*Econometrica Supplementary Material*

# SUPPLEMENT TO “OPTIMAL TAXATION AND R&D POLICIES”

(*Econometrica*, Vol. 90, No. 2, March 2022, 645–684)

UFUK AKCIGIT
Economics, University of Chicago, NBER, and CEPR

DOUGLAS HANLEY
Economics, University of Pittsburgh

STEFANIE STANTCHEVA
Economics, Harvard, NBER, and CEPR

## APPENDIX OA.1: PROOFS OF THE PROPOSITIONS IN THE MAIN TEXT

PROOF OF PROPOSITION 1: TAKING THE FOC of program $P$ in (8) with respect to $r_t(\theta^t)$ yields

$$\begin{aligned}\big[r_t(\theta^t)\big]:\ &\frac{1}{R}\mathbb{E}\left(\sum_{s=t+1}^{\infty}\left(\frac{1-\delta}{R}\right)^{s-t-1}\frac{\partial\tilde{Y}(\theta^s,\bar{q}_s)}{\partial q_s}\frac{\partial\lambda(\theta^{t+1})}{\partial r_t}\right)\\ &-\frac{1}{R}\mathbb{E}\left(\frac{1-F^1(\theta_1)}{f^1(\theta_1)}I_{1,t+1}(\theta^{t+1})\phi'_{t+1}\big(l(\theta^{t+1})\big)\frac{\lambda_\theta\lambda_r}{\lambda\lambda_l}[\rho_{\theta r}-\rho_{lr}]\right)\\ &-M'_t\big(r(\theta^t)\big)+\frac{1}{R}\mathbb{E}\left(\sum_{s=t+1}^{\infty}\left(\frac{1-\delta}{R}\right)^{s-t-1}\eta_s\frac{\partial\lambda(\theta^{t+1})}{\partial r_t}\right)=0.\end{aligned}$$

Using the definition of the R&D wedge as

$$s(\theta^t)=M'_t\big(r(\theta^t)\big)-\frac{1}{R}\mathbb{E}\left(\sum_{s=t+1}^{\infty}\left(\frac{1-\delta}{R}\right)^{s-t-1}\frac{\partial\pi_s(\theta^s)}{\partial q_s}\frac{\partial\lambda_{t+1}}{\partial r_t}\right)$$

to substitute for the marginal cost $M'_t(r_t(\theta^t))$ in the FOC, we obtain formula (10).

Taking the FOC with respect to $l_t(\theta^t)$ yields

$$\begin{aligned}\big[l_t(\theta^t)\big]:\ &\mathbb{E}\left(\sum_{s=t}^{\infty}\left(\frac{1-\delta}{R}\right)^{s-t}\frac{\partial\tilde{Y}(\theta^s,\bar{q}_s)}{\partial q_s}\frac{\partial\lambda(\theta^t)}{\partial l_t}\right)\\ &-\frac{1-F^1(\theta_1)}{f^1(\theta_1)}I_{1,t}(\theta^t)\frac{\partial}{\partial l_t}\left[\phi'_t\big(l_t(\theta^t)\big)\frac{\partial\lambda(\theta^t)/\partial\theta_t}{\partial\lambda(\theta^t)/\partial l_t}\right]\\ &-\phi'_t\big(l_t(\theta^t)\big)+\mathbb{E}\left(\sum_{s=t}^{\infty}\left(\frac{1-\delta}{R}\right)^{s-t}\eta_s\frac{\partial\lambda(\theta^t)}{\partial l_t}\right)=0.\end{aligned}$$

Ufuk Akcigit: uakcigit@uchicago.edu
Douglas Hanley: doughanley@pitt.edu
Stefanie Stantcheva: sstantcheva@fas.harvard.edu

