922. "Are inventions inevitable? A note on social evolution." *Political science quarterly*, 37(1): 83–98.

Peters, Michael. 2020. "Heterogeneous Markups, Growth, and Endogenous Misallocation." *Econometrica : journal of the Econometric Society*, 88(5): 2037–2073.

Ribeiro, Bernardo. 2025. "Growth with New and Old Technologies." Princeton University Working Paper.

Romer, Paul. 1992. "Two Strategies for Economic Development: Using Ideas and Producing Ideas." *The World Bank economic review*, 63–91.

Schumpeter, Joseph A. 1934. *The Theory of Economic Development. Routledge Classics*, Abingdon, Oxon; New York, NY:Routledge.

Sertkaya, Aylin, Trinidad Beleche, Amber Jessup, and Benjamin D. Sommers. 2024. "Costs of Drug Development and Research and Development Intensity in the US, 2000-2018." *JAMA Network Open*, 7(6): e2415445.

Sternitzke, Christian. 2013. "An Exploratory Analysis of Patent Fencing in Pharmaceuticals: The Case of PDE5 Inhibitors." *Research Policy*, 42(2): 542–551.

Taylor, David. 2015. "The Pharmaceutical Industry and the Future of Drug Development." In *Pharmaceuticals in the Environment*. The Royal Society of Chemistry.

Teng, Songyuan. 2025. "Learn Complementarity via Experiment: A Bayesian Theory of Combinatorial Innovation."

Tranchero, Matteo. 2024. "Finding Diamonds in the Rough: Data-Driven Opportunities and Pharmaceutical Innovation." *Academy of Management Proceedings*, 2024(1): 13751.

U. S. Government Accountability Office, NA. 2017. "Drug Industry: Profits, Research and Development Spending, and Merger and Acquisition Deals | U.S. GAO."

Uzzi, Brian, Satyam Mukherjee, Michael Stringer, and Ben Jones. 2013. "Atypical Combinations and Scientific Impact." *Science (American Association for the Advancement of Science)*, 342(6157): 468–472.

Wagner, Stefan, Christian Sternitzke, and Sascha Walter. 2022. "Mapping Markush." *Research Policy*, 51(10): 104597.

Weitzman, Martin L. 1998. "Recombinant Growth." *The Quarterly journal of economics*, 113(2): 331–360.

Youn, Hyejin, Deborah Strumsky, Luis M. A. Bettencourt, and José Lobo. 2015. "Invention as a Combinatorial Process: Evidence from US Patents." *Journal of the Royal Society interface*, 12(106): 20150272–.

Appendix

A Identification of Functional Groups

I identify functional groups (FGs) using the algorithm of Ertl (2017), which proceeds by (i) marking atoms that constitute chemically functional moieties, (ii) merging connected marked atoms into single FG, and (iii) recording the immediate carbon environment of each FG (aliphatic versus aromatic). I follow a Python implementation via the RDKit cheminformatics toolkit (Hall and Godin, 2017) that does not apply the optional generalization step of Ertl (2017).

A.1 Algorithm (Ertl, 2017)

(i) Mark all heteroatoms, i.e., any atom other than carbon and hydrogen.

(ii) Mark carbon atoms that are part of functional connectivity:

  (a) any carbon atom connected to any heteroatom by non-aromatic double or triple bond (C=X or C≡X, with X being heteroatom),

  (b) any carbon in a non-aromatic double or triple bond (C=C or C≡C),

  (c) acetal carbons: sp^3 carbons bonded to $\geq$ 2 atoms from {O, N, S} where those heteroatoms only have single bonds,

  (d) any carbon atom as part of oxirane, aziridine and thiirane rings; these rings have high reactivity and are thus traditionally considered to be functional groups.

(iii) Merge connected marked atoms into a single FG.

(iv) For each FG, attach any directly bonded unmarked carbon atoms as its environment; these carbons are not part of the FG itself.

B Model Extension: Link Creative Destruction to Growth

This appendix embeds a minimal macro block (consumer, final-good aggregator, product market Bertrand competition) into the baseline model in Section 4. Doing so makes the rate of creative destruction implied by firms' innovation decisions map directly into aggregate growth. The extended model yields a balanced growth path (BGP) in which firm innovation policy functions are constant over time while firm value, aggregate product quality, and wage rate grow over time at a common exponential rate.

B.1 Environment and Demand

Time and varieties. Continuous time $t \geq 0$. A unit mass of differentiated product lines indexed by $i \in [0, 1]$.

Household. A representative household has preferences

$$U = \int_0^\infty e^{-\rho t} \ln C_t \, dt, \qquad \rho > 0, \tag{B.1}$$

owns all firms, supplies one unit of labor inelastically, and trades a risk-free bond at rate r_t. The Euler equation is

$$\dot{C}_t / C_t = r_t - \rho. \tag{B.2}$$

Final-good aggregator. The final good aggregates line outputs $\{y_{i,t}\}$ via a log aggregator:

$$\ln Y_t = \int_0^1 \ln y_{i,t} \, di, \tag{B.3}$$

with the final good as numeraire (price normalized to 1). Standard demand under (B.3) equalizes revenue across lines:

$$p_{i,t} \, y_{i,t} = Y_t \quad \text{for all } i. \tag{B.4}$$

B.2 Production, Pricing, and Wages

Technology in each line. The leader in line i produces with linear technology

$$y_{i,t} = q_{i,t} \, \ell_{i,t}, \tag{B.5}$$

where $q_{i,t}$ is the line-specific leader quality and $\ell_{i,t}$ is labor hired in that line.

Quality ladder and Bertrand pricing. Competition is Bertrand on a quality ladder with common step size $\gamma > 1$.[35] The leader's quality is γ times the follower's, implying a constant markup γ. With wage w_t, the leader's price is

$$p_{i,t} = \frac{w_t}{q_{i,t}} \cdot \gamma. \tag{B.6}$$

[35] It is straightforward to allow heterogeneous step sizes, e.g., novel innovations with step size γ_b and combinatorial γ_c. This extension does not change the key implication: the growth rate equals the creative-destruction rate times a constant, namely a weighted average of (γ_b, γ_c). For simplicity, I abstract from this heterogeneity.

Per-line allocations and profits. Using (B.4) – (B.6),

$$y_{i,t} = \frac{Y_t}{p_{i,t}} = \frac{Y_t\, q_{i,t}}{\gamma\, w_t}, \quad \ell_{i,t} = \frac{y_{i,t}}{q_{i,t}} = \frac{Y_t}{\gamma\, w_t}. \tag{B.7}$$

Per-line operating profit is independent of i:

$$\pi_{i,t} = \left(p_{i,t} - w_t/q_{i,t}\right) y_{i,t} = \left(1 - \frac{1}{\gamma}\right) Y_t. \tag{B.8}$$

Aggregate quality and wage–quality link. Define the geometric-mean quality

$$\ln Q_t \equiv \int_0^1 \ln q_{i,t}\, di. \tag{B.9}$$

From (B.3) and (B.7),

$$\ln Y_t = \int_0^1 \ln q_{i,t}\, di + \ln\left(\frac{Y_t}{\gamma w_t}\right) = \ln Q_t + \ln\left(\frac{Y_t}{\gamma w_t}\right),$$

which implies the wage–quality link

$$w_t = \frac{Q_t}{\gamma}. \tag{B.10}$$

Production labor and output. Total production labor is

$$L_t^P = \int_0^1 \ell_{i,t}\, di = \frac{Y_t}{\gamma\, w_t}, \tag{B.11}$$

hence

$$Y_t = \gamma\, w_t\, L_t^P = Q_t\, L_t^P. \tag{B.12}$$

Therefore, the output growth rate decomposes as

$$g_{Y,t} \equiv \frac{\dot{Y}_t}{Y_t} = g_{Q,t} + g_{L^P,t}. \tag{B.13}$$

B.3 Innovation and Creative Destruction

Step size and improvements. An improvement event in line i replaces the incumbent leader and multiplies quality by γ:

$$q'_{i,t} = \gamma\, q_{i,t}, \qquad \gamma > 1. \tag{B.14}$$

Firm states and profits. A firm's state is (i, n) where i is knowledge/ingredient stock and n is the number of product lines it currently leads. Using (B.8) and (B.12), per-line operating

profit is

$$\pi_t \equiv \left(1 - \frac{1}{\gamma}\right) Y_t = \left(1 - \frac{1}{\gamma}\right) Q_t \, L_t^P. \tag{B.15}$$

A firm with n products earns $\pi_t \, n$.

