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

September 23, 2026

[*preliminary — 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 (DOE). The framework links public innovation to firm-level production, allowing us to estimate how the resulting gains are divided among firms, factor suppliers, and buyers of firms' output. To address the endogeneity of investment, we develop a political production function approach that isolates funding windfalls: residual variation in investment that cannot be predicted using requested investments. We combine these windfalls with measures of firms' exposure to DOE-funded technologies, and we estimate the effects of DOE-funded innovations. We find that public innovation raises revenues and input expenditures and lowers markups. Mapping these changes into surplus terms requires some assumptions, but provides a new view into the incidence of the value of public innovation.

Dugoua: e.dugoua@lse.ac.uk. Gerarden: gerarden@cornell.edu. Myers: kmyers@hbs.edu. Pless: jacq.pless@gmail.com. There are no financial relationships or other potential conflicts of interest that apply to the authors.

1. Introduction

Governments invest in R&D through a range of mechanisms, supporting innovation at public institutions, non-governmental organizations, and for-profit firms. A growing body of evidence suggests that these public investments generate large economic returns (Azoulay et al. 2019; Jones and Summers 2022a; Fieldhouse and Mertens 2026a,b). But who captures these gains from public innovation?

For private R&D, profit maximization provides a starting point for understanding which ideas firms will pursue and how their benefits are distributed (Acemoglu 2002, 2007; Acemoglu and Restrepo 2018). But public R&D serves a wider range of objectives, so theory provides less guidance on how these investments will ultimately affect the economy through firms' productivity, demand for inputs, and market power.

In this paper, we estimate how the gains from public innovation are split among input suppliers (i.e., labor, capital, R&D), shareholders, and consumers. The first ingredient is a research design that addresses the potential endogeneity of public investments. We develop a model of, what we term, *political production functions* and apply it to the U.S. Department of Energy (DOE). The key idea is that, while budget requests from the executive branch may be correlated with the productivity of those potential investments, there is some noise in final budget appropriations (due to the legislative branch) that is uncorrelated with productivity. Our model specifies the conditions necessary to isolate this exogenous variation in investment. Newly collected data on technology-level budget requests and appropriations at the DOE provide the data necessary.

The second ingredient is a framework for tracing the incidence of public innovation. We specify a model that places some, but not many, restrictions on the way in which public innovations can affect input suppliers, producers, and demand. The model motivates a series of simple regressions, and gives a set of parameters that must be estimated or calibrated in order to solve for the share of gains generated by public innovation (i.e., in value-add or surplus terms) that accrues to different groups in the economy. The surplus incidence metrics are the most policy-relevant as they directly describe welfare changes; however, solving for them requires the most restrictive assumptions.

The final ingredient is a way to map technology-level public innovation to firm-level production. Here, we leverage the recentered shift-share methodology of Borusyak and Hull (2023), which provides econometric guidance. For data, we rely on the Compustat dataset, which describes publicly-traded firms in the US and includes some of the key

variables. We supplement Compustat with data on firms' patents in order to locate firms in technology-space and measure their exposure to DOE's investments. Furthermore, thanks to recent work on markup estimation that also uses Compustat (De Loecker et al. 2020; Benkard et al. 2025), we can recover firms' markups, which allows us to engage with consumers' surplus.

Many aspects of our approach — the timing of public innovations, the estimation of windfall investments, the construction of patent exposure measures, assumptions about supply elasticities — require choices for which there are a range of plausible alternatives. Thus, we often report the full distribution of results across all of the plausible choices we can make. Overall, we find that public innovation raises revenues and input expenditures, and it decreases markups. In our preferred specification, a one-percent increase in exposure raises revenue by 0.48%. Expenditures on non-R&D labor rise by 0.46%, capital services by 0.62%, R&D by 0.37%, and intermediates by 0.48%. Markups fall by 0.05%. Each response is statistically distinguishable from zero at the 5% level under both clustered and randomization inference.

Focusing on the incidence of value add generated by public innovation, we generally find the following: labor receives about 38 cents of each additional dollar in value-add produced by the DOE's investments, capital suppliers receive 29 cents, R&D suppliers receive 12 cents, and shareholders receive 21 cents as profits. Labor receives the largest share in almost every specification. On average, compared with baseline (average) value-add shares, the additional value added has a similar labor share, a larger capital share, and smaller shares for R&D and profits.

The surplus calculation includes consumers' gains and deducts factor suppliers' opportunity costs (per their supply curves). We generally find the following: consumers receive 70 cents of each additional dollar of total surplus generated by the DOE's investments. Shareholders receive 16 cents, labor suppliers 7.3 cents, R&D suppliers 5.3 cents, and capital suppliers 1.8 cents. Consumers receive the largest share in every specification, although their marginal share is below their baseline (average) surplus shares. These estimates are only attainable through a range of restrictive assumptions and are most easily interpreted as partial equilibrium incidence metrics. Still, they provide a new view of who gains from public innovation.

Related Literature

First, we contribute to work measuring the returns to innovation and public R&D. A long tradition emphasizes that knowledge spillovers can make social returns exceed the private returns captured by innovators (Arrow 1962; Griliches 1979; Jones and Williams 1998; Bloom et al. 2013). Jones and Summers (2022b) quantify large economy-wide returns to innovation, while Fieldhouse and Mertens (2026a) estimate the productivity effects of federal non-defense R&D. Other work studies public research in settings including biomedical science, wartime research, the Space Race, and agriculture (Azoulay et al. 2019; Gross and Sampat 2023; Kantor and Whalley 2025; Akerman et al. 2026). We complement this work by studying how the gains are divided among the firms, input suppliers, and consumers connected to public innovation.

Second, we contribute to work on innovation rents and rent-sharing. Firm-specific innovation shocks can generate gains for both firms and their workers (Van Reenen 1996; Aghion et al. 2018; Kline et al. 2019; Alam and Diegmann 2026). A related literature studies how technological change, firm reallocation, and market power affect the division of income between labor and capital (Karabarbounis and Neiman 2014; Autor et al. 2020; De Loecker et al. 2020; Kehrig and Vincent 2021). We study these questions for publicly funded innovation and distinguish payments to capital suppliers from shareholders' residual profits. We also extend the accounting to consumers, whose gains are absent from measures of value added within firms.

Third, we contribute to the public finance literature on the incidence and welfare effects of government policy. Recent work develops methods for comparing policies using causal estimates of their benefits and fiscal costs (Hendren and Sprung-Keyser 2020; Finkelstein and Hendren 2020). Studies of business taxation show how policy effects are shared among workers, firm owners, and other suppliers (Suarez Serrato and Zidar 2016; Fuest et al. 2018). We apply this incidence perspective to public investment in innovation. Our framework distinguishes the distribution of additional income from the distribution of surplus after accounting for opportunity costs.

Fourth, we build on evidence that public grants and subsidies affect private R&D, innovation, and firm growth (Bronzini and Iachini 2014; Howell 2017; Azoulay et al. 2019; Myers and Lanahan 2022; Santoleri et al. 2024; Pless 2026). These firm responses are the starting point for our analysis. We show how a set of revenue, expenditure, and markup responses can be used to estimate the allocation of the resulting gains.

Methodologically, we develop a framework for using public investment variation when funding requests contain information about expected returns. The model makes explicit the conditions under which requests can control for this information, leaving residual investment variation suitable for identification. In our application, the residuals concern effective innovation stocks constructed from funding histories. We combine this technology-level variation with firms' nonrandom exposure using the approach of Borusyak and Hull (2023).

Finally, we contribute to work on energy innovation and mission-oriented public research (Howell 2017; Myers and Lanahan 2022; Dugoua and Gerarden 2025). The DOE supports research across technologies and institutions, making it a useful setting for studying the effects of public innovation beyond individual grant recipients. The distribution of these gains also matters for evaluating energy research policy, alongside its effects on technological progress, environmental quality, and energy security.

2. Political Production Functions

2.1. Setup

For notation, let x_t be flow of a variable at time t and X_t be the stock of that variable at time t (to include any discounting of historic flows).

The political production function involves the stock of a single input, where b is the flow and B 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:

$$B_{jt} = \delta \times B_{j(t-1)} + b_{j(t-\ell)} \,, \tag{1}$$

which assumes an implementation lag of ℓ years and a depreciation rate of $1 - \delta$; more complicated stock constructions are feasible. Technology-year output y_{jt} is produced according to the following exponential model:

$$y_{jt} = B_{jt}^{\beta} \times \exp\left(\omega_{jt} + \epsilon_{jt}\right) \quad \text{or, in logs} \quad \log\left(y_{jt}\right) = \beta \times \log\left(B_{jt}\right) + \omega_{jt} + \epsilon_{jt} \,, \tag{2}$$

where β is the output elasticity, $\omega_{jt} + \epsilon_{jt}$ is a composite productivity term, where ω_{jt} may be known to some 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 b 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. 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

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

First the executive requests an investment flow r_{jt}, which implies a requested stock per:

$$R_{j(t+1)} = \delta \times B_{jt} + r_{jt} \,. \tag{3}$$

The legislator determines the *actual* investment flows, b_{jt}, and realized stock, $B_{j(t+1)}$, 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 R_{jt} (the executive's choice) and B_{jt} (the legislator's choice).