© 2022 The Econometric Society https://doi.org/10.3982/ECTA15445

Transform the derivative of the envelope condition (using the notation $\phi_{lt}$ and $\phi_{ll,t}$ to denote, respectively, the first and second derivatives of $\phi$ with respect to $l$):

$$\frac{\partial}{\partial l_t}\left[\phi_{lt}\frac{\lambda_{\theta t}}{\lambda_{lt}}\right] = \left(\phi_{ll,t} - \phi_{lt}\frac{\lambda_{ll,t}}{\lambda_{lt}}\right)\frac{\lambda_{\theta t}}{\lambda_{lt}} + \phi_{lt}\frac{\lambda_{\theta l,t}}{\lambda_{lt}} = \frac{\phi_{lt}\lambda_{\theta t}}{\lambda_t}\left[\frac{\left(\phi_{ll,t} - \phi_{lt}\frac{\lambda_{ll,t}}{\lambda_{lt}}\right)}{\phi_{lt}}\frac{\lambda_t}{\lambda_{lt}} + \frac{\lambda_{\theta l,t}\lambda_t}{\lambda_{\theta t}\lambda_{lt}}\right]$$

$$= \frac{\phi_{lt}\lambda_{\theta t}}{\lambda_t}\left[\frac{1}{\varepsilon_{l,1-\tau}}\frac{\lambda_t}{\lambda_{lt}l_t} + \rho_{\theta l,t}\right] = \frac{\phi_{lt}\lambda_{\theta t}}{\lambda_t}\left[\frac{1}{\varepsilon_{l,1-\tau}}\frac{1}{\varepsilon_{\lambda l,t}} + \rho_{\theta l,t}\right].$$

Using the definition of the wedge $\tau(\theta^t)$ to substitute for $\phi_t'(l_t(\theta^t))$ yields the formula in the text.

PROOF OF IMPLEMENTATION RESULT: For every period, define the following objects:

$$D_s(\theta^{s-1}, \theta_s) = E\left(\sum_{t=s}^{\infty} I_{(s),t}\left(\frac{1}{R}\right)^{t-s}\frac{\partial v_t}{\partial \theta_t} \,\Big|\, \theta^s\right),$$

$$Q_s(\theta^{s-1}, \theta_s) = \int_{\underline{\theta}}^{\theta_s} D_s(\theta^{s-1}, q)\, dq,$$

where the expectation is explicitly conditioned on history $\theta^t$.

With a stochastic process such that the impulse response is independent of $\theta^t$ except through $\theta_1$ and $\theta_t$, we have that $I_{(s),t} = i(\theta_1, \theta_t, t)$ for some function $i()$. In addition, $\frac{\partial v_t}{\partial \theta_t} = \phi_t'(l_t(\theta^t))\frac{\frac{\partial \lambda(\theta^t)}{\partial \theta_t}}{\frac{\partial \lambda(\theta^t)}{\partial l_t}}$, so that

$$D_s(\theta^{s-1}, \theta_s) = E\left(\sum_{t=s}^{\infty}\left(\frac{1}{R}\right)^{t-s} i(\theta_1, \theta_t, t)\phi_t'(l_t(\theta^t))\frac{\partial \lambda(\theta^t)/\partial \theta_t}{\partial \lambda(\theta^t)/\partial l_t} \,\Big|\, \theta^s\right).$$

In the unrestricted mechanism, the transfers provided every period are

$$T_t(\theta^t) = Q_t(\theta^{t-1}, \theta_t) - \frac{1}{R}E_t(Q_{t+1}(\theta^t, \theta_{t+1})) + \phi_t(l_t(\theta^t)). \quad \text{(OA1)}$$

Given the time separable utility and the assumption on the impulse response functions, the transfer hence depends on $\lambda_t$, $r_{t-1}$, $\theta_t$, and $\theta_1$ (and, naturally, on age $t$). Denote it by $T_t^*(\lambda_t, r_{t-1}, \theta_t, \theta_1)$.

With the price subsidy in place, the total price faced by the monopolist is $\frac{Y(q,k)}{k}$. Hence, conditional on $q_t$, the monopolist maximizes social surplus from production and the choice will be a deterministic function of quality, denoted by $k_t(q_t)$. As a result, profits earned are a deterministic function of quality, denoted by $\pi_t(q_t)$.

Note that in period 1, since $r_0$ and $q_0$ are given and observed, the realization

$$q_1 = H(q_0, \lambda_1(l(\theta_1), r(\theta_0), \theta_1))$$

can be inverted to obtain $\theta_1$ (at the optimal allocation, under incentive compatibility) as long as for every $\theta_1$ there is a uniquely optimal $l(\theta_1)$. Hence, we will use conditioning on $q_1$ instead of $\theta_1$. Let $\Theta^t(q_1, r_{t-1}, q_{t-1})$ be the set of all histories (including $\theta_t$) that are consistent with $q_1$ in period 1, and $r_{t-1}$ and $q_{t-1}$. For each $\theta_t$ in this set, the optimal allocations and transfer are the same (independent of what exactly happened in the full past). Let $r_t^*(\theta)$, $l_t^*(\theta)$ be the optimal allocations given to each $\theta$ in this set (they are equal for each such $\theta$ by inspection of the wedge formulas at the optimum). The implied optimal quality is then $q_t^*(\theta) = q_{t-1} + \lambda_t(r_{t-1}, l_t^*(\theta), \theta)$.