R&D technologies and research labor. Innovation uses labor as the sole input. An incumbent in state (i, n) chooses Poisson arrival rate λ_b for novel innovation and λ_c for combinatorial innovation by hiring research workers:

$$\text{Novel:} \quad \ell_b(i, \lambda_b) \equiv \frac{A_b}{\phi_b} \lambda_b^{\phi_b} \, i^{-\zeta_b}, \qquad \phi_b > 1, \ \zeta_b > 0, \tag{B.16}$$

$$\text{Combinatorial:} \quad \ell_c(i, \lambda_c) \equiv \frac{A_c}{\phi_c} \lambda_c^{\phi_c} \, i^{-\zeta_c}, \qquad \phi_c > 1, \ \zeta_c > 0. \tag{B.17}$$

The corresponding flow costs are $w_t \, \ell_b(i, \lambda_b)$ and $w_t \, \ell_c(i, \lambda_c)$, respectively. A combinatorial success adds a product, $(i, n) \to (i, n+1)$. A novel success raises knowledge and adds a product, $(i, n) \to (i+1, n+1)$.

There is a unit continuum of potential entrants that undertake novel innovation with arrival rate η_t, hiring

$$\ell_e(\eta_t) \equiv \frac{A_e}{\phi_b} \eta_t^{\phi_b}, \tag{B.18}$$

at flow cost $w_t \, \ell_e(\eta_t)$. Aggregate research labor L_t^R sums innovation labor hired by incumbents and entrants, integrating over firm states.

Knowledge depreciation and product creative destruction. Let ξ denote the exogenous per-ingredient depreciation rate and μ denote the endogenous per-line destruction rate implied by aggregate innovation.

B.4 Firm problem

Incumbent HJB. Let $V_t(i, n)$ denote the value of an incumbent with knowledge $i \in \mathbb{N}_0$ and product count $n \in \mathbb{N}_0$. Given the flow profit π_t per product in (B.15), the R&D costs in (**??**)–(**??**), knowledge depreciation at rate ξ per ingredient, and product creative destruction at rate μ per

product, the HJB is

$$
\begin{aligned}
r_t\, V_t(i,n) \;=\; \max_{\lambda_b,\lambda_c} \Bigg\{ & \underbrace{\pi_t\, n}_{\text{flow profits}} \;+\; \underbrace{\dot{V}_t(i,n)}_{\text{capital gain}} \;-\; \underbrace{w_t\left[\frac{A_b}{\phi_b}\,\lambda_b^{\phi_b} i^{-\zeta_b} + \frac{A_c}{\phi_c}\,\lambda_c^{\phi_c} i^{-\zeta_c}\right]}_{\text{R\&D costs}} \\
& + \underbrace{\lambda_b\Big[V_t(i+1,n+1) - V_t(i,n)\Big]}_{\text{novel innovation}} + \underbrace{\lambda_c\Big[V_t(i,n+1) - V_t(i,n)\Big]}_{\text{combinatorial innovation}} \\
& - \underbrace{\xi i\,\Big[V_t(i,n) - V_t(i-1,n)\Big]}_{\text{ingredient depreciation}} \;-\; \underbrace{\mu_t n\,\Big[V_t(i,n) - V_t(i,n-1)\Big]}_{\text{creative destruction}} \Bigg\}.
\end{aligned}
\tag{B.19}
$$

The first line on the right-hand side collects the firm's instantaneous flow payoff net of innovation costs: each product yields flow profit π_t; the term $\dot{V}_t(i,n)$ is the capital gain from moving through time; and the innovation costs depend on firm's knowledge i, the chosen arrival rates (λ_b, λ_c), and wage w_t. The second and third lines capture the expected value changes from the firm's Poisson events. The second line reflects the expected gains from successful innovation: novel innovation, at rate λ_b, expands both the ingredient set and the product portfolio, and combinatorial innovation, at rate λ_c, adds a product using existing ingredients. The third line captures expected losses in firm value. Ingredients depreciate at a constant Poisson rate ξ, reducing the firm's capacity for future innovation, while creative destruction occurs at Poisson rate μ_t per product, displacing one product from the firm's portfolio.

Entrant problem. Entrants choose novel R&D intensity η_t and solve

$$
\max_{\eta} \; \underbrace{-\frac{A_e}{\phi_b}\eta^{\phi_b} w_t}_{\text{innovation costs}} \; + \underbrace{\eta\,[V_t(1,1) - V_t(0,0)]}_{\text{expected gains from entry}}. \tag{B.20}
$$

The FOC (if $\eta_t > 0$) is

$$
\eta_t = \left(\frac{V_t(1,1) - V_t(0,0)}{A_e w_t}\right)^{\frac{1}{\phi_b - 1}} \tag{B.21}
$$

B.5 Balanced Growth Path

I now restrict attention to the balanced-growth paths (BGP) of this economy, defined as paths along which all level variables grow at constant rate and firm distribution is stationary.

Balanced growth path definition A balanced growth path (BGP) consists of allocations of aggregate variables $(Y_t, Q_t, L_t^P, L_t^R, \mu_t, r_t, w_t)$, household choices $(C_t, y_{i,t})$, firm choices $(\lambda_{b,t}, \lambda_{c,t}, \eta_t)$, and cross-sectional firm distribution $\Phi_t(i,n)$, such that: (i) aggregate variables grow at a constant rate, (ii) consumers choose optimally to maximize utility, (iii) firms choose optimally to maximize profits, (iv) all markets clear, and (v) the cross-sectional distribution of firms is stationary.

I now characterize the BGP. Given the Euler equation (B.2) and constant growth of C_t, the interest rate r_t must be constant on a BGP. Because production labor L_t^P and research labor L_t^R add up to the unit aggregate labor endowment, it must be the case that (L_t^P, L_t^R) are constant, i.e., (L^P, L^R). I conjecture that μ_t is constant over time ($\mu_t = \mu, \forall t$) and verify this guess.

Incumbent value and policy functions. On a BGP, conjecture that firm value is linear in the product count and scales with aggregate quality,

$$V_t(i,n) = Q_t\Big[a(i) + b\,n\Big], \qquad \dot{V}_t(i,n) = g_Q\, V_t(i,n), \tag{B.22}$$

with $a(i)$ and b time invariant. Substituting (B.22) into the HJB (B.19), using

$$V_t(i,n{+}1){-}V_t(i,n) = Q_t b, \quad V_t(i{+}1,n{+}1){-}V_t(i,n) = Q_t\Big[a(i{+}1){-}a(i) + b\Big],$$

$$V_t(i,n){-}V_t(i{-}1,n) = Q_t\Big[a(i){-}a(i{-}1)\Big], \quad V_t(i,n){-}V_t(i,n{-}1) = Q_t b,$$

together with $\pi_t = \bar{\pi}\, Q_t$ where $\bar{\pi} \equiv (1 - 1/\gamma)L^P$ and $w_t = Q_t/\gamma$, and then dividing both sides by Q_t, yields the objective inside the maximization in (B.19) as

$$\begin{aligned}\mathcal{J}(i,n;\lambda_b,\lambda_c) = \bar{\pi}\, n + g_Q\left[a(i) + bn\right] - \frac{1}{\gamma}\left[\frac{A_b}{\phi_b}\lambda_b^{\phi_b} i^{-\zeta_b} + \frac{A_c}{\phi_c}\lambda_c^{\phi_c} i^{-\zeta_c}\right] \\ + \lambda_b \underbrace{\Big\{a(i{+}1) - a(i) + b\Big\}}_{\equiv\, \Delta_b(i)} + \lambda_c \underbrace{b}_{\equiv\, \Delta_c} \\ - \xi i \Big\{a(i) - a(i{-}1)\Big\} - \mu n\, b.\end{aligned} \tag{B.23}$$

Maximization over λ_b, λ_c. Taking first-order conditions of $\mathcal{J}$ with respect to λ_b and λ_c (and using $\phi_b, \phi_c > 1$ for interior optima) gives

$$0 = \frac{\partial \mathcal{J}}{\partial \lambda_b} = -\frac{1}{\gamma} A_b\, \lambda_b^{\phi_b - 1} i^{-\zeta_b} + \Delta_b(i) \quad \Longrightarrow \quad \lambda_b(i,n) = \left[\frac{\gamma\, \Delta_b(i)}{A_b}\, i^{\zeta_b}\right]^{\frac{1}{\phi_b - 1}}, \tag{B.24}$$

$$0 = \frac{\partial \mathcal{J}}{\partial \lambda_c} = -\frac{1}{\gamma} A_c\, \lambda_c^{\phi_c - 1} i^{-\zeta_c} + \Delta_c \quad \Longrightarrow \quad \lambda_c(i,n) = \left[\frac{\gamma\, \Delta_c}{A_c}\, i^{\zeta_c}\right]^{\frac{1}{\phi_c - 1}}. \tag{B.25}$$

