Scientific Input Costs and Budgetary Constraints in University Life Sciences Research

Connie Xu Ruby Zhang*

June 14, 2026

Abstract

The real value of a research grant depends not only on the dollars it provides but on the prices those dollars must pay for laboratory inputs. We study how principal investigators (PIs) absorb higher input prices when their grant budgets are fixed in dollars and their personnel commitments cannot be adjusted within the funding cycle. Most of a grant is committed in advance to multi-year personnel contracts, so consumables such as reagents and enzymes are the only margin a PI can adjust, even as the firms supplying them have consolidated. We assemble novel procurement records of 4.8 million purchases across 165 universities from 2010 to 2019, capturing what each PI buys, in what quantities, and at what prices. We study the 2014 merger of Thermo Fisher Scientific and Life Technologies as an exogenous shock to input prices. A difference-in-differences design estimates price increases of approximately 20 percent in affected product markets with little quantity response, indicating highly inelastic demand. We then trace these cost shocks through an author-year panel of 9,166 U.S. life sciences PIs. The average PI faces a 3.8 percent increase in consumables costs and publishes 2.0 percent fewer papers, an implied elasticity of output with respect to consumables costs of approximately -0.54. The decline appears concentrated in less-cited, lower-tier work, while output in the 15 highest-impact journals and citation-weighted output show no significant change. The real value of public research funding therefore depends on the prices science pays for its inputs, not only the dollars that fund it.

*__Xu:__ Harvard University. Email: conniexu@g.harvard.edu. __Zhang:__ Harvard University. Email: rzhang15@g.harvard.edu. We gratefully acknowledge support from the Lab for Economic Applications and Policy Award, John and Kiendl Gordon Economics Support Fund, the National Institute on Aging (T32-AG000186; Xu), the National Science Foundation (DGE #2140743; Xu), and the Alfred P. Sloan Foundation grant (#G-2024-22690; Xu).

1 Introduction

The return to public investment in science depends not only on the dollars a grant provides but on how much research those dollars can purchase. Federal funding for university life sciences research in the United States (U.S.) exceeds $30 billion each year, channeled primarily through the National Institutes of Health (NIH) and the National Science Foundation (NSF), and helps create much of the foundational knowledge behind new drugs and therapeutics. A grant is fixed in dollars when it is awarded and does not adjust when prices change, so as the prices of essential laboratory inputs rise, driven by consolidation among the firms that supply them, each grant dollar buys less research. This paper asks how principal investigators (PIs) absorb higher input prices when their grant budgets are fixed and their personnel obligations cannot be adjusted in the short run. We find that higher input prices reduce research output, with the decline appearing to fall on PIs' less-cited work rather than their highest-impact projects.

The economics of science literature has largely measured research output as a function of the funding researchers receive, not the prices they pay for the material inputs that research consumes. A large literature estimates how funding levels and the design of grants shape what researchers produce (Jones, 2009; Jacob and Lefgren, 2011; Li et al., 2017; Myers, 2020; Azoulay et al., 2021), while treating the cost of laboratory inputs as given and the input bundle itself as unobserved. We instead are able to observe the input bundle directly. Drawing on procurement records that capture, for each PI, what they buy, in what quantities, and at what prices, we bring in the supply side of the research budget that this literature has not been able to measure. We then show that the markets from which PIs procure their inputs govern the research a grant dollar can buy, using the merger as exogenous variation in input prices rather than as a welfare object in itself.

Two features of the university research setting make higher input prices especially likely to reduce output. On the demand side, grant budgets are committed in advance, primarily to multi-year personnel contracts that absorb most of an award, so consumables such as reagents and antibodies are the only margin a PI can adjust within an active grant. On the supply side, the market for these

consumables is concentrated among a few firms, including Thermo Fisher Scientific, Avantor, and MilliporeSigma, that negotiate prices privately and supply protocol-specific inputs with high switching costs, so a PI facing a price increase can neither benchmark it nor substitute away from it. Together these features make PI demand for consumables highly inelastic, so an increase in input prices passes through to the cost of research rather than being avoided through substitution.

We study these adjustments using novel purchase-level procurement records, assembled through public records requests, that resolve prices, quantities, and spending across thousands of life sciences inputs. To generate variation in input prices, we examine the 2014 merger of Thermo Fisher Scientific and Life Technologies Corporation, a $13.6 billion transaction that combined two of the largest suppliers of life sciences inputs and drew substantial antitrust scrutiny. The analysis proceeds in two steps, first measuring the size of the price shock and then tracing its effect on research output.

The first step quantifies the input price shock. Using a difference-in-differences design that compares treated product markets with matched control markets, we estimate that the merger raised prices in treated markets by approximately 20 percent. Spending rose by a comparable amount while quantities were essentially unchanged, which indicates that demand for these consumables is highly inelastic and that the price increase passed through almost entirely to spending. The higher cost of the same inputs must therefore be absorbed within a budget that does not adjust.

The second step traces how this cost shock reaches research output. We assemble an author-year panel of 9,166 U.S. life sciences PIs spanning 2010 to 2019 and construct, for each PI, an exposure measure equal to a spending-weighted average of the market-level price effects. Our procurement data link purchases to individual PIs at only a subset of universities, so we develop a textual matching procedure that imputes each PI's exposure measure from the similarity between their publication record and the records of PIs whose purchases we observe directly. Event-study estimates show a sustained decline in publication output among more-exposed PIs following the merger. For the average PI, the merger raised consumables costs by 3.8 percent and reduced total publication output by 2.0 percent, an elasticity of output with respect to input costs of approximately -0.54.

The decline does not appear to fall evenly across the quality distribution of output. Citation-weighted output shows no significant change and the estimated decline seems to be concentrated in lower-tier journals. We read this as suggestive evidence that more-exposed PIs accommodate higher input costs by scaling back their more marginal work rather than their highest-impact projects.

Taken together, these findings show that the real value of a research grant depends on the prices a PI pays for inputs, not only the dollars the award provides. The price of doing science, not only the budget for it, shapes how much science gets done.

2 Institutional Setting

University PIs operate their research labs under tight federal budgets whose only flexible margin is spending on consumable supplies. Those same supplies come from a few firms with the market power to raise prices, so the one area a PI can adjust is also the one where costs are rising. We study the effect of this tension on a PI's research output through a single large shock, Thermo Fisher Scientific's 2014 acquisition of Life Technologies.

2.1 University Research Budgets

A grant fixes a PI's budget in dollars at the time it is awarded and does not adjust within the funding cycle. Federal grants from the NIH and NSF fund the vast majority of academic life sciences research, and most PIs depend on them to pay for personnel, equipment, and supplies. These awards are typically three- to five-year commitments that are not renegotiated within the cycle (Stephan, 2012), so the grant sets the budget a PI must operate under.

Most of a grant's budget is committed before the research begins. Personnel (PI summer salary, postdocs, graduate students, and technicians) absorb 60–70% of a typical grant's direct costs and cannot be cut on short notice without breaking multi-year commitments to trainees (Stephan, 2012). Major instrumentation such as freezers and microscopes is lumpy, front-loaded, and frequently covered by institutional support. Consumables (reagents, plasticware, kits, antibodies, and enzymes) account for

the remaining 30–40% of direct costs.

Consumables are the only margin a PI can adjust within an active grant. Personnel cannot be dismissed in the short run, so when input prices rise the full burden of adjustment falls on this single flexible category, and paying more for the same inputs leaves fewer resources for everything else. A modest price increase therefore feeds through to research output, as the PI scales back projects, defers long-term hiring, or devotes more time to securing additional funding and less to running experiments.

2.2 Market Structure of the Life Sciences Supplier Industry

The life sciences supplier market concentrates pricing power in a handful of firms. Three structural features give them the ability to raise prices: high concentration, opaque pricing, and high switching costs.

Years of consolidation have left the life sciences supplier market dominated by three firms, Thermo Fisher Scientific, Avantor, and MilliporeSigma, each running integrated manufacturing and distribution arms that together supply the majority of life sciences equipment and consumables. The other firms in this space are smaller, specialized manufacturers such as New England Biolabs and Integrated DNA Technologies, which compete only in narrow segments like restriction enzymes or synthetic DNA oligonucleotides and often depend on the distribution channels of the three large firms to reach customers. A PI ordering inputs for their lab therefore buys, in most cases, from one of the three large suppliers, the largest being Thermo Fisher Scientific.

Prices in this market are also opaque. Each university negotiates a private contract that sets discounts off list price, so two universities can unknowingly pay very different prices for the same reagent. This lack of transparency extends within institutions as well, since PIs are able to negotiate separately with sales representatives causing prices to differ across PIs in the same department.

Even when cheaper alternatives exist, PIs face high switching costs at the bench. Experimental protocols are calibrated to specific reagents, instruments, and kits, so adopting a new supplier's product

requires revalidation that can take weeks or months and risks results that do not reproduce. The cost of price shopping compounds the cost of switching, as most PIs lack the time to compare suppliers and leave available alternatives unexplored. Together these frictions render demand for consumables highly inelastic and lock PIs to their suppliers.

Taken together, these three features leave a small number of suppliers facing buyers whose demand is inelastic. A firm that gains market power in this setting can therefore raise prices with little risk of losing customers.

2.3 The Thermo Fisher–Life Technologies Merger

We study one large shock to the supplier market, Thermo Fisher Scientific's acquisition of Life Technologies Corporation, which combined two of the largest suppliers of life sciences inputs into a single firm and raised concentration across the consumables markets they supplied.

Thermo Fisher Scientific, formed in 2006 from the merger of Thermo Electron and Fisher Scientific, was the largest supplier of scientific instruments, reagents, and consumables, with a distribution network reaching most U.S. university PIs. Life Technologies, created in 2008 from the combination of Invitrogen and Applied Biosystems, held leading positions in molecular biology and genomics reagents, including the enzymes, kits, and cell-culture products. Both ranked among the few full-line suppliers able to cover a wide range of laboratory inputs.

Thermo Fisher agreed to acquire Life Technologies for approximately $13.6 billion in April 2013 and completed the transaction on February 3, 2014, reinforcing its position as the largest supplier in the life sciences tools and reagents industry. The combined firm offered a broad catalog spanning instruments, reagents, and consumables, backed by a sales and distribution network reaching the large majority of academic laboratories. Antitrust authorities in the United States and the European Union reviewed the transaction before clearing it.

2.4 From Price Shock to Research Output

We formalize how the features of this setting can turn an input price increase into lower research output. A PI is endowed with a grant of size B, fixed in dollars over the funding cycle, since additional funding arrives only through new grants that take effect in later cycles. The PI produces research output $y = f(K, L)$ from capital K (reagents and consumables) and labor L (personnel), available at prices p_K and p_L, and allocates the grant to maximize output subject to the budget constraint,

$$\max_{K,L} f(K, L) \quad \text{s.t.} \quad p_K K + p_L L \leq B. \tag{1}$$

The PI chooses inputs to maximize output, but within the funding cycle there is little flexibility. Personnel are committed in advance, and the affected reagents have no close substitutes at the bench, so capital and labor are combined in fixed proportions, with each unit of research requiring a_K units of capital and a_L units of labor. The cost of producing one unit is then $p_K a_K + p_L a_L$, and the PI spends the whole grant on research, so output is

$$y = \frac{B}{p_K a_K + p_L a_L}. \tag{2}$$

The merger raises p_K and therefore the cost of each unit, so that $\partial y / \partial p_K = -a_K B / (p_K a_K + p_L a_L)^2 < 0$. A higher input price thus acts like a cut to the grant's real purchasing power, not a cost the PI can avoid by switching inputs.

3 Data, Measurement, and Sample Construction

Tracing the effect of input prices on research output requires linking three kinds of data: the prices paid for research inputs, the PIs who paid them, and the papers those PIs produced. We assemble the first two from purchase-level procurement data of public universities and the third from a panel of publication metadata.

3.1 Procurement Purchases

Our procurement data come from a commercial data vendor covering purchase records at 165 public universities and from Freedom of Information Act (FOIA) records we file directly with public universities. The FOIA records primarily serve to link a subset of purchases to individual PIs, and they additionally supplement the price records.

GovSpend. GovSpend provides 4.8 million life-sciences purchases across 165 public universities between 2010 and 2019, each a single line-item record reporting the item description, price, quantity, total spending, purchase date, supplier name, and purchasing university. GovSpend is a procurement data platform that aggregates publicly available purchase records from U.S. public entities through open records requests and partnerships with institutional purchasing portals, and we build our extract by focusing on public universities and applying over 1,500 keyword filters designed to capture life sciences items. These data, however, do not report the unit size of the item purchased, such as a 50ml versus a 100ml bottle of fetal bovine serum. Our price measure is therefore the average price per purchase rather than the price per unit, and our quantity measure is the number of purchases rather than the number of units. We return to this limitation in Section 4.3, where the absence of a noticeable quantity response to the merger rules out the leading concern that the observed price increases reflect shifts toward larger package sizes rather than genuine increases in unit prices.

FOIA Requests. We link purchases to individual PIs using records obtained by filing FOIA requests with U.S. public universities, since a PI's consumables spending composition determines the size of their cost increase and the GovSpend records do not identify the PI behind each purchase. We received usable records from five institutions: the University of Texas at Dallas (UT Dallas), the University of Texas at Austin (UT Austin), the University of Michigan (UMich), Texas Tech University (TTU), and East Carolina University (ECU). In addition to having purchase-level spending information, these data provide either the name of the responsible PI or the federal grant award number that funded the purchase. For purchases reported with only an award number, we identify the funded PI by

manually searching for the award number in databases including NIH RePORTER, NSF Award Search, USAspending.gov, and paper acknowledgments. Of the 396,227 purchases identified by award number alone, we are able to attribute 347,989 (92% in spending) purchases to a PI. We additionally obtained purchase records through a FOIA request to Oregon State University, which report item descriptions and prices but not spending or the responsible PI. These records therefore enter the price estimation alongside the GovSpend data but do not contribute to the PI-linked spending shares.

3.2 Classification of Items into Product Markets

Our measure of prices is within a *product market*, which we define as a set of inputs that are interchangeable at the lab bench. We classify purchases into these markets by building a machine learning classifier trained on a set of hand-labeled data and applying the classifier to our procurement data.

Labeled Training Data. Our training labels come from manually classifying the procurement records we obtained through FOIA requests from the life sciences departments at UT Dallas (2011–2024, N=61,882) and UMich (2010–2019, N=306,883).[1] As the largest full-line supplier, Thermo Fisher Scientific offers a product catalog that spans nearly the entire range of life sciences inputs, so we use its catalog hierarchy as the backbone taxonomy for our product markets and map items from smaller suppliers onto the closest Thermo Fisher category. For each unique item description–supplier pairing, we look the item up in its supplier's product catalog to identify what it is, then assign it to the narrowest Thermo Fisher category that applies. For example, this methodology separates *Taq DNA Polymerase*, *Hot Start DNA Polymerase*, and *High-Fidelity Hot Start DNA Polymerase* into distinct markets rather than a single *DNA Polymerase* category, matching their distinct use cases in polymerase chain reaction (PCR) protocols.[2] We also code each product market as either consumables or non-consumables. This

[1] UT Dallas covers the Departments of Biological Sciences and Chemistry & Biochemistry, while UMich covers the Departments of Molecular, Cellular and Developmental Biology, Biological Chemistry, and Pharmacology.