We now have to make the tax system such that allocations which do not arise in the Planner's solution are very unattractive to the agent. First, we can rule out allocations that never occur for any $\theta$ in $\Theta^t(q_1, r_{t-1}, q_{t-1})$ by making the transfer at points $q_t^*(\theta)$, $r_t^*(\theta)$ following $q_{t-1}$, $r_{t-1}$, $q_1$ highly negative. We can also directly rule out histories $q_{t-1}$ and $r_{t-1}$ which should never occur in the Planner's problem in the same way.

For all remaining consistent histories and for each $\theta$ in $\Theta^t(q_1, r_{t-1}, q_{t-1})$, the tax or transfer given as a function of the observables needs to be such that

$$T_t\big(q_t^*(\theta), r_t^*(\theta), q_{t-1}, r_{t-1}, q_1\big) + \pi_t\big(q_t^*(\theta)\big) = T_t^*\big(\lambda_t\big(r_{t-1}, l_t^*(\theta), \theta\big), r_{t-1}, \theta\big).$$

Consider the firm's choice. First, for given $r_{t-1}$, $q_{t-1}$, and $\theta_1$, the firm should rationally only select a pair $q_t^*, r_t^*$ that is consistent with some $\theta \in \Theta^t(q_1, r_{t-1}, q_{t-1})$ or else the transfer it receives would be very negative. For each $r_{t-1}$, $q_{t-1}$, and $\theta_1$, if the firm chooses $q_t^*(\theta)$ and $r_t^*(\theta)$ meant for type $\theta$ in the planner's problem, it receives the utility it would get from reporting to be type $\theta$ in the Planner's problem. By incentive compatibility, the firm will choose the allocation meant for its true type realization. *Q.E.D.*

## APPENDIX OA.2: WORKED EXAMPLE WITH CONSTANT MARKUPS

### *Production*

We can specialize the functional form to one that delivers constant markups. Let the cost of production be $C(k, \bar{q}) = \frac{k}{\bar{q}^{\zeta}}$, and the output as valued by consumers be $Y(q_t(\theta^t), k_t(\theta^t)) = \frac{1}{1-\beta} q_t(\theta^t)^{\beta} k_t(\theta^t)^{1-\beta}$. The demand function under a patent system that grants monopoly rights is then

$$p\big(q_t(\theta^t), k_t(\theta^t)\big) = q_t\big(\theta^t\big)^{\beta} k_t\big(\theta^t\big)^{-\beta}$$

and the quantity chosen by the monopolist is

$$k\big(q_t(\theta^t), \bar{q}_t\big) = \big[(1-\beta)\bar{q}_t^{\ \zeta}\big]^{\frac{1}{\beta}} q_t\big(\theta^t\big).$$

At the optimum, the price is a constant markup over marginal cost equal to

$$p(\bar{q}_t) = \frac{1}{(1-\beta)\bar{q}_t^{\ \zeta}}.$$

Profits are then given by

$$\pi\big(q_t(\theta^t), \bar{q}_t\big) = q_t\big(\theta^t\big)(1\text{-}\beta)^{\frac{1-\beta}{\beta}} \cdot \beta \cdot \bar{q}_t^{\ \zeta\frac{1-\beta}{\beta}}.$$

$Y(q_t(\theta^t), \bar{q}_t)$, the output from the private producer in the laissez-faire with a monopoly right, is

$$Y\big(q_t(\theta^t), \bar{q}_t\big) = Y\big(q_t(\theta^t), k\big(q_t(\theta^t), \bar{q}_t\big)\big) = \frac{1}{1-\beta} q_t\big(\theta^t\big)\big((1-\beta)\bar{q}_t^{\zeta}\big)^{\frac{1-\beta}{\beta}}.$$