Crucially, the marginal benefits $\Delta_b(i) = a(i{+}1) - a(i) + b$ and $\Delta_c = b$ do *not* depend on n, and the cost terms depend on $(i, \lambda_\cdot)$ but not on n. Hence the optimal policies obtained from (B.24)–(B.25)

are functions of i only:

$$\lambda_b(i) = \left[\frac{\gamma\left(a(i{+}1) - a(i) + b\right)}{A_b}\, i^{\zeta_b}\right]^{\frac{1}{\phi_b - 1}}, \qquad \lambda_c(i) = \left[\frac{\gamma\, b}{A_c}\, i^{\zeta_c}\right]^{\frac{1}{\phi_c - 1}}. \tag{B.26}$$

Substituting (B.26) back into the HJB and collecting terms yields the *stationary* Bellman equation in efficiency units:

$$\begin{aligned} r\left[a(i) + bn\right] = \bar{\pi}\, n + g_Q\left[a(i) + bn\right] - \frac{1}{\gamma}\left[\frac{A_b}{\phi_b}\lambda_b(i)^{\phi_b} i^{-\zeta_b} + \frac{A_c}{\phi_c}\lambda_c(i)^{\phi_c} i^{-\zeta_c}\right] \\ + \lambda_b(i)\Big\{a(i{+}1) - a(i) + b\Big\} + \lambda_c(i)\, b - \xi i\left\{a(i) - a(i{-}1)\right\} - \mu n\, b. \end{aligned} \tag{B.27}$$

Matching coefficients on n in (B.27) and using $r = \rho + g_Q$ (Euler equation) gives

$$(\rho + \mu)\, b = \bar{\pi} \qquad \Longrightarrow \qquad b = \frac{\bar{\pi}}{\rho + \mu}. \tag{B.28}$$

The i–block of (B.27) then delivers the one–dimensional recursion for $a(i)$:

$$\rho\, a(i) = -\frac{1}{\gamma}\left[\frac{A_b}{\phi_b}\lambda_b(i)^{\phi_b} i^{-\zeta_b} + \frac{A_c}{\phi_c}\lambda_c(i)^{\phi_c} i^{-\zeta_c}\right] + \lambda_b(i)\Big\{a(i{+}1) - a(i) + b\Big\} + \lambda_c(i)\, b - \xi i\left\{a(i) - a(i{-}1)\right\}. \tag{B.29}$$

Equations (B.26)–(B.29) verify the conjecture (B.22): the FOCs first deliver optimal λ_b, λ_c; these depend on i only (not on n or t); b is constant by (B.28); and $a(i)$ solves the stationary recursion (B.29). Labor market clearing condition thus reads

$$L^P + L^R \;=\; L^P + \frac{A_e}{\phi_b}\eta^{\phi_b} \;+\; \sum_{i,n}\left[\frac{A_b}{\phi_b}\lambda_b(i)^{\phi_b} i^{-\zeta_b} + \frac{A_c}{\phi_c}\lambda_c(i)^{\phi_c} i^{-\zeta_c}\right]\Phi(i,n) = 1,$$

where L^P is given by equation (B.11), and $\Phi(i,n)$ denotes the stationary firm distribution characterized by equation (B.33).

Entry rate and creative destruction. Let $\widetilde{b}^E \equiv a(1) - a(0) + b$ denote the efficiency–units value of creating a new firm with state $(1,1)$ relative to $(0,0)$. Using the first–order condition (B.21) and $V_t(1,1) - V_t(0,0) = Q_t\, \widetilde{b}^E$ with $w_t = Q_t/\gamma$,

$$\eta_t = \left(\frac{Q_t\, \widetilde{b}^E}{A_e\, w_t}\right)^{\frac{1}{\phi_b - 1}} = \left(\frac{\gamma\, \widetilde{b}^E}{A_e}\right)^{\frac{1}{\phi_b - 1}} \equiv \eta, \tag{B.30}$$

which is time invariant on a BGP. The aggregate improvement (product creative destruction) intensity per line is

$$\Lambda \;=\; \eta \;+\; \sum_{i,n} \Big[\lambda_b(i) + \lambda_c(i)\Big]\Phi(i,n), \tag{B.31}$$

and, by symmetry, each existing product faces displacement at the per–product hazard

$$\mu_t \equiv \Lambda \qquad \text{(constant over time)}, \tag{B.32}$$

verifying the guess of constant creative destruction along BGP.

Equilibrium firm distribution. Setting the Kolmogorov forward equation to zero gives, for all $i, n \in \mathbb{N}_0$,

$$\begin{aligned} 0 = {} & \underbrace{\lambda_b(i{-}1)\,\Phi(i{-}1,n{-}1)\,\mathbf{1}\{i{\geq}1,\; n{\geq}1\}}_{\text{novel inflow}} \;+\; \underbrace{\lambda_c(i)\,\Phi(i,n{-}1)\,\mathbf{1}\{n{\geq}1\}}_{\text{combination inflow}} \\ & + \quad \underbrace{\xi(i{+}1)\,\Phi(i{+}1,n)}_{\text{ingredient depreciation inflow}} \quad + \quad \underbrace{\mu(n{+}1)\,\Phi(i,n{+}1)}_{\text{creative destruction inflow}} \;+\; \underbrace{\eta\,\mathbf{1}\{i{=}1,\; n{=}1\}}_{\text{entry}} \\ & - \quad \underbrace{\Big[\lambda_b(i) + \lambda_c(i) + \xi i + \mu n\Big]\,\Phi(i,n)}_{\text{outflows}}. \end{aligned} \tag{B.33}$$

Intuitively, mass reaches state (i, n) via: (i) a novel success from $(i-1, n-1)$ at rate $\lambda_b(i-1)$; (ii) a combinatorial success from $(i, n-1)$ at rate $\lambda_c(i)$; (iii) ingredient depreciation from $(i+1, n)$ at rate $\xi(i+1)$; (iv) creative destruction from $(i, n+1)$ at rate $\mu(n+1)$; and (v) entry places mass at $(1, 1)$ at rate η. Symmetrically, firms exit (i, n) due to their own successes, depreciation, and product destruction.

Growth rate. Each improvement multiplies the upgraded line's quality by γ, so over $[t, t{+}dt]$ a fraction $\Lambda\, dt$ of lines experience a jump of size $\ln\gamma$ in log quality. Therefore

$$\ln Q_{t+dt} = \int_0^1 \ln q_{i,t+dt}\, di = \int_0^1 \ln q_{i,t}\, di \;+\; \Lambda\, dt\, \ln\gamma = \ln Q_t \;+\; \Lambda\, dt\, \ln\gamma. \tag{B.34}$$

Divide by dt and let $dt \to 0$:

$$\frac{\dot{Q}_t}{Q_t} \;=\; \Lambda\, \ln\gamma \;\equiv\; g_{Q,t}. \tag{B.35}$$

On a BGP, $g_Q = \Lambda \ln\gamma$ and, with $Y_t = Q_t L^P$ and $C_t = Y_t$ on the BGP, it follows that $g_Y = g_C = g_Q$, and the Euler equation pins down the interest rate $r = \rho + g_Q$.

By embedding a minimal macro block into the baseline model, I show that the rate of creative destruction implied by firms' innovation decisions maps directly into aggregate growth, scaled by

the innovation step size γ.

This extension also introduces a standard "standing on giants' shoulders" force: the private value of an innovation does not internalize the option value it creates for future innovators, so equilibrium innovation intensity can be inefficiently low relative to the social optimum (Klette and Kortum, 2002). This motivates the policy study in Section 6.

C Proofs on Firm Dynamic Problem

C.1 Proof of proposition 1

Proof of Proposition 1. Conjecture $V(i,n) = a(i) + b\,n$ with $b > 0$ and plug this guess into maintext equation (4.4). Under this conjecture,

$$\begin{aligned} V(i{+}1,n{+}1) - V(i,n) &= a(i{+}1) - a(i) + b, \\ V(i,n{+}1) - V(i,n) &= b, \\ V(i,n) - V(i{-}1,n) &= a(i) - a(i{-}1). \end{aligned}$$

FOCs for λ_b, λ_c. Taking derivatives of the right-hand side of maintext equation (4.4) with respect to (λ_b, λ_c) (using the R&D cost terms $\frac{A_b}{\phi_b} i^{-\zeta_b} \lambda_b^{\phi_b}$ and $\frac{A_c}{\phi_c} i^{-\zeta_c} \lambda_c^{\phi_c}$ with $\phi_b, \phi_c > 1$), we obtain

$$\begin{aligned} \Big[V(i{+}1,n{+}1) - V(i,n)\Big] - A_b i^{-\zeta_b} \lambda_b^{\phi_b - 1} &= 0, \\ \Big[V(i,n{+}1) - V(i,n)\Big] - A_c i^{-\zeta_c} \lambda_c^{\phi_c - 1} &= 0, \end{aligned}$$

hence, using the differences above,

$$\lambda_b(i) = \left(\frac{a(i{+}1) - a(i) + b}{A_b\, i^{-\zeta_b}} \right)^{\frac{1}{\phi_b - 1}}, \qquad \lambda_c(i) \;=\; \left(\frac{b}{A_c\, i^{-\zeta_c}} \right)^{\frac{1}{\phi_c - 1}}. \tag{C.1}$$