[2] In certain cases, the classifier cannot resolve a broad category into lab bench-interchangeable subcategories from the item description alone, so we retain a coarser market rather than impose splits based on functional substitutability. *Primary Antibodies*, for example, target different antigens and are not substitutes at the bench, yet their descriptions are too similar to separate reliably. This coarsening is concentrated in control markets, which we exclude from the candidate controls used to estimate the price effects, so it does not affect our treated-market estimates.

labeling effort yields 368,765 item descriptions mapped to 2,108 product markets.

Classification Algorithm. Our machine learning classifier operates in two steps. The first step separates consumables from non-consumables by combining deterministic keyword rules with a logistic regression on TF-IDF text representations of the item descriptions.[3] The keyword rules classify obvious cases directly, treating descriptions that contain *polymerase* or *antibody* as consumables and those that contain *publishing* or *furniture* as non-consumables, while the logistic regression classifies the remainder. The second step assigns each consumable to one of 2,108 product markets using centroid-based nearest-neighbor matching in TF-IDF space, refined with keyword override rules for item descriptions that contain unambiguous market identifiers. These override rules are particularly useful when two product markets share many common keywords. For example, items containing "taq" and "polymerase" but not "hot start" or "high-fidelity" are assigned to *Taq Polymerase* rather than *Hot Start Taq Polymerase* or *High-Fidelity Taq Polymerase*, regardless of cosine similarity score. We train both steps on an 80/20 train/validation split of the labeled UT Dallas and UMich data, supplemented for non-consumables coverage with California state procurement records (2012–2015) and Thermo Fisher's negotiated price list with New York State filtered to non-consumables purchases.

3.3 Research Outcomes

We measure each PI's output from annual publication counts, using PubMed to define the universe of life sciences publications and OpenAlex to attribute each publication to an individual PI.

PubMed Medical Subject Headings (MeSH) Extract. PubMed's Medical Subject Headings (MeSH), maintained by the National Library of Medicine, provide a curated taxonomy for biomedical literature that lets us define "life sciences publications" precisely. Our query covers the MeSH descriptors *Anatomy*, *Chemicals and Drugs*, *Organisms*, *Phenomena and Processes*, *Diseases*, *Anesthesia and Analgesia*, *Therapeutics*, and *Clinical Study*, which jointly capture basic, translational, and clinical life sciences

[3] We ran a BERT-based alternative and found that the TF-IDF text representations outperformed at substantially lower computational cost.

research. The predominant MeSH terms in the resulting sample include descriptors such as *Signal Transduction, HIV Infections, Neoplasms, Brain, Gene Expression Regulation,* and *RNA*, consistent with our goal of measuring output across the life sciences. We restrict to journal articles and exclude comments, case reports, technical notes, letters, and reviews.

OpenAlex. OpenAlex provides the consistent author and institutional identifiers we need to link each publication to a PI, which PubMed does not supply. We therefore supplement PubMed with OpenAlex metadata, drawing on a comprehensive open catalog of the global research system developed by the nonprofit OurResearch (Priem et al., 2022). OpenAlex aggregates and standardizes information from sources including PubMed, Microsoft Academic Graph, Crossref, and ORCID. It disambiguates authors with an algorithm that draws on names, publication records, citation patterns, and ORCID identifiers where available, and it standardizes institutional affiliations by parsing raw affiliation strings with a machine learning model.

3.4 Measuring PI Exposure to the Merger

Each PI's exposure to the merger is the measure that links the price shock to research output, capturing how much it raised the cost of the consumables that PI buys. The exposure measure is a shift-share index built from two pieces. The shares are each PI's pre-merger spending across product markets, which we observe directly from the FOIA records and capture the composition of their input bundle. The shifts are the market-level price effects of the merger, which we estimate in Section 4. A PI who devoted a large share of their pre-merger budget to the markets the merger hit hardest therefore has a high exposure measure.

We observe these shares directly for only the 204 PIs in the FOIA sample. For every other PI in the publication panel, we impute the exposure measure from publication text, under the assumption that PIs working on similar scientific topics use similar experimental protocols and therefore buy similar shares of inputs. For each such PI, we compare their pre-merger publication profile, built from titles, abstracts, and MeSH terms, against those of the FOIA PIs, and assign them the average exposure

measure of their closest matches. We defer the full construction of the exposure measure to Section 5.1, once the price effects are estimated.

3.5 Estimation Samples

Our analysis uses two estimation samples. A product market–university–year panel of our procurement data identifies the price effects in Section 4, and an annual PI-year panel of publications identifies the research output effects in Section 6.

Procurement Panel. We impose two restrictions on the procurement panel to ensure precise estimation of market-level price effects. To keep the panel's composition stable, we only keep the 76 universities that report consistently from 2010 to 2019, which leaves 1.2 million consumables purchases across 1,105 product markets after we apply our classifier to the GovSpend extract. To keep the market definitions accurate, we retain only the product markets the classifier assigns with high confidence, defined as at least 25 training examples and precision and recall of at least 0.8, because both fall sharply when training examples are scarce. The resulting sample covers 152 product markets, accounting for 55% of total consumables spending ($112 million). Tables 1, 2, and 3 report the top universities, suppliers, and product markets that account for the largest shares of consumables spending in this sample.

PI Panel. We define PIs as the last author on publications, following the convention in the life sciences that last authorship denotes the PI directing the project. We restrict the panel to authors with at least one publication before the merger year of 2014 and those with a positive exposure measure. Moreover, our procurement data come from public universities, so we restrict the institution sample to those classified as R1 or R2 under the Carnegie Classification. The resulting panel comprises 9,166 U.S. life sciences PIs observed annually from 2010 to 2019 at 136 universities, with an average of 1.8 publications per PI-year in the pre-merger period. The panel is balanced at the PI-year level, so a PI who publishes nothing in a given year gets assigned a publication count of zero rather than leaving the sample. Using k-means textual clustering on publication text, we assign each PI to one of five research areas. Most PIs are concentrated in molecular biology, biochemistry, and structural biology (41.8%), followed by

clinical and translational medicine (23.1%), with neuroscience (12.3%), microbiology, virology, and infectious disease (11.6%), and plant and environmental biology (11.3%) each near a tenth, covering the experimental approaches used across U.S. life sciences research.

4 Estimating Cost Shocks

The merger raised prices and spending in the affected product markets by approximately 20% with little quantity response. We estimate these effects with a difference-in-differences design that compares treated markets, those the European Commission identified as having competitive overlap between Thermo Fisher Scientific and Life Technologies, to matched control markets. The price effect varies widely across these markets, and this cross-market variation provides the identifying variation for the research-output analysis. The minimal quantity response indicates that demand for these consumables is highly inelastic, so the merger raised the cost of doing research rather than prompting substitution toward cheaper inputs.

4.1 Defining Treated Product Markets

We define a product market as treated if the European Commission identified pre-merger competitive overlap there, rather than building concentration measures from our own procurement data. A natural alternative would be to compute the Herfindahl-Hirschman Index (HHI) directly from purchasing records, since HHI captures pre-merger market concentration. Implementing this in our setting is complicated by the fact that GovSpend records supplier names as free-text fields rather than standardized identifiers, so the same firm often appears under multiple variants ("Thermo Fisher Scientific," "Thermo Fisher Sci.," and "ThermoFisher"). While manual cleaning of these fields helps, residual noise still propagates into any firm-level market share calculation and makes HHI a noisy proxy for market concentration. The European Commission, by contrast, defines markets from customer-substitution evidence and assesses competitive effects using reconstructed pre-merger market shares, transaction data, and the parties' internal documents. Thus, a product market is treated ($D_m^T = 1$) if it

appears in the European Commission's clearance decision (Case COMP/M.6944) as a market in which Thermo Fisher Scientific and Life Technologies Corporation had pre-merger competitive overlap. Markets not discussed in the decision are classified as control ($D_m^T = 0$).

The treated markets span the core molecular biology workflow, including cell culture media and sera (e.g., DMEM, fetal bovine serum) and nucleic-acid amplification enzymes and ready-to-use kits (PCR, qPCR, RT-PCR), while control markets cover consumables outside the merging firms' direct overlap, such as glassware, bulk lab chemicals, and pipette tips. Of the 1,105 consumables product markets in our labeled training data, 231 are treated under this definition, and restricting to the high-confidence estimation subset with at least 25 training examples and precision and recall above 0.8 leaves 46 treated and 106 control markets. Appendix B lists the full set of high-confidence markets.

4.2 Empirical Strategy

We use a difference-in-differences specification to estimate the merger's effect on three logged market-level outcomes: (1) average purchase price, (2) total quantity purchased, and (3) total spending. Together they identify both the size of the cost shock and the purchasing response to it. For each treated market m^*, we construct a set of matched controls of up to two markets using nearest-neighbor matching on pre-period log price trends and averages, sampling with replacement so that the same market can serve as a matched control for multiple treated markets. This matching process yields 51 distinct control markets across the 46 treated-market comparisons. Our difference-in-differences specification is as follows:

$$y_{jmt} = \beta_m \cdot D_m^T \cdot \text{Post}_t + \gamma_j + \delta_m + \lambda_t + \epsilon_{jmt}, \tag{3}$$

where y_{jmt} is the log outcome (price, quantity, or spending) in product market m purchased by university j in year t. The terms γ_j, δ_m, and λ_t are university, market, and year fixed effects. The indicator D_m^T equals one for treated markets, and $\text{Post}_t = \mathbf{1}\{t \geq 2014\}$ indicates the post-merger period. Observations are weighted by 2013 market spending, and standard errors are clustered at the market level. We estimate equation (3) separately for each of the 46 treated markets, producing market-

level price effects $\{\hat{\beta}_m^P\}_{m=1}^{46}$ and analogously $\{\hat{\beta}_m^Q\}$ and $\{\hat{\beta}_m^{PQ}\}$ for quantity and spending. Because these market-level estimates are noisily estimated, we apply empirical Bayes shrinkage to the $\hat{\beta}_m^P$ and use the shrunk estimates throughout the paper.

Identification rests on the common trends assumption that, absent the merger, treated and control markets would have followed the same counterfactual trajectory in each outcome. We assess the validity of this assumption by examining pre-trends in an equivalent event study specification shown in Section 4.3.

4.3 Price, Spending, and Quantity Results

Treated product markets experienced sharp, persistent price increases following the merger, while control prices stayed flat. Figure 1 presents the result in two panels, with panel (a) plotting the treated and control price trajectories separately and panel (b) the pooled difference between them, each normalized to the year before the merger. Treated prices jump upward in 2014 and climb steadily to roughly 20% above their pre-merger baseline while control prices show no movement, so the pooled difference in panel (b) traces out the same 20% effect, with flat pre-merger coefficients that lend support to the common trends assumption behind equation (3).

Despite these price increases, the quantity response was minimal, indicating that the higher prices passed through to spending rather than to a reduction in volume. Panel (a) of Figure 2 shows treated and control quantities moving in sync, and the pooled difference in panel (b) is statistically indistinguishable from zero throughout. This pattern is consistent with consumables being essential, protocol-critical inputs whose volumes are difficult to scale back. Both treated and control quantities dip in the first year after the merger, likely reflecting a temporary budget-wide pullback that year.

Spending decomposes as price times quantity, so with the quantity difference being relatively null the pooled spending difference in panel (b) of Figure 3 roughly matches the 20% price effect in Figure 1. In the separated treated and control trajectories shown in panel (a), control spending drops slightly more than treated spending in the years post-merger but this is a result of the shared quantity drop

and lack of price increase in control markets.

Together, these three series characterize demand for the treated consumables as highly inelastic. Treated quantities barely move relative to controls despite price increases of roughly 20%, so treated markets absorb nearly the entire increase as higher spending on the same essential inputs rather than substituting toward cheaper alternatives or reducing volumes. Table 4 reports the pooled difference-in-differences coefficient for each outcome, which is stable across alternative fixed-effects specifications. Appendix Figures A1, A2, and A3 reproduce all three event studies under a naive specification that compares treated markets to the full set of control markets without matching, and the resulting pre-merger trajectories are less comparable, which motivates our matched-control design.

The flat quantity response also rules out a measurement artifact in the price variable, that the 20% increase reflects a shift toward larger package sizes rather than a genuine rise in unit prices. GovSpend procurement records do not include unit quantities, so the 20% price increase could in principle reflect a compositional shift within product markets toward larger and more expensive package sizes (for example, from a 200-tip rack to a 960-tip rack) rather than a genuine increase in the unit price. Such a shift could arise because the merged firm changes the set of package sizes it offers after the merger. Any such compositional shift would mechanically change the quantity variable. A shift toward larger packages would reduce the total quantity purchased, while a shift toward smaller packages would increase it. Such a shift would have to occur in treated markets relative to controls, yet treated and control quantities move together throughout our time frame. The observed price increase therefore appears to reflect a genuine change in the unit price of a relatively stable bundle of goods rather than a shift in package size.

Behind this aggregate price effect lies substantial heterogeneity across product markets, the source of the identifying variation for our PI-level research output analysis. PIs differ in the composition of their input bundles, so a common set of market-level shocks generates different effects across PIs. Figure 4 plots the distribution of the shrunk market-specific price effects $\hat{\beta}_m^P$ against a placebo distribution that randomly shuffles treatment status across control markets, with the same empirical

Bayes shrinkage applied to both so that the comparison reflects genuine cross-market heterogeneity rather than sampling variation. The placebo distribution is tightly centered at zero, while the actual distribution is centered at 0.071 with a thicker right tail. This 0.071 is the mean of the unweighted distribution of market-level price effects $\hat{\beta}_m^P$ and is therefore smaller than the spending-weighted pooled effect of roughly 20%, which loads on the higher-spending treated markets. Both randomization inference and a Kolmogorov–Smirnov test reject the equality of our estimated treatment effects and the placebo treatment effects.

5 PI Exposure to the Cost Shock

Exposure to the cost shocks varies across PIs depending on their use of different experimental protocols and reliance on different bundles of consumables. This is the variation we use to identify the downstream effect on research outcomes. We measure it as a spending-weighted average of the market-level price effects, which will be larger for PIs who rely heavily on the markets with the largest price increases.

5.1 Shift-Share Exposure Construction

Exposure is a shift-share of the market-level price effects, observed directly for the 204 FOIA PIs from their pre-merger spending shares and imputed for the remaining panel PIs from the FOIA PIs whose published work is most similar.