Hence, the final good in the private market equilibrium is given by

$$Y_t = \int_{\Theta^t} Y\big(q_t(\theta^t), \bar{q}_t\big) P(\theta^t) = \int_{\Theta^t} \frac{1}{1-\beta} q_t(\theta^t) \big[(1-\beta)\bar{q}_t^{\,\zeta}\big]^{\frac{1-\beta}{\beta}} P(\theta^t)\, d\theta^t.$$

Conditional on a given quality $q_t(\theta^t)$, the production choice of the planner would be such that

$$k^*\big(q_t(\theta^t), \bar{q}_t\big) = \bar{q}_t^{\,\frac{\zeta}{\beta}} q_t(\theta^t) > k\big(q_t(\theta^t), \bar{q}_t\big).$$

*A Special Case With Very Simple Wedges*

We can impose additional restrictions to obtain particularly easy characterizations of the wedges. Assume the functional forms in Table II, but also assume the special case in which $\rho_{\theta r} = \rho_{rl} = 1$, so that the screening term in the R&D wedge is zero.

Let

$$B_e = 1 + \zeta\left(\frac{1-\beta}{\beta}\right),$$

$$B_m = \frac{2-\beta}{1-\beta},$$

and

$$G_t = H^1(\theta^1) p^{t-1}(1+\gamma)(1-\alpha).$$

Then, we can show that in this special case,

$$\frac{\tau_t}{1+G_t} = -\left(1 - \frac{1}{B_e}\right) - \frac{1}{B_e}\left(1 - \frac{1}{B_m}\right),$$

$$s_t = \left(1 - \frac{1}{B_e}\right) + \frac{1}{B_e}\left(1 - \frac{1}{B_m}\right),$$

and so the profit wedge $\tau_t$ depends only on time $t$ and the initial state $\theta_1$ and tends to a constant profit subsidy $-(1-\frac{1}{B_e}) - \frac{1}{B_e}(1-\frac{1}{B_m}) < 0$ over time. The net subsidy wedge is constant over time and type and equal to exactly $-\tau_t$. Both wedges are increasing in absolute value when the strength of the spillover ($\zeta$) increases.

## APPENDIX OA.3: EXTENSIONS

### OA.1. *Heterogeneity in Production Efficiency*

Suppose that firms are also heterogeneous in their production productivities, denoted by $\theta^p$, with realization $\theta_t^p$ and history $\theta^{p,t}$. For instance, production costs could be $C(k, \bar{q}_t, \theta_t^p)$. Allocations are now specified as functions of the full set of histories $(\theta^t, \theta^{p,t})$. If production productivity is observable, the planner will simply condition on it for each history of research productivities $\theta^t$. In fact, as long as quality $q$ and quantity $k$ are observable, the planner can perfectly infer $\theta^{p,t}$ from the observed production choices. Net output is then $\tilde{Y}(q_t(\theta^t, \theta^{p,t}), \bar{q}_t, \theta_t^p)$ and profits are $\pi(q_t(\theta^t, \theta^{p,t}), \bar{q}_t, \theta_t^p)$. Similarly to before, we can define $\Pi_t(\theta^t, \theta^{p,t}) := (\sum_{s=t}^{\infty} (\frac{1-\delta}{R})^{s-t} \frac{\partial \pi(q_s(\theta^s, \theta^{p,s}), \bar{q}_s, \theta_s^p)}{\partial q_s})$ as the marginal impact on future profit flows from an increase in quality. Let $Q_t(\theta^t, \theta^{p,t}) = \sum_{s=t}^{\infty}(\frac{1-\delta}{R})^{s-t}\frac{\partial \tilde{Y}(q_s(\theta^s, \theta^{p,s}), \bar{q}_s, \theta_s^p)}{\partial q_s}$ be the marginal impact of quality on future expected output net of production costs.