Pinning down b. Plug $V(i,n) = a(i) + bn$ into maintext equation (4.4) and collect the terms that are proportional to n on both sides. Using $V(i,n) - V(i,n{-}1) = b$, the creative destruction term contributes $-\mu n\, b$, while the flow-profit term is $\bar{\pi} n$. Equating coefficients of n yields

$$r\, b \;=\; \bar{\pi} - \mu\, b \quad \Rightarrow \quad b = \frac{\bar{\pi}}{r + \mu}.$$

Reduced Bellman equation for $a(i)$. Subtracting the linear-in-n component from maintext

equation (4.4), the value equation collapses to a one-dimensional recursion in i:

$$r\,a(i) = \max_{\lambda_b,\lambda_c}\Big\{ -\frac{A_b}{\phi_b} i^{-\zeta_b}\lambda_b^{\phi_b} - \frac{A_c}{\phi_c} i^{-\zeta_c}\lambda_c^{\phi_c} + \lambda_b\Big[a(i{+}1) - a(i) + b\Big] + \lambda_c\, b - \xi\, i\Big[a(i) - a(i{-}1)\Big]\Big\}. \tag{C.2}$$

Substitute the optimal policies (C.1) back into (C.2). Using the generic maximization $\max_{\lambda\geq 0}\{\lambda\Delta - \frac{A}{\phi} i^{-\zeta}\lambda^{\phi}\}$ with $\phi > 1$ gives the value $\frac{\phi-1}{\phi}\,(A\, i^{-\zeta})^{-1/(\phi-1)}\Delta^{\phi/(\phi-1)}$. Therefore,

$$r\,a(i) = \frac{\phi_b - 1}{\phi_b}\Big(A_b\, i^{-\zeta_b}\Big)^{-\frac{1}{\phi_b - 1}}\Big[a(i{+}1) - a(i) + b\Big]^{\frac{\phi_b}{\phi_b - 1}} + \frac{\phi_c - 1}{\phi_c}\Big(A_c\, i^{-\zeta_c}\Big)^{-\frac{1}{\phi_c - 1}} b^{\frac{\phi_c}{\phi_c - 1}} - \xi\, i\Big[a(i) - a(i{-}1)\Big]. \tag{C.3}$$

Conclusion. Equations (C.1)–(C.3), together with $b = \bar{\pi}/(r+\mu)$, verify that the conjectured form $V(i,n) = a(i) + bn$ is consistent with maintext equation (4.4), pin down the one-dimensional recursion for $a(i)$, and deliver the policy rules $\lambda_b(i), \lambda_c(i)$ that depend only on i. This proves the proposition. □

C.2 Proof of proposition 2

Proof of Proposition 2. Focus on a stationary equilibrium in which the cross-sectional distribution is time-invariant:

$$\dot{\Phi}_t(i) = 0 \quad \text{for all } i \geq 0,$$

and drop the time subscript to write $\Phi(i)$. Under stationarity, maintext equation (4.10) implies, for each interior state $i \geq 2$,

$$0 \;=\; \lambda_b(i{-}1)\,\Phi(i{-}1) \;+\; \xi\,(i{+}1)\,\Phi(i{+}1) \;-\; \Big[\lambda_b(i) + \xi\, i\Big]\Phi(i), \tag{C.4}$$

and, at the boundary $i = 1$ (where entry occurs at rate η),

$$0 \;=\; \eta \;+\; 2\xi\,\Phi(2) \;-\; \Big[\lambda_b(1) + \xi\Big]\Phi(1). \tag{C.5}$$

Link between i and $i{+}1$. Define the net probability current crossing the link $(i, i{+}1)$ (positive when mass flows from i to $i{+}1$) by

$$J(i) \;\equiv\; \lambda_b(i)\,\Phi(i) \;-\; \xi\,(i{+}1)\,\Phi(i{+}1), \qquad i \geq 0. \tag{C.6}$$

Using (C.6), the right-hand side of (C.4) can be written as $\lambda_b(i{-}1)\Phi(i{-}1) + \xi(i{+}1)\Phi(i{+}1) -$

$[\lambda_b(i) + \xi i]\Phi(i) = J(i-1) - J(i)$. Thus (C.4) is equivalent to

$$0 \;=\; J(i-1) - J(i), \qquad i \geq 2, \tag{C.7}$$

so $J(i)$ is constant in i on $\{1, 2, \dots\}$:

$$J(i) = J \quad \text{for all } i \geq 1. \tag{C.8}$$

Base case $i = 1$ (entry). Rewrite (C.5) using (C.6) at $i = 1$:

$$0 \;=\; \eta \;-\; \Big[\lambda_b(1)\Phi(1) - 2\xi\,\Phi(2)\Big] \;-\; \xi\,\Phi(1) \;=\; \eta \;-\; J(1) \;-\; \xi\,\Phi(1).$$

Using (C.8) we obtain

$$J \;=\; \eta \;-\; \xi\,\Phi(1). \tag{C.9}$$

Stationarity with finite mass and no sink at large i requires that there be no net drift of probability mass to infinity, hence the constant current must be zero: $J = 0$. Equation (C.9) then pins down the boundary mass

$$\Phi(1) = \frac{\eta}{\xi}. \tag{C.10}$$

Interior states $i > 1$. With $J = 0$, the definition (C.6) yields the pairwise balance across every adjacent link,

$$\lambda_b(i)\,\Phi(i) \;=\; \xi\,(i{+}1)\,\Phi(i{+}1), \qquad i \geq 1, \tag{C.11}$$

which is equivalent to the one-step forward recursion

$$\Phi(i{+}1) \;=\; \frac{\lambda_b(i)}{\xi\,(i{+}1)}\,\Phi(i), \qquad i \geq 1. \tag{C.12}$$

Conclusion. Setting $\dot{\Phi}_t(i) = 0$ in maintext equation (4.10) implies that the adjacent-state current $J(i)$ is constant across links; the boundary equation at $i = 1$ fixes $J = \eta - \xi\Phi(1)$, and the no-drift condition yields $J = 0$, which pins down $\Phi(1) = \eta/\xi$. The pairwise balance (C.11) then gives the recursion (C.12) for all $i > 1$, establishing the proposition. □

C.3 Proof of proposition 3

Proof of Proposition 3. Conjecture that the value function is linear in the number of products at each time t:

$$V_t(i, n) \;=\; a_t(i) \;+\; b_t\,n, \qquad b_t > 0,$$

and plug this guess into maintext equation (6.2). Under the conjecture,

$$\begin{aligned} V_t(i{+}1, n{+}1) - V_t(i,n) &= a_t(i{+}1) - a_t(i) + b_t, \\ V_t(i, n{+}1) - V_t(i,n) &= b_t, \\ V_t(i,n) - V_t(i{-}1,n) &= a_t(i) - a_t(i{-}1), \\ \dot{V}_t(i,n) &= \dot{a}_t(i) + \dot{b}_t\, n. \end{aligned}$$

FOCs for λ_b, λ_c. Differentiating the right-hand side with respect to (λ_b, λ_c) and using the cost exponents $\phi_b, \phi_c > 1$,

$$\begin{aligned} \Big[k_t V_t(i+1, n+1) + (1-k_t) V_t(i+1, n) - V_t(i,n)\Big] - A_b\, i^{-\zeta_b} \lambda_b^{\phi_b - 1} &= 0, \\ \Big[V_t(i, n+1) - V_t(i,n)\Big] - A_c\, i^{-\zeta_c} \lambda_c^{\phi_c - 1} &= 0, \end{aligned}$$

where $k_t \equiv \pi(I_t) + (1 - \pi(I_t))p$ is the probability that a novel success also yields a product at time t. Hence the optimal intensities are

$$\lambda_b(i) = \left(\frac{a_t(i{+}1) - a_t(i) + k_t b_t}{A_b\, i^{-\zeta_b}} \right)^{\frac{1}{\phi_b - 1}}, \qquad \lambda_c(i) = \left(\frac{b_t}{A_c\, i^{-\zeta_c}} \right)^{\frac{1}{\phi_c - 1}}. \tag{C.13}$$

Pinning down b_t. Collect the terms in maintext equation (6.2) that are proportional to n on the left-hand side and on the right-hand side. On the left-hand side, the coefficient on n is $r b_t$. On the right-hand side, the capital gain term contributes $\dot{b}_t\, n$ (using $\dot{V}_t(i,n) = \dot{a}_t(i) + \dot{b}_t\, n$), the flow-profit term contributes $\bar{\pi}\, n$ and the creative-destruction term contributes $-\mu_t\, n\, [V_t(i,n) - V_t(i, n{-}1)] = -\mu_t\, n\, b_t$. Equating the coefficients on n on the left-hand side and right-hand side gives

$$r\, b_t \;=\; \bar{\pi} + \dot{b}_t - \mu_t b_t,$$

which is the law of motion for b_t stated in the proposition.

