Who captures the value of public innovation?

Eugenie Dugoua
London School of Economics & CEPR

Todd Gerarden
Cornell University

Kyle R. Myers
Harvard University & NBER

Jacquelyn Pless
University of Toronto

June 17, 2026

[*In progress. Comments welcome.*]

Publicly funded innovations create social value, but the welfare consequences depend on who captures that value. This paper develops a framework for measuring the incidence of public innovation and applies it to R&D funding from the U.S. Department of Energy. The framework links public innovation to firm-level production allowing us to estimate how the gains generated by public R&D are allocated across workers, firms, and consumers. To address the endogeneity of investment, we develop a political production function approach that isolates funding windfalls: residual variation in congressional appropriations that cannot be predicted by budget requests from the executive branch. We combine these windfalls with measures of firms' exposure to DOE-funded technologies, and we estimate the effects the DOE's innovations. The estimates indicate that, within firms, the value from public innovation is directed slightly more towards capitalists than workers. Incorporating measures of demand, the estimates indicate that the surplus from public innovation is directed more towards consumers than firms. While we use the DOE as a laboratory for studying who captures the returns when governments invest in new technologies, understanding the incidence of public investment in energy technologies is substantively important in and of itself, as it is central to current debates over climate policy, industrial policy, national security, and technological competitiveness.

Dugoua: e.dugoua@lse.ac.uk. Gerarden: gerarden@cornell.edu. Myers: kmyers@hbs.edu. Pless: jacq.pless@gmail.com.

1. Introduction

Firms invest in R&D because new ideas are privately valuable. But these new ideas also create social value as surplus spills over to other firms, workers, and consumers (Arrow 1962; Griliches 1979; Jones and Williams 1998, 2000; Bloom et al. 2013). Governments also invest in R&D, subsidizing R&D at a range of public institutions, non-governmental organizations, and private firms. And a growing body of evidence suggests that these public investments generate large social returns (Azoulay et al. 2019; Jones and Summers 2022; Fieldhouse and Mertens 2026a,b). But who captures the gains from innovations?

For privately funded R&D, economic theory provides guidance on how the gains from innovation are distributed because profit-maximization leads firms to direct their R&D investments at a particular set of new ideas (e.g., Romer 1990; Grossman and Helpman 1991; Aghion and Howitt 1992; Acemoglu 2002; Acemoglu and Linn 2004; Acemoglu 2007; Acemoglu and Restrepo 2018). The case of public R&D is harder to predict. Governments invest in innovation with a wide range of missions and mechanisms, and without a full model of that process, it is unclear how the gains are distributed across capitalists, workers, and consumers.

In this paper, we estimate how the gains from public innovation are split. The first ingredient is a research design that isolates variation in public R&D investments that is not driven by the expected impact of those investments. We develop a method for estimating political production functions and apply it to data from the U.S. Department of Energy (DOE). The key idea is that, each year, the DOE budget requests reveal much of what the agency knows about the value of potential investments. With a set of assumptions about the budgeting process, we can use the budget requests to isolate funding windfalls: variation in actual budgets that cannot be predicted with requested budgets. Using this approach, we can identify windfall investments in key energy technologies on the scale of many millions of dollars per year, over the span of roughly three decades.

The second ingredient is a framework for tracing the incidence of public innovation. In general, public R&D investments can affect firms in many ways: changing productivity, factor demands, markups, etc. Our framework motivates a series of regressions that estimate the effect of public innovation on firms while making minimal restrictions on the channels and nature of these effects. The framework then shows how the reduced form elasticities from the regressions (e.g., the elasticity of firm spending on labor with respect to the stock of public innovation) can be combined with accounting identities to recover the distribution of gains from public innovation. We look both within firms using

measures of value add, and, perhaps more importantly, we look at how the total surplus generated by public innovations is distributed across workers, capitalists, and consumers. We use the Compustat dataset as our focal sample of firms. While this is a particular slice of the US economy, one of the many advantages of this dataset is that it is the exact data being used in recent work on markup estimation (e.g., De Loecker et al. 2020) and so we can lean on the frontier of empirical work and obtain firm-year specific estimates of markups. These markup estimates allow us to investigate changes in consumer surplus, an often overlooked aspect of innovation economics.

The final ingredient is a way to map technology-level public R&D windfalls to firms. We use firms' patenting records to measure where they sit in technology space. Thus, a firm is deemed more exposed to DOE funding in a technology when a larger share of its patent stock is in that same technology. Because these exposure shares may themselves be endogenous to public R&D investments, we use the recentered shift-share logic of Borusyak and Hull (2023).

We find that public innovation exposure raises firm activity across several margins. In general, a one-percent increase in exposure raises sales by about 0.50 percent. Labor compensation rises by about 0.47 percent, intermediate spending by about 0.51 percent, and R&D spending by about 0.34 percent. Markups fall by about 0.04 percent. These estimates are stable across choices about R&D-stock depreciation, implementation lags, first-stage controls, and first-stage learners. Conventional clustered standard errors and randomization inference give the same qualitative conclusion: exposure to DOE innovation expands firms and modestly lowers markups.

We use these estimates to recover the incidence of induced value added. Within the firm, public R&D generates gains for both workers and capital owners. Of each marginal dollar of value added created by public innovation, labor receives about 40 cents, R&D suppliers receive about 8 cents, and capital/profits receive about 50 cents. These marginal shares are close to average value-added shares in the sample, with a slight bias away from workers and R&D and towards capitalists. Thus, public R&D raises value added, but it does not sharply change the way value added is split inside the firm.

The total-surplus calculation adds consumers. Because public R&D lowers markups, some of the gain reaches buyers rather than staying inside the firm. Using the markup-based demand mapping, consumers receive roughly 70 percent of marginal total surplus (slightly more than their average in-sample surplus share), while capitalists, labor and R&D suppliers all receive slightly less than their average in-sample surplus share.

These estimates have limits. They are incidence estimates, not estimates of the aggregate social rate of return to public R&D. The sample is publicly traded firms that can be linked to patents and Compustat outcomes, not the whole economy. The consumer-surplus calculation depends on interpreting markups as demand-side wedges. The producer-surplus decomposition requires assumptions about factor supply elasticities. We show sensitivity to those assumptions. The central result is that public R&D creates gains that are spread across firms, input suppliers, and consumers, rather than being captured by a single group.

Related literature and contributions

Our paper speaks to several strands of literature. First, we contribute to the body of work measuring the social returns to innovation and public R&D. A long tradition in economics emphasizes that innovation generates spillovers and that the social returns to R&D may substantially exceed the private returns captured by inventors or firms (Arrow 1962; Griliches 1979; Jones and Williams 1998; Bloom et al. 2013). Recent work has made substantial progress quantifying these returns. Jones and Summers (2022) argue that the social returns to innovation are very large even under conservative assumptions, while Fieldhouse and Mertens (2026a) estimate large effects of federal non-defense R&D on private sector productivity and implied social returns to public R&D. Other work studies large public R&D episodes and institutions, including the Space Race, World War II research, NIH biomedical research, and agricultural research programs (Azoulay et al. 2019; Gross and Sampat 2023; Kantor and Whalley 2025; Akerman et al. 2026). These papers ask whether public R&D creates value and estimate aggregate or local productivity effects. We ask a complementary question to study the distributional consequences: conditional on public R&D creating value, who captures that value?

Second, we contribute to work on innovation rents, rent-sharing, and the distribution of gains from technological change. Evidence so far indicates that firm-specific productivity and innovation shocks generate rents that are shared unevenly within firms (Van Reenen 1996; Aghion et al. 2018; Kline et al. 2019; Alam and Diegmann 2026). For example, Kline et al. (2019) show that valuable patent allowances increase firm productivity and worker earnings, with workers capturing roughly 30 cents of each dollar of patent-induced surplus and gains concentrated among incumbent, higher-earning workers. More broadly, a growing literature links the decline in labor shares to technological change, reallocation toward superstar firms, rising markups, and changes in rent-sharing institutions (Karabarbounis and Neiman 2014; Autor et al. 2020; De Loecker et al. 2020; Kehrig and Vincent

2021). We build on this literature but shift the object of study from privately generated innovation rents to publicly generated innovation rents. This distinction matters because public R&D is financed by taxpayers and justified by social returns, yet those returns may ultimately accrue to workers through wages and employment, to capital owners through profits and firm value, or to consumers through lower markups and prices. In doing so, we connect the study of rent-sharing within innovative firms to the welfare and distributional evaluation of public innovation policy.

Third, we contribute to the broader public finance literature on the welfare evaluation and incidence of government policy. A large body of work studies who bears the burden of taxes and who benefits from government spending. Recent work has developed unified tools for comparing the welfare value of government policies by combining causal estimates of policy impacts with measures of willingness to pay and net fiscal costs (Hendren and Sprung-Keyser 2020; Finkelstein and Hendren 2020). A related incidence literature studies how the burdens and benefits of business taxation are distributed across workers, landowners, and firm owners, emphasizing that equilibrium responses in wages, rents, firm location, and profits are central to welfare analysis (Suarez Serrato and Zidar 2016; Fuest et al. 2018). We extend this incidence logic to innovation policy. Rather than asking who bears the burden of a tax or who receives a transfer, we ask who captures the gains from a public investment that creates new economic surplus. This distinction is important because public R&D is often justified by its ability to generate value beyond the direct recipient of government funds, especially when the direct recipient is in the private sector. Our framework therefore brings incidence analysis to a domain in which the central policy question is not only whether government spending creates value, but how that value is divided among workers, owners of capital, firms, and consumers.

Fourth, we contribute to the empirical literature on public R&D programs, grants, and private sector innovation. A growing set of quasi-experimental studies shows that public subsidies can stimulate private R&D, innovation, and firm growth (Bronzini and Iachini 2014; Howell 2017; Azoulay et al. 2019; Santoleri et al. 2024; Pless 2026). This literature has been central in establishing that public R&D and innovation subsidies affect firms' innovative and commercial outcomes. We focus on moving beyond estimating firm-level treatment effects to measuring the incidence of the gains those treatments generate. We therefore complement work on whether public R&D "works" by asking how the gains from public R&D are allocated across economic actors.

Methodologically, we provide a new framework for using political variation in public in-

vestment when budget requests are themselves endogenous to expected returns. Existing approaches often use appropriations shocks, funding rules, or grant score discontinuities. Our framework formalizes a distinct source of identifying variation: residual appropriations after conditioning on the information revealed by agency requests. This distinction matters because requests may embed agencies' private information about technological opportunity, while the remaining variation in appropriations can reflect legislative wedges, bargaining, or budgetary shocks that shift investment without shifting underlying productivity. In our setting, this approach yields technology-year DOE funding windfalls that can be linked to firms through predetermined technology exposure. More broadly, the framework provides a way to use political bargaining over public investment to estimate production effects in settings where funding requests themselves may contain information about expected returns.

Finally, we contribute to the literature on energy innovation, mission-oriented public R&D, and the energy transition. The DOE provides an especially useful setting for studying public innovation because its appropriations are large, technology-specific, and span the full innovation ecosystem, including universities, national laboratories, private firms, demonstration projects, and applied research programs. Prior work on energy and mission-oriented innovation has shown that public R&D can relax financial frictions, redirect technological change, and shape private sector innovation (Howell 2017; Myers and Lanahan 2022; Dugoua and Gerarden 2025). We use DOE as a laboratory for studying who captures the returns when governments invest in new technologies. However, the DOE setting is also substantively important in and of itself because public investment in energy technologies is central to current debates over climate policy, industrial policy, national security, and technological competitiveness.

2. Political production functions

2.1. Overview

For notation, let X_t be flow of a variable at time t and $\underline{X}_t$ be the stock of that variable at time t (to include any discounting of historic flows). Let $x_t \equiv \log(X_t)$.

The political production function involves the stock of a single input, where A is the flow and $\underline{A}$ is the stock. Investment occurs in production units indexed in the cross-section by j and in time by t. In our application, j will index technologies and t will index years.

Investment accumulates per:

$$\underline{A}_{jt} = \delta \times \underline{A}_{j(t-1)} + A_{j(t-\ell)} \,, \tag{1}$$

which assumes an implementation lag of ℓ years and a depreciation rate of $1-\delta$. Technology-year output Y_{jt} is produced according to the following exponential model:

$$Y_{jt} = \underline{A}_{jt}^{\beta} \times \exp\left(\omega_{jt} + \epsilon_{jt}\right) \quad \text{or, in logs} \quad y_{jt} = \beta \times \underline{a}_{jt} + \omega_{jt} + \epsilon_{jt} \,, \tag{2}$$

where β is the output elasticity and $\omega_{jt} + \epsilon_{jt}$ is a composite productivity term where ω_{jt} is known to (some) of the decision makers when making investments and ϵ_{jt} are i.i.d shocks. As is typically the case, the challenge of identifying β is that ω_{jt} may not equal zero and may not be observable to the econometrician. Thus, if investments A are (positively or negatively) correlated with ω, then naïve estimates of Equation (2) will be biased (upward or downward).

Below, we outline a model of political investment in production functions. Our goal is to identify the restrictions that can be placed on the investment process so that we can identify variation in investments that is orthogonal to productivity and use that to obtain an unbiased estimate of β and the productivity terms ($\omega + \epsilon$). This section avoids practical implementation details and discussion of our assumptions as applied to our specific empirical context. We highlight those points in the following section.