The executive's objective function is:

$$\max_{R_{jt}} \quad \mathbb{E}\left[p^E_{jt} \times y_{jt}\Big(B_{jt}\big(R_{jt}\big)\Big) - c^E_{j(t-1)} \times b_{j(t-1)}\big(R_{jt}\big) \mid \mathcal{I}^E_{j(t-1)} \right], \tag{4}$$

where p^E_{jt} are explicit or implicit prices that describe the executive's constant marginal value of output, c^E_{jt} is the executive's marginal cost of inputs, and $\mathcal{I}^E_{jt}$ is their information

set.[1] Likewise, the legislator's objective function is:

$$\max_{B_{jt}} \quad \mathbb{E}\Big[\, p^L_{jt} \times y_{jt}(B_{jt}) - c^L_{j(t-1)} \times b_{j(t-1)}(B_{jt}) \mid R_{jt}, \mathcal{I}^L_{j(t-1)} \Big] \tag{5}$$

where again p^L_{jt} are prices that describe their marginal value of output, c^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 requested log stock) and the legislator's optimal choice (i.e., the actual log stock given the executive's request), can be written as:

$$\log(R_{jt}) = \mathcal{R}_{jt}(\omega_{jt}\,;\, \mathcal{I}^E_{j(t-1)}) \tag{6}$$

$$\log(B_{jt}) = \mathcal{B}_{jt}\Big(\mathcal{R}_{jt}(\omega_{jt}\,;\, \mathcal{I}^E_{j(t-1)}), \omega_{jt}\,;\, \mathcal{I}^L_{j(t-1)}\Big)\,, \tag{7}$$

where the functions $\mathcal{R}_{jt}(\cdot)$ and $\mathcal{B}_{jt}(\cdot)$ 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.

Assumption: Stable game — *Conditional on some observables* $\mathbf{w}_{jt}$*, 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:

$$\log(R_{jt}) = \mathcal{R}(\omega_{jt}, \mathbf{w}_{jt}\,;\, \mathcal{I}^E_{j(t-1)}) \tag{8}$$

$$\log(B_{jt}) = \mathcal{B}\Big(\mathcal{R}(\omega_{jt}, \mathbf{w}_{jt}\,;\, \mathcal{I}^E_{j(t-1)}), \omega_{jt}, \mathbf{w}_{jt};\,;\, \mathcal{I}^L_{j(t-1)}\Big)\,, \tag{9}$$

[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., $B_{jt}(R_{jt})$); and in the legislator's objective function, we condition on the executives request R_{jt}.

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 $\mathbf{w}_{jt}$ 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 (i.e., prices p_L and/or costs c^L) vary conditional on (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, p_L or c^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:

$$\text{Var}\Big(\log\big(B_{jt}\big) \mid \log\big(R_{jt}\big), \mathbf{w}_{jt}\Big) > 0. \tag{10}$$

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:

$$\begin{aligned}
\log\big(B_{jt}\big) &= \mathcal{B}\Big(\mathcal{R}\big(\omega_{jt}, \mathbf{w}_{jt}\,;\, \mathcal{I}^E_{j(t-1)}\big), \omega_{jt}, \mathbf{w}_{jt}\,;\, \mathcal{I}^L_{j(t-1)}\Big) \\
&= \mathcal{B}\Big(\mathcal{R}\big(\omega_{jt}, \mathbf{w}_{jt}\,;\, \mathcal{I}^E_{j(t-1)}\big), \mathbf{w}_{jt}\,;\, \mathcal{I}^L_{j(t-1)}\Big) \\
&= \mathcal{B}\Big(\log\big(R_{jt}\big), \mathbf{w}_{jt}\,;\, \mathcal{I}^L_{j(t-1)}\Big)
\end{aligned} \tag{11}$$

This assumption allows requests to be endogenous to productivity. The executive may observe (or estimate) ω_{jt} and use that information when choosing investment requests, $r_{j(t-1)}$. However, the legislator cannot have more information on productivity than the executive.

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

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:

$$\phi\Big(\log(R_{jt}), \mathbf{w}_{jt}\Big) \equiv \mathbb{E}\Big[\log(B_{jt}) \mid \log(R_{jt}), \mathbf{w}_{jt}\Big] . \tag{12}$$

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

$$\eta_{jt} \equiv \log(B_{jt}) - \phi\Big(\log(R_{jt}), \mathbf{w}_{jt}\Big) , \tag{13}$$

where $\mathbb{E}\Big[\eta_{jt} \mid \log(R_{jt}), \mathbf{w}_{jt}\Big] = 0$ by construction. The *Stable game* assumption makes $\phi(\cdot)$ an object that can be estimated by pooling observations (.g., within years or technologies), while the continuity restriction justifies approximating it flexibly. The *Private wedges* assumption gives relevance by requiring that this residual component is not degenerate:

$$\mathrm{Var}\Big(\eta_{jt} \mid \log(R_{jt}), \mathbf{w}_{jt}\Big) > 0 . \tag{14}$$

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

$$\mathbb{E}\Big[\log(B_{jt}) \mid \log(R_{jt}), \mathbf{w}_{jt}, \omega_{jt}\Big] = \mathbb{E}\Big[\log(B_{jt}) \mid \log(R_{jt}), \mathbf{w}_{jt}\Big] . \tag{15}$$

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):

$$\begin{aligned}\mathbb{E}\Big[\eta_{jt} \mid \log(R_{jt}), \mathbf{w}_{jt}, \omega_{jt}\Big] &= \mathbb{E}\Big[\log(B_{jt}) \mid \log(R_{jt}), \mathbf{w}_{jt}, \omega_{jt}\Big] - \phi\Big(\log(R_{jt}), \mathbf{w}_{jt}\Big) \\ &= 0 \,.\end{aligned} \tag{16}$$

Then, by the law of iterated expectations:

$$\begin{aligned}\mathbb{E}\big[\eta_{jt} \mid \omega_{jt}\big] &= \mathbb{E}\Big[\mathbb{E}\Big[\eta_{jt} \mid \log(R_{jt}), \mathbf{w}_{jt}, \omega_{jt}\Big] \mid \omega_{jt}\Big] \\ &= 0 \,.\end{aligned} \tag{17}$$

Therefore:

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

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 $\Big(\log(B_{jt}), \log(R_{jt}), \mathbf{w}_{jt}\Big)$, we can estimate $\phi(\cdot)$ using, for instance, semi-parametric or machine-learning methodologies and construct:

$$\widehat{\eta}_{jt} = \log(B_{jt}) - \widehat{\phi}\Big(\log(R_{jt}), \mathbf{w}_{jt}\Big) \,. \tag{19}$$

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

$$\log(B_{jt}) = \alpha \times \widehat{\eta}_{jt} + u_{jt} \tag{20}$$

$$\log(y_{jt}) = \beta \times \log(B_{jt}) + \big(\omega_{jt} + \epsilon_{jt}\big) \,, \tag{21}$$

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 $\mathbf{w}_{jt}$ used in the estimation of windfalls.

We report conventional 2SLS standard errors as a benchmark and approximate randomization-inference p-values following Borusyak and Hull (2023). The latter require an additional assignment assumption: conditional on requested stocks, controls, the sample, and potential outcomes under the null, the joint distribution of windfall shocks is unchanged by the permutations we use. This is stronger than the conditional mean restriction in Equation (15).

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 B_0 denote the predetermined investment stock at the time of the executive's request. Let R denote the new requested stock and B the realized stock. The associated requested and realized flows are: $r = R - B_0$ and $b = B - B_0$, respectively.

We focus on the interior where $R > B_0$ and $B > B_0$ and normalize the price of inputs to one for both parties, with no lags or discounting of investment. Following the notation above, uppercase R and B denote stocks, lowercase r and b denote flows, and logarithms are written explicitly. The production function is:

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

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:

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

Thus:

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

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

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:

$$\max_B \quad p_L \times \sqrt{B} \times \exp(\widehat{\omega}) - \left(B - B_0\right).$$

The legislator's optimal realized stock is:

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

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:

$$\max_{\widehat{\omega}} \quad \mathbb{E}_\nu\Big[\sqrt{B^*(\widehat{\omega}, \nu)} \times \exp(\omega) - \big(B^*(\widehat{\omega}, \nu) - B_0\big)\Big].$$

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

$$\max_{\widehat{\omega}} \quad \frac{\exp(\omega + \widehat{\omega})}{2} - \frac{\left(1+\sigma_\nu^2\right) \times \exp(2 \times \widehat{\omega})}{4} + B_0.$$

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:

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

Therefore the log requested stock is:

$$\log\big(R^*(\omega)\big) = 2 \times \omega - \log(4) - \log\left(1+\sigma_\nu^2\right),$$

which is monotone in productivity. For any observed requested stock, the legislators inference is thus given by[4]:

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

After observing R and ν, the legislator chooses:

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

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

Taking logs gives:

$$\log(B) = \log(R) + \log\Big(\frac{(1+\nu)^2}{1+\sigma_\nu^2}\Big).$$

With our identification strategy described above in mind, define:

$$\phi(\log(R)) \equiv \mathbb{E}[\log(B) \mid \log(R)].$$

Then:

$$\phi(\log(R)) = \log(R) + \mathbb{E}_{\widetilde{\nu}}\Big[\log\Big(\frac{(1+\widetilde{\nu})^2}{1+\sigma_\nu^2}\Big)\Big],$$

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

$$\eta = \log(B) - \phi(\log(R)) = \log\Big(\frac{(1+\nu)^2}{1+\sigma_\nu^2}\Big) - \mathbb{E}_{\widetilde{\nu}}\Big[\log\Big(\frac{(1+\widetilde{\nu})^2}{1+\sigma_\nu^2}\Big)\Big].$$

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:

$$\mathbb{E}[\eta \mid \log(R), \omega] = 0.$$

Therefore:

$$\mathbb{E}[\eta \times \omega] = \mathbb{E}[\omega \times \mathbb{E}[\eta \mid \log(R), \omega]] = 0.$$

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

Our budget data come from the DOE's annual *Congressional Budget Justification* documents. For each program and fiscal year, we record the executive's funding request and the amount Congress appropriated. We collect these amounts by hand, excluding administra-

tive costs and funding devoted solely to deployment.[5] Requests are available for 1989–2020 and appropriations for 1987–2020. We use the appropriations from 1987–1988 when constructing innovation stocks, but begin estimating windfalls in 1989, when we first observe requests.