Reduced Bellman equation for $a_t(i)$. Subtract the linear-in-n component from maintext equation (6.2). Using the optimal policies in (C.13), the remaining one-dimensional problem in i is

$$r\, a_t(i) - \dot{a}_t(i) = \max_{\lambda_b, \lambda_c} \left\{ -\frac{A_b}{\phi_b}\, i^{-\zeta_b} \lambda_b^{\phi_b} - \frac{A_c}{\phi_c}\, i^{-\zeta_c} \lambda_c^{\phi_c} + \lambda_b \Big[a_t(i{+}1) - a_t(i) + k_t b_t\Big] + \lambda_c\, b_t - \xi\, i \Big[a_t(i) - a_t(i{-}1)\Big] \right\}. \tag{C.14}$$

Generically, maximization $\max_{\lambda \geq 0} \{\lambda \Delta - \frac{A}{\phi} i^{-\zeta} \lambda^{\phi}\}$ with $\phi > 1$ gives the value $\frac{\phi - 1}{\phi}\, (A\, i^{-\zeta})^{-1/(\phi-1)} \Delta^{\phi/(\phi-1)}$.

Applying this expression to the two R&D margins in (C.14) yields the reduced recursion

$$r\,a_t(i) - \dot{a}_t(i) = \frac{\phi_b - 1}{\phi_b}\left(A_b\, i^{-\zeta_b}\right)^{-\frac{1}{\phi_b - 1}}\Big[a_t(i{+}1) - a_t(i) + k_t b_t\Big]^{\frac{\phi_b}{\phi_b - 1}} \;+\; \frac{\phi_c - 1}{\phi_c}\left(A_c\, i^{-\zeta_c}\right)^{-\frac{1}{\phi_c - 1}}\Big[b_t\Big]^{\frac{\phi_c}{\phi_c - 1}} \\ -\;\xi\, i\Big[a_t(i) - a_t(i{-}1)\Big]. \tag{C.15}$$

Conclusion. The expressions (C.13), the law of motion $r\,b_t = \bar{\pi} + \dot{b}_t - \mu_t b_t$, and the one-dimensional recursion (C.15) verify that the conjecture $V_t(i,n) = a_t(i) + b_t n$ is consistent with maintext equation (6.2), determine b_t and $a_t(i)$, and imply that the optimal policies $\lambda_b(i)$ and $\lambda_c(i)$ depend on i but not on n. □

D Robustness Checks

D.1 R&D-to-net-income ratio

In the baseline calibration (Section 5), I target a moderate R&D-to-net-income ratio of 0.8. This section recalibrates the model to a higher ratio of 1.0 and shows that it fits both targeted moments and untargeted patterns well and implies a similar contribution of combinatorial innovation. Results are reported in Tables D.1 – D.3 and Figure D.1.

Table D.1: Target higher R&D-to-net-income ratio: Calibration and parameter choice.

Panel A: Parameters calibrated externally					
Parameter	Symbol	Value	Parameter	Symbol	Value
interest rate	r	2%	R&D curvature	(ϕ_b, ϕ_c)	2
Panel B: Parameters calibrated through moment matching					
Parameter	Symbol	Value	Moments	Data	Model
novel R&D, scale	A_b	55	pharma R&D / net income	1.0	1.0
comb. R&D, scale	A_c	75	share of within-firm comb. drugs	38%	35%
entry novel R&D, scale	A_e	200	entrants / active firms	26%	21%
novel R&D knowledge elas.	ζ_b	0.50	ing. share, top 20% firms	56%	55%
comb. R&D knowledge elas.	ζ_c	1.38	value: novel / combinatorial	1.6	1.6
depreciation rate	ξ	0.15	ingredient hazard rate	0.15	0.15

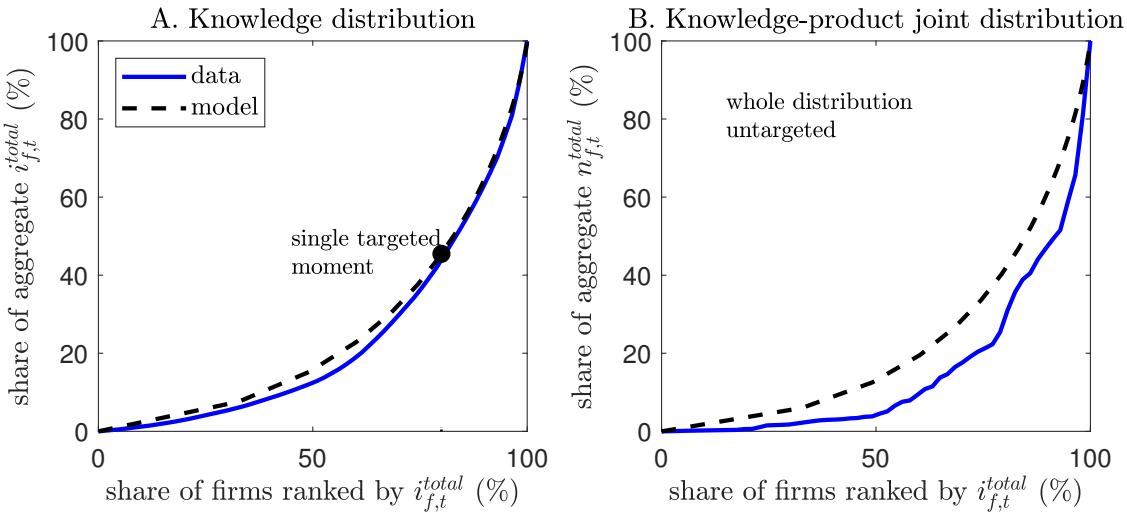

Note: This figure reports the model implications by targeting a higher R&D-to-net-income ratio. Variable definitions are the same as Figure 6. Parameter values are reported in Table D.1.

Figure D.1: Target higher R&D-to-net-income ratio: Knowledge and product concentration

Table D.2: Target higher R&D-to-net-income ratio: Heterogeneous innovation across firms.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$n^{new}_{f,t}$	$i^{new}_{f,t}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$	$n^{new}_{f,t}$	$i^{new}_{f,t}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$
knowledge, $\ln(i^{total}_{f,t})$	0.29	0.18	-0.06	0.06	0.44	0.27	-0.09	0.09
	(0.00)	(0.01)	(0.00)	(0.00)	(0.02)	(0.02)	(0.01)	(0.01)
product, $\ln(n^{total}_{f,t})$					-0.13	-0.08	0.03	-0.03
					(0.02)	(0.02)	(0.01)	(0.01)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.02	0.01	–		0.02	0.01	–	
observations	92,267	92,267	54,482		92,267	92,267	54,482	

Note: This table reports the model implications by targeting a higher R&D-to-net-income ratio. Variable definitions are the same as Table 7. Parameter values are reported in Table D.1.

Table D.3: Target higher R&D-to-net-income ratio: Contribution of combinatorial innovation.

	Baseline	No recombination	Δ (level)	Δ (%)
Product displacement dynamics				
aggregate creative-destruction rate	0.201	0.148	−0.053	−26%
incumbent-driven novel displacement	0.088	0.101	+0.013	+15%
incumbent-driven recomb. displacement	0.071	0	−0.071	−100%
entry-driven displacement	0.042	0.047	+0.005	+12%
Concentration				
top 20% (by knowledge) share of products	59%	54%	−5%	−8%

Note: This table reports the model implications by targeting a higher R&D-to-net-income ratio. Variable definitions are the same as Table 8. Parameter values are reported in Table D.1.

D.2 Value difference between novel and combinatorial innovation

In the baseline calibration (Section 5), I target a value ratio of novel to combinatorial innovation at 1.6. This section recalibrates the model to 1.4 and 1.8 and shows that it fits both targeted moments and untargeted patterns well and implies a similar contribution of combinatorial innovation. Results are reported in Tables D.4 – D.9 and Figures D.2 – D.3.

Table D.4: Target higher value difference: Calibration and parameter choice.