Then, the optimal profit wedge can be set for each history $(\theta^t, \theta^{p,t})$ and satisfies

$$\tau(\theta^t, \theta^{p,t}) = -\mathbb{E}\left(\sum_{s=t}^{\infty}\left(\frac{1-\delta}{R}\right)^{s-t}\eta_s\right)\frac{\partial \lambda_t}{\partial l_t} - \mathbb{E}\big(Q_t(\theta^t, \theta^{p,t}) - \Pi_t(\theta^t, \theta^{p,t})\big)\frac{\partial \lambda_t}{\partial l_t}$$
$$+ \frac{1-F^1(\theta_1)}{f^1(\theta_1)} I_{1,t}(\theta^t)\frac{\phi_t' \lambda_{\theta t}}{\lambda_t}\left[\frac{1}{\varepsilon_{l,1-\tau}}\frac{1}{\varepsilon_{\lambda l,t}} + \rho_{\theta l,t}\right]$$

and the optimal R&D subsidy is given by

$$s(\theta^t, \theta^{p,t}) = \mathbb{E}\left(\sum_{s=t+1}^{\infty}\left(\frac{1-\delta}{R}\right)^{s-t-1}\eta_s\frac{\partial \lambda_{t+1}}{\partial r_t}\right)$$
$$+ \mathbb{E}\left(\big(Q_{t+1}(\theta^{t+1}, \theta^{p,t+1}) - \Pi_{t+1}(\theta^{t+1}, \theta^{p,t+1})\big)\frac{\partial \lambda(\theta^{t+1})}{\partial r_t}\right)$$
$$+ \frac{1}{R}\mathbb{E}\left(\frac{1-F^1(\theta_1)}{f^1(\theta_1)} I_{1,t+1}(\theta^{t+1})\phi_{t+1}'\big(l(\theta^{t+1})\big)\frac{\lambda_\theta \lambda_r}{\lambda \lambda_l}(\rho_{lr} - \rho_{\theta r})\right).$$

The productivity differences only enter the monopoly valuation term, as they only affect how effectively each firm can transform the quality into output. As a result, productivity differences in production do not really change the previous results.

More generally, any additional heterogeneity that is observable can be treated in a similar way, by conditioning the optimal policies on it. The problem becomes much more complicated if there is additional unobservable heterogeneity that is correlated with research productivity $\theta$. Already in much simpler static settings without spillovers, Rochet and Choné (1998) showed that, with two-dimensional heterogeneity, there are barely any general results. Incorporating non-trivial two-dimensional heterogeneity in a dynamic model with spillovers like this one (and being able to estimate it) would be an important big step for future research.

Empirically, we do not let this additional observable heterogeneity (such as production sector, technology sector, or business-cycle induced effects) contaminate the results and filter it out from the variables thanks to fixed effects before computing our data moments. What could be quite interesting for future research would be to actually specifically estimate the model and simulate differentiated optimal policies, allowing explicitly for different sectors, different technology classes, or different parts of the business cycles.

## OA.2. *Different Types of Observable R&D Investments*

Suppose that there are several types of observable R&D investments that firms can make, denoted by $r^1, \ldots, r^j, \ldots, r^J$. A natural interpretation would be the investments in different technology classes.

The step size is determined as a function of the observable R&D investments, unobservable R&D effort, and firm research productivity:

$$\lambda_t = \lambda_t(r_{t-1}^1, \ldots, r_{t-1}^j, \ldots, r_{t-1}^J, l_t, \theta_t).$$

We can define the Hicksian complementarity of each R&D type with firm effort and research productivity as

$$\rho^j_{\theta r,t} := \frac{\frac{\partial^2 \lambda_t}{\partial r^j_{t-1}\partial \theta_t}\lambda_t}{\frac{\partial \lambda_t}{\partial \theta_t}\frac{\partial \lambda_t}{\partial r^j_{t-1}}} \quad \text{and} \quad \rho^j_{lr,t} := \frac{\frac{\partial^2 \lambda_t}{\partial r^j_{t-1}\partial l_t}\lambda_t}{\frac{\partial \lambda_t}{\partial l_t}\frac{\partial \lambda_t}{\partial r^j_{t-1}}}.$$

Different types of R&D investments can have very different complementarity profiles with R&D effort and firm type (or, equivalently, their exposure to risk as embodied by the stochastic type). Some investments may generate returns with high certainty, regardless of the type realization, while others may only yield returns when firms are particularly good or in period of good realizations of the stochastic type.

Let the subsidy on investment $r^j_t$ be denoted by $s^j(\theta^t)$. At the optimum, formula (10) holds separately for each type of R&D investment wedge $s^j(\theta^t)$. The wedge $s^j(\theta^t)$ will be increasing in the effect of investment $j$ on the step size (in the Pigouvian correction term), as well as in the relative complementarity of that investment to unobservable R&D effort relative to its complementarity with respect to firm research productivity, $\rho^j_{\theta l} - \rho^j_{\theta r}$.