2.2. Investment game setup

The investment game has two stages: first, an *executive* receives some common and private information about the production environment and submits a requested investment level to a *legislator*; second, after receiving the request, the legislator receives some common and private information about the production environment and decides on the final investment level, after which stock accumulation and production occurs mechanically as described in Equations (1) and (2). For simplicity, we assume implementation lags of $\ell = 1$.

Formally, first the executive requests an investment flow R_{jt}, which implies a requested stock per:

$$\underline{R}_{j(t+1)} = \delta \times \underline{A}_{jt} + R_{jt} \,. \tag{3}$$

The legislator determines the *actual* investment flows, A_{jt}, and realized stock, $\underline{A}_{jt}$, per Equation 1. Thus, even while the agents may practically choose investment flows since

costs are paid on the flow basis, we can evaluate their implicit choices of investment stocks $\underline{R}_{j(t+1)}$ (the executive's choice) and $\underline{A}_{j(t+1)}$ (the legislator's choice).

The executive's objective function is:

$$\max_{\underline{R}_{j(t+1)}} \quad \mathbb{E}\Big[\nu^E_{jt} \times Y_{j(t+1)}\left(\underline{A}_{j(t+1)}\left(\underline{R}_{j(t+1)}\right)\right) - \zeta^E_{jt} \times A_{jt}\left(\underline{R}_{j(t+1)}\right) \mid \mathcal{I}^E_{jt}\Big], \tag{4}$$

where ν^E_{jt} are explicit or implicit prices that describe their constant marginal value of output, ζ^E_{jt} is their marginal cost of inputs, and $\mathcal{I}^E_{jt}$ is their information set.[1] Likewise, the legislator's objective function is:

$$\max_{\underline{A}_{j(t+1)}} \quad \mathbb{E}\Big[\nu^L_{jt} \times Y_{j(t+1)}\left(\underline{A}_{j(t+1)}\right) - \zeta^L_{jt} \times A_{jt}(\underline{A}_{j(t+1)}) \mid \underline{R}_{j(t+1)}, \mathcal{I}^L_{jt}\Big] \tag{5}$$

where ν^L_{jt} are prices that describe their marginal value of output, ζ^L_{jt} is their marginal cost of inputs, and $\mathcal{I}^L_{jt}$ is their information set. These objective functions are written to reflect the fact that the executive's request can signal information and, in turn, influence the legislator's decision.[2]

In general, the executive's optimal choice (i.e., their implicit requested log stock) and the legislator's optimal choice (i.e., the actual log stock given the executive's request), can be written as:

$$\underline{r}_{j(t+1)} = \mathcal{R}_{jt}(\omega_{j(t+1)}; \mathcal{I}^E_{jt}) \tag{6}$$

$$\underline{a}_{j(t+1)} = \mathcal{A}_{jt}(\mathcal{R}_{jt}(\omega_{j(t+1)}; \mathcal{I}^E_{jt}), \omega_{j(t+1)}; \mathcal{I}^L_{jt}), \tag{7}$$

where the functions $\mathcal{R}_{jt}(\cdot)$ and $\mathcal{A}_{jt}(\cdot)$ describe each party's equilibrium policy functions given the structure of the game involving these investment decisions.

2.3. Key assumptions

We impose three key restrictions on the investment game and information structure to facilitate the identification of "noise" in investment — investment uncorrelated with productivity — that we will use to identify the focal output elasticity β.

[1] More general pricing or utility functions over output are feasible in our framework.

[2] In the executive's objective function, we write the realized stock of investment as a function of the requested stock (i.e., $\underline{A}_{jt}(\underline{R}_{jt})$); and in the legislator's objective function, we condition their beliefs on the executives request $\underline{R}_{jt}$.

Assumption: Stable game — *Conditional on some observables W_{jt}, the equilibrium request and investment rules are common across production units and continuous in their arguments.* Given this assumption, the policy functions can be written as:

$$\underline{r}_{j(t+1)} = \mathcal{R}(\omega_{j(t+1)}, W_{jt}; \mathcal{I}^E_{jt}) \tag{8}$$

$$\underline{a}_{j(t+1)} = \mathcal{A}(\mathcal{R}(\omega_{j(t+1)}, W_{jt}; \mathcal{I}^E_{jt}), \omega_{j(t+1)}, W_{jt}; \mathcal{I}^L_{jt}) \,, \tag{9}$$

and can be empirically estimated with standard techniques. Note that, unlike Equations (6) and (7), these policy functions are not indexed by jt. This rules out the case in which each production unit has its own unrestricted equilibrium policy functions, which is what would otherwise prevent us from separating investment noise from production-unit-specific equilibria. The restriction is still relatively flexible because W_{jt} can include technology and year indicators, so the common policy functions still accommodate j- and t-specific shifters.

Assumption: Private wedges or investment noise — *Either (i) legislator-side payoff primitives C_L and/or P_L vary conditional on $(\underline{r}, W)$, or (ii) there is idiosyncratic noise in the final investments.* In the absence of noise, this requires some payoff primitive affecting the legislator's final choice, C_L or P_L, to differ from the executive's expectations when the request is chosen. To the extent there is truly noise in the final investment (i.e., idiosyncratic forces that lead to sub-optimal investments by the legislator), then variation in private wedges is not necessary to generate windfalls.[3] Either of these conditions ensures that:

$$\text{Var}(\underline{a}_{j(t+1)} \mid \underline{r}_{j(t+1)}, W_{jt}) > 0. \tag{10}$$

Assumption: Request sufficiency — *The legislator has no private information about productivity beyond what can be inferred from the executive's request.* Equivalently, after conditioning on the requested stock and observables, the legislator's decision is orthogonal to productivity. Given this assumption, we can simplify the legislator's policy function:

$$\begin{aligned} \underline{a}_{j(t+1)} &= \mathcal{A}(\mathcal{R}(\omega_{j(t+1)}, W_{jt}; \mathcal{I}^E_{jt}), \omega_{j(t+1)}, W_{jt}; \mathcal{I}^L_{jt}) \\ &= \mathcal{A}(\mathcal{R}(\omega_{j(t+1)}, W_{jt}; \mathcal{I}^E_{jt}), W_{jt}; \mathcal{I}^L_{jt}) \\ &= \mathcal{A}(\underline{r}_{j(t+1)}, W_{jt}; \mathcal{I}^L_{jt}) \end{aligned} \tag{11}$$

[3] For example, an experiment could randomize final investment decisions to generate variation conditional on requests that is idiosyncratic by construction.

This assumption allows requests to be endogenous to productivity. The executive may observe (or estimate) $\omega_{j(t+1)}$ and choose higher requested stocks for higher-productivity technologies. However, the legislator cannot have more information on productivity than the executive.

2.4. Windfall identification

The three key assumptions deliver the conditions needed to use windfall stock shocks as an instrument: relevance, exclusion, and exogeneity. The *Stable game* assumption implies that there is a common conditional mean function mapping requested log stocks and observables into realized log stocks, which allows us to define and estimate windfall shocks. The *Private wedges* assumption implies relevance: realized log stocks vary around their request-predicted conditional mean. The *Request sufficiency* assumption implies exclusion and exogeneity: after conditioning on requested log stocks and observables, the residual variation in realized log stocks is not information about productivity and is therefore orthogonal to the structural productivity term in the outcome equation. We formalize this below.

Let us index the request-predicted stock, the windfall, and output by the common period t in which the realized stock is dated. Define the request-predicted conditional mean of realized log stocks as:

$$\phi(\underline{r}_{jt}, W_{jt}) \equiv \mathbb{E}[\underline{a}_{jt} \mid \underline{r}_{jt}, W_{jt}] \,. \tag{12}$$

Then define the windfall stock shock as the residual component of realized log stocks:

$$\eta_{jt} \equiv \underline{a}_{jt} - \phi(\underline{r}_{jt}, W_{jt}) \,, \tag{13}$$

where $\mathbb{E}[\eta_{jt} \mid \underline{r}_{jt}, W_{jt}] = 0$ by construction. The *Stable game* assumption makes $\phi(\cdot)$ a common object that can be estimated by pooling observations, while the continuity restriction justifies approximating it flexibly. The *Private wedges* assumption gives relevance by requiring that this residual component is not degenerate:

$$\mathrm{Var}(\eta_{jt} \mid \underline{r}_{jt}, W_{jt}) > 0 \,. \tag{14}$$

On the exclusion restriction, the *Request sufficiency* assumption can be stated as the conditional mean restriction:

$$\mathbb{E}[\underline{a}_{jt} \mid \underline{r}_{jt}, W_{jt}, \omega_{jt}] = \mathbb{E}[\underline{a}_{jt} \mid \underline{r}_{jt}, W_{jt}] \,. \tag{15}$$

This condition allows requested stocks to be endogenous to productivity. It only requires that, after conditioning on requested stocks and observables, the remaining variation in realized stocks is not additional information about productivity. Using (12), (13), and (15):

$$\begin{aligned}\mathbb{E}[\eta_{jt} \mid \underline{r}_{jt}, W_{jt}, \omega_{jt}] &= \mathbb{E}[\underline{a}_{jt} \mid \underline{r}_{jt}, W_{jt}, \omega_{jt}] - \phi(\underline{r}_{jt}, W_{jt}) \\ &= 0 \, .\end{aligned} \tag{16}$$

Then, by the law of iterated expectations:

$$\begin{aligned}\mathbb{E}[\eta_{jt} \mid \omega_{jt}] &= \mathbb{E}\Big[\mathbb{E}[\eta_{jt} \mid \underline{r}_{jt}, W_{jt}, \omega_{jt}] \mid \omega_{jt}\Big] \\ &= 0 \, .\end{aligned} \tag{17}$$

Therefore:

$$\mathbb{E}[\eta_{jt} \times \omega_{jt}] = 0 \, . \tag{18}$$

Thus, requested stocks may be correlated with productivity, but windfall stock shocks are not. The identifying variation is the component of realized log stocks that remains after conditioning on requested log stocks and observables.

With data on $(\underline{a}_{jt}, \underline{r}_{jt}, W_{jt})$, we can estimate $\phi(\cdot)$ using, for instance, semi-parametric or machine-learning methodologies and construct:

$$\widehat{\eta}_{jt} = \underline{a}_{jt} - \widehat{\phi}(\underline{r}_{jt}, W_{jt}) \, . \tag{19}$$

We can then use the estimated windfall stock shock, $\widehat{\eta}_{jt}$, as an instrument in a 2SLS equation:

$$\underline{a}_{jt} = \alpha \times \widehat{\eta}_{jt} + u_{jt} \tag{20}$$

$$y_{jt} = \beta \times \underline{a}_{jt} + (\omega_{jt} + \epsilon_{jt}) \, , \tag{21}$$

where we may include covariates in this 2SLS regression as well, but note that the covariates included here need not be the same as the covariates W_{jt} used in the estimation of windfalls.

Because $\widehat{\eta}_{jt}$ is generated from an estimated first-stage conditional mean, conventional robust 2SLS standard errors do not account for the sampling error in $\widehat{\phi}(\cdot)$ and should be read as a lower bound on the true uncertainty. We therefore report them as a benchmark but treat randomization inference, based on the assignment model in the spirit of Borusyak and Hull (2023), as our preferred inference procedure. The randomization procedure

holds requested stocks, controls, and support fixed, permutes the estimated residual stock shocks within the relevant assignment cells, and recomputes the 2SLS estimate under each draw. This inference procedure targets the variation isolated by the model: realized-stock variation conditional on requested stocks.

2.5. A strategic communication microfoundation

This subsection gives a simple microfoundation for the empirical model above, which is in the spirit of models of strategic communication (e.g., Kartik et al. 2007). Suppress all subscripts and consider a cross-sectional version of the model. Let $\underline{A}_0$ denote the predetermined investment stock at the time of the executive's request. Let $\underline{R}$ denote the new requested stock and $\underline{A}$ the realized stock. The associated requested and realized flows are:

$$R = \underline{R} - \underline{A}_0, \qquad A = \underline{A} - \underline{A}_0.$$

We focus on the interior where $\underline{R} > \underline{A}_0$ and $\underline{A} > \underline{A}_0$ and normalize the price of inputs to one for both parties, $C_E = C_L = 1$, with no lags or discounting of investment. Following the notation above, lowercase variables denote logs, so $\underline{a} \equiv \log(\underline{A})$ and $\underline{r} \equiv \log(\underline{R})$. The production function is:

$$Y(\underline{A}; \omega) = \sqrt{\underline{A}} \times \exp(\omega), \qquad \text{or in logs,} \qquad y(\underline{a}; \omega) = \frac{1}{2} \times \underline{a} + \omega.$$

The executive has known preferences with marginal utility of output normalized to one:

$$P_E = 1.$$

The legislator's marginal utility of output is uncertain at the time of the executive's signal:

$$P_L = 1 + \nu > 0, \qquad \mathbb{E}[\nu] = 0, \qquad \mathrm{Var}(\nu) = \sigma_\nu^2, \qquad \nu \perp \omega.$$

Thus:

$$\mathbb{E}[P_L] = 1, \qquad \mathbb{E}\left[P_L^2\right] = 1 + \sigma_\nu^2.$$

The timing of the game is as follows: (*i*) the executive observes productivity ω and chooses a requested stock $\underline{R}$; (*ii*) the legislator observes $\underline{R}$, makes an inference $\widehat{\omega}(\underline{R})$, observes their private value shock ν, and chooses the realized stock $\underline{A}$.

The legislator observes ν before choosing the realized stock, but has no direct knowledge

of productivity. The executive observes productivity ω and knows the distribution of ν, but does not observe the realization of ν when choosing the request.