DOE frequently changes its program names and how it groups research activities. We assign the budget lines to 14 technology categories that we follow over time: coal, natural gas, oil, carbon capture and storage, nuclear fission, solar, wind, geothermal, bioenergy, hydropower and marine energy, grid systems, energy storage, hydrogen and fuel cells, and batteries and electric vehicles. These are also the categories we use to measure firms' exposure in Section 6. Where the documents provide enough detail, we separate technologies reported within the same program. Where they do not, we retain combined categories.[6]

For appropriations, we use the "actual" or "adjusted" amounts reported in later budget documents when available. These figures account for rescissions, reallocations, and the use of prior-year balances. For requests, we use the most recent restatement of each year's figure. Some funding supports both research and other activities, so we must make assumptions about how much to count as R&D. This is particularly relevant for the 2009 American Recovery and Reinvestment Act, which funded manufacturing and deployment alongside research.[7]

The unit of observation is a technology j in fiscal year t. We convert all amounts to 2024 dollars using the GDP deflator. Table 1 summarizes the annual funding flows. Average requests are roughly $150 million per technology-year, compared with average appropriations of roughly $160 million.

[5] For early fossil-energy programs, we also use Department of the Interior budget justifications. We obtain the documents from HathiTrust, https://www.govinfo.gov/app/details/CHRG-106hhrg63101/CHRG-106hhrg63101, and the DOE's budget archive at https://www.energy.gov/cfo/listings/budget-justification-supporting-documents.

[6] For example, we separate wind from solar and fuel-cell research from fossil-energy programs when the funding amounts are reported separately. We keep hydropower and marine energy together because DOE combined them under "Water Power" in 2010–2013, and batteries and electric vehicles together because they merged under "Vehicle Technologies" in 2013. We also group hydrogen and fuel cells because these research programs overlap.

[7] We count 20% of "Transportation Electrification" funding and 30% of "Advanced Battery Manufacturing" funding as R&D. We assign the remaining funds to manufacturing and deployment.

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

Note: Reports annual flows in millions of dollars, using matched requests and appropriations from 1989–2020.

3.2. Timing Public Innovation

Public investment in innovative R&D can take many years to affect private production. For example, Fieldhouse and Mertens (2026a) find that the productivity response to non-defense R&D spending builds gradually and is statistically significant at horizons of eight to fifteen years. In the case of biomedical research, Azoulay et al. (2019) document substantial delays between NIH funding and subsequent patents citing the research it supports, with these links continuing to emerge more than a decade after funding. The traditional stock of investments described in the previous section was useful for illustrating our identification strategy, and may be a useful model for more traditional investments, but it may not be a credible model of innovation. Here, we describe an alternative approach that allows investment in public innovation to have a more drawn out effect.

Let $g_h(\lambda)$ describe the contribution of an investment to the effective public innovation stock h years later. For technology j, this stock is:

$$B_{jt}(\lambda) = \sum_{h=0}^{\infty} g_h(\lambda) \times b_{j(t-h)} \,, \tag{22}$$

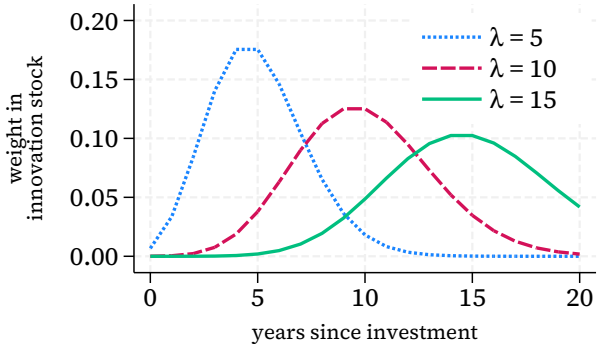
where b_{jt} is the annual flow of real appropriations. We use a poisson timing assumption which implies the poisson weights are:

$$g_h(\lambda) = \exp(-\lambda) \times \frac{\lambda^h}{h!} \,. \tag{23}$$

These weights are nonnegative and sum to one. Normalizing the weights to sum to one sets the scale of the stock; it does not impose a dollar return to public investment. They allow an investment's contribution to rise as research develops and reaches private producers, then decline as that research becomes less useful. For the empirical analysis, this profile replaces the fixed implementation lag and geometric depreciation in Equation (1).

Given this poisson timing assumption, the parameter λ describes the mean horizon

FIGURE 1. Poisson Production of the Stock of Public Innovation



Note: Shows the Poisson timing profiles.

at which investment yields output relevant for the stock of public innovation; it also determines the dispersion, since the poisson variance equals its mean. We consider $\lambda \in \{5, 10, 15\}$ to assess how our results depend on the assumed timing. These alternatives are motivated by the evidence of long delays described above. As a descriptive check, fitting scaled Poisson cumulative distributions to digitized curves from these studies gives mean horizons of about 7 years for the productivity response in Fieldhouse and Mertens (2026a) and 9 years for the pooled grant-to-patent linkage curve in Azoulay et al. (2019). These fits help motivate our timing assumptions; they do not directly estimate the weights in DOE's innovation stock.[8] Figure 1 illustrates these alternative timing assumptions.

3.3. 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

[8]We assume that investment in each technology begins in its first observed funding year, setting all earlier flows to zero. We construct stocks from observed funding using the Poisson weights above.

for that year to a new "Unconventional Fossil Energy Technologies" program. The impetus appears to have come from legislators representing oil- and gas-producing states, rather than from any new assessment of the technology's scientific promise;.

Nuclear fission. 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. Nuclear-fission appropriations fell below requests in 1997–2000 and exceeded them in 2001–2006. These annual funding gaps are distinct from the estimated stock windfalls below, which also account for funding history and the other predictors in our empirical model. The DOE may well have anticipated this support and shaded its nuclear requests accordingly, so part of the wedge may have been predictable. But recall, predictability *per se* 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.

3.4. Legislative Windfalls: Estimation and Results

We use the public innovation stocks constructed above to estimate legislative windfalls. Following the investment game in Section 2, the requested stock is based on past appropriations and the the executive's request:

$$R_{jt}(\lambda) = \sum_{h=1}^{\infty} g_h(\lambda) \times b_{j(t-h)} + g_0(\lambda) \times r_{jt}\,, \tag{24}$$

where r_{jt} is the annual flow of requested funding. Requested and realized stocks therefore contain the same funding history and differ in the current year's request and appropriation. We construct both stocks under each timing assumption, $\lambda \in \{5, 10, 15\}$, and suppress λ below to simplify notation.

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

$$\log\big(B_{jt}\big) = \phi\Big(\log\big(R_{jt}\big), \mathbf{w}_{jt}\Big) + \eta_{jt}\,. \tag{25}$$

The estimated windfall is the difference between realized log stocks and their predicted value:

$$\widehat{\eta}_{jt} = \log(B_{jt}) - \widehat{\phi}\left(\log(R_{jt}), \mathbf{w}_{jt}\right). \tag{26}$$

This removes predictable differences between requests and appropriations, including any tendency for the executive to systematically request more or less than Congress ultimately provides.

For each timing assumption, we consider three nested sets of predictors: (*i*) requested log stocks and year fixed effects; (*ii*) adds a cubic spline in requested log stocks and fixed effects for presidential administration interacted with clean-versus-dirty technology status; (*iii*) adds interactions of requested log stocks and their spline terms with presidential administration indicators and with the administration-by-clean/dirty indicators. The spline knots are placed at the 25th, 50th, and 75th percentiles of requested log stocks.

We estimate each specification using four approaches: lasso with plug-in penalties, adaptive lasso, lasso with cross-validation, and elastic net with cross-validation. For the cross-validated estimates, we use the one-standard-error rule.[9]

Averaging across all of these sets of specification choices, the standard deviation of $\widehat{\eta}_{jt}$ is 0.013, implying that a one-standard-deviation windfall corresponds to roughly a 1.3% change in realized DOE stock relative to predicted stock. A positive one-standard-deviation windfall corresponds to an effective innovation stock about $1.4 million above its predicted value, evaluated at each observation's predicted stock and averaged across observations and specifications. The corresponding average predicted stock is $110 million.

Figure 2 shows the windfalls by technology over time, after residualizing on year fixed effects and averaging across our plausible specifications.

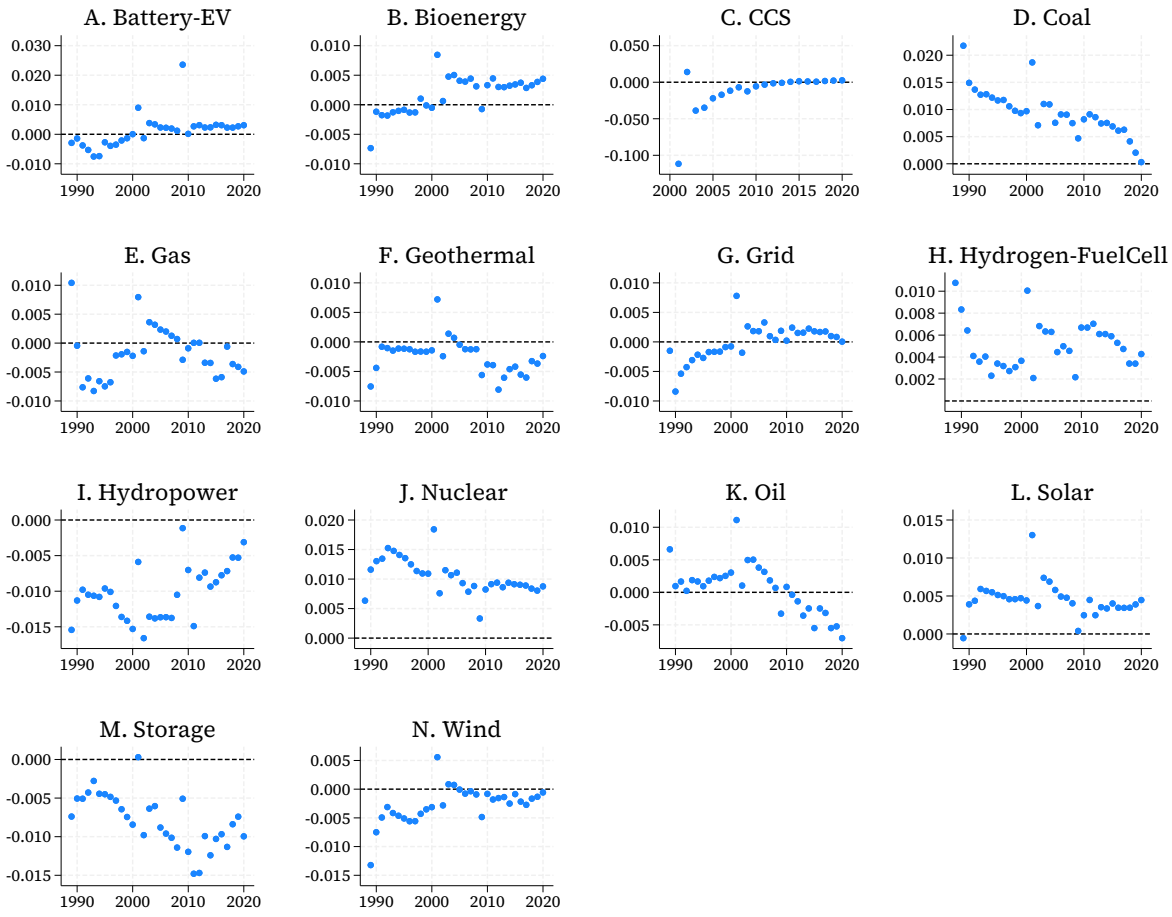
4. A Model of the Private Returns to Public Innovation