The lesson is that, while it is optimal to subsidize investments with higher externalities at a higher rate, it is not as beneficial if these investments are also highly sensitive to the firm productivity and firm research productivity is unobservable.

## OA.3. *Different Externalities From Different Types of Research*

It is also possible to directly incorporate different externalities from each type of R&D investments by letting the cost function be decreasing in each aggregate investment type:

$$C(k, \bar{q}^1, \ldots, \bar{q}^J) \quad \text{with } \bar{q}^j = \int_{\Theta^t} q^j_t(\theta^t)\, d\theta^t \quad \text{and} \quad q^j_t(\theta^t) = q^j_t(\theta^{t-1})(1-\delta) + \lambda^j_t(r^j_{t-1}, l_t, \theta_t).$$

This is important in order to be able to speak to the very different spillovers from different types of research such as basic and applied research. Basic research may only add little to the total quality of a firm's product, but if its effect on the costs of production of other firms is important, it will suffer from a large under-investment in the laissez-faire, as highlighted in Akcigit, Hanley, and Serrano-Velarde (2021), and will warrant a large Pigouvian correction.

At the firm level, the (single) product quality is given by

$$q_t = (1-\delta)q_{t-1} + \sum_{j=1}^{J} \lambda^j_t(r^j_{t-1}, l_t, \theta_t).$$

We have to impose $j$ consistency constraints in the partial program in each period $t$, each with multiplier $\eta^j_t$. Formula (10) then tells us that R&D investments with the highest spillovers (highest $\eta^j_t = \int_{\Theta^t} \frac{\partial \tilde{Y}^*(q_t(\theta^t), \bar{q}^1_t, \ldots, \bar{q}^J_t)}{\partial \bar{q}^j_t} P(\theta^t)\, d\theta^t$) are the ones that should be subsidized most (bearing in mind that their complementarities with effort and firm research productivity may dampen the benefits from subsidizing them).

## APPENDIX OA.4: COMPUTATIONAL APPENDIX

### OA.1. *Computational Procedure*

All code is written in standard Python 3, and depends only on common numerical and scientific modules such as numpy, scipy, pandas, statsmodels, patsy, and matplotlib. The parameter estimation and optimal policy calculations are done using either the Nelder–Mead algorithm or simulated annealing.

Because of the staggered nature of research spending and firm effort decisions, we find the optimal decisions for a log-uniform grid of possible $(\theta_t, \theta_{t+1})$ values. In addition, in the case of the optimal mechanism, one also tracks the initial type $\theta_1$, as this bears on the constraints imposed by informational limitations.

When solving for both the optimal mechanism and the linear tax equilibrium outcome, the solution method is constructed as a fixed point problem on the path of $\bar{q}$. Because $\bar{q}$ evolves according to a firm's research decisions and these decisions are made based on expectations that condition on the future path of $\bar{q}$, the decisions made by firms are in a sense both forward and backward looking.

Given a certain candidate path for $\bar{q}$, we can find the optimal choices for research spending and firm effort (for either the firm or the planner), which itself amounts to solving a one-dimensional equation for each point in the type space in each time period. Using these decisions, one can construct an updated path for $\bar{q}$. When this process reaches a fixed point, we have found the equilibrium path for $\bar{q}$. In practice, as the equations characterizing firm choices are analytical but not closed form, it is more efficient to formulate the problem as a fixed point over both the path of $\bar{q}$ and firm choices for $r$ and $l$ for each type. Updating is then done only using the $M'(r)$ and $\phi'(l)$ terms in the first-order conditions. Additionally, it is useful to dampen the updating process to avoid any numerical instabilities.

Moving to nonlinear policies considerably complicates matters. In this case, the relevant state space of the firm must include the actual value of $q$. As a result, we must track the joint distribution of $q_t$, $\theta_t$, and $\theta_{t+1}$. Conceptually, the convergence process and criterion are similar to the linear case, but the run time is much longer. The advantage is that we can entertain tax and subsidy policies that are arbitrary (differentiable) functions of firm profit and R&D investment.