By backward induction, given a request, an inferred productivity $\widehat{\omega}$, and a realized value $P_L = 1 + \nu$, the legislator solves:

$$\max_{\underline{A}} \quad P_L \times \sqrt{\underline{A}} \times \exp(\widehat{\omega}) - (\underline{A} - \underline{A}_0).$$

The legislator's optimal realized stock is:

$$\underline{A}^*(\widehat{\omega}, \nu) = \frac{(1+\nu)^2 \times \exp(2 \times \widehat{\omega})}{4}.$$

Because the request has no direct payoff except through the legislator's inference, the executive's problem can be written as choosing the inference $\widehat{\omega}$ that maximizes their payoff. Without any lying penalty, the executive solves:

$$\max_{\widehat{\omega}} \quad \mathbb{E}_\nu\left[\sqrt{\underline{A}^*(\widehat{\omega}, \nu)} \times \exp(\omega) - (\underline{A}^*(\widehat{\omega}, \nu) - \underline{A}_0)\right].$$

Substituting in the legislator's rule into the executive's expectation gives:

$$\max_{\widehat{\omega}} \quad \frac{\exp(\omega + \widehat{\omega})}{2} - \frac{\left(1 + \sigma_\nu^2\right) \times \exp(2 \times \widehat{\omega})}{4} + \underline{A}_0.$$

The first-order condition implies:

$$\widehat{\omega}^*(\omega) = \omega - \log\left(1 + \sigma_\nu^2\right),$$

which, in turn, yields the mapping the executive uses to choose their request given their observation of productivity:

$$\underline{R}^*(\omega) = \frac{\exp(2 \times \omega)}{4 \times \left(1 + \sigma_\nu^2\right)}.$$

Therefore the log requested stock is:

$$\underline{r}^*(\omega) = 2 \times \omega - \log(4) - \log\left(1 + \sigma_\nu^2\right),$$

which is monotone in productivity. For any observed requested stock, the legislators

inference is thus given by[4]:

$$\widehat{\omega}(\underline{R}) = \frac{1}{2} \times \log\left(\frac{4 \times \underline{R}}{1 + \sigma_{\nu}^2}\right).$$

After observing $\underline{R}$ and ν, the legislator chooses:

$$\underline{A} = \frac{(1+\nu)^2}{1+\sigma_{\nu}^2}\underline{R}.$$

Taking logs gives:

$$\underline{a} = \underline{r} + \log\left(\frac{(1+\nu)^2}{1+\sigma_{\nu}^2}\right).$$

With our identification strategy described above in mind, define:

$$\phi(\underline{r}) \equiv \mathbb{E}[\underline{a} \mid \underline{r}].$$

Then:

$$\phi(\underline{r}) = \underline{r} + \mathbb{E}_{\widetilde{\nu}}\left[\log\left(\frac{(1+\widetilde{\nu})^2}{1+\sigma_{\nu}^2}\right)\right],$$

where $\widetilde{\nu}$ is an independent draw from the same distribution as ν. The residual log-stock shock is:

$$\eta = \underline{a} - \phi(\underline{r}) = \log\left(\frac{(1+\nu)^2}{1+\sigma_{\nu}^2}\right) - \mathbb{E}_{\widetilde{\nu}}\left[\log\left(\frac{(1+\widetilde{\nu})^2}{1+\sigma_{\nu}^2}\right)\right].$$

This expression shows the source of identifying variation directly. Conditional on the requested stock, residual variation in realized log stocks is generated only by the legislator's private value shock ν. If ν has no uncertainty, then $\eta = 0$.

Because $\nu \perp \omega$ and ν is realized after the request:

$$\mathbb{E}[\eta \mid \underline{r}, \omega] = 0.$$

Therefore:

$$\mathbb{E}[\eta \times \omega] = \mathbb{E}[\omega \times \mathbb{E}[\eta \mid \underline{r}, \omega]] = 0.$$

Adding a direct lying penalty changes the mapping from productivity ω to the requested

[4]This decoding rule should be read as the action-relevant signaling rule rather than literal Bayesian inversion of a separating request rule.

stock $\underline{R}$, because it changes the executive's optimal induced inference. However, as long as the penalty either (*i*) does not depend on the legislator's private value shock ν, and/or (*ii*) is orthogonal to productivity, it still yields exogenous variation in investment.

3. R&D investment at the US Department of Energy

3.1. Budget data: Requests and appropriations

We assemble the budget panel from the Department of Energy's annual *Congressional Budget Justification* documents, supplemented for early-year fossil-fuel programs by the corresponding Department of the Interior justifications. For each program and fiscal year, these documents report both the funding the agency requested and the amount Congress ultimately appropriated. Because the justifications are largely unstructured PDFs whose formats and program labels change over time, we manually extracted and harmonized the line-item data, retaining R&D activities and excluding administrative program direction and pure deployment. The historical documents were obtained from the HathiTrust Digital Library, the `govinfo.gov` repository,[5] and the DOE Office of the Chief Financial Officer's website (for fiscal years 2005 through 2023).[6]

Our request series runs from fiscal year 1989 through 2023 and our appropriations series from fiscal year 1987 through 2022; because the funding windfall we exploit is the gap between an appropriation and its corresponding request, the estimation panel effectively begins in fiscal year 1989, the first year in which DOE budget requests are available. As a further check on our hand-collected figures, we will cross-validate the aggregate of our series against additional published sources such as the American Association for the Advancement of Science (AAAS) *Historical Trends in Federal R&D*, which independently compiles federal R&D budget authority by agency, to verify that the sum of our technology-level DOE appropriations tracks an external benchmark for total DOE R&D.

We harmonize program lines across years into a consistent set of 14 technology categories indexed by j: fossil energy (coal, natural gas, oil, and carbon capture and storage); nuclear fission; renewables (solar, wind, geothermal, bioenergy, and hydropower and marine energy); grid systems; energy storage; hydrogen and fuel cells; and batteries and electric vehicles. This taxonomy is also used to measure firm exposure in Section 6.

[5] https://www.govinfo.gov/app/details/CHRG-106hhrg63101/CHRG-106hhrg63101

[6] https://www.energy.gov/cfo/listings/budget-justification-supporting-documents

Harmonizing the budget lines across years is non-trivial because DOE program structures and labels change frequently. We track each technology at the most granular level the documents allow and split lines that bundle distinct technologies: for example, "Wind Energy" was at times reported as a sub-category of "Solar Energy" in early requests, which we separate into independent series, and we likewise carve "Fuel Cell" research out of the Coal and Natural Gas ("Fossil Energy") lines wherever the budget detail permits. Conversely, when DOE reporting collapses two programs we cannot un-merge them, so we treat them as a single category throughout: hydropower and marine energy, which DOE combined into one "Water Power" line for fiscal years 2010–2013, and batteries and electric vehicles, which merged under "Vehicle Technologies" in 2013. We also treat hydrogen and fuel cells as one category given their substantial programmatic overlap. On the appropriations side, we prefer the ex-post "actual" or "adjusted" figures reported in later budget documents over the initial "enacted" values, since the former reflect rescissions, reallocations, and the use of prior-year balances; for requests, we use the most recent restatement of each fiscal year's figure to capture the final presidential request. Finally, we treat the one-time American Recovery and Reinvestment Act (ARRA, 2009) surge conservatively: because its appropriations blended R&D with manufacturing and deployment, our baseline attributes only 20% of the "Transportation Electrification" funds and 30% of the "Advanced Battery Manufacturing" funds to R&D, assigning the remainder to non-research deployment and manufacturing support.

We construct a technology-year panel of DOE budget activity from the annual request and appropriation records. The unit of observation is a technology j in fiscal year t. All dollar values are deflated to 2024 dollars using the GDP deflator.

Let R_{jt} denote the requested flow of investment into each technology-year per the agency's budget request, and A_{jt} the actual appropriation flow. We use these flows to construct requested and actual investment stocks, exactly as in Equations (1) and (3):

$$\underline{R}_{jt} = \delta \times \underline{A}_{j(t-1)} + R_{j(t-\ell)}, \qquad \underline{A}_{jt} = \delta \times \underline{A}_{j(t-1)} + A_{j(t-\ell)}$$

Table 1 summarizes the DOE budget panel. Average requested investment flows are about $147 million per technology-year, while average appropriation flows are about $164 million. The fact that appropriations generally exceed requests suggests that the nature of the strategic game between the two branches often yields a particular kind of equilibrium where executives "under-request". However, we would emphasize that our identification strategy does not assume that all of the difference between the request and appropriation

is uncorrelated with productivity; rather, the appropriations that are *unpredictable* given requests will be the focal variation we try to isolate.

TABLE 1. Technology-year investment flows

	Mean	SD	p25	p50	p75
Request flow	147.1	176.6	35.0	90.0	192.3
Appropriation flow	163.5	194.9	50.7	107.9	196.9

Notes: The table reports request and appropriation flows for the DOE technology-year budget panel. Amounts are deflated to real 2024 dollars and reported in millions.

3.2. Legislative windfalls: Case studies

The identifying content of our design rests most importantly on the *Request sufficiency* assumption of Section 2: conditional on the executive's request and observable political controls, residual appropriations reflect legislative wedges that are orthogonal to a technology's underlying productivity. Before turning to estimation, we make this assumption concrete with two episodes in which the gap between requested and appropriated funding was driven by idiosyncratic political circumstances rather than by new information about a technology's promise.

Unconventional fossil energy. Our first episode involves a program the executive branch did not request at all. In its fiscal year 2010 budget request, the DOE proposed no dedicated funding for unconventional (oil and gas) fossil R&D, yet Congress appropriated $20 million for that year to a new "Unconventional Fossil Energy Technologies" program. The impetus came from legislators representing oil- and gas-producing states rather than from any new departmental assessment of the technology's scientific promise; the gap between the (zero) request and the $20 million appropriation is therefore a clean example of a legislative wedge orthogonal to DOE's own opportunity ranking.

Nuclear fission under Senator Domenici. Our second episode is a more persistent, institutionally driven wedge. Senator Pete Domenici of New Mexico was the senior Republican on the Senate Energy and Water Appropriations Subcommittee from 1997 to 2006, chairing it in the years his party held the Senate majority. New Mexico hosts the Los Alamos and Sandia National Laboratories, whose large nuclear programs gave Domenici a strong constituency interest in elevating nuclear funding regardless of its relative scientific merit in any given year. Consistent with this, nuclear-fission appropriations increasingly exceeded the levels DOE requested over his tenure: after an initial period in the late 1990s in which appropriations ran somewhat below requests, the windfall turned positive and remained

so through the 2000s. DOE may well have anticipated this support and shaded its nuclear requests accordingly, so part of the wedge is predictable. Predictability, however, is not what threatens our design; correlation with productivity is. Because Domenici's support reflects New Mexico's stake in its national laboratories rather than any new assessment of nuclear's promise, these above-request appropriations are plausibly orthogonal to the productivity of nuclear R&D, which is what our identifying assumption requires.

3.3. Legislative windfalls: Empirical estimates

To obtain our windfall estimates, we estimate the conditional mean of realized stocks using:

$$\log(\underline{A}_{jt}) = \phi\left(\log\left(\underline{R}_{jt}\right), W_{jt}\right) + \eta_{jt}\,, \tag{22}$$

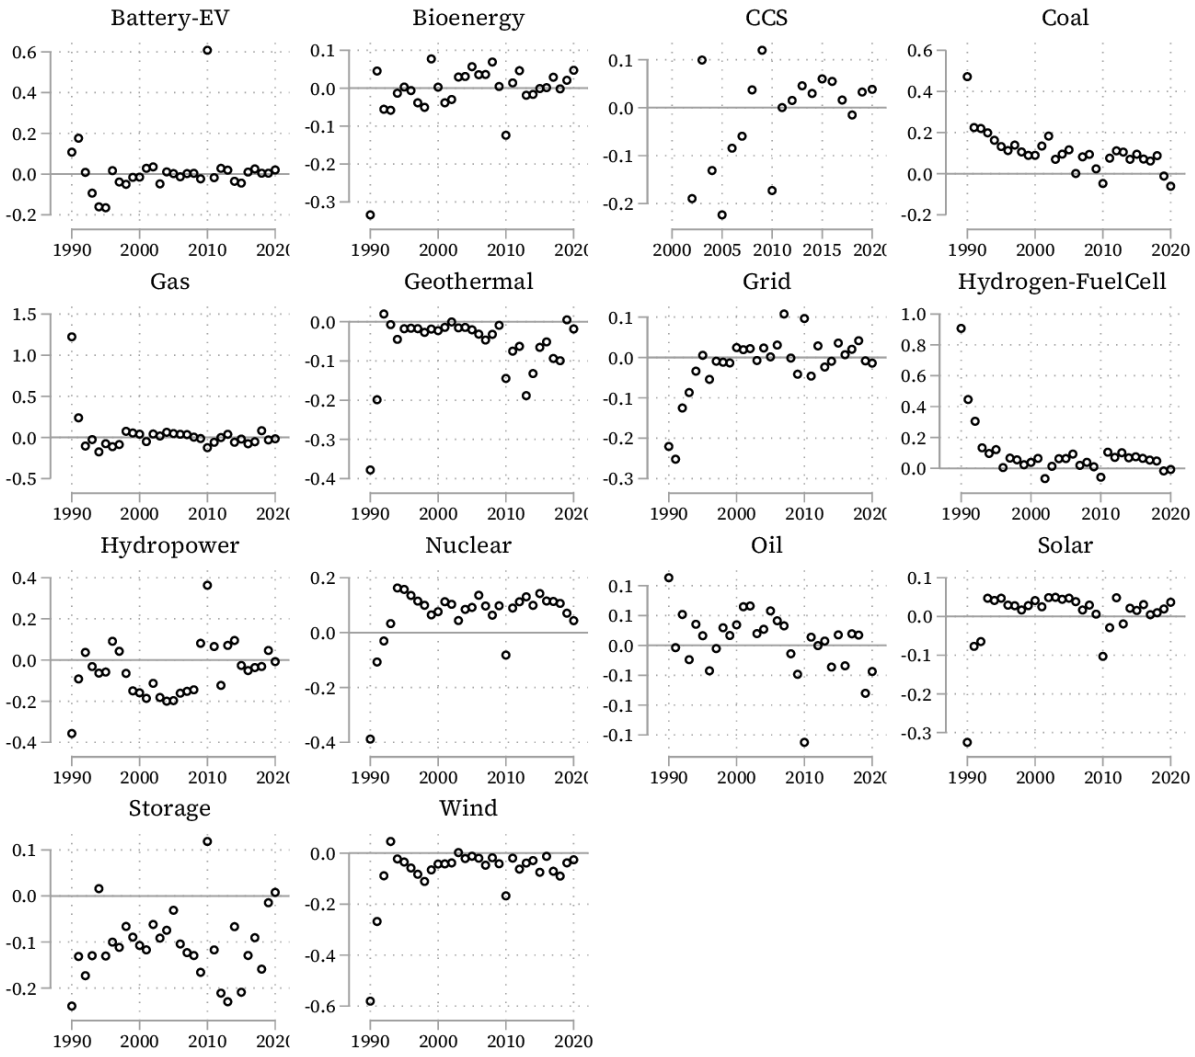
where residual $\widehat{\eta}_{jt}$ is the funding windfall. We explore the following specification choices when estimating Equation (22):