4.1. Factor Supply, Production, and Output Demand

Let firm i in year t produce physical output q_{it} at price $p_{q,it}$. The firm uses labor l_{it}, capital k_{it}, intermediate inputs m_{it}, and research and development rd_{it}, which are mutually exclusive and measured here as flows (e.g., capital services with a rental cost). We let $x \in \{l, k, rd\}$ index the three focal factors and $\mathbf{x} = (l, k, rd)$ be the vector of all three focal

[9] For elastic net, we consider mixing parameters of 0.1, 0.25, 0.5, 0.75, and 1.

FIGURE 2. Windfalls across Technologies and Time



Note: Panels (A.)–(N.) show windfalls by technology, in log points. Vertical scales differ.

factors. The physical production function is:

$$q_{it} = \Omega_{it}\left(\mathbf{x}_{it}, m_{it}\right) . \tag{27}$$

Let $p_{x,it}$ be the price of input x so total expenditures on each factor are:

$$c_{x,it} = p_{x,it} \times x_{it} , \tag{28}$$

and likewise for intermediates m. Factor supply schedules are isoelastic:

$$p_{x,it} = \sigma_{x,it} \times x_{it}^{1/\varepsilon_{x,it}} , \tag{29}$$

where $\sigma_{x,it} > 0$ is a supply shifter and $\varepsilon_{x,it} > 0$ is the factor's supply elasticity.

Each firm faces residual isoelastic demand for its output given by the inverse demand function:

$$p_{q,it} = \sigma_{q,it} \times q_{it}^{-1/\varepsilon_{q,it}} \quad \text{where } \varepsilon_{q,it} > 1\,, \tag{30}$$

where $\sigma_{q,it} > 0$ is a demand shifter and $\varepsilon_{q,it}$ is the absolute value of the demand elasticity. Finite supply and demand elasticities allow for monopsony and monopoly power, respectively.

4.2. Value-add and Surplus Metrics

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

$$\text{VA}_{it} \equiv p_{q,it} \times q_{it} - c_{m,it}\,. \tag{31}$$

Value added comprises the claims of labor, R&D, capital, and shareholders' profits. A factor x's share of value added at each firm is also an accounting expression:

$$\text{S}^{\text{VA}}_{x,it} \equiv \frac{c_{x,it}}{\text{VA}_{it}} = \frac{p_{x,it} \times x_{it}}{p_{q,it} \times q_{it} - c_{m,it}}\,. \tag{32}$$

To solve for economic surplus metrics, note that for each input x, the isoelastic supply curve in (29) implies the firm's input-suppliers' producer surplus is: For surplus accounting, we interpret inverse input supply as marginal opportunity cost and inverse output demand as marginal willingness to pay; thus, we measure surplus relative to zero trade. For each factor x, the isoelastic schedules imply that suppliers' producer surplus is:

$$\begin{aligned} \text{PS}_{x,it} &\equiv c_{x,it} - \int_0^{x_{it}} \sigma_{x,it} \times u^{1/\varepsilon_{x,it}}\, du \\ &= \frac{1}{\varepsilon_{x,it} + 1} \times c_{x,it}\,. \end{aligned} \tag{33}$$

We cannot trace input-output links to intermediate-input suppliers, so their surplus lies outside our accounting boundary. Letting firm profits be $\Pi_{it} \equiv p_{q,it} \times q_{it} - c_{m,it} - \sum_x c_{x,it}$, we define the combined producer surplus of the firm and its focal factor suppliers as:

$$\begin{aligned} \text{PS}_{it} &\equiv \Pi_{it} + \sum_x \text{PS}_{x,it} \\ &= \text{VA}_{it} - \sum_x \frac{\varepsilon_{x,it}}{\varepsilon_{x,it} + 1} \times c_{x,it}\,. \end{aligned} \tag{34}$$

A factor x's share of this producer surplus is:

$$\mathbb{S}^{\text{PS}}_{x,it} \equiv \frac{\text{PS}_{x,it}}{\text{PS}_{it}} = \frac{\frac{1}{\varepsilon_{x,it}+1} \times c_{x,it}}{\text{PS}_{it}}, \tag{35}$$

and likewise for the other factors. For consumers, defined here as the direct buyers of the firm's output, isoelastic demand implies that their surplus is:

$$\begin{aligned} \text{CS}_{it} &\equiv \int_0^{q_{it}} \sigma_{q,it} \times u^{-1/\varepsilon_{q,it}} du - p_{q,it} \times q_{it} \\ &= \frac{p_{q,it} \times q_{it}}{\varepsilon_{q,it} - 1}, \end{aligned} \tag{36}$$

which can be used to solve for total surplus ($\text{PS}_{it} + \text{CS}_{it}$) and factors' shares thereof.

Letting MC_{it} denote a firms' marginal cost of production (accounting for firms' input supply schedules), the isoelastic residual demand curve above implies that the profit-maximizing price for firm i is:

$$p_{q,it} \times \left(1 - \frac{1}{\varepsilon_{q,it}}\right) = \text{MC}_{it}, \qquad \mu_{it} \equiv \frac{p_{q,it}}{\text{MC}_{it}} = \frac{\varepsilon_{q,it}}{\varepsilon_{q,it} - 1} > 1. \tag{37}$$

Here μ_{it} is the firm's markup over marginal cost. Consumer surplus can therefore be written as:

$$\text{CS}_{it} = p_{q,it} \times q_{it} \times \left(\mu_{it} - 1\right). \tag{38}$$

4.3. Empirical Model: Firm-level Production

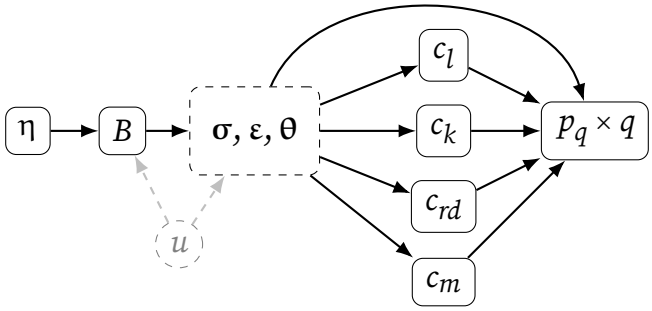
In general, publicly funded innovation can affect a factor's share of value added (or surplus) through three sets of channels:

- product demand level ($\sigma_{q,it}$) and / or elasticity ($\varepsilon_{q,it}$);
- factor productivity (level shifters in $\mathbb{Q}_{it}$), returns-to-scale or substitutability (curvature parameters in $\mathbb{Q}_{it}$);
- factor supply level ($\sigma_{x,it}$) or elasticity ($\varepsilon_{x,it}$).

Let $\boldsymbol{\sigma}$ and $\boldsymbol{\varepsilon}$ collect the supply and demand shifters and elasticities, and let θ collect the production parameters. Let B_{it} denote the stock of publicly funded R&D relevant to firm i in year t, η denote exogenous variation in public innovation, and u denote other unobserved forces that may affect both public innovation and market conditions. Figure 3 illustrates how public innovation affects firm outcomes through these primitives and

firms' input choices.

FIGURE 3. Public Innovation and Firm-level Outcomes



For each positive outcome y_{it}, let $\beta_y \equiv d\log(y_{it})/d\log(B_{it})$ denote its total elasticity with respect to public innovation. These responses allow input quantities, input prices, output, and product prices to adjust. Our empirical specification assumes a common elasticity for each outcome across firms and years:

$$\log(y_{it}) = \alpha_y + \beta_y \times \log(B_{it}) + \Gamma_y \times \mathbf{z}_{it} + u_{yit}\,, \tag{39}$$

where $y_{it} \in \{p_{q,it} \times q_{it}, c_{x,it}, c_{m,it}, \mu_{it}\}$ and $\mathbf{z}_{it}$ contains identifying covariates, including fixed effects. We estimate each outcome equation separately, instrumenting public innovation using the exogenous funding variation described in Section 6.

Limitations of Input-conditional Regressions. Our objective is to recover total responses to public innovation, including changes in firms' input choices. To illustrate why conditioning on these choices can introduce bad controls, consider:

$$\log(\text{VA}_{it}) = \alpha + \gamma_l \times \log(c_{l,it}) + \gamma_k \times \log(c_{k,it}) + \xi \times \log(B_{it}) + u_{it}\,. \tag{40}$$

The coefficient ξ describes the relationship between public innovation and value added conditional on labor and capital expenditures. If public innovation changes factor supply, production, or demand, firms may adjust these expenditures. To the extent that these adjustments contribute to the change in value added, conditioning on them removes part of the total effect of interest. Such conditioning can also introduce bias because expenditures respond to unobserved determinants of value added. Consequently, ξ need not recover the total value-added elasticity.

We therefore exclude contemporaneous input expenditures when estimating the revenue response. When a given input expenditure is the outcome, we exclude other input

expenditures, which may be jointly affected by public innovation and unobserved firm conditions.