To generate simulated moments for parameter estimation, we simulate a large number of firms ($2^{15} = 32{,}768$) for the entirety of their life cycle and compute various statistics on this panel of simulated data. All of the moments are relatively straightforward to calculate, with the notable exception of the coefficients for the spillover regression (M8) and the R&D-cost elasticity regression (M9), which are used to identify the externality parameter and various cost elasticities.

For the spillover regression (M8), we actually re-solve and re-simulate the model for a variety of different scenarios in which innovations contribute an additional boost to average productivity $\bar{q}$, which we interpret as innovation spillovers between firms. We perform this exercise for a variety of boost parameters centered around unity (the baseline model value). We interpret each simulated economy as representing a particular industry with a particular level of innovation spillovers. This mimics the exogenous variation used to identify the spillovers in the Bloom, Schankerman, and Van Reenen (2013) paper. Using this variation, we then run a regression of firm sales on the amount of research spending undertaken by the firm as well as the average research spending by all firms in that time period and industry. We then match this to an analogous regression run by Bloom, Schankerman, and Van Reenen (2013).

Similarly, for the R&D-cost regression (M9), we simulate a variety of economies having different values of the R&D cost parameters $\kappa_r$ centered at the baseline value. These differences can represent actual differences in cost, or alternatively, differences in R&D subsidy levels or tax credits. We then run a firm-level regression across time and industry of R&D investment on the level of $\kappa_r$.

To generate estimates for the standard errors of our parameter estimates, we take 100 draws from the distribution induced by our data moment means and variances, fully re-estimate the parameters of our model for each of these draws, then report the standard deviation of these estimates. Because some of our data moments (in particular, moments M8 and M9) come from different sources, it is not clear what the interpretation of off-diagonal elements would be. A natural choice is to set them to zero, using a diagonal matrix for the data moment standard errors.

## OA.2. *Ex Post Verification Procedure*

To perform the ex post verification, we start with the allocations under truth-telling in the optimal mechanism, $\lambda(\theta^t)$, $r(\theta^t)$, and $T(\theta^t)$ (where the transfers $T(\theta^t)$ are constructed following (OA1)). These allocations are defined for all histories $\theta^t$ which could arise along the equilibrium path by the optimal mechanism—thus, any history $\theta^t$ that can never arise given the distribution of stochastic shocks is ruled out (with, for instance, infinitely negative transfers $T(\theta^t)$).

For every history $\theta^{t-1}$, we can compute the allocations that would be assigned to an agent of type $\theta$ who reports $\theta'$ (not necessarily truthfully) among the feasible types in the space $\Theta$ at time $t$. Under any report $\theta'$, the agent will be assigned the allocations $\lambda(\theta^{t-1}, \theta')$, $r(\theta^{t-1}, \theta')$ and $T(\theta^{t-1}, \theta')$, which are meant for the "true" type $(\theta^{t-1}, \theta')$. The agent whose true type realization is $\theta$ chooses the report $\theta'$ that will maximize his expected discounted payoff, which is

$$\max_{\theta'} T(\theta^{t-1}, \theta') - \phi\big(\lambda(\theta^{t-1}, \theta')/w\big(r_{t-1}(\theta^{t-1}), \theta\big)\big) + \frac{1}{R}\int \omega\big(\theta^{t-1}, \theta', \theta_{t+1}\big) f^{t+1}(\theta_{t+1}|\theta).$$

The ex post verification consists in checking whether the agent will, in fact, choose $\theta' = \theta$ (i.e., report his true type) when faced with the set of allocations that can arise for *any* type at the optimum. Note that this amounts to checking that the global incentive constraints are satisfied at the optimal allocations derived using the first-order approach.

## REFERENCES

AKCIGIT, U., D. HANLEY, AND N. SERRANO-VELARDE (2021): "Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth," *The Review of Economic Studies*, 88 (1), 1–43. [6]

BLOOM, N., M. SCHANKERMAN, AND J. VAN REENEN (2013): "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81 (4), 1347–1393. [7]

ROCHET, J.-C., AND P. CHONÉ (1998): "Ironing, Sweeping, and Multidimensional Screening," *Econometrica*, 66 (4), 783–826. [5]

*Co-editor Giovanni L. Violante handled this manuscript.*

*Manuscript received 19 June, 2017; final version accepted 19 August, 2021; available online 27 September, 2021.*