- *Discounting δ and lags ℓ.* We vary $\delta \in \{1, 0.9, 0.8\}$ and $\ell \in \{0, 1\}$
- *Vector of controls W.* We consider three nested control sets: (*i*) $\log(\underline{R}_{jt})$ and year fixed effects; (*ii*) adds a cubic spline in $\log(\underline{R}_{jt})$ and fixed effects for presidential administration interacted with clean-versus-dirty technology status; (*iii*) adds interactions of $\log(\underline{R}_{jt})$ and its spline terms with presidential administration indicators and with the administration-by-clean/dirty indicators.
- *Machine-learning estimation methods.* We estimate $\phi(\cdot)$ using regularized linear prediction using (*i*) lasso with cross-validation using the one-standard-error rule, or (*ii*) elastic net with cross-validation using the one-standard-error.

Averaging across all of these sets of specification choices, the standard deviation of $\widehat{\eta}_{jt}$ is 0.174, implying that a one-standard-deviation windfall corresponds to roughly a 18.9% change in realized DOE stock relative to predicted stock. Or, in terms of investment levels, a one-standard-deviation windfall corresponds to a technology receiving roughly $31.0 million more or less than expected, which can be benchmarked off the average observed technology-year investment level of $163.5 million.

Figure 1 shows the windfalls by technology over time, after residualizing on year fixed effects and averaging across our plausible specifications. Since the productivity of these investments is truly unobservable, we cannot test our key assumption that these windfalls are uncorrelated with productivity. However, Figure 1 illustrates that, for most technologies, there is no persistent trend in the windfalls; except for what might be referred to as a regression-to-the-mean trend whereby most technologies have particularly large (relative)

FIGURE 1. Windfalls across technologies and time



Notes: The figure plots technology-year windfall estimates, $\widehat{\eta}_{jt}$. Windfalls are residualized on year fixed effects before plotting.

windfalls in the early days of R&D funding at the DOE in our data.

4. A model of the private returns to public innovation

4.1. Setup: Production, supply, and demand

Let firm i in year t produce physical output Q_{it} at price P_{Qit}. The firm uses labor L_{it}, capital K_{it}, intermediate inputs M_{it}, and research and development R_{it}. We set aside R&D as a separate factor throughout the following exposition, because it will behave exactly as

labor does. We let $X \in \{L, K, M, R\}$ index the four factors. The physical production function is:

$$Q_{it} \equiv Q_{it}\left(L_{it}, K_{it}, M_{it}, R_{it}\right) \ . \tag{23}$$

Let P_{Xit} be the price of input X so total expenditures on each factor are:

$$C_{Xit} \equiv P_{Xit} \times X_{it} \ , \tag{24}$$

Input supply schedules are isoelastic:

$$P_{Xit} = \sigma_{Xit} \times X_{it}^{1/\varepsilon_{Xit}} \ , \tag{25}$$

where σ_{Xit} are supply shifters and ε_{Xit} is the supply elasticity. Each firm faces residual isoelastic demand for its output given by the inverse demand function:

$$P_{Qit} = \sigma_{Dit} \times Q_{it}^{-1/\varepsilon_{Dit}} \ , \tag{26}$$

where σ_{Dit} is a demand shifter and ε_{Dit} is the demand elasticity. In input and output markets, a finite ε implies firms have monopsony or monopoly power.

4.2. Value and surplus metrics

A firm's value added is an accounting concept given by revenue less intermediate inputs:

$$VA_{it} \equiv P_{Qit} \times Q_{it} - C_{Mit} \ . \tag{27}$$

Because intermediate inputs exclude R&D, value added comprises the claims of labor, R&D, and capital-and-profits. Labor's share of value added at each firm is also an accounting expression:

$$\mathcal{S}_{Lit}^{VA} \equiv \frac{C_{Lit}}{VA_{it}} = \frac{P_{Lit} \times L_{it}}{P_{Qit} \times Q_{it} - P_{Mit} \times M_{it}} \ . \tag{28}$$

To solve for economic surplus metrics, note that for each input X, the isoelastic supply curve in (25) implies the firm's input-suppliers' producer surplus is:

$$PS_{Xit} \equiv C_{Xit} - \int_0^{X_{it}} \sigma_{Xit} \times u^{1/\varepsilon_{Xit}} du = \frac{1}{\varepsilon_{Xit} + 1} \times C_{Xit} \ . \tag{29}$$

Since we do not fully observe firm-to-firm input-output links in our data, we work with a value-added boundary of producer surplus. Letting profits be given by $\Pi_{it} \equiv P_{Qit} \times Q_{it} -$

$C_{Lit} - C_{Kit} - C_{Mit}$, we define firm-wide producer surplus as:

$$PS_{it} \equiv \Pi_{it} + PS_{Lit} + PS_{Kit} + PS_{Rit} = VA_{it} - \frac{\varepsilon_{Lit}}{\varepsilon_{Lit}+1} \times C_{Lit} - \frac{\varepsilon_{Kit}}{\varepsilon_{Kit}+1} \times C_{Kit} \,. \tag{30}$$

Thus, the labor share of a firm's producer surplus is:

$$\mathcal{S}_{Lit}^{PS} \equiv \frac{PS_{Lit}}{PS_{it}} = \frac{\frac{1}{\varepsilon_{Lit}+1} \times C_{Lit}}{PS_{it}} \,, \tag{31}$$

and likewise for the other factors. With isoelastic demand, consumers' surplus is:

$$CS_{it} \equiv \frac{P_{Qit} \times Q_{it}}{\varepsilon_{Dit} - 1} \,, \tag{32}$$

which can be used to solve for total surplus ($PS_{it} + CS_{it}$), and labor's and consumers' shares thereof.

4.3. Empirical model and shares of value-add

In general, publicly funded innovation can affect a factor's share of value added (or surplus) through three sets of channels:

- product demand level (σ_{Dit}) or elasticity (ε_{Dit});
- factor productivity (level shifters in $\mathcal{Q}_{it}$), returns-to-scale or substitutability (curvatures in $\mathcal{Q}_{it}$);
- factor supply level (σ_{Xit}) or elasticity (ε_{Xit}).

Let $\{\sigma, \varepsilon, \theta\}$ be the set of parameters that govern input supply, production, and output demand functions, and B_{it} be a measure of public innovation (i.e., the stock of investments up to year t in R&D relevant to firm i). A general empirical model of the reduced-form channels through which public innovation can affect producers and consumers is illustrated as a DAG in Figure 2.

Focusing first on our value-add estimands of interest, we want to know the following objects:

- labor's share of the marginal value added generated by public innovation
- the semi-elasticity of labor's share of value added with respect to public innovation

The first object answers the question: for each additional dollar of value-add generated by public innovation, how many cents are paid to labor? The second object answers a

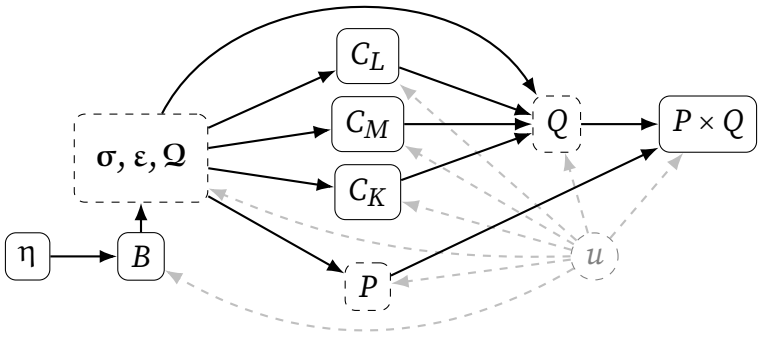


FIGURE 2. Public innovation, structural primitives, and firm outcomes

different question: do firms become more or less labor-intensive (in value-added terms) due to public innovation?

Labor's share of the marginal value-add generated by public innovation is:

$$\mathcal{M}^{VA}_{Lit} \equiv \frac{dC_{Lit} \,/\, d\log(B_{it})}{dVA_{it} \,/\, d\log(B_{it})} = \frac{C_{Lit} \times \beta_{C_L}}{P_{Qit} \times Q_{it} \times \beta_{PQ} - C_{Mit} \times \beta_{C_M}}\,, \tag{33}$$

where we use the notation that $\beta_Y \equiv d\log(Y_{it})/d\log(B_{it})$.[7] The derivatives in Equation (33), which we need to estimate, are total derivatives. Figure 2 illustrates this logic: public innovation affects the primitives of the market, and these primitives in turn affect input expenditures, prices, quantities, and revenue.

More generally, for any factor X whose compensation is a direct claim on value added (i.e., labor and R&D) the share of marginal value added it captures is

$$\mathcal{M}^{VA}_{Xit} \equiv \frac{dC_{Xit} \,/\, d\log(B_{it})}{dVA_{it} \,/\, d\log(B_{it})} = \frac{C_{Xit} \times \beta_{C_X}}{P_{Qit} \times Q_{it} \times \beta_{PQ} - C_{Mit} \times \beta_{C_M}}\,, \tag{34}$$

while capital and profits are treated jointly as a single residual claimant, $\mathcal{M}^{VA}_{\Pi it} \equiv 1 - \mathcal{M}^{VA}_{Lit}$. Because this claim is defined residually, the value-added incidence measures require no estimate of the capital-expenditure elasticity β_{C_K}. These three objects—labor, R&D, and capital/profits—are reported in Figure 4.

The estimands above involve all directed paths from B to each accounting component. For

[7] Furthermore, the semi-elasticity of labor's share of value with respect to public innovation is: $\frac{dS^{VA}_{Lit}}{d\log(B_{it})} = S^{VA}_{Lit} \times \left(\beta_{C_L} - \frac{P_{Qit} \times Q_{it} \times \beta_{PQ} - C_{Mit} \times \beta_{C_M}}{VA_{it}}\right)$.

example, the effect of B on revenues $P_{Qit} \times Q_{it}$ includes paths through product demand, productivity, input expenditures, output quantity, and price. Conditioning on variables in this pathway (e.g., input expenditures) would block some of these paths. Furthermore, it could open non-causal paths because these variables are jointly determined by public innovation B and unobserved shocks u. Therefore, the relevant elasticities should not condition on other outcome components or mechanisms.

The regressions implied by this model are separate reduced-form equations for each component. For each outcome $Y_{it} \in \{P_{Qit} \times Q_{it}, C_{Lit}, C_{Kit}, C_{Mit}, \mu_{it}\}$, we estimate β_Y from:

$$\log(Y_{it}) = \alpha + \beta_Y \times \log(B_{it}) + \Gamma_Y \times Z_{it} + u_{Yit}\,, \tag{35}$$

where Z_{it} may contain covariates; however, those covariates should be included only because of identification concerns regarding the focal parameter(s) β_Y. These covariates should not be, or be correlated with, the proximal outcomes (i.e., input levels) for reasons we discuss next.