4.4. Incidence of Marginal Value-add

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

$$\mathcal{M}^{\text{VA}}_{x,it} \equiv \frac{dc_{x,it} \,/\, d\log(B_{it})}{d\text{VA}_{it} \,/\, d\log(B_{it})} = \frac{c_{x,it} \times \beta_{c_x}}{p_{q,it} \times q_{it} \times \beta_{pq} - c_{m,it} \times \beta_{c_m}} \,. \tag{41}$$

Shareholders receive the residual of value added after payments to the focal factors. Their share of marginal value added is therefore:

$$\mathcal{M}^{\text{VA}}_{\Pi,it} \equiv \frac{d\Pi_{it} \,/\, d\log\left(B_{it}\right)}{d\text{VA}_{it} \,/\, d\log\left(B_{it}\right)} = 1 - \sum_{x} \mathcal{M}^{\text{VA}}_{x,it} \,. \tag{42}$$

Thus, the same revenue and expenditure responses determine the incidence of marginal value added across factors and shareholders, without requiring a separate regression for profits. The marginal shares defined here and below require a nonzero total response. They sum to one across the included claimants but need not lie between zero and one, since individual claimants can experience gains or losses.

4.5. Incidence of Marginal Producer Surplus

Allowing factor supply elasticities to respond to public innovation, the change in factor x's producer surplus is:

$$\frac{d\text{PS}_{x,it}}{d\log\left(B_{it}\right)} = \text{PS}_{x,it} \times \left(\beta_{c_x} - \frac{\varepsilon_{x,it}}{\varepsilon_{x,it}+1} \times \beta_{\varepsilon_x}\right), \tag{43}$$

where β_{ε_x} is the elasticity of the factor's supply elasticity with respect to public innovation. Expenditure responses alone do not identify this parameter. We therefore calibrate supply elasticities and maintain the following restriction.

Assumption: Fixed Factor Supply Elasticities. Public innovation does not change factor supply elasticities. It may change supply shifters $\sigma_{x,it}$, as well as production and demand conditions. 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.

Under this assumption, a factor's share of marginal producer surplus is:

$$\begin{aligned}\mathcal{M}^{\text{PS}}_{x,it} &\equiv \frac{d\text{PS}_{x,it} \,/\, d\log\left(B_{it}\right)}{d\text{PS}_{it} \,/\, d\log\left(B_{it}\right)} \\ &= \frac{\frac{1}{\varepsilon_{x,it}+1} \times c_{x,it} \times \beta_{c_x}}{p_{q,it} \times q_{it} \times \beta_{pq} - c_{m,it} \times \beta_{c_m} - \sum_x \frac{\varepsilon_{x,it}}{\varepsilon_{x,it}+1} \times c_{x,it} \times \beta_{c_x}} \,.\end{aligned} \tag{44}$$

For each calibration, we assign each factor a common supply elasticity across firms and years. We consider:

- Labor supply elasticity: $\varepsilon_l \in \{1.5, 3, 6\}$,
- Capital supply elasticity: $\varepsilon_k \in \{5, 10, \infty\}$,
- R&D supply elasticity: $\varepsilon_{rd} \in \{0, 1, 2\}$.

As $\varepsilon_x \searrow 0$, the fraction of expenditures accruing as supplier surplus approaches one; as $\varepsilon_x \nearrow \infty$, it approaches zero.

Shareholders' producer surplus is their profits. Their share of marginal producer surplus is therefore:

$$\mathcal{M}^{\text{PS}}_{\Pi,it} \equiv \frac{d\Pi_{it} \,/\, d\log\left(B_{it}\right)}{d\text{PS}_{it} \,/\, d\log\left(B_{it}\right)} = 1 - \sum_x \mathcal{M}^{\text{PS}}_{x,it} \,. \tag{45}$$

Holding the estimated revenue and expenditure responses fixed, this share depends on the calibrated supply elasticities.

4.6. Incidence of Marginal Consumer and Total Surplus

The markup elasticity β_μ supplies the additional response needed to measure consumer gains. Differentiating Equation (38) gives:

$$\frac{d\text{CS}_{it}}{d\log\left(B_{it}\right)} = p_{q,it} \times q_{it} \times \left(\left(\mu_{it} - 1\right) \times \beta_{pq} + \mu_{it} \times \beta_\mu\right) . \tag{46}$$

The two terms capture the revenue and markup contributions to the consumer-surplus response. A negative markup response reduces consumer gains for a given revenue

response. Adding the producer- and consumer-surplus responses gives:

$$\frac{d\mathrm{TS}_{it}}{d\log\left(B_{it}\right)} = p_{q,it} \times q_{it} \times \mu_{it} \times \left(\beta_{pq} + \beta_{\mu}\right) - c_{m,it} \times \beta_{c_m} - \sum_{x} \frac{\varepsilon_{x,it}}{\varepsilon_{x,it}+1} \times c_{x,it} \times \beta_{c_x} \,. \tag{47}$$

The shares of marginal total surplus received by factor suppliers, consumers, and shareholders are, respectively:

$$\begin{aligned} \mathcal{M}^{\mathrm{TS}}_{x,it} &\equiv \frac{d\mathrm{PS}_{x,it} \,/\, d\log\left(B_{it}\right)}{d\mathrm{TS}_{it} \,/\, d\log\left(B_{it}\right)}\,, \\ \mathcal{M}^{\mathrm{TS}}_{\mathrm{CS}it} &\equiv \frac{d\mathrm{CS}_{it} \,/\, d\log\left(B_{it}\right)}{d\mathrm{TS}_{it} \,/\, d\log\left(B_{it}\right)}\,, \\ \mathcal{M}^{\mathrm{TS}}_{\Pi,it} &\equiv 1 - \sum_{x} \mathcal{M}^{\mathrm{TS}}_{x,it} - \mathcal{M}^{\mathrm{TS}}_{\mathrm{CS}it}\,. \end{aligned} \tag{48}$$

4.7. Aggregate Incidence and Returns

The marginal incidence measures above describe how the gains from innovation are divided at each firm. For each specification, we evaluate the same proportional increase in public innovation at every firm-year. We sum each claimant's gains and divide by the summed increase in value added or surplus. This gives each claimant's share of the combined gains, rather than an average of firm-level shares. Interpreting these combined gains as aggregate effects requires further consideration.

Our regressions (Eq. 39) measure how firms respond to public innovation. For instance, they allow input choices and output prices to adjust, but they do not identify aggregate effects. For example, one firm may expand partly by taking business from its competitors. We therefore interpret the resulting estimates as partial-equilibrium responses. To obtain a more general-equilibrium interpretation, we must assume that any aggregate effects of public innovation yield the same incidence shares that we estimate. One model consistent with this assumption is that positive and negative externalities generated by public innovation occur between firms with the same baseline technologies and market power. Here, positive externalities (e.g., knowledge spillovers) can increase aggregate gains, while congestion in implementing innovation can reduce them, but neither changes the shares.[10]

[10] Of course, arbitrary business stealing, changes in relative input prices, or reallocation among firms with different incidence shares can generate discrepancies between the partial-equilibrium sort of estimates we can obtain and the general-equilibrium estimates of interest.

5. Firm-level Data

5.1. Compustat Data

Our firm-level analysis requires firm-year measures of revenues, input expenditures, and industry classifications over a long horizon. The Compustat dataset is one of the few non-administrative datasets that meets these requirements. Compustat, which reports annual financial information for publicly traded firms, provides the sales and expenditure measures needed to construct the objects in Section 4. Using Compustat also allows us to link firms' to their patenting activity and, through those patents, to DOE-funded technologies, a process we describe in the next Section. Compustat firms are substantially larger than the typical US firm and account for a much larger share of employment and output than their share of the firm population (Dinlersoz et al. 2024).

The unit of observation is a firm-year. We convert monetary variables to real 2024 US dollars using the GDP deflator. Our estimation sample requires usable two-digit NAICS information, positive revenue and all four expenditure measures, and a markup greater than one. After matching firms to their patent exposure, we retain observations with positive exposure measures under every specification and drop industry-year cells containing only one firm. We use this common sample for every outcome and specification; Table 2 summarizes it.

5.2. Revenue and Input Expenditures

In general, we engage with the Compustat data using prevailing approaches in the literature. We adapt the labor-imputation approach of Hartman-Glaser et al. (2019) and the capital-cost measures of De Loecker et al. (2020) and Benkard et al. (2025) to construct the objects required by our model. Table 2 summarizes the key variables.

Revenues, $p_{q,it} \times q_{it}$. We measure revenue, using Compustat sales (`sale`). This records the total amount customers pay for the firm's output. We do not separately observe the firm's output price or the physical quantity it sells; our revenue regressions allow both to respond to public innovation.

R&D expenditures, $c_{rd,it}$. We measure R&D expenditures using reported research and development expenses (`xrd`), which we treat as an approximation to the annual cost of composite research services. These expenditures include employee compensation, capital costs, and purchased inputs.