A PI's exposure measures the percentage increase in the cost of their pre-merger consumables bundle, formed by weighting each market's estimated price effect by the PI's pre-merger spending share there:

$$\text{Exposure}_i = \sum_m s_{im} \hat{\beta}_m^P \tag{4}$$

where s_{im} denotes PI i's pre-merger spending share in product market m and $\hat{\beta}_m^P$ is the empirical Bayes shrunk merger-induced price effect in market m from equation (3). We compute the spending shares

relative to the PI's total pre-merger spending across the high-confidence markets defined in Section 4.1. Only treated markets have an estimated $\hat{\beta}_m^P$, so the shares that enter this shift-share measure will not sum up to one but to the PI's treated-market spending share within the set of high-confidence markets. An exposure value of 0.1 corresponds to a 10% increase in PI i's consumables costs. Variation across PIs comes from two sources: (1) differences in the composition of consumables used (the s_{im} shares), which reflect each PI's research focus and experimental protocols, and (2) differences across markets in the size of the merger-induced price effect (the $\hat{\beta}_m^P$ shifts). PIs who devoted larger pre-merger budget shares to markets that experienced larger price increases receive higher exposure values.

Observed Exposure from FOIA Data. For PIs at universities where we obtained FOIA data linking purchases to individual PIs, we calculate the exposure measure directly using observed pre-merger spending shares. We draw on FOIA records from UT Dallas, UT Austin, UMich, TTU, and ECU, keeping PIs who have both sufficient pre-2013 purchasing records and pre-2013 publications. We further restrict the sample to life scientists by applying k-means clustering to PIs' publication text and dropping any PI who does not fall into a life-science-relevant cluster. The resulting 204 PIs span the same five research areas as the full PI panel.

Imputed Exposure Using Textual Matching. For the 8,962 non-FOIA PIs, we impute exposure from publication-text similarity, relying on the link between a PI's published topics and their likely input bundle. The underlying assumption is that PIs working on similar scientific topics use the same experimental protocols, and therefore have comparable spending compositions across consumables product markets.

Imputed exposure is a similarity-weighted average of the observed exposures of a PI's five most textually similar FOIA PIs. We measure similarity as the cosine similarity between TF-IDF representations of pre-merger publication text, which we form for each PI from titles, abstracts, and MeSH

keywords, and combine the neighbors' exposures with weights proportional to that similarity:

$$\widehat{\text{Exposure}}_i = \sum_{k \in \mathcal{N}_5(i)} w_{ik} \cdot \text{Exposure}_k \tag{5}$$

where $\mathcal{N}_5(i)$ denotes the set of five nearest FOIA PIs to target PI i and w_{ik} are cosine similarity weights normalized to sum to one. We choose to include five neighbors by validation on a held-out test set. For each PI who has coauthored with a FOIA PI, we predict their exposure from their independent publication text, excluding any jointly authored papers with a FOIA PI, and compare the prediction to the observed exposure of their FOIA coauthor. We find that imputation accuracy peaks at five neighbors and declines on either side, which motivates the choice of $\mathcal{N}_5(i)$ in equation (5). Figure 5 compares the distribution of the imputed exposure measure across the full PI sample to the observed exposure distribution among FOIA PIs.

6 Effects on Research Output

More-exposed PIs publish fewer papers after the merger than less-exposed PIs, and the lost papers appear to come from their lower-tier, less-cited work.

6.1 Empirical Strategy

Our empirical design treats each PI's exposure to the merger as a continuous treatment and traces the effect on research output with an event study that compares the publication trajectories of more- and less-exposed PIs around the 2014 merger. The merger generates exogenous shifts in the prices of the treated product markets, and a PI's exposure measure aggregates those shifts through the PI's pre-merger consumables bundle, so PIs who rely on the hardest-hit markets receive the largest cost shocks.

Our event-study specification interacts each PI's continuous exposure measure with relative-year indicators centered on the merger year, and the coefficients on those interactions are the parameters

of interest:

$$y_{it} = \mu_i + \lambda_t + \sum_{n=-G}^{N} \beta_n \left[D_{i,t-n} \times \text{Exposure}_i \right] + \sum_{n=-G}^{N} \gamma_n \left[D_{i,t-n} \times S_i \right] + \epsilon_{it} \tag{6}$$

where y_{it} is our measure of research output (i.e., raw publication counts) for PI i in year t. The terms μ_i and λ_t are PI and year fixed effects, $D_{i,t-n}$ are relative-year dummies centered on 2014, and Exposure_i is the imputed exposure measure constructed in Section 5.1. The term $S_i = \sum_m s_{im}$ is the sum of PI i's pre-merger spending shares across the treated markets, which we interact with the same relative-year indicators. We normalize the effect at $t = -1$ to zero, so that all β_n are interpreted relative to the year immediately before the merger.

The exposure measure combines two sources of variation, and our empirical design isolates one of them by controlling for the sum of shares. Since the spending shares are defined over the treated markets alone, they sum to S_i, the overall share of a PI's pre-merger budget exposed to treated markets, rather than to one. A PI's exposure is therefore high for either of two reasons. The PI may have devoted a large share of the budget to treated markets to begin with, captured by a high S_i, or, holding S_i fixed, the PI may have concentrated that budget in the treated markets whose prices rose the most. Only the second reflects the cost shock of interest. The first does not, because S_i is not assigned as good as at random. A high S_i identifies a reagent-intensive PI, and reagent intensity is correlated with field, lab scale, and baseline research trajectory, all of which affect output independently of the merger. Left uncontrolled, exposure would in part proxy for the type of PI rather than the cost shock the PI faced. Following Borusyak et al. (2022), we control for S_i directly. Holding S_i fixed, two PIs are then compared only on which treated markets they relied on and how sharply those markets were hit, which is the variation our design is meant to use. In the event study of equation (6), the PI fixed effects absorb the level of S_i, and the interactions of S_i with the relative-year indicators remove any differential trend associated with overall exposure intensity. The coefficients β_n are thereby identified from this composition-driven variation alone.

A coefficient of $\beta_n = -0.10$ in a post-merger year means a PI with an exposure measure of 1.0, a

100% increase in consumables costs, would publish 0.1 fewer papers that year, while the average PI, with an exposure measure of around 0.038, will only face a small fraction of that. Identification comes from variation in the *degree* of treatment across PIs, since every PI is exposed to the merger to some extent. PI fixed effects absorb time-invariant differences in baseline productivity, and year fixed effects capture aggregate publication trends common to all PIs.

We report the reduced-form effect of the exposure measure on output rather than a structural elasticity of output with respect to consumables spending. Recovering that elasticity would require relating each PI's exposure to the realized increase in their consumables costs, but PI-linked spending is observed only for the FOIA sample, too few to estimate that relationship on the full panel with any precision. We therefore report reduced-form effects throughout and treat the implied output-cost elasticity as an interpretive benchmark rather than a structural estimand.

Our identification here relies on a common-trends assumption that, absent the merger, the publication trajectories of more- and less-exposed PIs would have evolved in parallel, conditional on the PI and year fixed effects. Our exposure measure is a shift-share of market-level price effects, so we place the identifying assumption on exogenous shifts rather than on exogenous shares. This requires that the merger-induced price effects $\hat{\beta}_m^P$ are unrelated to the counterfactual research trajectories of the PIs who bought in those markets. This is plausible because the size of a market's price increase reflects how concentrated that market was before the merger, a supply-side feature that should be unrelated to the research output of the PIs who buy those inputs. Thermo Fisher acquired Life Technologies to combine two product portfolios for reasons orthogonal to where high-output PIs concentrate. The event study in Section 6.2 offers some supporting evidence, with pre-merger coefficients that are individually statistically insignificant and show no systematic trend by exposure, although the individual estimates are noisy and the pre-period is not perfectly flat.

6.2 Reduced-Form Effects on Total Publication Output

More-exposed PIs publish modestly fewer papers after the merger compared to less-exposed PIs, implying an output-cost elasticity of approximately -0.54.

The event-study estimates in panel (a) of Figure 6 are noisy and the pre-period is not flat, but more-exposed PIs publish modestly fewer papers after the merger. We therefore summarize the effect with a difference-in-differences estimate in Table 5 that pools the pre- and post-merger years, regressing annual publication counts on the interaction of post-merger timing with PI exposure and including PI and year fixed effects. Controlling for the treated-market share, which isolates the variation in the price shifts from overall treated-market intensity, the coefficient is -0.976 compared with -1.186 without the control. Panel (b) of the same figure shows the corresponding binned scatter.

For the average PI, the merger's 3.8% cost increase implies a decline of about 0.037 papers per year, a 2.0% reduction from the baseline rate of 1.822 papers per year and an implied output-cost elasticity of approximately -0.54. This follows from multiplying the average exposure measure of 0.038 by the magnitude of the difference-in-differences coefficient, 0.038×0.976, and dividing the resulting decline by the cost increase. This is an elasticity with respect to consumables costs, not the full grant budget, and should not be read as a grant-level elasticity. Consumables are only 30 to 40 percent of direct costs, and rescaling it to the grant would require assuming that a dollar lost to higher input prices reduces output like a dollar removed from any other budget line, an assumption our design does not allow for.

In dollar terms, an additional $20,541 in annual consumables costs corresponds to one fewer publication per PI per year. We obtain this from our manually classified UT Dallas and UMich records, which most accurately link consumables spending to individual PIs. Figure 7 shows that the average PI spends roughly $20,000 per year on consumables, with some spending upwards of $100,000, so the average 3.8% increase of about $760 per year corresponds to the 0.037 papers per year decline.

6.3 Composition of the Output Response

The output decline appears to come out of the PI's least-cited work, while their highest-impact output shows no significant change. We establish this with a citation-weighted measure of output, which is expressed in the same units as the raw publication count and can therefore be compared with it directly. The cost shock therefore appears to operate through a change in the composition of research output.

Our primary composition measure is a citation-weighted count of output, which substitutes for the raw count in the event study of equation (6). We build it by quantifying each author's "effective" contribution to a given paper (Azoulay et al., 2010, 2009) and weighting each paper by its average annual citation count since publication. The weights are normalized so that the citation-weighted measures sum across all authors to the total number of papers in our dataset. This keeps the measure in units of papers and makes it directly comparable to the publication raw count from Section 6.2, so the gap between the two outcomes isolates the average citation impact of the papers a PI forgoes. As an alternative exercise, we also count each PI's publications in the 15 highest-impact life-sciences journals.[4]

While the total number of papers declines, as established in Section 6.2, the estimated effect on citation-weighted output is statistically indistinguishable from zero throughout the post-merger period as shown in Figure 8. Since the two measures share the same units, the contrast between a negative count effect and a citation-weighted effect near zero indicates that the papers a PI forgoes carry below-average citations, that is, that the decline falls on the PI's least-cited work. Appendix Figure A4 shows the same for a count of publications in the top 15 journals, with an effect also indistinguishable from zero.

This pattern is consistent with a shift in the composition of output rather than a uniform contraction. A PI who simply did less of everything would show proportional declines across the quality

[4]The 15 journals are *Cell*, *Nature*, *Science*, *Nature Cell Biology*, *Nature Genetics*, *Nature Medicine*, *Nature Biotechnology*, *Nature Chemical Biology*, *Nature Neuroscience*, *Neuron*, *Cell Stem Cell*, *PLOS One*, *Oncogene*, *Journal of Biological Chemistry*, and *The FASEB Journal*. The list comes from Chandra and Xu (2025), where leading scientists were asked what they viewed as the top journals in their field.

distribution, whereas the estimates instead concentrate the decline away from the highest-impact margin, the pattern expected if more-exposed PIs protect their highest-impact projects. Less-cited, lower-tier work still carries substantial scientific value, so a decline concentrated there should not be mistaken for a loss of unimportant research.

6.4 Robustness

Our reduced-form estimate is robust to the two features of the empirical design, the incomplete shares underlying the exposure measure and the imputation of exposure for PIs outside the FOIA sample.

As discussed in Section 6.1, the spending shares underlying the exposure measure sum to a PI's treated-market share rather than to one, and the baseline specification controls for that share directly. An alternative is to renormalize the shares to sum to one, so that the exposure measure becomes a weighted average of the price shifts across only the treated markets a PI uses. Column (1) of Appendix Table A1 reports this specification, and the estimate remains negative and significant.

Exposure is imputed from publication-text similarity for the PIs outside the FOIA sample, and this measurement error is smaller for PIs whose matches are closer. We restrict to the most similar 40 percent of all imputed PIs, which leaves 7,954 of the 9,166 panel members. The estimate remains negative, significant, and if anything slightly larger in magnitude in column (2) of Appendix Table A1, consistent with classical measurement error attenuating the effect toward zero.

7 Discussion

An input price increase imposed on a largely fixed grant operates as a hidden tax on public research funding, lowering the real value of an award. The economics of science literature has focused on the funding side, whether budgets are adequate and how awards should be structured, but research capacity depends on costs as well. With funding largely fixed, a change in input market structure alters the purchasing power of a grant and therefore the quantity of science it produces.

7.1 The Fixed-Budget Channel

Higher input prices reduce research output because the grant is fixed in dollars. Section 4.3 shows that the quantity of affected reagents purchased does not fall after the merger, so the price increase arrives as a higher cost for the same inputs. A PI's grant budget is divided among personnel, spending in affected markets, and spending elsewhere, so with the budget and personnel fixed, the higher cost of the affected inputs forces spending in other markets, and the experiments it supports, to fall. The cost increase therefore surfaces as lower output rather than as a reduction in the affected reagents themselves. Consistent with this, Section 6.3 shows the decline concentrates in a PI's least-cited, lower-tier work rather than spreading proportionally, and Appendix Figure A5 shows average team size does not change, so the cost is paid in forgone projects rather than fewer people.

We confirm this directly at the university level, where the whole input budget is observed and where exposure is built from each university's own pre-merger procurement rather than imputed from publication text. We regress a university's total consumables purchasing on its exposure to the merger, absorbing university and year fixed effects and controlling for treated-market concentration so that the estimate isolates the price-shift component of exposure. Figure 9 shows that more-exposed universities bought fewer units of consumables overall after the merger. A university given extra funding to cover the higher prices would have held total purchasing steady, so the decline is the signature of a budget that does not expand to absorb the shock.

7.2 Relation to Prior Work on Funding and Cost Shocks

Our results connect to a literature on how funding and cost shocks shape scientific output, and they differ from it on the margin that adjusts. The shocks studied in that work change who does science, whereas ours changes how much a given PI produces. Tham et al. (2024) and Ganguli and Murciano-Goroff (2025) study a funding interruption and a higher minimum wage for research assistants, each of which pushes researchers out of science. Our input price shock instead leaves personnel in place, with average team size unchanged after the merger as Appendix Figure A5 shows, and it lowers output

per researcher rather than the number of researchers.