4.4. Limitations of input-conditional regressions

It is useful to contrast the approach motivated by the model above — estimate unconditional effects of B on inputs and outputs to obtain the β parameters — with a common alternative. A typical approach to estimating the value produced by technological shocks is to estimate regressions of the form:

$$\log(P_{Qit} \times Q_{it} - C_{Mit}) = \alpha + \gamma_L \times \log(C_{Lit}) + \gamma_K \times \log(C_{Kit}) + \xi \times \log(B_{it}) + u_{it}\,, \tag{36}$$

which conditions on input expenditures (and may include other controls) and focuses on the estimate of the ξ parameter as a description of the value-add due to the innovation in B. Intuitively, if the innovation increases revenues conditional on inputs, the producer must be more productive. However, Equation (36) implicitly assumes that *none* of the value-add produced by public innovation is captured by labor because, with C_L included as a control, $\left.\frac{\partial C_L}{\partial \log(B)}\right|_{C_L} = 0$.[8]

[8] But does the *magnitude* of ξ estimated via Equation (36) describe the elasticity of value-add with respect to the innovation? In general, no. As a function of ξ estimated from Equation (36), the elasticity of value-add with respect to public innovation is: $\frac{d\log VA_{it}}{d\log(B_{it})} = \xi + \gamma_L \times \beta_{C_L} + \gamma_K \times \beta_{C_K}$, recalling that β_{C_X} is the true elasticity of expenditures on input X with respect to B. Thus, if innovation affects input levels, $dVA_{it} \,/\, d\log(B_{it}) = \xi$ requires either a knife-edge cancellation (i.e., $\gamma_L \times \beta_{C_L} + \gamma_K \times \beta_{C_K} = 0$), or it must be the case that public innovation does not affect input levels (i.e., $\beta_{C_L} = \beta_{C_K} = 0$).

4.5. Producer surplus and public innovation

Before solving for labor's share of the marginal surplus produced by public innovation, it is useful to focus first on the semi-elasticity of labor's surplus. In general, if public innovation can affect expenditures, profits, *and* the elasticities of supply, then the change in labor's surplus is:

$$\begin{aligned}\frac{\partial PS_{Lit}}{\partial \log(B_{it})} &= \frac{\partial \left(\frac{1}{\varepsilon_{Lit}+1} \times C_{Lit}\right)}{\partial \log(B_{it})} \\ &= PS_{Lit} \times \left(\beta_{C_L} - \frac{\varepsilon_{Lit}}{\varepsilon_{Lit}+1} \times \beta_{\varepsilon_L}\right),\end{aligned} \tag{37}$$

which requires estimating β_{C_L} and β_{ε_L}, where the latter is the elasticity of labor's supply elasticity with respect to public innovation. This parameter, β_{ε_L}, is exceptionally difficult to estimate as it requires exogenous variation both in public R&D *and* labor demand, since it is effectively an interaction term in the labor supply equation. We therefore leave that pursuit to future work and instead impose two assumptions on factor supply.

Assumption: No input supply elasticity effects — *Public innovation has no effect on factor supply elasticities.*. This assumption—that public innovation shifts factor supply levels but not curvatures—parallels the treatment of supply elasticities as structural primitives in the task-based technology literature (e.g., Acemoglu and Restrepo 2022): technological change influences production functions and the demand for certain types of workers, but it does not rotate the underlying factor supply curves.

Assumption: Supply Elasticities. First, we assume capital is supplied perfectly elastically, $\varepsilon_{Kit} \to \infty$. Under this assumption capital suppliers retain no surplus, since $PS_{Kit} = C_{Kit}/(\varepsilon_{Kit}+1) \to 0$, and the entire capital-side claim on producer surplus collapses to economic profit Π_{it}. As in the value-added accounting, we therefore treat capital and profits as a single residual claimant. Second, we set labor and R&D input supply elasticities to fixed baseline values motivated by the literature. In the baseline implementation, we set the firm-facing labor supply elasticity to

$$\varepsilon_{Lit} = \varepsilon_L = 3\,.$$

This value is meant to capture an upward-sloping residual labor supply curve faced by the firm, not the aggregate Frisch or Hicksian labor supply elasticity. For R&D inputs, our

baseline assigns R&D expenditure fully to R&D suppliers:

$$PS_{Rit} = C_{Rit}.$$

In the isoelastic-supply notation above, this is equivalent to setting

$$\varepsilon_{Rit} = \varepsilon_R = 0,$$

so that the R&D opportunity-cost component is zero and the entire R&D expenditure response is treated as compensation to R&D suppliers. This is a deliberately transparent benchmark: it treats R&D labor, contractors, and specialized research inputs as scarce factors that capture the induced R&D spending. We assess sensitivity to this assumption by varying ε_R over a broader range.

With these restrictions, labor's share of the marginal producer surplus generated by public innovation is:

$$\mathcal{M}_{Lit}^{PS} \equiv \frac{dPS_{Lit} \,/\, d\log(B_{it})}{dPS_{it} \,/\, d\log(B_{it})} = \frac{\frac{1}{\varepsilon_{Lit}+1} \times C_{Lit} \times \beta_{C_L}}{\Pi_{it} \times \beta_{\Pi} + \frac{1}{\varepsilon_{Lit}+1} \times C_{Lit} \times \beta_{C_L} + \frac{1}{\varepsilon_{Rit}+1} \times C_{Rit} \times \beta_{C_R}}, \tag{38}$$

Because $\Pi_{it} = P_{Qit} \times Q_{it} - C_{Mit} - C_{Lit} - C_{Kit}$ is a derived quantity, $\Pi_{it} \times \beta_{\Pi}$ is pinned down by the other elasticities.[9]

4.6. Consumer surplus and public innovation

Throughout, we interpret the firm-level markup estimated in Section 5 as the demand-side wedge, $\mu_{it} = \varepsilon_{Dit}/(\varepsilon_{Dit} - 1)$, so that the estimated markup elasticity β_{μ} maps directly into changes in the residual demand elasticity; this is what allows us to recover a consumer-side component without observing firm-level prices and quantities separately. The semi-elasticity of consumer surplus with respect to public innovation is:

$$\begin{aligned} \frac{\partial CS_{it}}{\partial \log(B_{it})} &= P_{Qit} \times Q_{it} \times \left(\frac{\beta_{PQ}}{\varepsilon_{Dit} - 1} - \frac{\varepsilon_{Dit} \times \beta_{\varepsilon_D}}{(\varepsilon_{Dit} - 1)^2} \right) \\ &= P_{Qit} \times Q_{it} \times \left((\mu_{it} - 1) \times \beta_{PQ} + \mu_{it} \times \beta_{\mu} \right), \end{aligned} \tag{39}$$

[9] Likewise, the semi-elasticity of labor's share of producer surplus is $\frac{\partial \mathcal{S}_{Lit}^{PS}}{\partial \log(B_{it})} = \mathcal{S}_{Lit}^{PS} \times \left(\beta_{C_L} - \frac{1}{PS_{it}} \times \left(\Pi_{it} \times \beta_{\Pi} + \frac{1}{\varepsilon_{Lit}+1} \times C_{Lit} \times \beta_{C_L} \right) \right)$.

where the second line expresses the effect in terms of markups $\mu_{it} = \varepsilon_{Dit}/(\varepsilon_{Dit} - 1)$. Thus, the semi-elasticity of total surplus with respect to public innovation is:

$$\begin{aligned}\frac{dTS_{it}}{d\log(B_{it})} = P_{Qit} \times Q_{it} \times \left((\mu_{it} - 1) \times \beta_{PQ} + \mu_{it} \times \beta_{\mu}\right) + P_{Qit} \times Q_{it} \times \beta_{PQ} \\ - C_{Mit} \times \beta_{C_M} - \frac{\varepsilon_{Lit}}{\varepsilon_{Lit} + 1} \times C_{Lit} \times \beta_{C_L} - C_{Kit} \times \beta_{C_K},\end{aligned} \tag{40}$$

which ultimately yields our expressions for the shares of marginal total surplus generated by public innovation that are appropriated by labor, R&D, capital/profits, and consumers:

$$\mathcal{M}^{TS}_{Lit} \equiv \frac{dPS_{Lit} \,/\, d\log(B_{it})}{dTS_{it} \,/\, d\log(B_{it})}, \tag{41}$$

$$\mathcal{M}^{TS}_{CSit} \equiv \frac{dCS_{it} \,/\, d\log(B_{it})}{dTS_{it} \,/\, d\log(B_{it})}, \tag{42}$$

$$\mathcal{M}^{TS}_{\Pi it} \equiv 1 - \mathcal{M}^{TS}_{Lit} - \mathcal{M}^{TS}_{CSit}. \tag{43}$$

4.7. Distribution versus level of returns

Two distinct objects are of interest: the *level* of the total returns to public innovation and their *distribution* across claimants. The level — the general-equilibrium social return per dollar of public R&D — is difficult to recover in an applied micro setting, because public innovation generates positive and negative externalities both within the estimation sample (across more- and less-exposed firms) and beyond it, and a design that conditions on controls and compares differentially exposed firms is not immune to these critiques. Credibly estimating the level requires a macroeconomic model (e.g., Fieldhouse and Mertens 2026a). The distribution, by contrast, raises only a question of representativeness: the shares we estimate describe how a marginal dollar of induced value added or surplus is split among claimants. Our aim is to provide distributional measures that can be used to scale separately estimated total returns to public R&D.

5. Firm-level data

5.1. Compustat sample and production variables

Our firm-level analysis is built from Compustat because the model requires firm-year measures of sales, input expenditures, and accounting variables for publicly traded firms that can be linked to patenting activity and therefore to exposure to DOE-funded technolo-

gies. These firms are not a random sample of the U.S. economy. They are large, formal, publicly traded producers with standardized financial statements, and they are precisely the firms for which patent-based exposure to public innovation can be measured consistently over time. We keep firm-years with usable industry information from either NAICS or SIC. Monetary Compustat variables are converted to real 2024 dollars using the GDP deflator.

We use `sale`, rather than broader revenue measures, because it is the closest Compustat object to the model's revenue, $P_{Qit} \times Q_{it}$; however, firm-years with missing, zero, or negative sales are not in the regression sample.

Labor expenditures are based on Compustat `xlr`, the labor and related expense item. Because many firms do not report `xlr`, we impute missing labor bills from reported labor costs among similar firms. For firm-years with observed `xlr` and employment, we compute labor cost per employee and predict this average cost using additive industry and year fixed effects, moving from NAICS4 to NAICS3 to NAICS2 and filling missing values from the most detailed prediction available.

Intermediate expenditures are designed to measure non-labor, non-capital, non-R&D operating input purchases. We begin with cost of goods sold and SG&A. Because Compustat can include R&D inside SG&A, we define adjusted SG&A as $\texttt{xsga}^{*}_{it} = \texttt{xsga}_{it}$ when $\texttt{xsga}_{it} \geq \texttt{xrd}_{it}$, and $\texttt{xsga}^{*}_{it} = \texttt{xsga}_{it} + \texttt{xrd}_{it}$ when $\texttt{xsga}_{it} < \texttt{xrd}_{it}$. Intermediate expenditures are then $C_{Mit} = (\texttt{cogs}_{it} + \texttt{xsga}^{*}_{it} - \widehat{\texttt{xlr}}_{it} - \texttt{xrd}_{it})$. Negative reported COGS, SG&A, or R&D components are not allowed to add negative costs to this object; missing components generally leave the constructed intermediate-cost variable missing. Firm-years with missing or nonpositive C_{Mit} are excluded from the regressions.

R&D expenditures are measured directly by Compustat `xrd`, so that $C_{Rit} = \texttt{xrd}_{it}$. Because we treat R&D as a distinct value-added claimant (Section 4), `xrd` is removed from intermediate expenditures, as above, and is reported as its own factor; firm-years with missing or negative `xrd` are excluded from the regressions.

Table 2 reports summary statistics for the common-positive accounting sample, which spans 1982–2016 and comprises 1,105 firms and 18,595 firm-years, before treatment-support and fixed-effect singleton restrictions; imposing those restrictions yields the firm-level regression sample of roughly 920 firms and 11,000 firm-years. Dollar variables are in millions of real 2024 dollars.

The median firm-year has sales of about $2.2 billion, labor expenditures of about $497

TABLE 2. Compustat Firm-Year Summary Statistics

	N	Mean	SD	p25	p50	p75
Sales	18,595	14,103.2	39,742.3	416.6	2,155.1	9,813.7
Labor expenditures	18,595	2,272.1	4,924.4	101.4	496.8	2,089.7
Intermediate expenditures	18,595	8,969.5	29,844.9	183.3	1,052.9	5,174.0
R&D expenditures	18,595	578.2	1,612.3	15.7	62.3	303.6
Markup	18,595	1.46	0.69	1.08	1.26	1.61

Notes: The table reports firm-year summary statistics for the common-positive accounting sample used in the firm-level log-linear regressions: sales, labor expenditures, intermediate expenditures, R&D expenditures, and markups must all be positive and finite. Dollar variables are real 2024 dollars and are reported in millions. The sample contains 1,105 firms and 18,595 firm-years over 1982–2016.

million, intermediate expenditures of about $1.1 billion, and R&D expenditures of about $62 million.

5.2. Markups and demand elasticities

We estimate firm-year markups following De Loecker et al. (2020). Under the residual-demand model of Section 4, a markup μ_{it} implies a residual demand elasticity $\varepsilon_{Dit} = \mu_{it}/(\mu_{it} - 1)$; these markup estimates are therefore the link that lets the surplus calculations add a consumer-side component without observing firm-level prices and physical quantities separately. The mapping requires $\mu_{it} > 1$, which holds throughout our sample. Specifically, we follow Benkard et al. (2025) closely, thanks to their publicly available codebase.[10] For each firm-year, the markup is $\mu_{it} = \theta^{\texttt{cogs}}_{g(i),t} \times \texttt{sale}_{it}/\texttt{cogs}_{it}$, where $\theta^{\texttt{cogs}}_{g,t}$ is an industry-year output elasticity for cost-of-goods-sold (COGS).[11] Table 2 summarizes the markup estimates and we refer the reader to Benkard et al. (2025) for more on variation in markups in the Compustat sample.

[10] See: https://web.stanford.edu/ lanierb/.