Labor expenditures, $c_{l,it}$. We first construct total-company labor expenditures using reported labor and related expenses (`xlr`). Many firms do not report this item, so we impute missing labor expenditures using employment and the labor expenses reported by other firms. Specifically, we predict log labor expense per employee using industry and year effects, exponentiate the prediction, and multiply it by the firm's number of employees (`emp`, converted from thousands to workers). We begin with four-digit NAICS industries and use broader industry groups when the more detailed prediction is unavailable. Reported labor expenses are retained when available. Because this total includes research employees, we subtract the labor expenditures assigned to R&D.[11]

Capital (service) expenditures, $c_{k,it}$. Capital expenditures (i.e., implicit rental rates times the capital quantity stock) measure the annual cost of using capital. Purchases of new equipment do not directly measure this cost, since equipment can provide services for many years. We first construct the total-company capital-service bill as $\rho_t \times K_{it}$, where K_{it} is gross property, plant, and equipment (`ppegt`) in real dollars. Following De Loecker et al. (2020), we use the annual user-cost series supplied with the replication files of Benkard et al. (2025). Its construction subtracts inflation from the nominal federal funds rate and adds a 12% allowance for depreciation and risk. The rate varies by year and is common to all firms and industries. We then subtract the capital services assigned to R&D.[12]

Intermediate expenditures, $c_{m,it}$. We begin with operating expenditures, measured as cost of goods sold (`cogs`) plus selling, general, and administrative expenses (`xsga`), adjusted for the treatment of R&D. We subtract total-company labor expenditures and the 35% of R&D expenditures assigned to purchased research inputs: $c_{m,it} = \texttt{cogs}_{it} + \texttt{xsga}^{*}_{it} - c^{\text{total}}_{l,it} - 0.35 \times c_{rd,it}$. Here, $\texttt{xsga}^{*}_{it}$ adjusts SG&A for the treatment of R&D.[13] The labor deduction uses the total-company measure, before removing research employees. Our accounting checks show that sales less COGS and unadjusted SG&A correspond to operating

[11] Compustat does not separately identify R&D factors, so we assume that 60% of R&D expenditures pay for the firm's own research employees. This allocation assumption is motivated by federal statistics (National Center for Science and Engineering Statistics 2019).

[12] As in the case of R&D labor, Compustat does not separately identify R&D capital, so we assume that 5% pay for capital services used in research. Reported depreciation accounts for approximately 3.5% of the same expenditure total (National Center for Science and Engineering Statistics 2019), but we choose a 5% capital-service share to allow for a return on capital as well as depreciation. This distinction follows the economic accounting described by Howells et al. (2018).

[13] We set $\texttt{xsga}^{*}_{it} = \texttt{xsga}_{it}$ when $\texttt{xsga}_{it} \geq \texttt{xrd}_{it}$ and $\texttt{xsga}^{*}_{it} = \texttt{xsga}_{it} + \texttt{xrd}_{it}$ otherwise. This convention treats R&D as included in reported SG&A in the first case and adds it to the operating-cost total in the second.

income before reported depreciation. We therefore do not separately subtract reported depreciation when constructing intermediates.

Profits, Π_{it}. Profits here are the residual defined by the model: $\Pi_{it} = p_{q,it} \times q_{it} - c_{l,it} - c_{k,it} - c_{rd,it} - c_{m,it}$.

TABLE 2. Compustat Sample Descriptive Statistics

	N	Mean	SD	p25	p50	p75
Revenue	8,334	16,719.1	37,106.3	722.1	3,079.3	14,622.8
Non-R&D labor	8,334	2,474.4	4,972.4	120.5	596.8	2,469.2
Non-R&D capital services	8,334	1,433.7	4,302.7	34.9	193.7	1,040.2
Non-R&D intermediates	8,334	10,450.0	25,835.7	392.8	1,664.2	8,158.6
Composite R&D services	8,334	891.3	2,062.9	27.1	108.4	591.3
Profits	8,334	1,469.8	4,611.0	8.4	176.8	1,025.9
Markup	8,334	1.52	0.66	1.15	1.31	1.62

Note: Reports the regression sample for 1989–2016. Dollar amounts are in millions.

5.3. Markups and Demand Elasticities

We estimate firm-year markups following the methodology of De Loecker et al. (2020). Under the residual-demand model of Section 4, a markup μ_{it} implies a residual demand elasticity $\varepsilon_{q,it} = \mu_{it}/(\mu_{it} - 1)$; 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.

Specifically, we follow Benkard et al. (2025) closely, thanks to their publicly available code-base.[14] 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).[15] Table 2 summarizes the markup estimates and we refer the reader to De Loecker et al. (2020) and Benkard et al. (2025) for more on variation in markups in the Compustat sample.

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

[15] 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 log investment, lagged log capital, and squared lagged log capital, instrumenting current log COGS with lagged log COGS. Capital in this step is deflated gross PP&E, and investment is $i_{it} = K_{it} - 0.9 \times K_{i(t-1)}$.

6. Firms' Non-random Exposure to Exogenous Public Innovation

6.1. Empirical Model: Firm-level Production and Technology-level Innovation

The model in Section 4 motivated the reduced-form regression of Equation (39):

$$\log(y_{it}) = \alpha_y + \beta_y \times \log(B_{it}) + \Gamma_y \times \mathbf{z}_{it} + u_{yit} \,,$$

for the focal outcomes $y_{it} \in \{p_{qit} \times q_{it},\, c_{xit},\, c_{mit},\, \mu_{it}\}$.

The problem is that our source of plausibly exogenous variation in public innovation is not at the level of *firm-year* but rather at the *technology-year*. 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:

$$B_{it} = \sum_j \chi_{ijt} \times B_{jt} \,, \tag{49}$$

where χ_{ijt} measures firm i's exposure to technology j in year t. We construct these χ weights from firms' patenting records, as described below. Substituting this exposure measure into Equation (39) gives:

$$\log(y_{it}) = \alpha_y + \beta_y \times \log\Big(\sum_j \chi_{ijt} \times B_{jt}\Big) + \Gamma_y \times \mathbf{z}_{it} + u_{yit} \,, \tag{50}$$

The control vector $\mathbf{z}$ accounts for differences in firms' total technology exposure and for common shocks. We always include $\log\big(\sum_{j\in\mathcal{J}_t} \chi_{ijt}\big)$, where $\mathcal{J}_t$ contains the technologies with observed funding requests in year t. We also consider three sets of fixed effects: year effects, year and two-digit NAICS industry effects, and industry-by-year effects.

Notably, we do not need an exogenous measure of χ_{ijt} (i.e., that is uncorrelated with productivity), since we can leverage the methodology of Borusyak and Hull (2023). Recall from Section 2 that our technology-level windfall estimates are:

$$\widehat{\eta}_{jt} = \log(B_{jt}) - \widehat{\phi}\left(\log(R_{jt}), \mathbf{w}_{jt}\right) \,.$$

Thus, the technology-year level stock of public innovation can be written as:

$$B_{jt} = \exp\left(\widehat{\phi}\left(\log(R_{jt}), \mathbf{w}_{jt}\right) + \widehat{\eta}_{jt}\right) \,. \tag{51}$$

Following Borusyak and Hull (2023), we can generate an instrument for firm-level public innovation exposure, B_{it}, that isolates only the variation in technology-level public innovation, B_{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:

$$B_{jt}^{(s)} = \exp\left(\widehat{\phi}\left(\log(R_{jt}), \mathbf{w}_{jt}\right) + \widehat{\eta}_{jt}^{(s)}\right), \tag{52}$$

We use $S = 1{,}000$ draws, randomly reassigning estimated windfalls across the technology-year observations used in windfall estimation. Stocks without an estimated windfall remain at their observed values. Holding firms' exposure weights fixed, we recompute firm-year exposure under each draw and, following Borusyak and Hull (2023), construct the "recentered" instrument:

$$\text{IV}_{it} = \log\left(\sum_j \chi_{ijt} \times B_{jt}\right) - \frac{1}{S} \times \sum_{s=1}^{S} \log\left(\sum_j \chi_{ijt} \times B_{jt}^{(s)}\right). \tag{53}$$

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.[16]

6.2. Exposure Metric and Patent Data

We use the patent record 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 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 $1 - \delta_N$ and lag ℓ_N, the

[16] We use the same simulated treatments that generated the recentered instrument. Define the simulated recentered instruments $\text{IV}_{it}^{(s)} = \log\left(\sum_j \chi_{ijt} \times B_{jt}^{(s)}\right) - \frac{1}{S} \times \sum_{s=1}^{S} \log\left(\sum_j \chi_{ijt} \times B_{jt}^{(s)}\right)$. 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 $T_{\text{obs}} = \left|\sum_{it} \text{IV}_{it} \times \tilde{u}_{it}\right|$, where $\tilde{u}_{it}$ is the residualized outcome under the null. For each draw we compute $T_s = \left|\sum_{it} \text{IV}_{it}^{(s)} \times \tilde{u}_{it}\right|$, and the randomization-inference p-value is $\frac{1 + \sum_{s=1}^{S} \mathbf{1}\{T_s \geq T_{\text{obs}}\}}{S+1}$.

technology-specific patent stock is

$$N_{ijt}(\delta_N, \ell_N) = \delta_N \times N_{ij,t-1}(\delta_N, \ell_N) + n_{ij,t-\ell_N}\,, \tag{54}$$

and the total patent stock 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

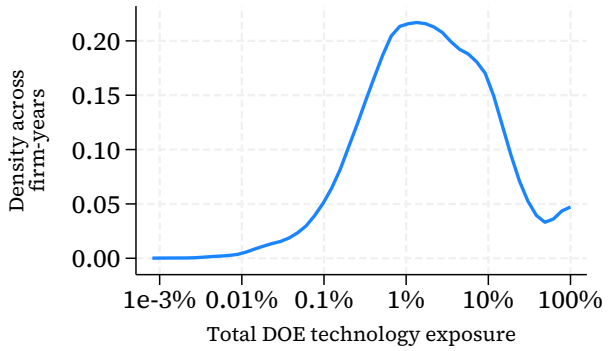
$$\chi_{ijt} \equiv \frac{N_{ijt}(\delta_N, \ell_N)}{N_{it}(\delta_N, \ell_N)} \quad \text{if } N_{it}(\delta_N, \ell_N) > 0, \tag{55}$$

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 defined in Section 3, the shares need not sum to one; this is the "missing share" problem highlighted by Borusyak et al. (2022). We account for differences in total exposure using the log of the sum of shares over the observed request categories, as described above.

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, we use a detailed dictionary-based text classification approach. We construct an energy dictionary containing over 400 technology-specific terms spanning different energy technologies.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.

We explore $\delta_N \in \{0.7, 0.8, 0.9\}$ and $\ell_N \in \{0, 1, 2\}$. Using $\delta_N = 0.8$ and $\ell_N = 1$, Figure 4 shows the distribution of total exposure across firm-years in our estimation sample, summing shares across all 14 DOE technologies. At the median, these technologies account for roughly 1.8 percent of a firm's depreciated patent stock. Some firms concentrate nearly all of their patent stock in these technologies.