Panel A: Parameters calibrated externally

Parameter	Symbol	Value	Parameter	Symbol	Value
interest rate	r	2%	R&D curvature	(ϕ_b, ϕ_c)	2

Panel B: Parameters calibrated through moment matching

Parameter	Symbol	Value	Moments	Data	Model
novel R&D, scale	A_b	120	pharma R&D / net income	0.8	0.8
comb. R&D, scale	A_c	160	share of within-firm comb. drugs	38%	38%
entry novel R&D, scale	A_e	750	entrants / active firms	26%	21%
novel R&D knowledge elas.	ζ_b	0.55	ing. share, top 20% firms	56%	56%
comb. R&D knowledge elas.	ζ_c	1.69	value: novel / combinatorial	1.8	1.8
depreciation rate	ξ	0.15	ingredient hazard rate	0.15	0.15

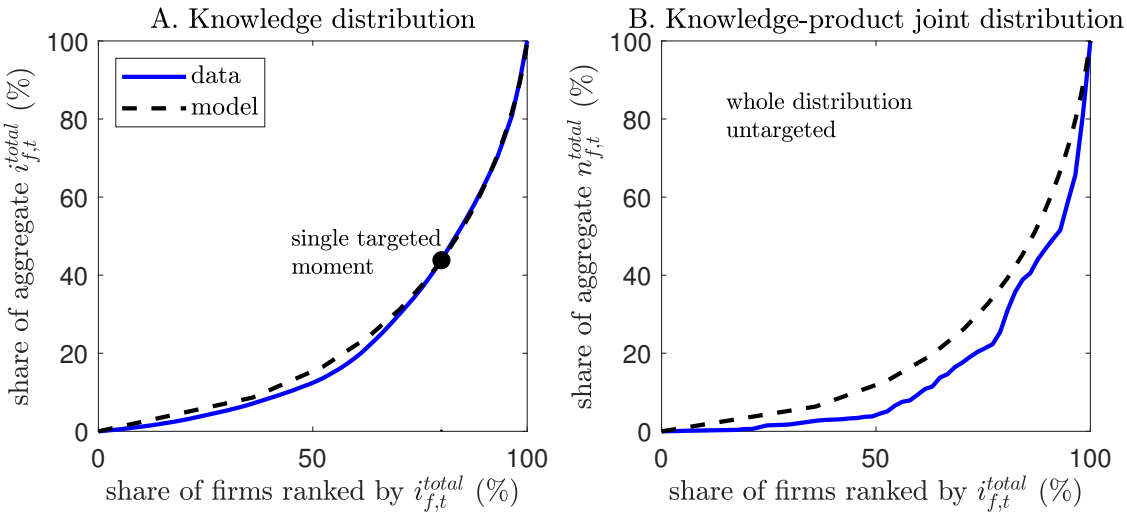

Note: This figure reports the model implications by targeting a higher value difference between novel and combinatorial innovation. Variable definitions are the same as Figure 6. Parameter values are reported in Table D.4.

Figure D.2: Target higher value difference: Knowledge and product concentration

Table D.5: Target higher value difference: Heterogeneous innovation across firms.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$n_{f,t}^{new}$	$i_{f,t}^{new}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$	$n_{f,t}^{new}$	$i_{f,t}^{new}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$
knowledge, $\ln(i_{f,t}^{total})$	0.42	0.26	-0.08	0.08	0.55	0.33	-0.10	0.10
	(0.01)	(0.01)	(0.00)	(0.00)	(0.02)	(0.02)	(0.01)	(0.01)
product, $\ln(n_{f,t}^{total})$					-0.10	-0.06	0.02	-0.02
					(0.02)	(0.02)	(0.01)	(0.01)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.04	0.01	–		0.04	0.01	–	
observations	83,398	83,398	46,939		83,398	83,398	46,939	

Note: This table reports the model implications by targeting a higher value difference between novel and combinatorial innovation. Variable definitions are the same as Table 7. Parameter values are reported in Table D.4.

Table D.6: Target higher value difference: Contribution of combinatorial innovation.

	Baseline	No recombination	Δ (level)	Δ (%)
Product displacement dynamics				
aggregate creative-destruction rate	0.097	0.066	−0.031	−32%
incumbent-driven novel displacement	0.039	0.042	+0.003	+8%
incumbent-driven recomb. displacement	0.037	0	−0.037	−100%
entry-driven displacement	0.021	0.023	+0.002	+10%
Concentration				
top 20% (by knowledge) share of products	62%	54%	−8%	−13%

Note: This table reports the model implications by targeting a higher value difference between novel and combinatorial innovation. Variable definitions are the same as Table 8. Parameter values are reported in Table D.4.

Table D.7: Target lower value difference: Calibration and parameter choice.

Panel A: Parameters calibrated externally

Parameter	Symbol	Value	Parameter	Symbol	Value
interest rate	r	2%	R&D curvature	(ϕ_b, ϕ_c)	2

Panel B: Parameters calibrated through moment matching

Parameter	Symbol	Value	Moments	Data	Model
novel R&D, scale	A_b	80	pharma R&D / net income	0.8	0.8
comb. R&D, scale	A_c	100	share of within-firm comb. drugs	38%	36%
entry novel R&D, scale	A_e	700	entrants / active firms	26%	21%
novel R&D knowledge elas.	ζ_b	0.25	ing. share, top 20% firms	56%	52%
comb. R&D knowledge elas.	ζ_c	1.09	value: novel / combinatorial	1.4	1.5
depreciation rate	ξ	0.15	ingredient hazard rate	0.15	0.15

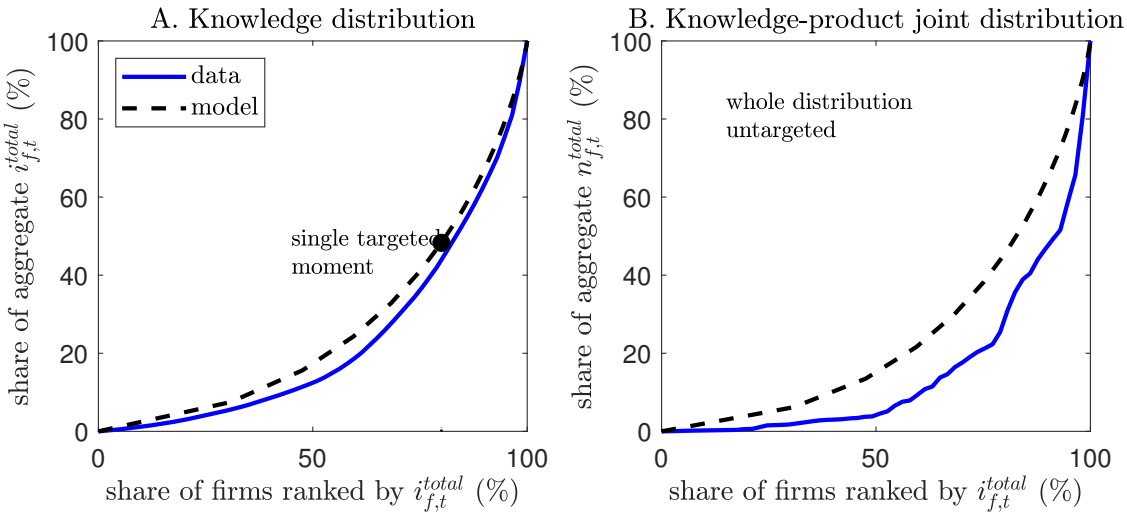

Note: This figure reports the model implications by targeting a lower value difference between novel and combinatorial innovation. Variable definitions are the same as Figure 6. Parameter values are reported in Table D.7.

Figure D.3: Target lower value difference: Knowledge and product concentration

Table D.8: Target lower value difference: Heterogeneous innovation across firms.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$n^{new}_{f,t}$	$i^{new}_{f,t}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$	$n^{new}_{f,t}$	$i^{new}_{f,t}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$
knowledge, $\ln(i^{total}_{f,t})$	0.21	0.12	-0.05	0.05	0.30	0.17	-0.07	0.07
	(0.00)	(0.01)	(0.00)	(0.00)	(0.02)	(0.02)	(0.01)	(0.01)
product, $\ln(n^{total}_{f,t})$					-0.08	-0.05	0.02	-0.02
					(0.01)	(0.02)	(0.01)	(0.01)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.01	0.00	–		0.01	0.00	–	
observations	100,825	100,825	63,399		100,825	100,825	63,399	

Note: This table reports the model implications by targeting a lower value difference between novel and combinatorial innovation. Variable definitions are the same as Table 7. Parameter values are reported in Table D.7.

Table D.9: Target lower value difference: Contribution of combinatorial innovation.

	Baseline	No recombination	Δ (level)	Δ (%)
Product displacement dynamics				
aggregate creative-destruction rate	0.108	0.078	−0.030	−28%
incumbent-driven novel displacement	0.047	0.054	+0.007	+15%
incumbent-driven recomb. displacement	0.039	0	−0.039	−100%
entry-driven displacement	0.022	0.024	+0.002	+9%
Concentration				
top 20% (by knowledge) share of products	55%	52%	−3%	−5%

Note: This table reports the model implications by targeting a lower value difference between novel and combinatorial innovation. Variable definitions are the same as Table 8. Parameter values are reported in Table D.7.

D.3 R&D curvatures

In the baseline calibration (Section 5), I externally set R&D curvature parameters (ϕ_b, ϕ_c) to 2, which implies an elasticity of patents with respect to R&D expenditures of 0.5. This section recalibrates the model to 0.4 and 0.6 and shows that it fits both targeted moments and untargeted patterns well and implies a similar contribution of combinatorial innovation. Results are reported in Tables D.10 – D.15 and Figures D.4 – D.5.