[11] We estimate $\theta^{cogs}_{g,t}$ in rolling five-year NAICS2 windows. We drop firm-years with missing, zero, or negative sales or COGS, trim the sales-to-COGS ratio at the 1st and 99th percentiles within year, drop negative SG&A, trim SG&A-to-sales when SG&A is reported, and combine NAICS 48 and 49. Within each rolling industry-year window, we estimate a production function with log sales on log COGS, log capital, lagged investment, lagged capital, and lagged capital squared, instrumenting current log COGS with lagged log COGS. Capital in this step is deflated gross PP&E, and investment is $I_{it} = K_{it} - 0.9K_{i,t-1}$.

6. Firms' non-random exposure to exogenous public innovation

6.1. Empirical model

The model in Section 4 motivated the reduced-form regression of Equation (35):

$$\log(Y_{it}) = \alpha_Y + \beta_Y \times \log(B_{it}) + \Gamma_Y \times Z_{it} + u_{Yit} \, ,$$

for the focal outcomes $Y_{it} \in \{P_{Qit} \times Q_{it}, C_{Lit}, C_{Kit}, C_{Mit}, C_{Rit}, \mu_{it}\}$.

The problem is that our source of plausibly exogenous variation in public innovation is not at the level of *firm-year* (per B_{it}) but rather *technology-year* (per $\underline{A}_{jt}$). We therefore need to map technology-specific investments into firm-specific exposure metrics.

We use a "shift-share" formulation to express the stock of public innovation relevant to each firm-year:

$$B_{it} = \sum_j \chi_{ijt} \times \underline{A}_{jt} \, , \tag{44}$$

which implies our focal estimating equation:

$$\log(Y_{it}) = \alpha_Y + \beta_Y \times \log\left(\sum_j \chi_{ijt} \times \underline{A}_{jt}\right) + \Gamma_Y \times Z_{it} + u_{Yit} \, . \tag{45}$$

The control vector Z_{it} serves two purposes, both tied to identification of β_Y rather than to the mechanisms of interest. First, because our exposure shares need not sum to one (the "missing shares" problem discussed below), we always include the firm-year total covered share, $\log \sum_j \chi_{ijt}$, so that firms with more of their patenting inside the DOE technology space are not mechanically assigned greater exposure. Second, we include SIC2 × year fixed effects to absorb industry-by-year common shocks. Consistent with the model's guidance (Section 4), Z_{it} deliberately excludes input levels and other proximal outcomes, which would block the causal paths we aim to estimate.

Notably, we do not need an exogenous measure of χ_{ijt} (i.e., one uncorrelated with productivity), since we can leverage recent advances in non-random exposure to exogenous shocks (Borusyak and Hull 2023). Recall that our windfall estimates are

$$\widehat{\eta}_{jt} = \log(\underline{A}_{jt}) - \widehat{\phi}\Big(\log(\underline{R}_{jt}), W_{jt}\Big) \, ,$$

so that the technology-year stock of public innovation can be written as

$$\underline{A}_{jt} = \exp\Big(\widehat{\phi}\Big(\log(\underline{R}_{jt}), W_{jt}\Big) + \widehat{\eta}_{jt}\Big) \ .$$

Following Borusyak and Hull (2023), we generate an instrument for B_{it} that leverages only the variation in $\underline{A}_{jt}$ due to the windfalls $\widehat{\eta}_{jt}$. For each of $s = 1, \ldots, S$ permutation draws from the jt-level empirical distribution of $\widehat{\eta}_{jt}$, we construct a counterfactual stock

$$\underline{A}^{(s)}_{jt} = \exp\Big(\widehat{\phi}\Big(\log(\underline{R}_{jt}), W_{jt}\Big) + \widehat{\eta}^{(s)}_{jt}\Big) \ , \tag{46}$$

and recompute firm-year exposure under each draw, which, in the language of Borusyak and Hull (2023), yields a "recentered" instrument:

$$IV_{it} = \log\left(\sum_j \chi_{ijt} \times \underline{A}_{jt}\right) - \frac{1}{S} \times \sum_{s=1}^{S} \log\left(\sum_j \chi_{ijt} \times \underline{A}^{(s)}_{jt}\right) \ . \tag{47}$$

Importantly, while this approach allows the exposure terms χ to be endogenous, it remains important to control for the "missing share" problem highlighted by Borusyak et al. (2022) when $\sum_j \chi_{ijt} < 1$. Because our measure of χ generically leaves shares short of one, we always include $\log \sum_j \chi_{ijt}$ as a firm-year control, as noted above.

Inference uses both conventional clustered standard errors and randomization inference. For tradition's sake, we report standard errors clustered at the firm level. More in line with our identification strategy, and following Borusyak and Hull (2023), we also report randomization-inference results.[12]

6.2. Exposure metric and patent data

We use patenting to locate firms in technology space and measure the exposure object χ. The idea is that a firm's past patenting reveals the set of technologies in which it has scientific, engineering, organizational, or commercial absorptive capacity. DOE funding in technology j is therefore more relevant for firm i in year t when the firm has recently

[12] We use the same simulated treatments that generated the recentered instrument. Define the simulated recentered instruments $IV^{(s)}_{it} = \log\left(\sum_j \chi_{ijt} \times \underline{A}^{(s)}_{jt}\right) - \frac{1}{S} \times \sum_{s=1}^{S} \log\left(\sum_j \chi_{ijt} \times \underline{A}^{(s)}_{jt}\right)$. For the sharp null $\beta_Y = 0$, we residualize the outcome with respect to the same controls and fixed effects used in the regression and compute the score statistic $T_{\text{obs}} = |\sum_{it} IV_{it} \times \tilde{u}_{it}|$, where $\tilde{u}_{it}$ is the residualized outcome under the null. For each draw we compute $T_s = \left|\sum_{it} IV^{(s)}_{it} \times \tilde{u}_{it}\right|$, and the randomization-inference p-value is $\frac{1+\sum_{s=1}^{S} \mathbf{1}\{T_s \geq T_{\text{obs}}\}}{S+1}$.

patented more in technology j, relative to its overall patenting activity.

Let n_{ijt} be firm i's effective patent flow in technology j and year t, and let n_{it} be its total effective patent flow in year t. Given a patent-stock depreciation δ_N and lag ℓ_N, the technology-specific patent stock is

$$\underline{N}_{ijt}(\delta_N, \ell_N) = \delta_N \times \underline{N}_{ij,t-1}(\delta_N, \ell_N) + n_{ij,t-\ell_N}\,, \tag{48}$$

and the total patent stock $\underline{N}_{it}$ is defined analogously over all of the firm's patents, whether or not they relate to in-sample technologies. The firm-technology exposure share is

$$\chi_{ijt} \equiv \frac{\underline{N}_{ijt}(\delta_N, \ell_N)}{\underline{N}_{it}(\delta_N, \ell_N)} \quad \text{if } \underline{N}_{it}(\delta_N, \ell_N) > 0, \tag{49}$$

and $\chi_{ijt} = 0$ when the total patent stock is zero. Because the denominator is the firm's *total* effective patent stock while the numerator is restricted to the DOE technology categories $\mathcal{J}$ defined in Section 3, the shares need not sum to one. This is the "missing shares" issue raised above, and it motivates the $\log \sum_j \chi_{ijt}$ control: firm-years with a smaller covered patent share have less measured exposure to the DOE technology space even before variation in DOE funding is considered.

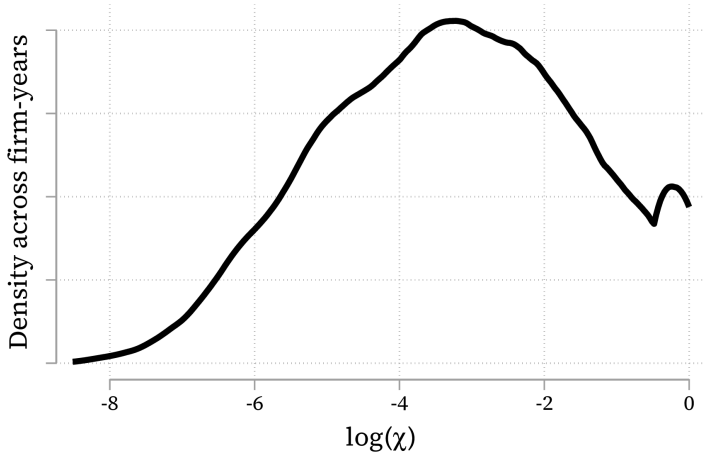
We use patent data drawn from the PATSTAT Global database (Spring 2022 edition) to construct firms' exposure shares, restricting attention to patent families that contain an application filed at the U.S. Patent and Trademark Office (USPTO). To locate each patent in the DOE technology space $\mathcal{J}$, we use a detailed dictionary-based text classification approach. We construct an energy dictionary containing over 400 technology-specific terms spanning different energy technologies.[13] We transform patent titles and abstracts into n-grams and search for the dictionary terms, assigning each family to one or more of the categories in $\mathcal{J}$. To link patents to firms, we use the firm-to-patent concordance of Dyevre and Seager (2023), which maps USPTO patents to Compustat firms, assigning each patent to the ultimate corporate owner active in Compustat at the time of the patent application. These linked data allow us to construct the firm-technology-year patent counts n_{ijt} from which the exposure shares χ_{ijt} are built.

Figure 3 plots the firm-year distribution of the sum of the exposure shares across technologies, in logs. The peak of the distribution is around -3, where roughly 5 percent of a firm's patents in a given year are classified as within a DOE technology. There is also a

[13] Appendix Table A1 provides a complete list of energy terms.

mass of firms that specialize in energy technologies, with DOE energy technology shares that sum to nearly one (zero in log terms).

FIGURE 3. Exposure metrics, χ



Notes: The figure plots the kernel density of log firm-year patent-exposure shares, $\log \sum_j \chi_{ijt}$, for the baseline patent-stock specification with one-year lag and depreciation parameter 0.9.

7. Results

We estimate the reduced-form elasticities of Equation (35) and map them, through the model of Section 4, into measures of who captures the value that public innovation creates. Three findings stand out. Public-innovation exposure raises every margin of firm activity in close to fixed proportion. It lowers markups. And the value created is shared: within the firm, workers and the owners of capital each take a large slice, and once the fall in markups reaches buyers, consumers capture most of the total surplus.

7.1. Elasticities

Table 3 reports the elasticity of each firm variable with respect to public-innovation exposure, averaged over the 80 specifications that vary the depreciation and lag parameters, the control set, and the first-stage learner. A one-percent increase in exposure raises sales by about 0.50%, labor compensation by 0.47%, intermediate spending by 0.51%, and R&D by 0.34%, and it lowers the markup by about 0.04%. Every estimate is significant under both clustered and randomization inference, and the first stage is strong. The instrumented estimates exceed their OLS counterparts (0.39, 0.36, 0.40, and 0.25) by a

modest margin, consistent with mild attenuation and/or weak negative selection of DOE R&D appropriations on firm prospects.

TABLE 3. Elasticities with respect to public innovation

	Sales (1)	Labor (2)	Inter-mediates (3)	R&D (4)	Markups (5)
Panel A: OLS					
Coefficient	0.393	0.361	0.399	0.249	-0.042
Mixture SE	(0.094)	(0.088)	(0.093)	(0.082)	(0.014)
p-value	0.000	0.000	0.000	0.002	0.003
Panel B: IV					
Coefficient	0.502	0.469	0.509	0.336	-0.042
Mixture SE	(0.098)	(0.093)	(0.097)	(0.093)	(0.017)
p-value	0.000	0.000	0.000	0.000	0.015
RI p-value	0.001	0.001	0.001	0.004	0.012
Specifications averaged	80	80	80	80	80
Mean firm-year observations	11,078	11,078	11,078	11,078	11,078

Notes: The table averages across all completed specifications available when the reporting script is run. Point estimates are equal-weight averages across specifications. Mixture standard errors use $\frac{1}{S}\sum_s[\widehat{SE}_s^2 + (\hat{\beta}_s - \bar{\beta})^2]$, combining within-specification uncertainty and between-specification dispersion. Reported conventional p-values are two-sided normal tests using the mixture standard error. Mean RI p-values are equal-weight averages of specification-level randomization-inference p-values.