The magnitude of our estimate lets us read an input price increase as a reduction in the real value of a grant. Tabakovic and Wollmann (2019) study the opposite shock and find that a one percent funding windfall raises output by about 0.31 percent. Their elasticity is taken against total funding while ours is against consumables costs, so we compare only qualitatively. Both shocks act on the high-impact frontier of a PI's work. A funding windfall expands it, raising highly cited output the most, whereas our cost shock protects it, concentrating the decline in lower-tier papers.

The composition of the work that is lost depends on how the shock reaches the PI. Babina et al. (2023) find that federal funding cuts move universities toward applied, patentable output as PIs turn to private funders. Our shock raises input prices rather than withdrawing funds, so PIs have no alternative funder to turn to in the short run, and the work they forgo comes from their lower-tier rather than their highest-quality output. The two settings therefore speak to different dimensions of composition. Babina et al. (2023) focus on the division between basic and applied research whereas we focus on the quality of the work regardless of where it falls on that division.

7.3 Aggregate Implications for Scientific Output

Scaled across the 9,166 life sciences PIs in our sample, the cost-shock channel implies on the order of 340 fewer publications per year against roughly $7 million in additional consumables spending. This scaling extends the response to differential exposure across the sample, so it gauges the size of the cost-shock channel rather than the total effect of the merger. These higher prices do not cost the funding agency anything more. Under a largely fixed budget the same dollars are redirected out of research and toward the suppliers charging them, so public money meant to produce science instead accrues to supplier revenue, and the burden is paid in forgone work.

To the extent the forgone work comes from less-cited and lower-tier journals, it still carries real value. It feeds much of the medical innovation that public funding produces, since the patents behind new drugs and therapeutics draw on a wide range of funded work rather than only its most prominent

papers. Li et al. (2017) find that only about 10 percent of NIH grants yield a patent directly, yet roughly 30 percent produce research later cited by a private-sector patent, so the commercial value of public funding runs largely through a broad base of papers rather than the few that patent directly.

8 Conclusion

The real value of a research grant depends not only on the dollars it provides but on the prices those dollars must pay. This paper has shown that supplier consolidation that raises the price of laboratory inputs imposes a cost on the science that public funding is intended to produce. Drawing on purchase-level procurement data and the 2014 merger of Thermo Fisher Scientific and Life Technologies, we find that the merger raised consumables prices by approximately 20 percent in affected markets, that demand for these inputs is highly inelastic, and that more-exposed PIs subsequently publish less. The implied elasticity of publication output with respect to consumables costs is approximately -0.54. This elasticity is measured against consumables costs rather than the full grant, so it is not directly comparable to estimates of how output responds to grant funding. It captures how output responds when input prices reduce the purchasing power of an award whose dollar size is held fixed.

The decline also appears not to fall evenly across a PI's body of work. With the grant budget largely fixed and the affected inputs costly to substitute, a PI facing higher prices may absorb the shock by forgoing marginal projects while continuing their highest-impact work. Consistent with this, the estimated decline appears concentrated in lower-tier journals where citation-weighted output and output in the most visible journals show no significant change, a pattern suggestive of reprioritization rather than a uniform contraction.

Scaled across the 9,166 life sciences PIs in our sample, the per-PI effect implies on the order of 340 fewer publications per year and roughly $7 million in additional consumables spending. Under a largely fixed budget, these higher prices are research resources redirected toward suppliers, so an input price increase operates as a hidden tax on public research funding. The forgone work lies in the less-cited portion of the research distribution, itself a meaningful source of downstream medical

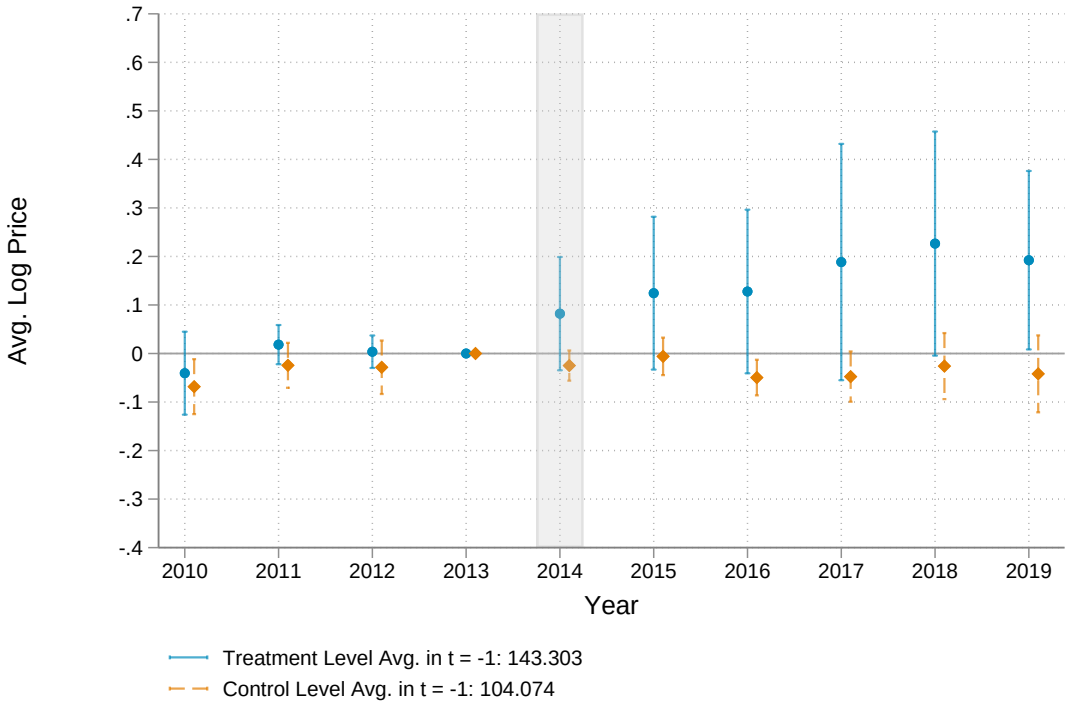
innovation (Li et al., 2017; Azoulay et al., 2019), and this tax has gone largely unmeasured because the prices scientists pay are difficult to observe.

These findings point to where the cost of rising input prices can be addressed. The real value of a grant depends on the prices a PI pays for inputs, yet the prevailing system of opaque, individually negotiated contracts gives PIs and their institutions little ability to benchmark prices or to bring scale to their negotiations. Although individual labs make many of their own purchases, the scale needed to negotiate sits at the institutional level. University purchasing offices and cross-institution consortia could centralize negotiation, share price information, and pool volume across labs, recovering some of the pricing power that concentrated suppliers now hold and increasing the research that each grant dollar buys. Operating through input prices rather than award amounts, this channel remains largely unobserved by the agencies that fund research, which leaves the institutions that purchase the inputs as the natural locus of intervention.

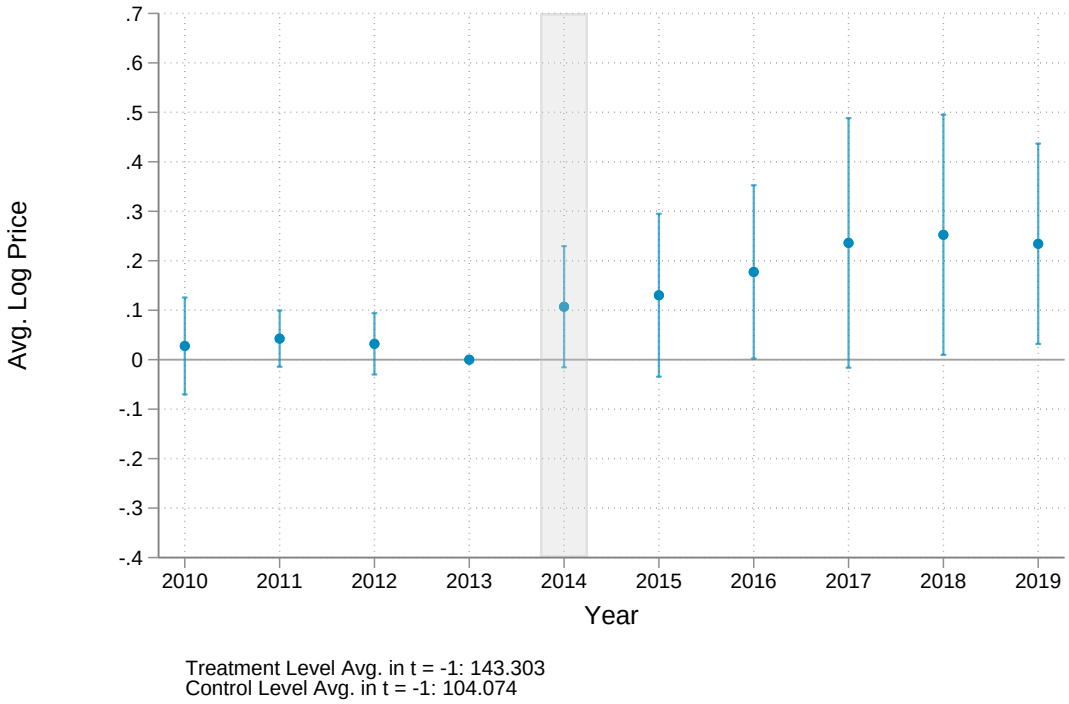
Several questions remain open. We observe publications but not the experiments underlying them, so the path from forgone purchases to forgone projects is inferred rather than directly traced. The welfare cost of the lost research also depends on its long-run scientific and commercial value, which our sample horizon only begins to capture. Scientists may also differ in how well they weather these shocks, since some PIs can substitute across suppliers, redesign protocols around cheaper inputs, or absorb higher prices without cutting projects, so skill in managing inputs is itself a dimension of research productivity alongside the generation of ideas. Quantifying these costs, identifying which scientists weather input shocks best and what enables them to do so, and assessing how far these effects recur across the many input markets that supply modern science are natural directions for future research.

Figure 1. Pooled Merger-Induced Price Effects

(a) Treated and control series



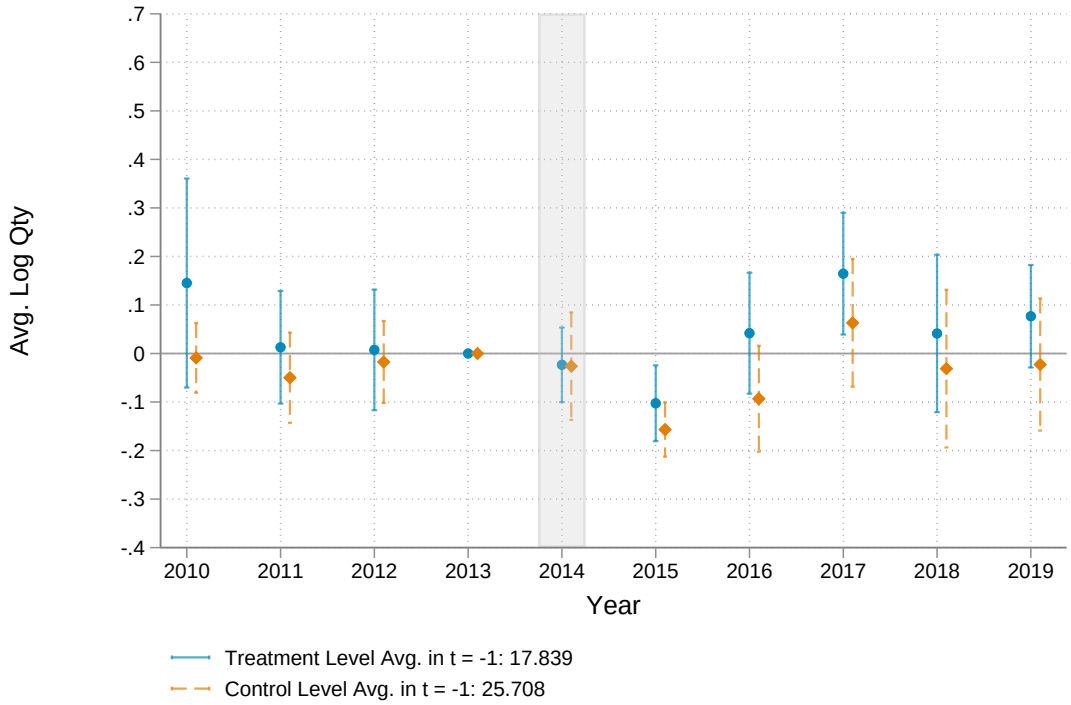
(b) Pooled difference



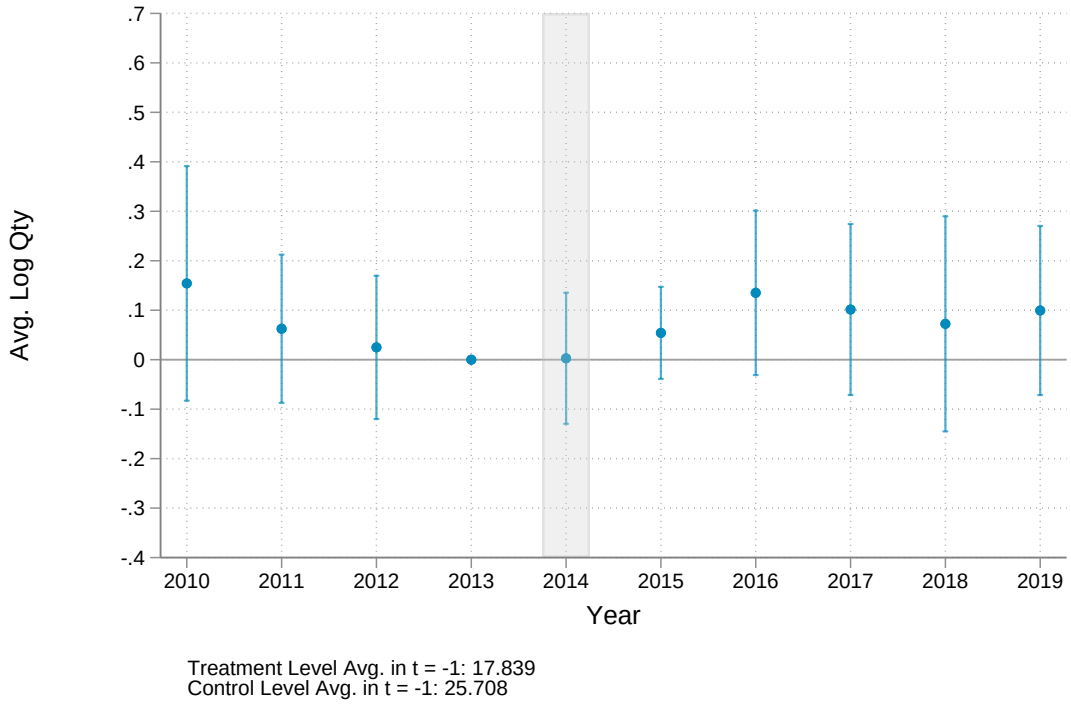
Notes: Event study estimates of average log prices around the Thermo Fisher-Life Technologies merger. Panel (a) plots treated product markets (blue) and their matched controls (orange) separately. Panel (b) plots the pooled difference between the two groups. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Figure 2. Pooled Merger-Induced Quantity Effects