FIGURE 4. Exposure Metric, χ



Note: Shows the density of log exposure, with percentage labels.

7. Results

7.1. A Note on Reporting

Our analysis requires several choices for which there are plausible alternatives. Table 3 collects these choices. We consider all combinations: approximately 970 regression specifications for value-added incidence and approximately 26,000 combinations of regression choices and factor supply elasticities for surplus incidence. All use the same firm-year sample, and we give each combination equal weight when summarizing the results. We first report a preferred regression specification, then summarize incidence across these combinations. Our approach to reporting distributions of point estimates is motivated by the uncertainty we have over the many calibrations we must make in order to convert our regression coefficient estimates into welfare-relevant metrics.

7.2. Elasticities

We begin with OLS and IV estimates of the focal elasticities in Table 4. We use the middle values of the timing parameters in Table 3, the most extensive windfall predictor set, cross-validated elastic net, and industry-by-year fixed effects.

In the IV estimates, a one-percent increase in public-innovation exposure raises revenue by about 0.48%. Expenditures on non-R&D labor rise by 0.46%, capital services by 0.62%, R&D by 0.37%, and intermediates by 0.48%. Markups fall by about 0.05%. Capital-service expenditures have the largest estimated response, while R&D expenditures have the smallest. These expenditure responses allow both input quantities and input prices to adjust.

TABLE 3. Parameters and Estimation Choices

Parameter or choice	Symbol	Values or alternatives
Panel A. Choices affecting the regression estimates		
mean horizon of public innovation	λ	5 , 10 , or 15 years
windfall covariates	$\mathbf{w}$	three sets, described in Section 3.4
windfall estimation method	-	four methods, described in Section 3.4
patent stock depreciation rate	$1-\delta_N$	0.1 , 0.2 , or 0.3
patent lag, years	ℓ_N	0, 1, or 2
firm-regression fixed effects	$\mathbf{z}$	year, year + industry, or year × industry
Panel B. Economic parameters for surplus incidence		
Labor supply elasticity	ε_l	1.5, 3, or 6
Capital supply elasticity	ε_k	5, 10, or ∞
R&D supply elasticity	ε_{rd}	0, 1, or 2

The IV estimates exceed their OLS counterparts in absolute value for every outcome. All IV responses are statistically distinguishable from zero at the 5% level under both clustered and randomization inference; the largest randomization-inference p-value is 0.016. The cluster-robust first-stage F statistic is approximately 6,200.

7.3. Incidence of Marginal Value-add

Figure 5 reports how the additional value added is divided across claimants. For each specification, we evaluate the same proportional increase in public-innovation exposure at each firm-year, sum the resulting gains, and divide each claimant's gains by the total increase in value added. Averaging these shares across specifications, each additional dollar of value added allocates about 38 cents to labor, 29 cents to capital suppliers, 12 cents to R&D suppliers, and 21 cents to shareholders as profits.

Labor receives the largest share in almost every specification. Its average marginal share is close to its 39-cent share of each dollar of existing value added. Capital suppliers' average marginal share exceeds their 23-cent share of existing value added. On average, the additional value added therefore has a similar labor share and a larger capital share, with smaller shares for R&D and profits. The point estimates imply positive gains for every claimant in every specification.

TABLE 4. Elasticities with respect to public innovation

	$p \times q$ (1)	c_l (2)	c_k (3)	c_{rd} (4)	c_m (5)	μ (6)
Panel A: OLS						
Coefficient (β)	0.301	0.295	0.420	0.197	0.304	-0.042
Clustered SE	(0.095)	(0.088)	(0.105)	(0.101)	(0.097)	(0.011)
p-value	0.002	<0.001	<0.001	0.051	0.002	<0.001
Panel B: IV						
Coefficient (β)	0.480	0.464	0.623	0.370	0.485	-0.050
Clustered SE	(0.099)	(0.092)	(0.110)	(0.107)	(0.101)	(0.012)
p-value	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
RI p-value	0.002	0.002	0.001	0.016	0.002	0.001
First-stage F (KP)	6,236.8	6,236.8	6,236.8	6,236.8	6,236.8	6,236.8
Firm-year observations	8,334	8,334	8,334	8,334	8,334	8,334
Firms / clusters	799	799	799	799	799	799

Note: Reports estimates controlling for log total patent exposure. RI uses 1,000 permutations; KP is the cluster-robust first-stage F statistic.

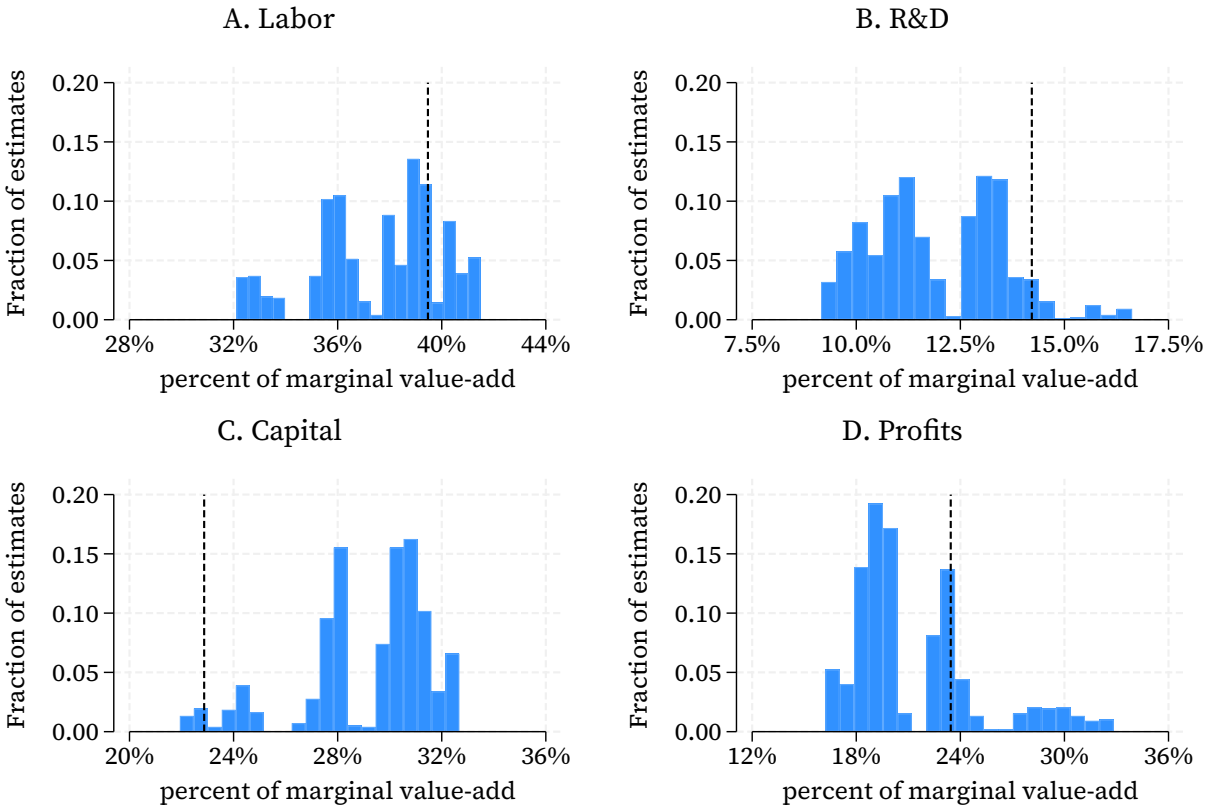
7.4. Incidence of Marginal Surplus

Figure 6 adds consumer gains and deducts factor suppliers' opportunity costs. Averaging across all regression choices and factor supply elasticities, consumers receive an average of 70 cents of each additional dollar of total surplus. Shareholders receive 16 cents, labor suppliers 7.3 cents, R&D suppliers 5.3 cents, and capital suppliers 1.8 cents.

Consumers receive the largest share in every specification. Under our demand model, revenue growth raises consumer surplus. Lower markups partly offset this gain in almost every regression specification, but consumer gains remain positive throughout. Consumers' marginal share is also below their share of existing surplus when both are evaluated using the same factor supply elasticities, so producers collectively receive a larger fraction of the additional surplus than of existing surplus. The dashed lines use the middle supply elasticities as a common reference.

Capital's small surplus share is conditional on the assumed capital supply elasticity. Most additional payments compensate suppliers for the opportunity cost of providing capital services; only the remainder enters supplier surplus. With perfectly elastic capital supply, all additional payments cover opportunity costs and capital suppliers receive no additional surplus. Shareholders' profit gains enter both accounts in full.