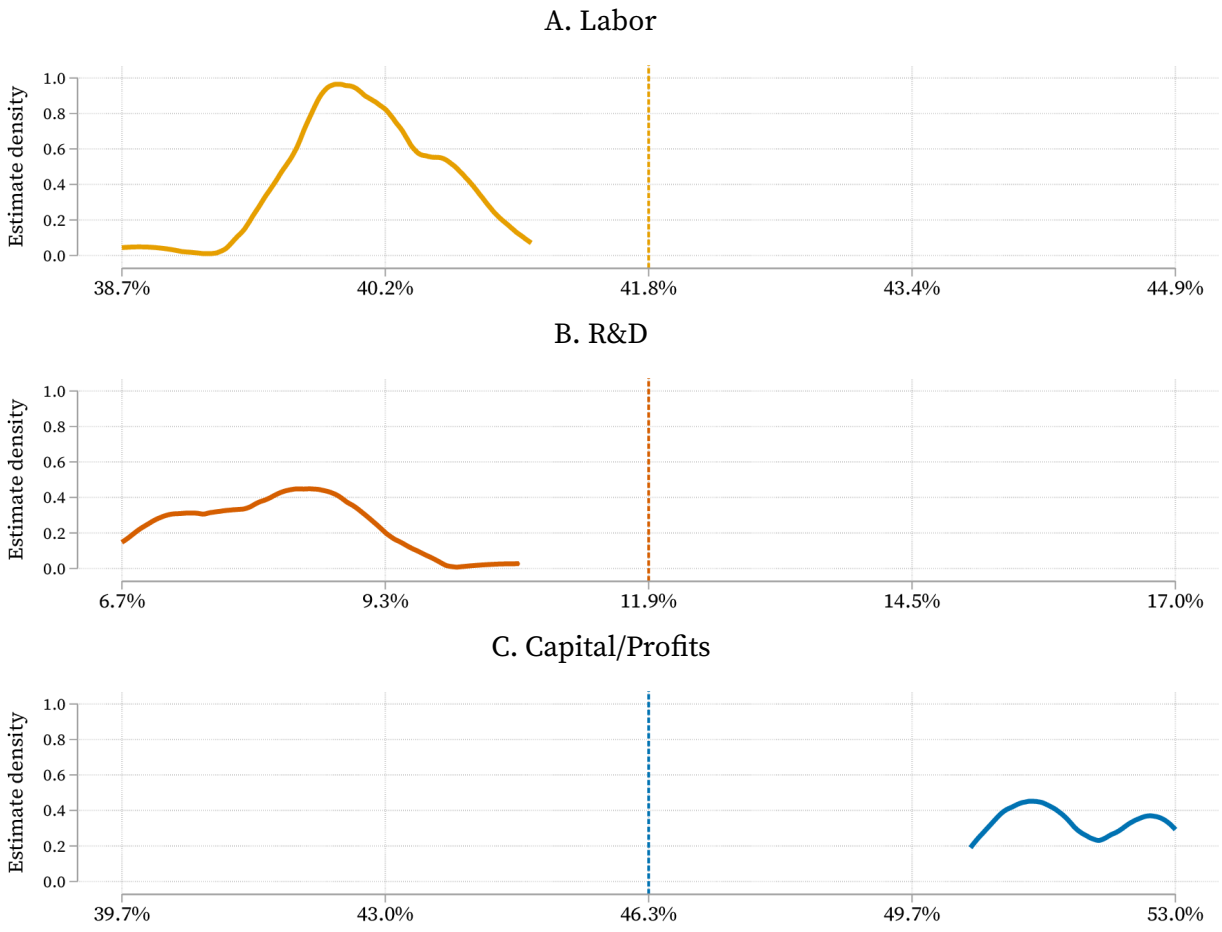
7.2. The incidence of induced value added

Figure 4 converts these elasticities into the share of marginal value added (spurred by the DOE's investments) that accrues to each claimant in the firm. Of every additional dollar of value added that public innovation generates, labor receives about 40 cents, the owners of capital about 51 cents, and R&D about 8 cents. These marginal shares are close to the corresponding average shares in our sample — 42, 46, and 12 cents — so public R&D leaves labor's share of value added essentially unchanged and tilts the margin only modestly toward capital.

7.3. The incidence of induced surplus

While value-added metrics are less demanding to estimate, as accounting metrics they do not reflect surplus splits and omit any gains or losses for consumers. On the latter, the elasticity estimates indicate that public R&D lowers markups, shifting some (marginal)

FIGURE 4. Share of marginal value add from public innovation, $\mathcal{M}^{VA}$



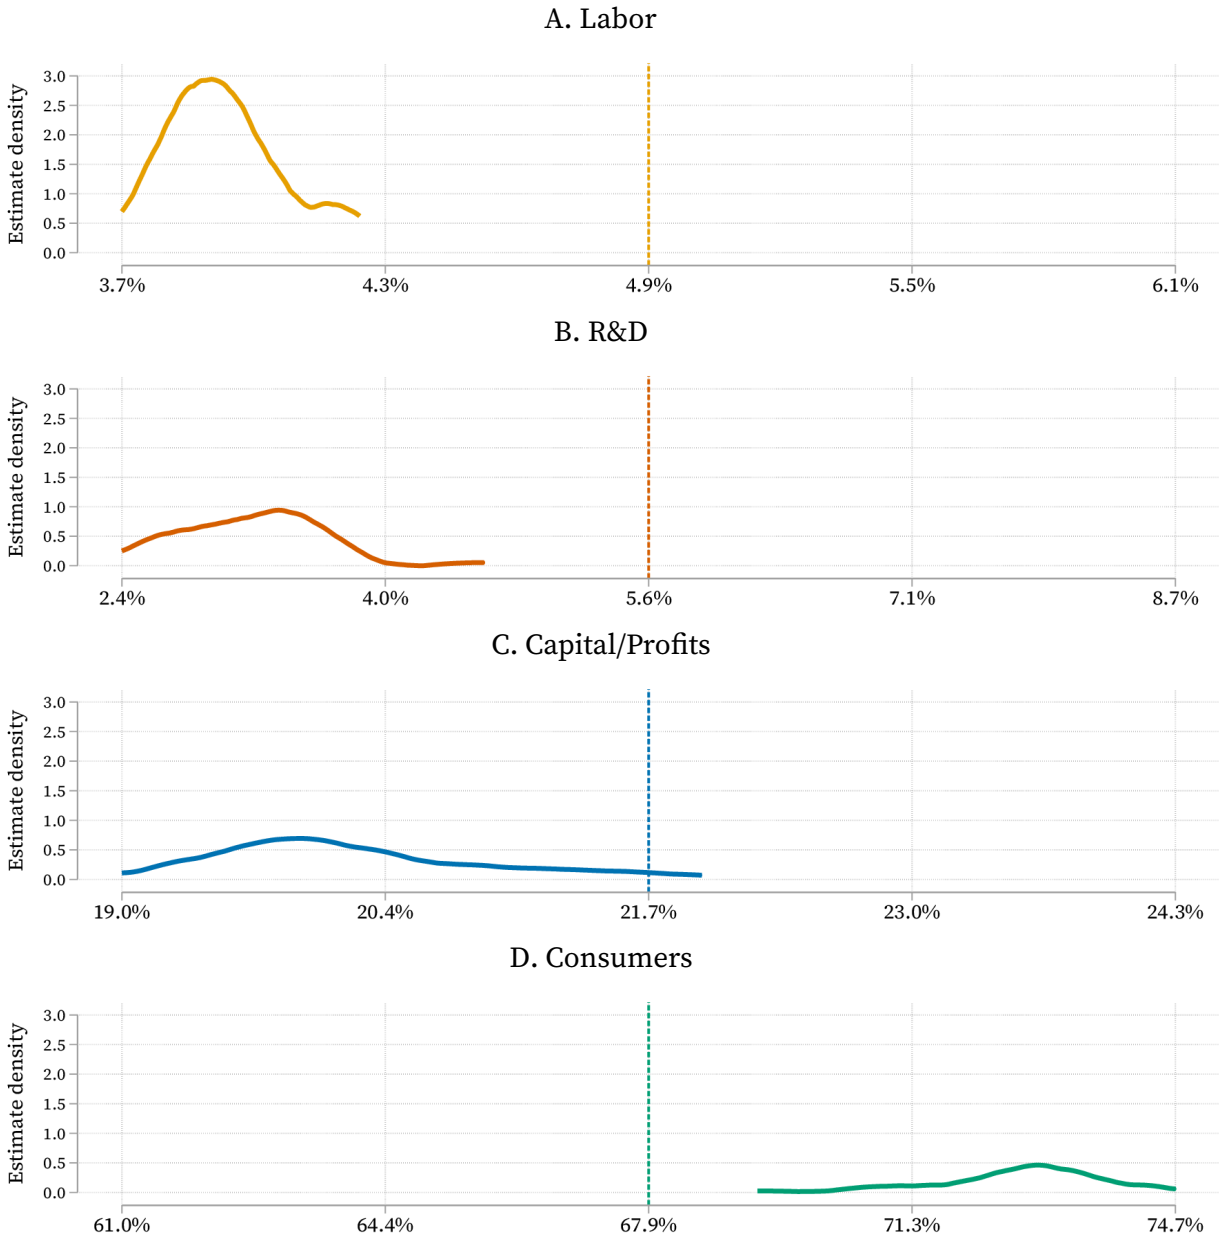
Notes: Each panel plots a kernel density across completed specifications of the party's IV-implied share of value added created by public innovation. Panels use plot-specific bandwidths and a common y-axis scale. Each x-axis is centered on the party's sample mean share of value added, shown by the vertical line.

surplus to consumers. Figure 5 reports the division of marginal *total* surplus. Consumers capture roughly 70%, the owners of capital about 20%, labor about 4%, and R&D about 3%. Two forces put consumers first: markups in our sample sit around 1.3, so the consumer-surplus base is large, and public R&D widens that base by pushing markups down.

7.4. Robustness

The estimates are stable across the specification grid. The sign and rough magnitude of every elasticity are invariant to the depreciation rate, the lag length, the control set, and the first-stage learner. Randomization-inference p-values, which respect the design of the windfall instrument, deliver the same conclusions as conventional clustered errors, and replacing the SIC2-by-year fixed effects with additive industry and year effects leaves the estimates unchanged.

FIGURE 5. Share of marginal surplus from public innovation, $\mathcal{M}^{TS}$



Notes: Each panel plots a kernel density across completed specifications of the party's IV-implied share of total surplus created by public innovation. Panels use plot-specific bandwidths and a common y-axis scale. Each x-axis is centered on the party's sample mean share of total surplus, shown by the vertical line.

8. Discussion

This paper asks who captures the value created when the government funds innovation, and focuses on answering this question in the context of energy technology R&D financed by the U.S. Department of Energy and the effects on large, publicly listed U.S. companies. First, we find public-innovation scales firms up roughly proportionally, raising sales and inputs by about 0.5% per one-percent increase. On a value-add basis, the marginal gains from this expansion accrue to capital-owners / shareholders at a slightly larger rate than average. On the margin, the within-firm division of value added tilts toward capital and profits: capital takes about 52 cents of each induced dollar against an average share of 46, while labor (40 cents against 42) and R&D (8 against 12) each take somewhat less. Once the fall in markups reaches buyers, the economy-wide division tilts further outward. Consumers capture about 73 percent of the induced total surplus against an average share of 68, and the firm and its factors together capture correspondingly less, with labor's share falling proportionally more than capital's.

To summarize, public R&D does not leave the distribution of value untouched, but neither does it overturn it: the marginal dollar resembles the average dollar, tilted a few cents toward capital inside the firm and a few points toward consumers overall. Against the concern that innovation increasingly concentrates gains among the owners of capital, the within-firm tilt is present but small, and it coexists with a larger movement toward consumers that the concern usually leaves out.

Three limitations bound these conclusions. We identify the distribution of the gains, not their level: the social rate of return to public R&D requires variation that a firm-level comparison does not contain, and the shares we estimate are meant to be combined with a separately estimated level rather than to stand in for it. The estimation sample is large, publicly traded firms, for which patent-based exposure can be measured but which are not the economy as a whole. And the consumer result, the largest reallocation we find, rests on interpreting estimated markups as demand-side wedges and on the isoelastic-demand mapping; we read it as evidence of meaningful pass-through rather than as a precise figure, while the producer-side split between labor and capital depends on assumed factor-supply elasticities.

These limitations point to the work we see as most valuable next. Because pass-through is at once the largest and the least certain part of our estimates, identifying the consumer-side response directly (e.g., from variation in prices or demand rather than from markups inferred off production functions) would be valuable. The reduced-form effects also invite

a closer look at mechanism: how public funding yields trained scientists and new ideas before it appears in firm accounts, and how the incidence measured among public firms extends to the rest of the economy.

Finally, the political-production-function design is not specific to our setting. Wherever a government funds production and records both what an agency requests and what legislators appropriate, the gap between the two could supply quasi-experimental variation in public investment. And looking beyond government settings, it is likely that many organizations and investors record similar data on planned and realized investment, which may be amenable to our methodology.

References

Acemoglu, Daron, "Directed technical change," *The review of economic studies*, 2002, *69* (4), 781–809.

___, "Equilibrium bias of technology," *Econometrica*, 2007, *75* (5), 1371–1409.

___ and Joshua Linn, "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry," *The Quarterly Journal of Economics*, 2004, *119* (3), 1049–1090.

___ and Pascual Restrepo, "The race between man and machine: Implications of technology for growth, factor shares, and employment," *American economic review*, 2018, *108* (6), 1488–1542.

___ and ___, "Tasks, Automation, and the Rise in US Wage Inequality," *Econometrica*, 2022, *90* (5), 1973–2016.

Aghion, P., U. Akcigit, A. Hyytinen, and O. Toivanen, "On the Returns to Invention within Firms: Evidence from Finland," *AEA Papers and Proceedings*, 2018, *108*, 208–12.

Aghion, Philippe and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992, *60* (2), 323–351.

Akerman, A., J. Moscona, H.S. Pellegrina, and K. Sastry, "Public R&D Meets Economic Development: Embrapa and Brazil's Agricultural Revolution," *NBER Working Paper 34213*, 2026.

Alam, A. and A. Diegmann, "Patents, Firm Rents, and Worker Compensation: Causal Evidence from Quasi-Random Patent Allocation," *CESifo Working Paper No. 12666*, 2026.

Arrow, Kenneth, "Economic Welfare and the Allocation of Resources for Invention," in "The Rate and Direction of Inventive Activity: Economic and Social Factors," Princeton University Press, 1962, pp. 609–626.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," *The Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Azoulay, Pierre, Joshua S Graff Zivin, Danielle Li, and Bhaven Sampat, "Public R&D investments and private-sector patenting: Evidence from NIH funding rules," *Review of Economic Studies*, 2019, *86* (1), 117–152.