(a) Treated and control series



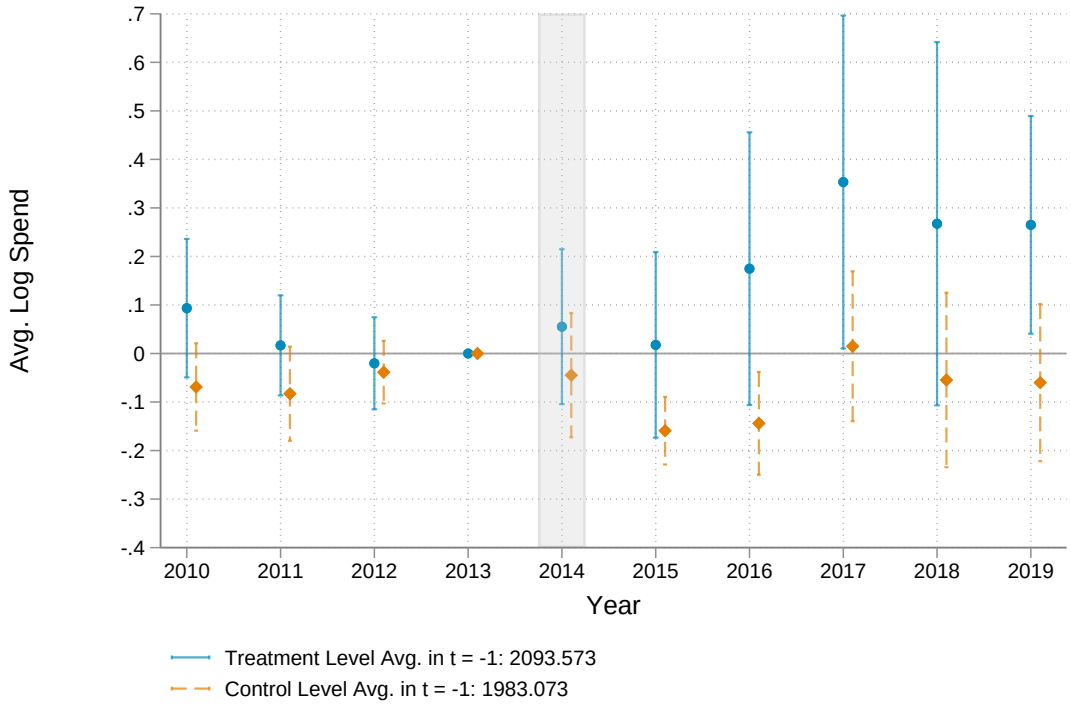
(b) Pooled difference



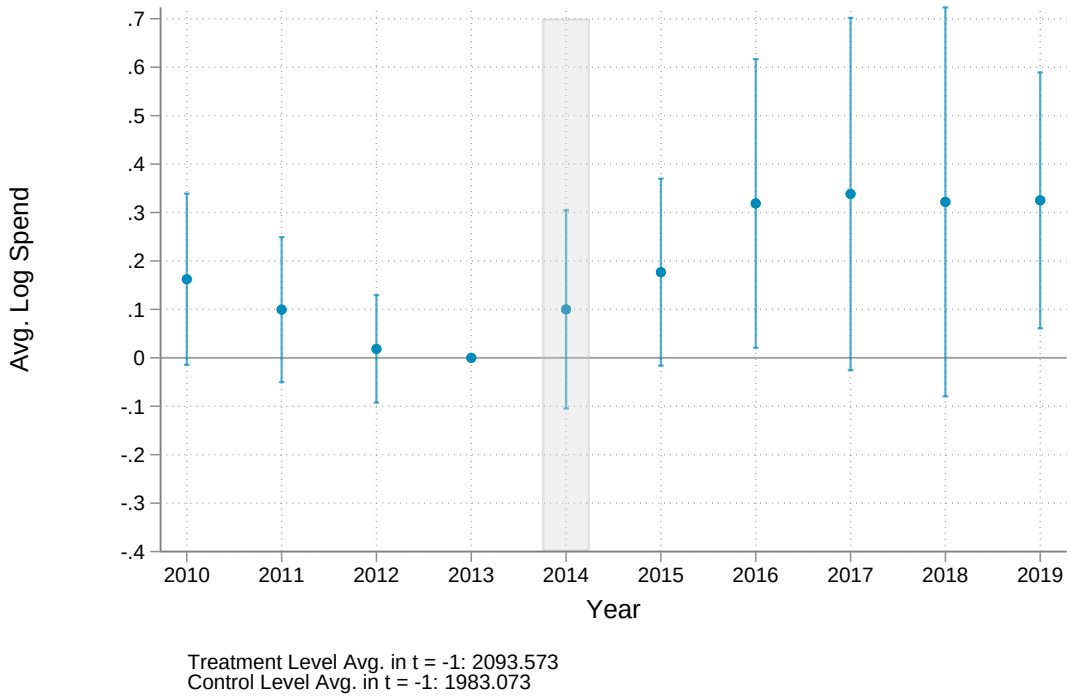
Notes: Event study estimates of average log quantity around the merger. Panel (a) plots treated product markets (blue) and their matched controls (orange) separately. Panel (b) plots the pooled difference between the two groups. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Figure 3. Pooled Merger-Induced Spending Effects

(a) Treated and control series

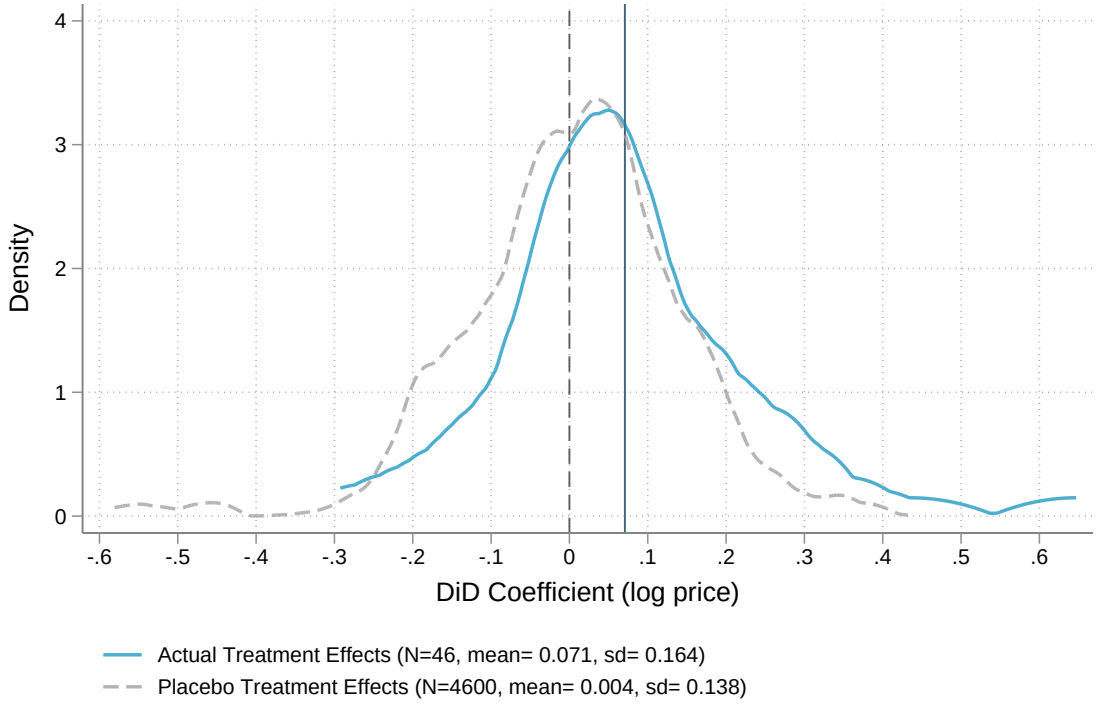


(b) Pooled difference



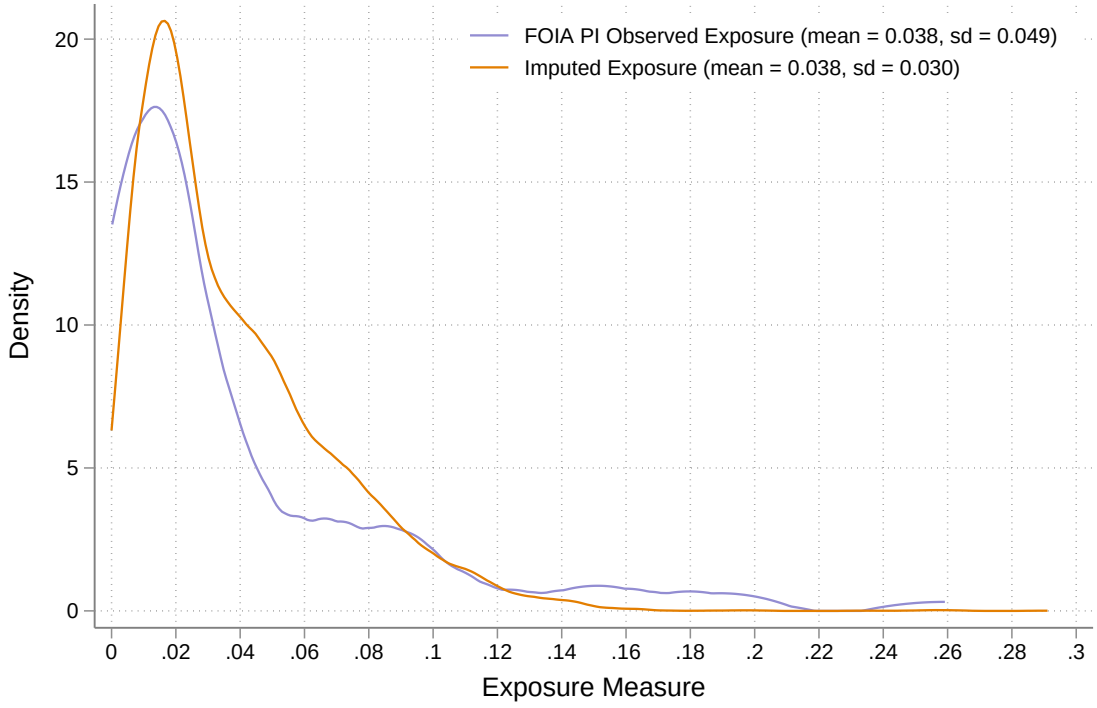
Notes: Event study estimates of average log spending around the merger. Panel (a) plots treated product markets (blue) and their matched controls (orange) separately. Panel (b) plots the pooled difference between the two groups. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Figure 4. Market-Level Merger-Induced Price Effects



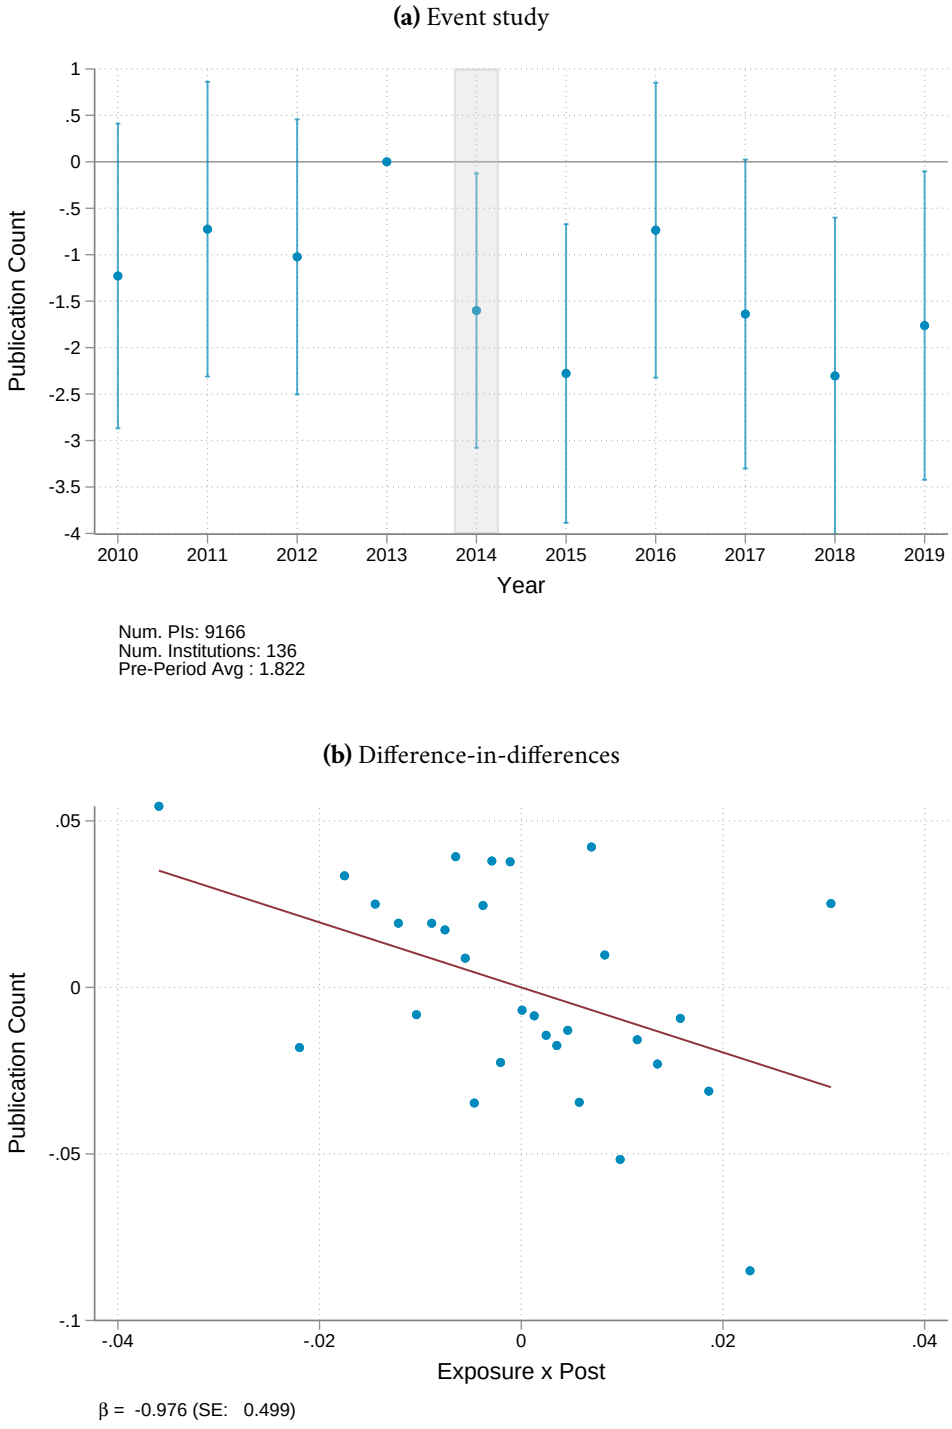
Notes: The figure overlays the unweighted distribution of market-specific difference-in-differences price effects $\hat{\beta}_m^P$ across the 46 treated markets (solid) against a placebo distribution constructed by randomly assigning treatment status across control markets (dashed).