Table D.10: Target lower R&D curvature: Calibration and parameter choice.

Panel A: Parameters calibrated externally

Parameter	Symbol	Value	Parameter	Symbol	Value
interest rate	r	2%	R&D curvature	(ϕ_b, ϕ_c)	1.7

Panel B: Parameters calibrated through moment matching

Parameter	Symbol	Value	Moments	Data	Model
novel R&D, scale	A_b	100	pharma R&D / net income	0.8	0.8
comb. R&D, scale	A_c	150	share of within-firm comb. drugs	38%	36%
entry novel R&D, scale	A_e	600	entrants / active firms	26%	21%
novel R&D knowledge elas.	ζ_b	0.40	ing. share, top 20% firms	56%	59%
comb. R&D knowledge elas.	ζ_c	1.30	value: novel / combinatorial	1.6	1.7
depreciation rate	ξ	0.15	ingredient hazard rate	0.15	0.15

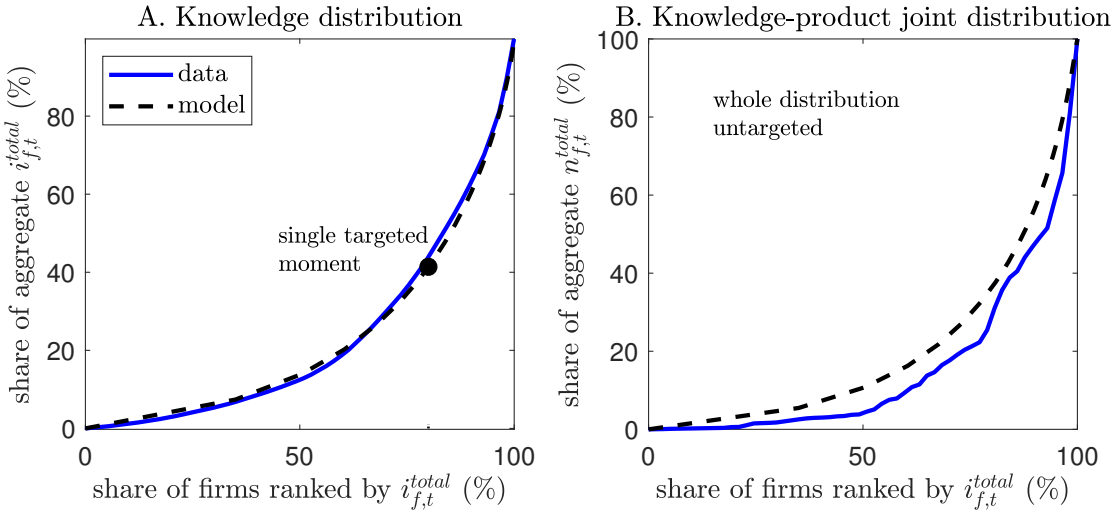

Note: This figure reports the model implications by targeting a lower R&D curvature. Variable definitions are the same as Figure 6. Parameter values are reported in Table D.10.

Figure D.4: Target lower R&D curvature: Knowledge and product concentration

Table D.11: Target lower R&D curvature: Heterogeneous innovation across firms.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$n_{f,t}^{new}$	$i_{f,t}^{new}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$	$n_{f,t}^{new}$	$i_{f,t}^{new}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$
knowledge, $\ln(i_{f,t}^{total})$	0.49	0.31	-0.09	0.09	0.57	0.36	-0.10	0.10
	(0.01)	(0.01)	(0.00)	(0.00)	(0.03)	(0.03)	(0.01)	(0.01)
product, $\ln(n_{f,t}^{total})$					-0.07	-0.05	0.01	-0.01
					(0.02)	(0.02)	(0.01)	(0.01)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.05	0.02	–		0.05	0.02	–	
observations	80,509	80,509	43,699		80,509	80,509	43,699	

Note: This table reports the model implications by targeting a lower R&D curvature. Variable definitions are the same as Table 7. Parameter values are reported in Table D.10.

Table D.12: Target lower R&D curvature: Contribution of combinatorial innovation.

	Baseline	No recombination	Δ (level)	Δ (%)
Product displacement dynamics				
aggregate creative-destruction rate	0.046	0.033	−0.013	−28%
incumbent-driven novel displacement	0.021	0.021	+0.000	+0%
incumbent-driven recomb. displacement	0.017	0	−0.017	−100%
entry-driven displacement	0.010	0.012	+0.002	+20%
Concentration				
top 20% (by knowledge) share of products	64%	55%	−9%	−14%

Note: This table reports the model implications by targeting a lower R&D curvature. Variable definitions are the same as Table 8. Parameter values are reported in Table D.10.

Table D.13: Target higher R&D curvature: Calibration and parameter choice.

Panel A: Parameters calibrated externally

Parameter	Symbol	Value	Parameter	Symbol	Value
interest rate	r	2%	R&D curvature	(ϕ_b, ϕ_c)	2.5

Panel B: Parameters calibrated through moment matching

Parameter	Symbol	Value	Moments	Data	Model
novel R&D, scale	A_b	70	pharma R&D / net income	0.8	0.8
comb. R&D, scale	A_c	90	share of within-firm comb. drugs	38%	39%
entry novel R&D, scale	A_e	300	entrants / active firms	26%	19%
novel R&D knowledge elas.	ζ_b	0.43	ing. share, top 20% firms	56%	52%
comb. R&D knowledge elas.	ζ_c	1.55	value: novel / combinatorial	1.6	1.7
depreciation rate	ξ	0.15	ingredient hazard rate	0.15	0.15

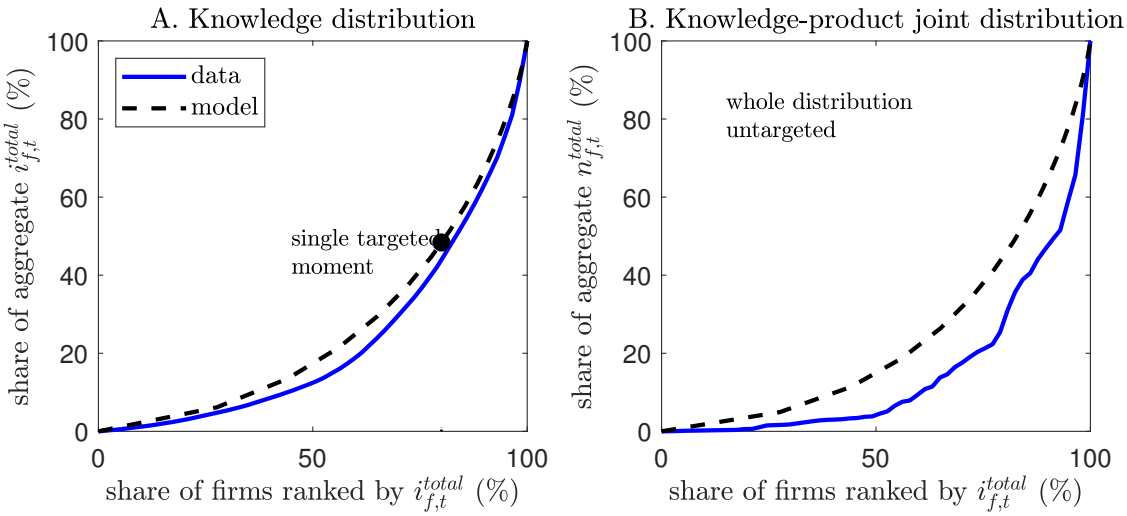

Note: This figure reports the model implications by targeting a higher R&D curvature. Variable definitions are the same as Figure 6. Parameter values are reported in Table D.13.

Figure D.5: Target higher R&D curvature: Knowledge and product concentration

Table D.14: Target higher R&D curvature: Heterogeneous innovation across firms.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$n_{f,t}^{new}$	$i_{f,t}^{new}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$	$n_{f,t}^{new}$	$i_{f,t}^{new}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$
knowledge, $\ln(i_{f,t}^{total})$	0.20	0.12	-0.04	0.04	0.29	0.15	-0.07	0.07
	(0.00)	(0.01)	(0.00)	(0.00)	(0.02)	(0.02)	(0.01)	(0.01)
product, $\ln(n_{f,t}^{total})$					-0.08	-0.03	0.02	-0.02
					(0.01)	(0.02)	(0.01)	(0.01)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.01	0.00	–		0.01	0.00	–	
observations	84,918	84,918	55,575		84,918	84,918	55,575	

Note: This table reports the model implications by targeting a higher R&D curvature. Variable definitions are the same as Table 7. Parameter values are reported in Table D.13.

Table D.15: Target higher R&D curvature: Contribution of combinatorial innovation.