Benkard, C Lanier, Nathan H Miller, and Ali Yurukoglu, "The Rise of Market Power and the Macroeconomic Implications: Comment," *Mimeo*, 2025.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying technology spillovers and product market rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

Borusyak, Kirill and Peter Hull, "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, 2023, *91* (6), 2155–2185.

___, ___, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *The Review of Economic Studies*, 2022, *89* (1), 181–213.

Bronzini, R. and E. Iachini, "Are Incentives for R&D Effective? Evidence from a Regression Discontinuity Approach," *American Economic Journal: Economic Policy*, 2014, *6* (4), 100–134.

Dugoua, Eugenie and Todd D. Gerarden, "Induced Innovation, Inventors, and the Energy Transition," *American Economic Review: Insights*, March 2025, *7* (1), 90–106.

Dyevre, Arnaud and Oliver Seager, "70 Years of Patents Matched to Firms: Methodology and Insights About Firm Heterogeneity," Working Paper 2023.

Fieldhouse, Andrew and Karel Mertens, "The Returns to Government R&D: Evidence from US Appropriations Shocks," *The American Economic Review*, 2026.

Fieldhouse, Andrew J and Karel Mertens, "The social returns to public R&D," *Entrepreneurship and Innovation Policy and the Economy*, 2026, *5* (1), 9–39.

Finkelstein, A. and N. Hendren, "Welfare Analysis Meets Causal Inference," *Journal of Economic Perspectives*, 2020, *34* (4), 146–67.

Fuest, C., A. Peichl, and S. Siegloch, "Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany," *The American Economic Review*, 2018, *108* (2), 393–418.

Griliches, Zvi, "Issues in Assessing the Contribution of Research and Development to Productivity Growth," *The Bell Journal of Economics*, 1979, *10* (1), 92–116.

Gross, Daniel P. and Bhaven N. Sampat, "America, Jump-Started: World War II R&D and the Takeoff of the US Innovation System," *The American Economic Review*, 2023, *113* (12), 3323–56.

Grossman, Gene M and Elhanan Helpman, "Quality ladders in the theory of growth," *The review of economic studies*, 1991, *58* (1), 43–61.

Hendren, N. and B. Sprung-Keyser, "A Unified Welfare Analysis of Government Policies," *The Quarterly Journal of Economics*, 2020, *135* (3), 1209–1318.

Howell, Sabrina, "Financing innovation: Evidence from R&D grants," *American Economic Review*, 2017, *107* (4), 1136–64.

Jones, Benjamin F. and Lawrence H. Summers, "A Calculation of the Social Returns to Innovation," in "Innovation and Public Policy," University of Chicago Press, 2022, pp. 13–60.

Jones, Charles I and John C Williams, "Measuring the social return to R&D," *The Quarterly Journal of Economics*, 1998, *113* (4), 1119–1135.

___ and ___, "Too much of a good thing? The economics of investment in R&D," *Journal of Economic Growth*, 2000, *5* (1), 65–85.

Kantor, Shawn and Alexander Whalley, "Moonshot: Public R&D and Growth," *The American Economic Review*, 2025, *115* (9), 2891–2925.

Karabarbounis, Loukas and Brent Neiman, "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani, "Credulity, Lies, and Costly Talk," *Journal of Economic Theory*, 2007, *134* (1), 93–116.

Kehrig, M. and N. Vincent, "The Micro-Level Anatomy of the Labor Share Decline," *The Quarterly Journal of Economics*, 2021, *136* (2), 1031–1087.

Kline, P., N. Petkova, H. Williams, and O. Zidar, "Who Profits from Patents? Rent-Sharing at Innovative Firms," *The Quarterly Journal of Economics*, 2019, *134* (3), 1343–1404.

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, "The Rise of Market Power and the Macroeconomic Implications," *The Quarterly Journal of Economics*, 2020, *135* (2), 561–644.

Myers, Kyle R. and Lauren Lanahan, "Estimating Spillovers from Publicly Funded R&D: Evidence from the US Department of Energy," *The American Economic Review*, 2022, *112* (7), 2393–2423.

Pless, Jacquelyn, "Are Complementary Policies Substitutes? Evidence from R&D Subsidies in the United Kingdom," *American Economic Journal: Economic Policy*, 2026, *18* (1), 92–126.

Romer, Paul M, "Endogenous technological change," *Journal of Political Economy*, 1990, *98* (5, Part

2), S71–S102.

Santoleri, Pietro, Andrea Mina, Alberto Di Minin, and Irene Martelli, "The causal effects of R&D grants: Evidence from a regression discontinuity," *The Review of Economics and Statistics*, 2024, *106* (6), 1495–1510.

Serrato, J.C. Suarez and O. Zidar, "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous FIrms," *The American Economic Review*, 2016, *106* (9), 2582–2624.

Van Reenen, J., "The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies," *The Quarterly Journal of Economics*, 1996, *111* (1), 195–226.

Supplemental Appendix

TABLE A1. Energy Term Dictionary

Battery

KWRD(battery)
anode battery
battery + electrolyte
battery anode
battery cathode
battery cell
battery discharge
battery electrode
battery module
battery pack
battery performance
battery rechargability
battery recharge
battery storage
cathode battery
discharge battery
electrochemical battery
lead acid battery
li battery
li ion battery
lithium + electrolyte
lithium battery
lithium ion battery
lithium sulfur battery
rechargability battery
recharge battery
rechargeable battery
redox flow battery
sulfur li battery
sulfur lithium battery

BiomassWaste

bio gas
bio oil
bio waste
biochar
biodiesel
bioenergy
bioenergy + CLASS(energy)
bioenergy + synthetic fuel
bioenergy + thermochemical conversion
bioethanol
biofuel
biogas
biomass + CLASS(energy)
biomass + producer gas
biomass + synthetic fuel
biomass + thermochemical conversion
biomass energy
biooil
biorefinery
cellulosic ethanol
lignocellulosic biomass
syngas + biochar
syngas + biofuel
syngas + biogas
syngas + biomass
syngas + ethanol
waste energy
waste heat
waste incineration
waste-to-energy

Building

KWRD(energy) + KWRD(building)
air leak + building
air leakage + building

CarbonCapture

carbon capture
carbon dioxide + electrochemical reduction
carbon dioxide + electrolysis reduction
carbon dioxide + photoelectrolysis reduction
carbon dioxide + reduction + catalysis
carbon dioxide + reduction + catalyst
carbon dioxide + reduction + catalytic
carbon dioxide + reduction + photocatalysis
carbon dioxide + reduction + photocatalyst
carbon dioxide + reduction + photocatalytic
carbon dioxide capture
carbon dioxide disposal
carbon dioxide sequestration
carbon dioxide storage
carbon sequestration
carbon storage
co2 + electrochemical reduction
co2 + electrolysis reduction
co2 + photoelectrolysis reduction
co2 + reduction + catalysis
co2 + reduction + catalyst
co2 + reduction + catalytic
co2 + reduction + photocatalysis
co2 + reduction + photocatalyst
co2 + reduction + photocatalytic
co2 capture
co2 sequestration
co2 storage

Coal

KWRD(coal)
KWRD(coal) + KWRD(hydrogen)
asphaltite
briquette
browncoal
carbonaceous + slurry
carbonaceous shale
clean coal
coal + combustion
coal + deposit
coal + slurry
coal shale
coal to gas
coal to liquid
coke oven
coking
disintegrating peat
hardcoal
integrated gasification combined-cycle
lignite
metaanthracite
ortholignite
peat + energy
peat + power
peat drying
solid fuel
syngas + coal
working-up + peat

ElectricVehicle

charging station
electric mobility
electric propelled vehicle
electric regenerative braking
electric vehicle
electrical propulsion units
electrodynamic brake

ElectricityGeneral

electric power

electric utility
electrical grid
electrical load
electricity market
electricity transmission
heat pump
load forecasting
power conversion
power generation
power system
power transformer
transmission line

EnergyEfficiency

energy conservation
energy efficiency
energy efficient
energy recuperation
energy saving

EnergyGeneral

clean energy
energy consumption
energy conversion
energy demand
energy harvesting
energy management
energy system
energy use
heat transfer
thermal conductivity
thermal management

EnergyStorage

KWRD(capacitor) + KWRD(energy)
KWRD(storage) + KWRD(energy)
capacitor + li
capacitor + li ion
capacitor + lithium
electrical double layer
electrochemical double layer
energy storage
flywheel
supercapacitor

FossilFuel

acid fracturing
carbonaceous
combustion furnace
fossil fuel
hydraulic fracking
hydraulic fracturing
hydrocarbon
kiln + engineering
liquid fuel
oil and gas
proppant fracturing
steam accumulator
steam engine
steam generation
unconventional reservoir
well productivity

FuelCell

fuel cell

Fusion

blanket + fusion
breeding blanket + fusion
fusion blanket
fusion chamber
fusion energy
fusion power
fusion reactor
inertial confinement + fusion
inertial electrostatic confinement + fusion
madison symmetric torus + fusion
magnetic confinement + fusion
magnetohydrodynamics + plasma + CLASS(energy)
reversed field pinch + fusion
tokamak + fusion

Geothermal

geothermal

HybridVehicle

hybrid propulsion system
hybrid vehicle
regenerative braking

Hydro

hydro electric
hydro energy
hydro plant
hydro power
hydroelectric
hydropower

Hydrogen

KWRD(biomass) + KWRD(syngas)
KWRD(hydrogen)
KWRD(hydrogen) + KWRD(coal)
KWRD(hydrogen) + KWRD(natural gas)
electrolysis + hydrogen
hydrogen evolution reaction
hydrogen fuel
hydrogen pipeline
hydrogen production
hydrogen storage
hydrogen tank
integrated gasification combined-cycle + syngas
steam methane reforming + biochar
steam methane reforming + biodiesel
steam methane reforming + biofuel
steam methane reforming + biogas
steam methane reforming + biomass
steam methane reforming + coal
steam methane reforming + ethanol
steam methane reforming + natural gas
steam reforming + biochar
steam reforming + biodiesel
steam reforming + biofuel
steam reforming + biogas
steam reforming + biomass
steam reforming + coal
steam reforming + ethanol
steam reforming + natural gas
syngas + biochar
syngas + biodiesel
syngas + biofuel
syngas + biogas
syngas + biomass
syngas + coal
syngas + ethanol
syngas + natural gas
water electrolysiselectrolysis of water
water splitting + hydrogen

InternalCombustion

KWRD(engine)
combustion + diesel
combustion + engine
combustion + gasoline
combustion engine
compression-ignition + engine
diesel + engine
gasoline + engine
spark ignition + engine

Manufacturing

KWRD(energy) + KWRD(manufacturing)

Marine

hydrokinetic
in-stream energy
marine energy
marine hydrokinetic energy
marine hydrokinetic power
marine power
marine renewable energy

marine renewable power
ocean thermal energy
ocean thermal power
salinity gradient energy
salinity gradient power
tidal energy
tidal power
tidal turbine
wave energy
wave power

NaturalGas

LNG
LPG
boiler + combustion
borehole + gas
clathrate hydrates
clathrates
coal bed methane
coal mine methane
coal seam gas
coalbed gas
coalbed methane
conventional + gas
drilling + gas
enhanced recovery + gas
fracking + gas
fracturing + gas
gas deposit
gas hydrates
gas to liquid
gas turbine
gaseous fuel
hydrates + gas
liquefied petroleum gas
methane + combustion
natural gas
natural gas + combustion
offshore platform + gas
pipeline + gas
seismography + gas
shale gas
synthetic natural gas
unconventional + gas
well productivity + gas
wellbore + gas

Nuclear

boiling water + reactor
generation I + reactor
generation II + reactor
generation III + reactor
generation IV + reactor
generationmolten salt + reactor
heavy water + reactor
light-water-moderated + reactor
liquid metal cooled + reactor
molten salt + nuclear
next generation + reactor
nuclear accident
nuclear energy
nuclear engineering
nuclear fission
nuclear fuel
nuclear power
nuclear reactor
nuclear reprocessing
nuclear safety
pressurized water + reactor
radioactive waste
supercritical water + reactor

Oil

bitumen
bituminous sand
borehole + oil
butane
butane + combustion
cetane
conventional oil
crude bitumen
crude oil
diesel
drilling + oil
enhanced oil recovery
enhanced recovery + oil
fracking + oil
fracturing + oil
gasoline
jet fuel
kerosene
octane
offshore platform + oil
oil + combustion
oil chalk
oil cracking
oil deposit
oil exploration
oil refinery
oil refining
oil sand
oil shale
oil well
petroleum engineering
petroleum oil
pipeline + oil
propane
propane + combustion
seismography + oil
shale oil
tar sand
unconventional oil
well productivity + oil
wellbore + oil

OtherRenewables

alternative energy
renewable energy

SmartGrid

distributed energy
microgrid
power grid
smart grid
smart power grid

Solar

KWRD(solar) + KWRD(energy)
photocatalysis + CLASS(energy)
photovoltaic
solar + stirling engine
solar cell
solar energy
solar panel
solar plant
solar power
solar thermal energy
solar tower

Wind

wind energy
wind farm
wind power
wind turbine

Notes: When the keyword is a combination of terms (e.g., “lead acid + battery”), we search for joint occurrences, that is, a dissertation would be classified as Battery if it contains the phrase “lead acid” and “battery” even if there are not immediately next to each other.