Figure 5. Observed and Imputed PI Exposure Distributions



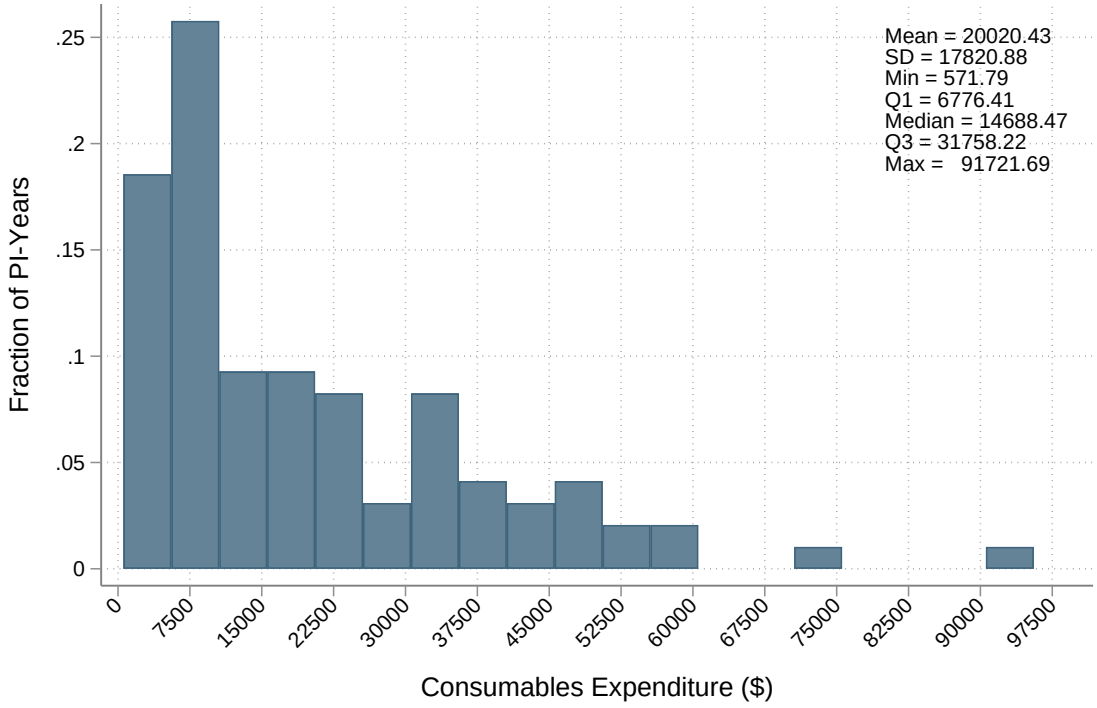
Notes: The purple line shows the distribution of directly observed exposure among 204 FOIA PIs. The orange line shows the distribution of imputed exposure for the full research panel of 8,962 PIs.

Figure 6. Effect of Merger Exposure on Publication Counts



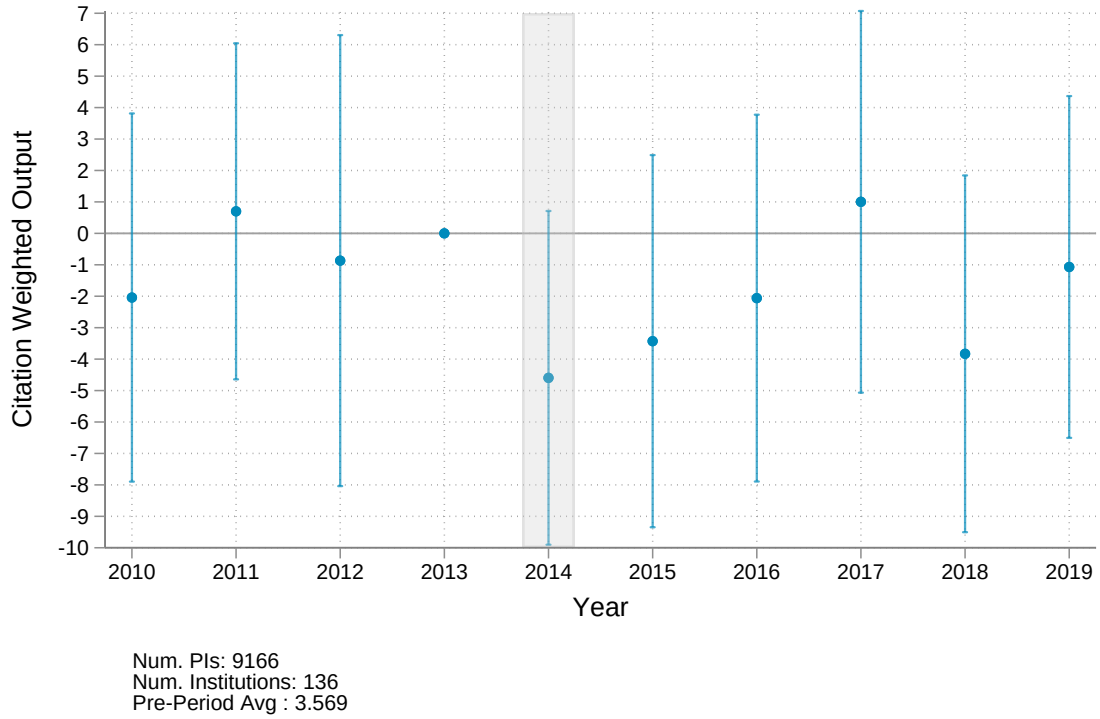
Notes: Panel (a) plots event study estimates from equation (4), the coefficients on the interaction of PI exposure with relative year indicators, with the vertical gray band indicating the merger year (2014) and coefficients normalized to zero at $t = -1$. Panel (b) is a binned scatter of annual publication counts against the $\text{Post}_t \times \text{Exposure}_i$ interaction, after partialling out PI and year fixed effects and the treated-market share control. Each point is a bin of PI-year observations. The slope of the fitted line is the difference-in-differences coefficient, $\beta = -0.976$ (SE 0.499).

Figure 7. Distribution of Annual PI Consumables Spending



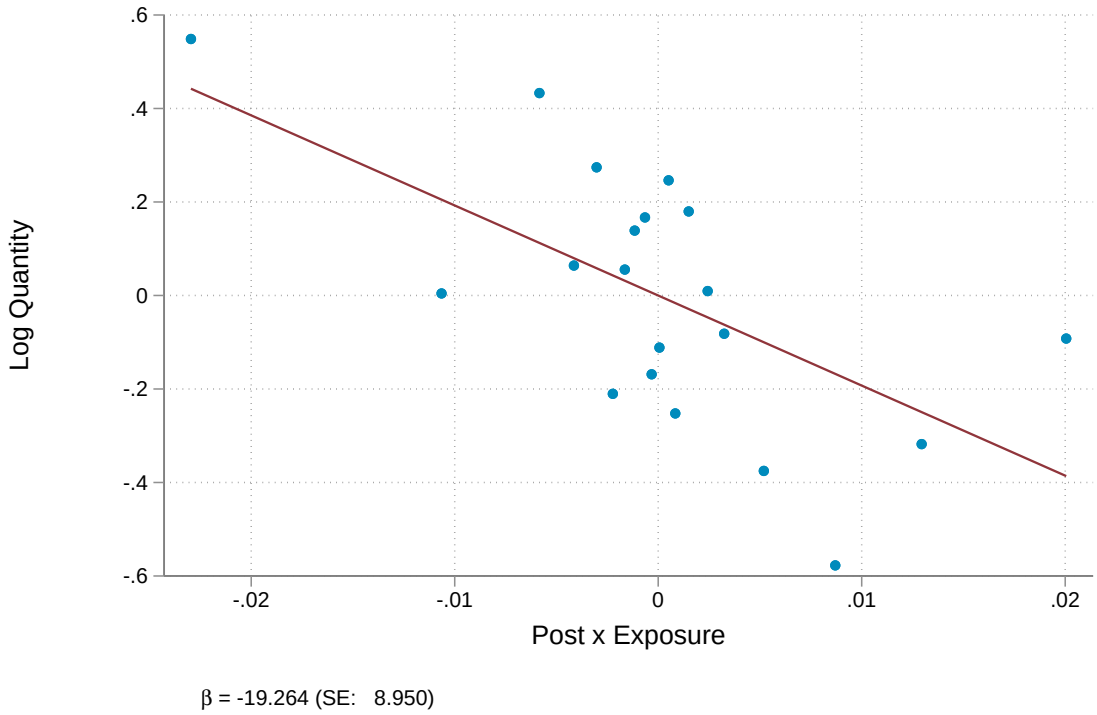
Notes: The figure plots the distribution of annual consumables spending per PI, computed from the manually classified UT Dallas and UMich FOIA records that link purchases to individual PIs.

Figure 8. Effect of Merger Exposure on Citation-Weighted Output



Notes: Event study estimates from equation (4), plotting coefficients on the interaction of PI exposure with relative year indicators. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Figure 9. University Total Consumables Purchasing and Merger Exposure



Notes: Binned scatter of the log of total consumables quantity purchased by a university, summed across all product markets, against the $\text{Post}_t \times \text{Exposure}_j$ interaction, after partialling out university and year fixed effects and the treated-market share control $\text{Post}_t \times S_j$. University exposure is the university-level analogue of the PI exposure measure, observed directly from procurement records. The slope is the difference-in-differences coefficient per unit of exposure, −19.3 (SE 8.95), with standard errors clustered at the university.

Table 1. Top Universities by Consumables Spending

University	Total Spending ($)	% of Total
University of Florida	27,828,927	24.6%
University of Michigan at Ann Arbor	15,229,709	13.5%
Georgia Institute of Technology	11,339,028	10.0%
University of Kentucky	6,232,692	5.5%
Wayne State University	5,611,603	5.0%
Texas A&M University	4,927,740	4.4%
Texas Tech University	4,808,370	4.3%
East Carolina University	4,565,450	4.0%
Florida State University	4,307,242	3.8%
Washington State University	3,935,903	3.5%

Notes: The ten universities with the highest total consumables spending in the high-confidence estimation sample. Spending is summed over all purchases in the 152 high-confidence product markets, 2010–2019. Shares are computed relative to total consumables spending across all universities in the sample.

Table 2. Top Suppliers by Consumables Spending

Supplier	Total Spending ($)	% of Total
Thermo Fisher Scientific	46,058,726	40.7%
VWR	15,595,166	13.8%
Life Technologies	8,232,818	7.3%
MilliporeSigma	6,789,028	6.0%
Abcam	3,276,299	2.9%
Bio-Rad Laboratories	2,700,172	2.4%
Cell Signaling Technology	2,534,807	2.2%
Qiagen	2,225,760	2.0%
Santa Cruz Biotechnology	1,466,029	1.3%
USA Scientific	1,349,498	1.2%

Notes: The ten suppliers with the highest total consumables spending in the high-confidence estimation sample. Spending is summed over all purchases in the 152 high-confidence product markets, 2010–2019. Shares are computed relative to total consumables spending across all suppliers in the sample. Thermo Fisher Scientific and Life Technologies are the two merging firms.

Table 3. Top Product Markets by Consumables Spending

Product Market	Total Spending ($)	% of Total
Primary Antibodies	19,117,698	16.9%
Pipette Tips	6,237,983	5.5%
US FBS	5,783,318	5.1%
Nitrile Gloves	5,175,748	4.6%
Serological Pipettes	4,735,295	4.2%
Vials	3,319,824	2.9%
Dye-Based qPCR Systems	3,005,386	2.7%
Synthetic DNA Oligonucleotide – Desalted	2,729,039	2.4%
Petri Dishes	2,334,389	2.1%
Cell Culture Plates	2,307,216	2.0%

Notes: The ten product markets with the highest total consumables spending in the high-confidence estimation sample. Spending is summed over all purchases in each market, 2010–2019, and shares are computed relative to total consumables spending across all markets in the sample.

Table 4. Pooled Merger Effects on Prices, Quantities, and Spending

	Log Price			Log Quantity			Log Spending		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post × Treated	0.164*	0.162*	0.153	0.023	0.050	0.001	0.199*	0.225*	0.167
	(0.099)	(0.096)	(0.096)	(0.063)	(0.064)	(0.063)	(0.111)	(0.121)	(0.112)
University FE	✓			✓			✓		
Market FE	✓		✓	✓		✓	✓		✓
Year FE	✓	✓		✓	✓		✓	✓	
University × Market FE		✓			✓			✓	
University × Year FE			✓			✓			✓
Observations	19,388	19,388	19,330	19,388	19,388	19,330	19,388	19,388	19,330

Notes: Each column reports the pooled difference-in-differences coefficient on $D_m^T \times \text{Post}_t$ from a single regression of the column's log outcome on treatment, estimated on the high-confidence treated and matched-control markets, 2010–2019. Columns (1), (4), and (7) are the baseline specification with university, market, and year fixed effects. Columns (2), (5), and (8) use university-by-market and year fixed effects. Columns (3), (6), and (9) use university-by-year and market fixed effects. Standard errors clustered at the market level in parentheses. These are unshrunk pooled estimates and differ from the empirical Bayes market-level price effects $\hat{\beta}_m^P$ used elsewhere in the paper. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5. Effect of Merger Exposure on Publication Output

	(1) Without share control	(2) With share control
Exposure measure	−1.186***	−0.976*
	(0.453)	(0.499)
Treated-market share S_i		−0.179
		(0.159)
Pre-period mean	1.822	1.822
Observations	91,660	91,660

Notes: Difference-in-differences estimates of the effect of merger exposure on annual publication counts, from a panel of 9,166 PIs observed 2010–2019 with PI and year fixed effects. The reported coefficient is on $\text{Post}_t \times \text{Exposure}_i$. Column (1) omits the treated-market share control; column (2) adds it following Borusyak et al. (2022) and is our preferred specification. Standard errors clustered at the PI level in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

References

Pierre Azoulay, Waverly Ding, and Toby Stuart. THE IMPACT OF ACADEMIC PATENTING ON THE RATE, QUALITY AND DIRECTION OF (PUBLIC) RESEARCH OUTPUT. *The Journal of Industrial Economics*, 57(4):637–676, December 2009. ISSN 00221821, 14676451. doi: 10.1111/j.1467-6451.2009.00395.x. URL https://onlinelibrary.wiley.com/doi/10.1111/j.1467-6451.2009.00395.x.