FIGURE 5. Share of Marginal Value-add from Public Innovation, $\mathcal{M}^{VA}$



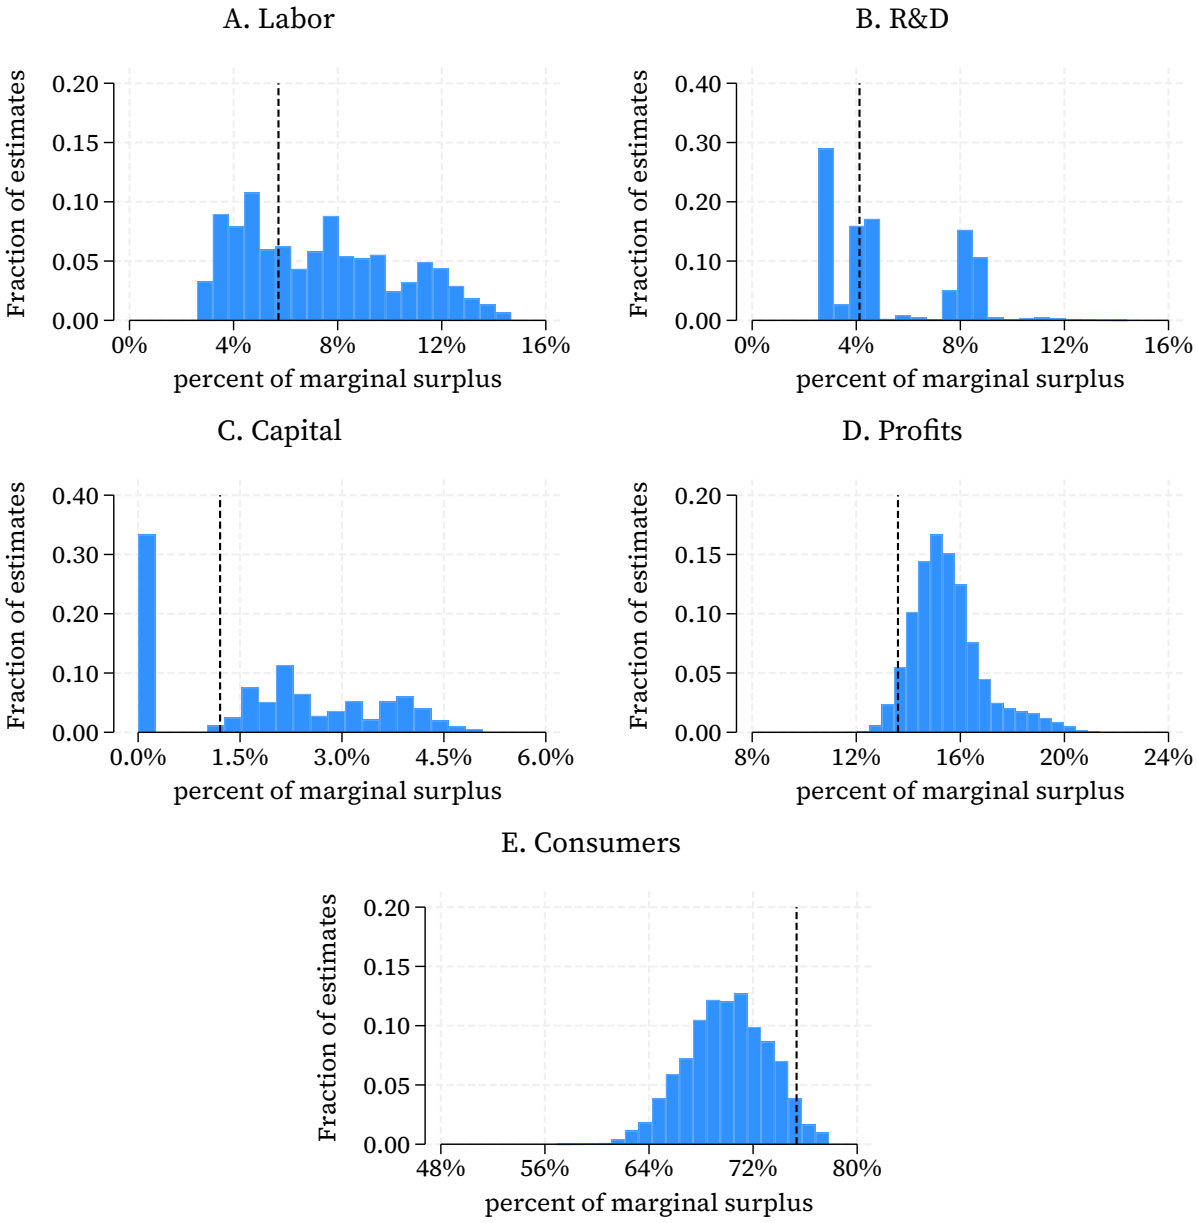
Note: Panel (A.) shows labor's share; panels (B.)–(D.) show R&D, capital, and profits. All regression choices vary. Dashed lines mark shares of existing value added.

7.5. Robustness

Revenue and input expenditures rise in every regression specification, while markups fall in almost every specification. Consumers' share of marginal surplus ranges from 57% to 78%, while labor's share ranges from 2.6% to 15%. Consumers receive the largest share throughout, but the ranking of the other claimants varies. These ranges reflect all regression choices and factor supply elasticities in Table 3.

Statistical support varies across specifications. In particular, randomization inference does not reject a zero R&D response at the 5% level in approximately 270 of the approximately 970 regression specifications. These cases occur under the middle and longest public-innovation timing assumptions.

FIGURE 6. Share of Marginal Surplus from Public Innovation, $\mathcal{M}^{TS}$



Note: Panel (A.) shows labor's share; panels (B.)–(E.) show R&D, capital, profits, and consumers. All regression and supply-elasticity combinations are included. Dashed lines mark shares of existing surplus at the middle supply elasticities.

8. Discussion

Publicly funded research can be commercially valuable while much of its benefit accrues to workers, suppliers, and buyers. In our estimates, exposure to DOE-funded innovation raises firms' revenues and input expenditures, and profits increase despite lower markups.

Labor receives the largest share of additional value added in almost every specification, while buyers receive most of the additional surplus under our maintained model. Shareholders gain, but receive only about one-sixth of the modeled surplus on average. A firm's profit response therefore captures only part of the economic value of the public knowledge it uses.

This pattern raises a question about the kinds of research that public and private funders choose. Models of directed technical change emphasize how prospective profits shape the research that firms undertake (Acemoglu 2002). A firm can find an innovation worth pursuing even if it captures a small fraction of its social benefits, provided its own expected returns cover its costs. Our estimated profit share therefore does not tell us whether private firms would have financed the same research, or whether privately financed innovation would benefit different groups. Comparing public and private research in related technological areas, using common measures of benefits, would help establish whether differences in incidence arise from project selection, access to knowledge, or firms' subsequent development of that knowledge.

The capital result illustrates the difference between paying for an input and generating a gain for its supplier. Capital suppliers receive an average of 29 cents of each additional dollar of value added, but 1.8 cents of each additional dollar of surplus once we include buyers and deduct suppliers' opportunity costs. Under the supply elasticities we consider, most additional capital payments compensate suppliers for providing additional services. Research expenditures raise a related issue: more spending can reflect both additional research activity and higher prices for scarce skills (Goolsbee 1998). For workers, an important distinction is whether public innovation raises the pay of existing employees or creates opportunities for new workers, a question closely related to evidence on rent-sharing at innovative firms (Kline et al. 2019). Our expenditure data do not separate these margins. Studying wage and hiring responses across technologies and labor markets with different supplies of relevant skills would help assess when research funding and policies that expand the scientific workforce are complements (Bloom et al. 2019).

The large buyer share directs attention to the product-market changes that follow public research. Innovation may make existing goods cheaper to produce, improve their quality, create new goods, or change competition among suppliers. Our financial outcomes do not distinguish these channels, and a lower markup does not by itself establish that prices have fallen. Under our isoelastic demand interpretation, lower markups imply more elastic demand and less buyer surplus per dollar of revenue. Revenue growth more than

offsets this effect in our estimates, so buyers gain, although their share of marginal surplus remains below their share of existing surplus. Linking funding shocks to product prices, quantities, and choices among competing products would help separate these mechanisms and assess the buyer-surplus estimates under more flexible demand models.

Household incidence requires following these gains to the people who ultimately receive them. Workers, shareholders, and buyers are overlapping groups, and some buyers in our accounting are themselves firms. A large labor or buyer share therefore does not establish that public innovation disproportionately benefits lower-income households. Combining employer–employee links with ownership and expenditure data would reveal whether the same households gain through several channels or whether different channels reach different parts of the income distribution.

Connecting this benefit incidence to the incidence of financing would allow us to ask whether the households receiving the gains also bear the costs of public research. The relevant comparison is with the economic burden of the marginal taxes used to finance additional investment, including their behavioral effects. Together with an estimate of aggregate benefits and evidence on the additional tax revenue and changes in public spending generated by innovation, this would permit a group-specific decomposition of innovation's marginal value of public funds (i.e., Hendren and Sprung-Keyser 2020; Finkelstein and Hendren 2020). These comparisons could differ across both research programs and financing instruments: the same innovation may have different distributional consequences when financed through taxes on earnings, consumption, or capital income. Incidence would thus help evaluate both the research being funded and the way society pays for it.

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 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.

___ **, Philippe Aghion, Leonardo Bursztyn, and David Hemous**, "The environment and directed technical change," *American Economic Review*, 2012, *102* (1), 131–66.

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.

Akerman, A., J. Moscona, H.S. Pellegrina, and K. Sastry, "Public R&D Meets Economic Development: Embrapa and Brazil's Agricultural Revolution," *Mimeo*, 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, John Van Reenen, and Heidi Williams, "A toolkit of policies to promote innovation," *Journal of Economic Perspectives*, 2019, *33* (3), 163–84.

___ **, 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.

Budish, Eric, Benjamin N Roin, and Heidi Williams, "Do firms underinvest in long-term research? Evidence from cancer clinical trials," *American Economic Review*, 2015, *105* (7), 2044–85.

Dinlersoz, Emin, Sebnem Kalemli-Ozcan, Henry Hyatt, and Veronika Penciakova, "Leverage over the Firm Life-Cycle, Firm Growth, and Aggregate Fluctuations," *Mimeo*, 2024.

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.

Goolsbee, Austan, "Does Government R&D Policy Mainly Benefit Scientists and Engineers?," *The American Economic Review*, 1998, *88* (2), 298–302.

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.

Hartman-Glaser, Barney, Hanno Lustig, and Mindy Z Xiaolan, "Capital share dynamics when firms insure workers," *The Journal of Finance*, 2019, *74* (4), 1707–1751.

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.

III, Thomas F. Howells, Edward T. Morgan, Kevin B. Barefoot, Louis E. Feagans, Teresa L. Gilmore, and Chelsea K. Smith-Nelson, "Preview of the 2018 Comprehensive Update of the Industry Economic Accounts," *Survey of Current Business*, August 2018, *98* (8).

Jones, Benjamin F and Larry Summers, "A Calculation of the Social Returns to Innovation," in Austan Gooslbee and Benjamin F Jones, eds., *Innovation and Public Policy*, 2022.

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.

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.

National Center for Science and Engineering Statistics, "Business R&D and Innovation: 2016," Detailed Statistical Tables NSF 19-318, National Science Foundation, Alexandria, VA May 2019.

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.

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.

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