	Baseline	No recombination	Δ (level)	Δ (%)
Product displacement dynamics				
aggregate creative-destruction rate	0.386	0.267	−0.119	−31%
incumbent-driven novel displacement	0.164	0.190	+0.026	+16%
incumbent-driven recomb. displacement	0.149	0	−0.149	−100%
entry-driven displacement	0.077	0.073	+0.004	+5%
Concentration				
top 20% (by knowledge) share of products	55%	52%	−3%	−5%

Note: This table reports the model implications by targeting a higher R&D curvature. Variable definitions are the same as Table 8. Parameter values are reported in Table D.13.

E Additional Results

Table E.16: Number of patented drugs by firm knowledge: 1990–2010.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$n_{f,t}^{new,1}$	$n_{f,t}^{new,2}$	$n_{f,t}^{new,3}$	$n_{f,t}^{new,4}$	$n_{f,t}^{new,1}$	$n_{f,t}^{new,2}$	$n_{f,t}^{new,3}$	$n_{f,t}^{new,4}$
knowledge, $\ln(i_{f,t}^{total})$	0.98	0.55	1.87	1.29	0.77	0.81	1.20	0.53
	(0.11)	(0.05)	(0.16)	(0.16)	(0.42)	(0.19)	(0.50)	(0.44)
size, $\ln(n_{f,t}^{total})$					0.13	-0.16	0.40	0.47
					(0.25)	(0.11)	(0.27)	(0.29)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.46	0.30	0.79	0.49	0.46	0.30	0.79	0.50
observations	265	265	265	265	265	265	265	265

Note: The sample spans 1990 – 2010 and is split into four five-year periods. All columns use Poisson pseudo–maximum likelihood specification with period fixed effects: $E[y_{f,t}|X_{f,t}] = \exp(\alpha + \beta_i \ln(i_{f,t}^{total}) + \beta_n \ln(n_{f,t}^{total}) + \gamma_t)$. The dependent variables are the number of new drugs patented by firm f in period t classified into four types: types 1, 2, 3, and 4 respectively refer to novel, across-firm combination, within-firm combination, and refinement, as defined in the note of Figure 3. Columns (1) - (2) include firm knowledge only ($\ln(i_{f,t}^{total})$), where $i_{f,t}^{total}$ is measured as the number of distinct functional groups used in f's patented drugs by the start of period t. Columns (3) - (4) include both firm knowledge and size ($\ln(i_{f,t}^{total}), \ln(n_{f,t}^{total})$), where $n_{f,t}^{total}$ is measured as its total number of patented drugs by the start of period t. Firms with no new patented drugs during the period are treated as exits and thus dropped. Standard errors are clustered at firm level and reported in parentheses.

Table E.17: Number of new functional groups by firm knowledge: 1990–2010.

	(1) $i_{f,t}^{new,market}$	(2) $i_{f,t}^{new,firm}$	(3) $i_{f,t}^{new,market}$	(4) $i_{f,t}^{new,firm}$
knowledge, $\ln(i_{f,t}^{total})$	1.00	0.21	0.87	0.41
	(0.09)	(0.04)	(0.38)	(0.15)
size, $\ln(n_{f,t}^{total})$			0.08	-0.13
			(0.23)	(0.09)
period fixed effect	✓	✓	✓	✓
R^2	0.45	0.08	0.45	0.09
observations	265	265	265	265

Note: The sample spans 1990 – 2010 and is split into four five-year periods. All columns use Poisson pseudo–maximum likelihood specification with period fixed effects: $E[y_{f,t}|X_{f,t}] = \exp(\alpha + \beta_i \ln(i_{f,t}^{total}) + \beta_n \ln(n_{f,t}^{total}) + \gamma_t)$. The dependent variables count the number of functional groups newly adopted in drugs patented by firm f over period t. Columns (1) and (3) count functional groups that are new to the market; columns (2) and (4) count groups that are new to firm f but previously used by other firms. Columns (1) - (2) include firm knowledge only ($\ln(i_{f,t}^{total})$), where $i_{f,t}^{total}$ is measured as the number of distinct functional groups used in f's patented drugs by the start of period t. Columns (3) - (4) include both firm knowledge and size ($\ln(i_{f,t}^{total}), \ln(n_{f,t}^{total})$), where $n_{f,t}^{total}$ is measured as its total number of patented drugs by the start of period t. Firms with no new patented drugs during the period are treated as exits and thus dropped. Standard errors are clustered at firm level and reported in parentheses.

Table E.18: Innovation outcomes and specialization by firm knowledge: random benchmark.

	(1)	(2)	(3)	(4)	(5)
	$n_{f,t}^{new}$	$i_{f,t}^{new}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}^{across}$	$\%\text{comb}_{f,t}^{within}$
knowledge, $\ln(i_{f,t}^{total})$	0.06	0.13	0.03	-0.05	0.02
	(0.14)	(0.13)	(0.03)	(0.03)	(0.01)
size, $\ln(n_{f,t}^{total})$	0.60	0.23	-0.02	-0.02	0.04
	(0.11)	(0.09)	(0.02)	(0.02)	(0.01)
period fixed effect	✓	✓	✓	✓	✓
R^2	0.63	0.34		–	
observations	265	265		265	

Note: This table reports placebo results from a simulated panel. I hold firms' pre-1990 drug portfolios fixed and fix the functional groups universe to those observed by 2010. For 1990 – 2010, for each observed new drug patented by firm f with k distinct functional groups, I generate one simulated drug with k distinct functional groups drawn at random (without replacement) from the 2010 functional group universe. I then re-estimate the specifications in Tables 2 and 3 using the simulated data. Variable definitions and empirical specifications follow those tables. Results are stable across repeated simulations.

Table E.19: Model-implied regression coefficients for innovation outcomes: $(i_{f,t}, n_{f,t})$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$n^{new}_{f,t}$	$i^{new}_{f,t}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$	$n^{new}_{f,t}$	$i^{new}_{f,t}$	$\%\text{novel}_{f,t}$	$\%\text{comb}_{f,t}$
knowledge, $\ln(i_{f,t})$	0.90	0.60	-0.16	0.16	0.91	0.61	-0.16	0.16
	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)
product, $\ln(n_{f,t})$					0.00	-0.01	-0.00	0.00
					(0.01)	(0.01)	(0.00)	(0.00)
period fixed effect	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.08	0.03	–		0.09	0.03	–	
observations	86,818	86,818	50,644		77,964	77,964	46,304	

Note: The simulated dataset is constructed to match the empirical dataset in time frame, variable definitions, and sample structure, but with a larger simulated population of 50,000 firms for improved precision. Columns (1) – (2) use PPML specification (5.1). Columns (3) – (4) use fractional logit specification (see equation 5.2). Columns (5) – (6) use PPML specification (5.1) but include the portfolio size control $\ln(n_{f,t})$. Columns (7) – (8) use fractional logit specification (5.2) but include the portfolio size control $\ln(n_{f,t})$. The dependent variables are the number of new drugs patented over the subsequent five-year period for columns (1) and (5); the number of new ingredients used in those drugs for columns (2) and (6); the fraction of those drugs being novel for columns (3) and (7); and the fraction of those drugs being combinatorial for columns (4) and (8). Knowledge $\ln(i_{f,t})$ and product portfolio $\ln(n_{f,t})$ are measured as the log-transformations of firm state variables (i, n).

Table E.20: Innovation specialization by firm knowledge: 1990–2010

	(1)	(2)	(3)	(4)
	$\%\text{novel}_{f,t}$ + $\%\text{comb}_{f,t}^{across}$	$\%\text{comb}_{f,t}^{within}$	$\%\text{novel}_{f,t}$ + $\%\text{comb}_{f,t}^{across}$	$\%\text{comb}_{f,t}^{within}$
knowledge, $\ln(i_{f,t}^{total})$	-0.15	0.15	−0.09	0.13
	(0.01)	(0.01)	(0.05)	(0.04)
size, $\ln(n_{f,t}^{total})$			−0.04	0.01
			(0.03)	(0.02)
period fixed effect	✓	✓	✓	✓
observations	265		265	

Note: The sample spans 1990–2010 and is split into four five-year periods. All columns estimate a fractional multinomial logit specification with period fixed effects, analogous to equations (3.1 – 3.2), except that two outcome categories are pooled into a single one. The dependent variables are shares of newly patented drugs. Columns (1) and (3) pool "novel" ($\%\text{novel}_{f,t}$) and "across-firm combination" ($\%\text{comb}_{f,t}^{across}$) into a single category, i.e., drugs that require obtaining external knowledge and thus correspond to novel innovation in the baseline model (Section 4). Columns (2) and (4) report the share of "within-firm combination" ($\%\text{comb}_{f,t}^{within}$), i.e., drugs that combine internal knowledge and thus correspond to combinatorial innovation in the baseline model. The omitted base category is refinement. Knowledge and size are defined as in Table 2. Firms with no new patented drugs during the period are treated as exits and thus dropped. Standard errors are clustered at firm level and reported in parentheses.