Pierre Azoulay, Joshua S. Graff Zivin, and Jialan Wang. Superstar Extinction. *The Quarterly Journal of Economics*, 125(2):549–589, May 2010. ISSN 0033-5533. doi: 10.1162/qjec.2010.125.2.549. URL https://doi.org/10.1162/qjec.2010.125.2.549.

Pierre Azoulay, Joshua S. Graff Zivin, Danielle Li, and Bhaven N. Sampat. Public R&D investments and private-sector patenting: evidence from NIH funding rules. *The Review of economic studies*, 86 (1):117–152, 2019. ISBN: 0034-6527.

Pierre Azoulay, Wesley H. Greenblatt, and Misty L. Heggeness. Long-term effects from early exposure to research: Evidence from the NIH "Yellow Berets". *Research Policy*, 50(9):104332, November 2021. ISSN 0048-7333. doi: 10.1016/j.respol.2021.104332. URL https://www.sciencedirect.com/science/article/pii/S0048733321001311.

Tania Babina, Alex Xi He, Sabrina T Howell, Elisabeth Ruth Perlman, and Joseph Staudt. Cutting the Innovation Engine: How Federal Funding Shocks Affect University Patenting, Entrepreneurship, and Publications*. *The Quarterly Journal of Economics*, 138(2):895–954, May 2023. ISSN 0033-5533. doi: 10.1093/qje/qjac046. URL https://doi.org/10.1093/qje/qjac046.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, 89(1):181–213, January 2022. ISSN 0034-6527. doi: 10.1093/restud/rdab030. URL https://doi.org/10.1093/restud/rdab030.

Amitabh Chandra and Connie Xu. Person and Place Effects in Scientific Discovery, July 2025. URL https://www.nber.org/papers/w33996.

Ina Ganguli and Raviv Murciano-Goroff. Exposure to Science and Scientific Careers: Evidence from Minimum Wage Increases and University Lab Employment, September 2025. URL https://www.nber.org/papers/w34244.

Brian A. Jacob and Lars Lefgren. The impact of research grant funding on scientific productivity. *Journal of Public Economics*, 95(9):1168–1177, October 2011. ISSN 0047-2727. doi: 10.1016/j.jpubeco.2011.05.005. URL https://www.sciencedirect.com/science/article/pii/S0047272711000776.

Benjamin F. Jones. The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder? *The Review of Economic Studies*, 76(1):283–317, January 2009. ISSN 0034-6527. doi: 10.1111/j.1467-937X.2008.00531.x. URL https://doi.org/10.1111/j.1467-937X.2008.00531.x.

Danielle Li, Pierre Azoulay, and Bhaven N. Sampat. The applied value of public investments in biomedical research. *Science*, 356(6333):78–81, 2017. ISBN: 0036-8075.

Kyle Myers. The Elasticity of Science. *American Economic Journal: Applied Economics*, 12(4):103–134, October 2020. ISSN 1945-7782. doi: 10.1257/app.20180518. URL https://www.aeaweb.org/articles?id=10.1257/app.20180518.

Jason Priem, Heather Piwowar, and Richard Orr. OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts, June 2022. URL http://arxiv.org/abs/2205.01833. arXiv:2205.01833 [cs].

Paula Stephan. *How Economics Shapes Science*. Harvard University Press, 2012. ISBN 978-0-674-04971-0. URL https://www.jstor.org/stable/j.ctt2jbqd1.

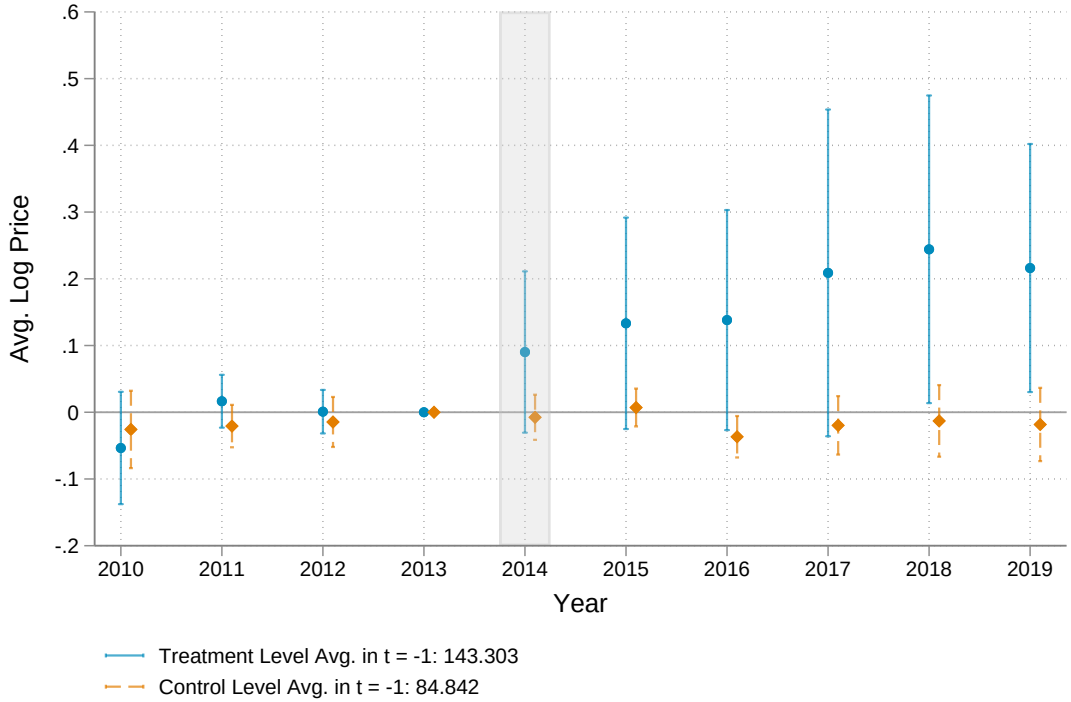
Haris Tabakovic and Thomas G. Wollmann. The impact of money on science: Evidence from unexpected NCAA football outcomes. *Journal of Public Economics*, 178:104066, October 2019. ISSN 0047-2727. doi: 10.1016/j.jpubeco.2019.104066. URL https://www.sciencedirect.com/science/article/pii/S0047272719301276.

Wei Yang Tham, Joseph Staudt, Elisabeth Ruth Perlman, and Stephanie D. Cheng. Scientific Talent Leaks Out of Funding Gaps, February 2024. URL http://arxiv.org/abs/2402.07235. arXiv:2402.07235 [econ].

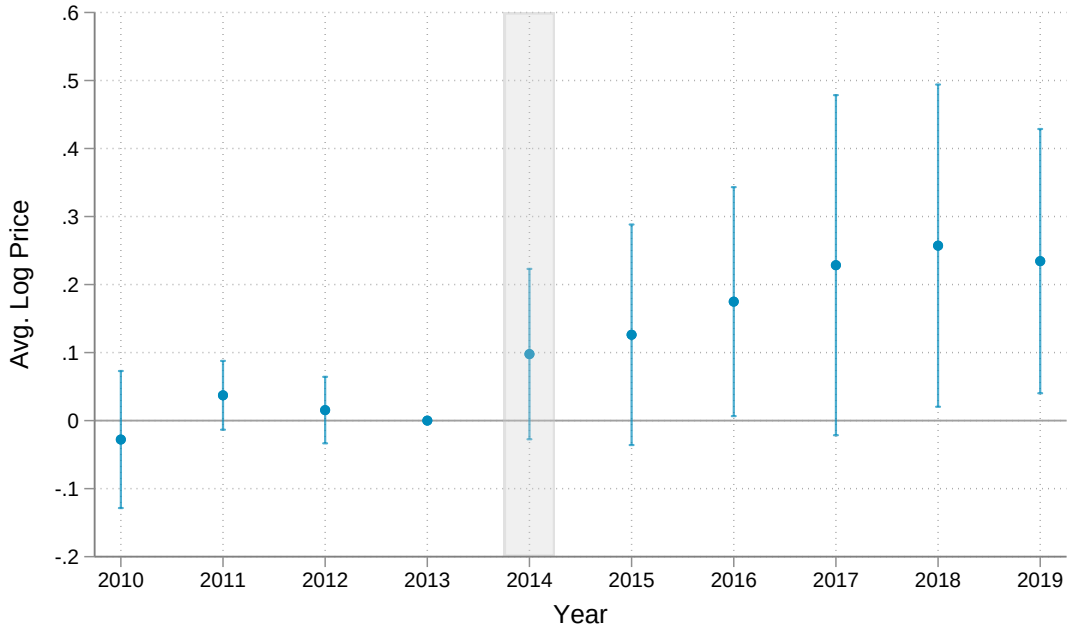
A Appendix Figures and Tables

Figure A1. Pooled Merger-Induced Price Effects, Naive Specification

(a) Treated and control series



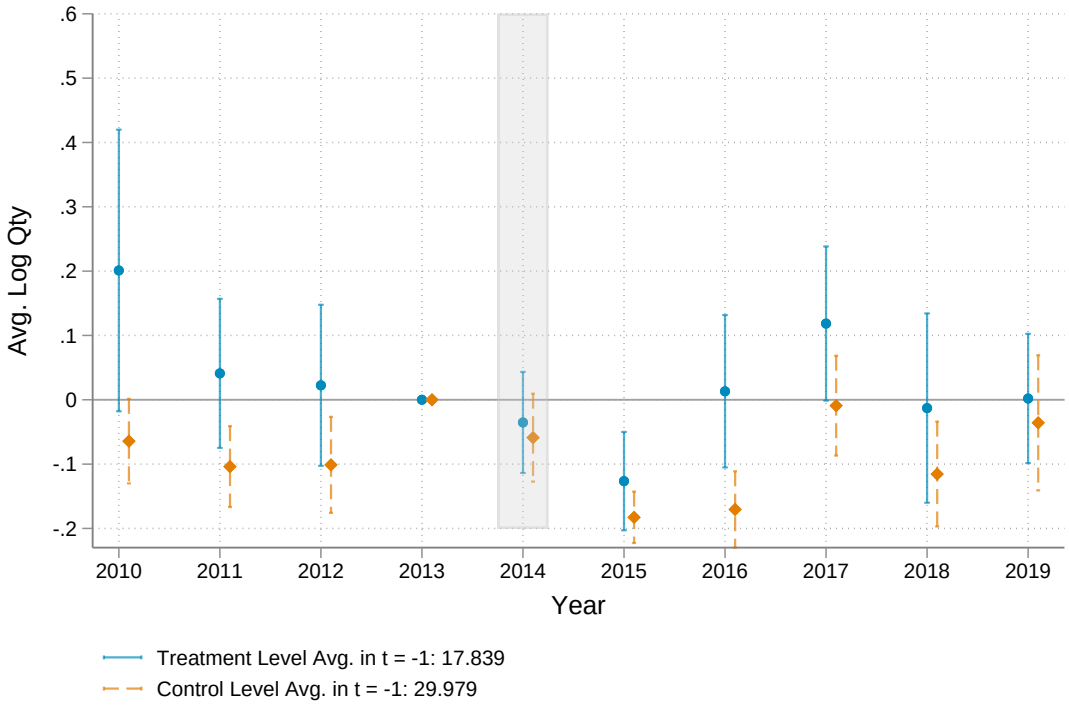
(b) Pooled difference



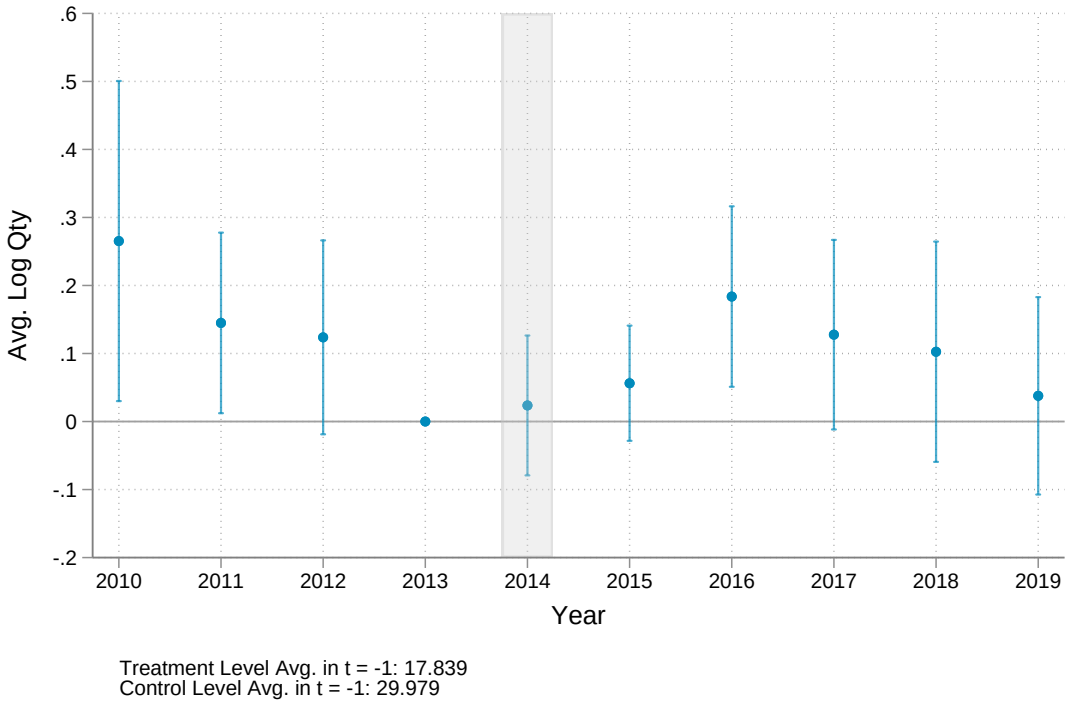
Notes: Naive event study estimates of average log prices around the Thermo Fisher-Life Technologies merger. Panel (a) plots treated product markets (blue) and the full set of unmatched control markets (orange) separately. Panel (b) plots the pooled difference between the two groups. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Figure A2. Pooled Merger-Induced Quantity Effects, Naive Specification

(a) Treated and control series



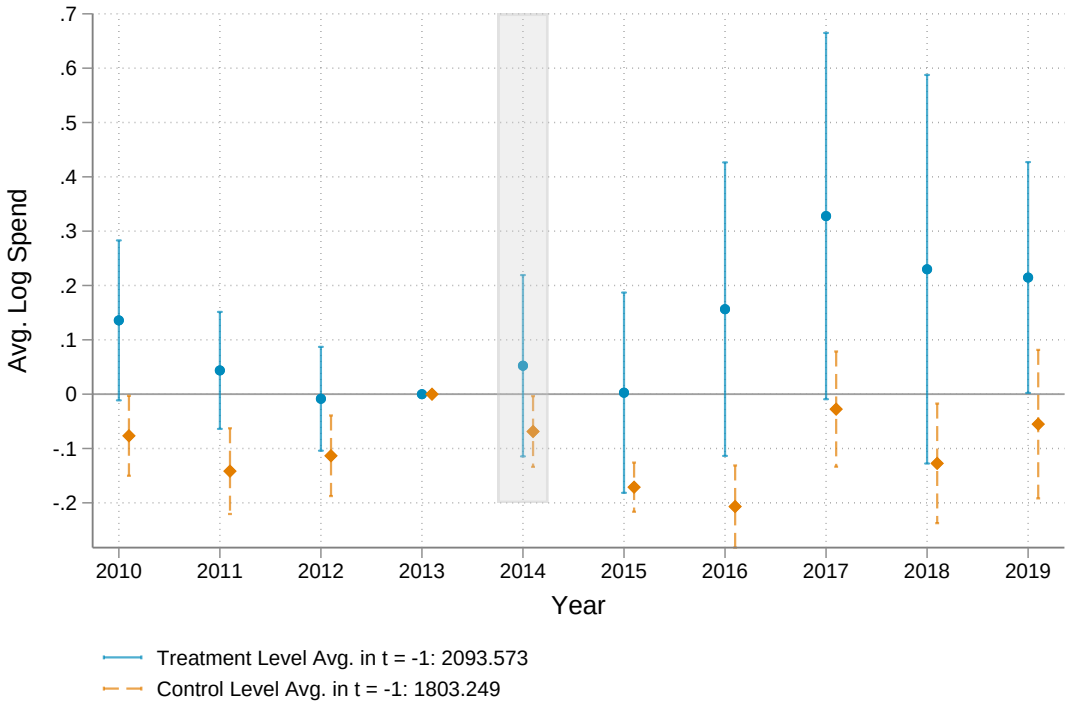
(b) Pooled difference



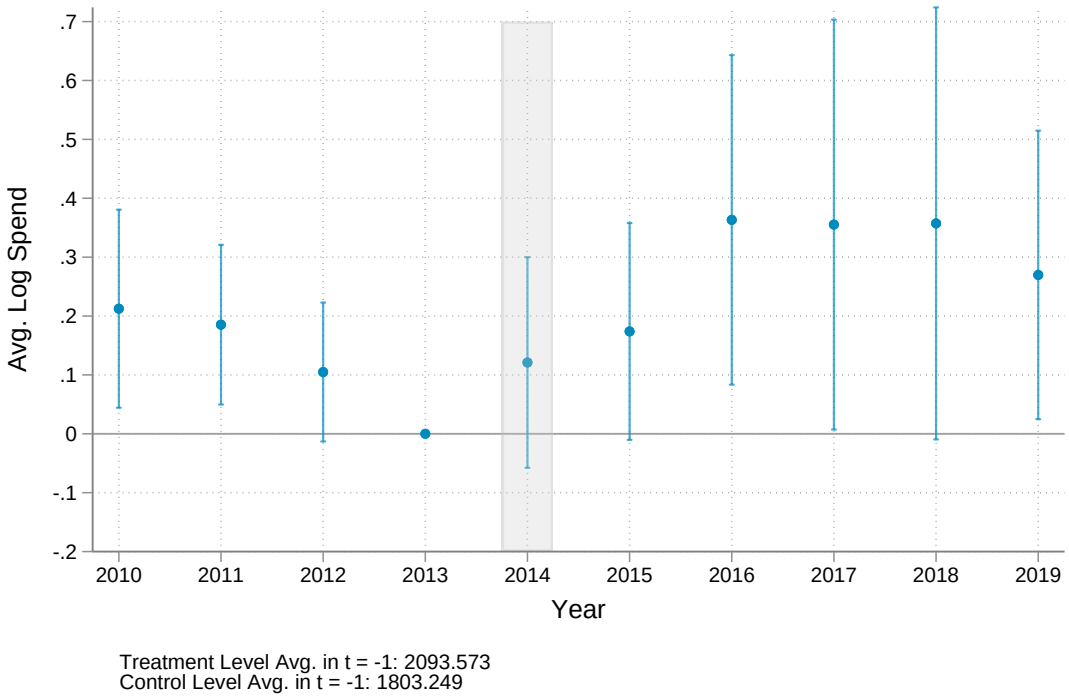
Notes: Naive event study estimates of average log quantity around the merger. Panel (a) plots treated product markets (blue) and the full set of unmatched control markets (orange) separately. Panel (b) plots the pooled difference between the two groups. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Figure A3. Pooled Merger-Induced Spending Effects, Naive Specification

(a) Treated and control series

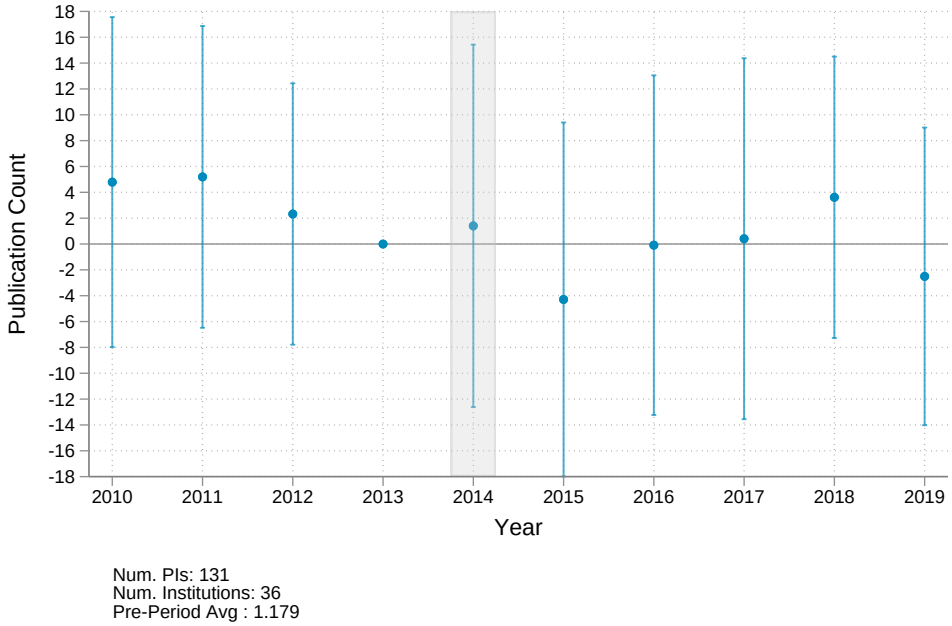


(b) Pooled difference



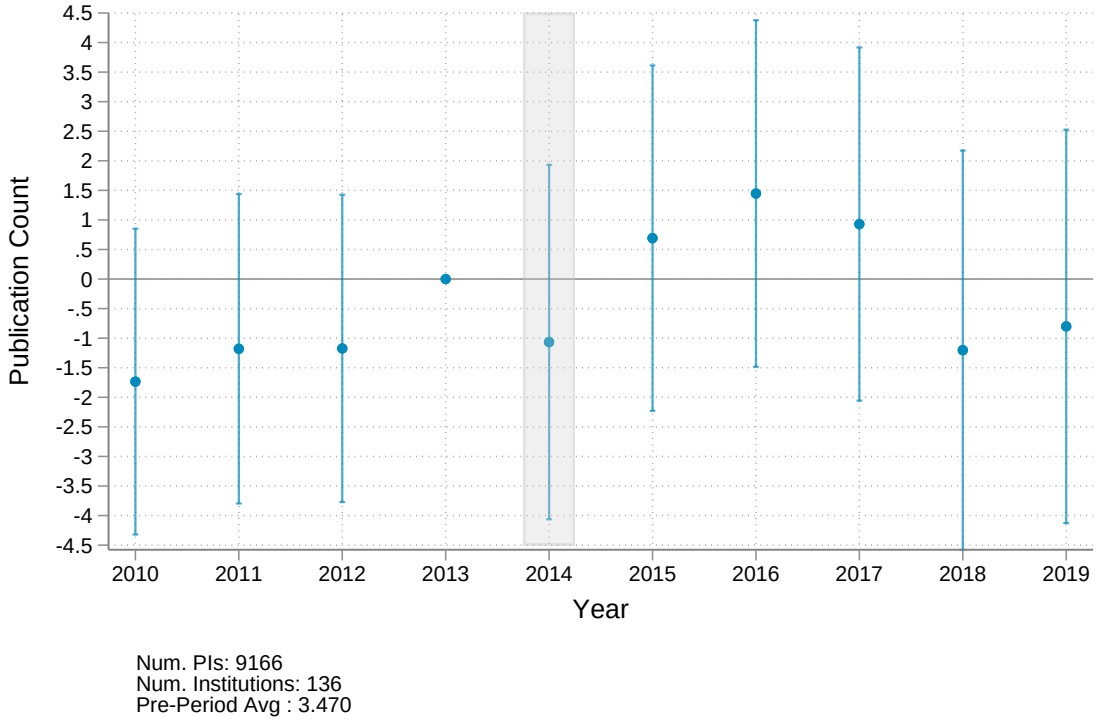
Notes: Naive event study estimates of average log spending around the merger. Panel (a) plots treated product markets (blue) and the full set of unmatched control markets (orange) separately. Panel (b) plots the pooled difference between the two groups. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Figure A4. Effect of Merger Exposure on Publication Counts, by Journal Tier



Notes: Event study estimates from equation (4) estimated separately for two outcomes, the count of a PI's publications in the 15 highest-impact journals and the count in all remaining journals. Journal tiers are fixed using pre-merger impact, as described in Section 3.3. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Figure A5. Effect of Merger Exposure on Average Number of Coauthors



Notes: Event study estimates from equation (4), plotting coefficients on the interaction of PI exposure with relative year indicators. The vertical gray band indicates the merger year (2014). Coefficients are normalized to zero at $t = -1$.

Table A1. Reduced-Form Robustness: Match Quality and Share Normalization

	(1) Normalized	(2) High similarity
Exposure measure	−0.720***	−1.017**
	(0.184)	(0.504)
Treated-market share S_i	—	0.002
		(0.130)
Observations	9,555	7,954

Notes: Reduced-form estimates under two variations of the specification in Table 5, column 2, each observation a PI. Column (1) normalizes the shift-share shares to sum to one, so S_i is constant across PIs and is omitted. Column (2) restricts to panel PIs in the most similar 40 percent of all imputed PIs, which leaves 7,954 of the 9,166 panel members, and includes the treated-market share S_i as a control following Borusyak et al. (2022). Standard errors in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B High-Confidence Markets

This appendix lists the 152 product markets that were considered as high-confidence through our machine learning classifier. This lists consists of 46 treated markets and 106 control markets. A market is classified as high-confidence if it has at least 25 training examples and a precision and recall above

0.8. Treated markets are those the European Commission identified as having competitive overlap in its clearance decision (Case COMP/M.6944), and control markets fall outside the merging firms' direct overlap. Within each group, the markets are organized into broad categories for readability.

Treated Markets

- **Cell culture media, sera, and supplements.** DMEM, DMEM/F-12, RPMI, U.S.-origin fetal bovine serum, donkey serum, goat serum, horse serum, Dulbecco's phosphate-buffered saline (DPBS), Hanks' balanced salt solution (HBSS), cell-culture dissociation reagents, tryptone, and yeast extract.
- **Nucleic-acid amplification and modifying enzymes.** Taq polymerases, high-fidelity DNA polymerase, hot-start PCR systems, reverse transcriptase, first-strand cDNA synthesis systems, DNase I, RNase inhibitors, Klenow fragment, T4 polynucleotide kinase, and site-directed mutagenesis systems.
- **Nucleic-acid purification kits.** Column-based plasmid miniprep, midiprep, and maxiprep kits, gel DNA extraction kits, microbial DNA purification kits, and RNA stabilization reagent.
- **Gel electrophoresis and protein blotting.** Acrylamide-bis solution, pre-cast Bis-Tris gels, pre-cast TBE gels, pre-cast Tris-glycine gels, pre-cast Tris-tricine gels, MES-SDS running buffer, Laemmli sample buffer, LDS sample buffer, nitrocellulose blotting membranes, PVDF blotting membranes, pre-stained protein molecular-weight ladders, unstained DNA ladders, nucleic-acid gel stains, protein gel stains, and chemiluminescent substrates.
- **General buffers and blocking reagents.** Phosphate-buffered saline (PBS), Tris-EDTA (TE) buffer, and bovine serum albumin.

Control Markets

- **Bulk chemicals, acids, and solvents.** Acetone, chloroform, ethyl acetate, isopropanol, methanol, tetrahydrofuran (THF), trifluoroacetic acid (TFA), xylenes, phenols, glycerol, glycine, nitric acid, sulfuric acid, boric acid, hydrogen peroxide, bleach, ammonium chloride, ammonium persulfate, lithium chloride, magnesium chloride, potassium carbonate, potassium phosphate, sodium acetate, sodium azide, sodium bicarbonate, sodium chloride, sodium citrate, sodium hydroxide, sodium sulfate, and zinc chloride.
- **Buffering agents.** HEPES buffer, MES buffer, and MOPS buffers.
- **Antibiotics and selection reagents.** Ampicillin, carbenicillin, chloramphenicol, hygromycin, penicillin-streptomycin, puromycin, zeocin, and IPTG.
- **Detergents, lysis, and reducing reagents.** Digitonin, SDS (sodium dodecyl sulfate), Triton X-100, Tween detergents, DTT, TCEP, and protease inhibitor cocktails.
- **Stains, substrates, and detection kits.** Apoptosis detection kits, bioluminescence assay kits, colorimetric substrates (DAB, MTT, and TMB), histology and cell stains, phalloidin conjugates, viability stains, western blot stripping buffers, and PAP pens.
- **Separation media and other molecular-biology reagents.** Agarose, dNTPs, GST-tag affinity resins, size-exclusion resins, spin desalting columns, and dialysis cassettes.

- **Cell-culture and microbiology reagents.** Carbon dioxide, laminin, polylysines, thrombin, dehydrated bacterial broth, and inoculating loops and needles.
- **Labware, plasticware, and glassware.** Cell culture dishes, flasks, plates, and tubes, coverslips, microscope slides, cell scrapers, cell strainers, cryovials, cuvettes, beakers, boiling flasks, dewar flasks, disposable pipettes, serological pipettes, pipette tips, extraction thimbles, film and foil wrapping, filter papers, funnels, graduated cylinders, lab coats, laboratory spatulas, latex gloves, media bottles, microcentrifuge tubes, PCR tube strips, petri dishes, spray bottles, square glass bottles, stir bars, storage jars, vials, wash bottles, weigh boats, and weighing papers.