

Designing Dynamic Subsidies to Spur Adoption of New Technologies*

Ashley Langer[†] and Derek Lemoine[‡]

July 9, 2018

We analyze the efficient subsidy for durable good technologies. We theoretically demonstrate that a policymaker faces a tension between intertemporally price discriminating by designing a subsidy that increases over time and taking advantage of future technological progress by designing a subsidy that decreases over time. Using new empirical estimates of household preferences for residential solar in California, we show that the efficient subsidy increases over time. The regulator's spending increases more than sixfold when households anticipate future technological progress and future subsidies.

JEL: H21, H23, H71, Q48

Keywords: subsidy, dynamics, price discrimination, expectations, foresight, technological change, durable good, solar

*We are grateful to the University of Arizona Renewable Energy Network for supporting this work. Kyle Wilson, Max Rosenthal, and Sanguk Nam provided valuable research assistance. We thank Gautam Gowrisankaran, Erin Mansur, Stan Reynolds, and Mo Xiao for helpful discussions. We also thank participants at the 2016 POWER conference, the 2016 Summer Conference of the Association of Environmental and Resource Economists, the 2017 Stanford Institute for Theoretical Economics, North Carolina State University, the University of British Columbia, and the University of North Carolina/Duke.

[†]University of Arizona; alanger@email.arizona.edu

[‡]University of Arizona and NBER; dlemoine@email.arizona.edu

1 Introduction

Policymakers commonly subsidize adoption of new durable good technologies. The U.S. government pays hospitals to adopt electronic medical record systems and car buyers to choose electric vehicles. The U.S. and other countries have paid farmers to install more efficient irrigation systems, firms to build renewable energy projects, and households to replace their aging appliances. And many U.S. states pay homeowners to install solar panels. These subsidies often evolve over time according to an announced schedule. Policymakers may employ these subsidies to increase adoption, but these subsidies also interact with other goals such as limiting public spending: two subsidy trajectories that eventually achieve the same level of adoption could have very different implications for the public purse. Yet the design of an efficient subsidy schedule has thus far remained an open policy question.

We investigate how a regulator should design a durable good subsidy that seeks to achieve a target level of adoption by a certain date.¹ Potential adopters have heterogeneous, private values for the technology. The regulator announces how the subsidy will evolve over time, values cumulative adoption, and dislikes spending public funds. Given that regulators often want to achieve suboptimal targets, whether for political or budgetary reasons, we remain agnostic about the welfare implications of the adoption target and refer to the “efficient” subsidy as the one that maximizes the regulator’s objective subject to the target.² We theoretically disentangle the forces that determine whether the efficient subsidy increases or decreases over time, and we combine the theoretical analysis with new empirical estimates of household preferences for rooftop solar electric systems in order to understand the relative importance of these different forces in a prominent real-world application.

We formally show that the regulator has conflicting interests when designing the efficient subsidy schedule. Of particular note, the regulator wants to offer a low initial subsidy to induce adoption by consumers with a particularly high willingness-to-pay for the technology and wants to increase the subsidy over time so as to then obtain adoption from consumers with lower willingness-to-pay. Within a period, the regulator cannot discriminate between consumers because all adopters receive the same subsidy, but the dynamic nature of the

¹We formally demonstrate that our setting is equivalent to one in which the regulator has a fixed budget instead of a fixed adoption target.

²We refer to our derived subsidy trajectories as “efficient” rather than “optimal” because the policy may not maximize welfare, and we avoid “cost-effective” because the regulator’s objective need not reduce to minimizing the spending required to meet the target.

subsidy enables intertemporal price discrimination. By using a low subsidy early and a high subsidy later, the regulator avoids “over-subsidizing” consumers who would be willing to adopt the technology at a low subsidy level and thereby reduces the total cost of achieving a given level of adoption. This price discrimination channel is strong when there are a lot of inframarginal consumers who would adopt even at low subsidy levels and is weak when most consumers are on the margin.³ However, when consumers anticipate future subsidies, some consumers will simply wait for the later, higher subsidies.⁴ Consumers’ expectations therefore constrain the regulator’s ability to intertemporally price discriminate and force the regulator to commit to a flatter subsidy schedule.⁵

In contrast, anticipated improvements in technology usually favor a declining subsidy. As technology improves, more people want to adopt the technology for a given subsidy. If the regulator offered households the same subsidy as in the case without technological change, then later periods would see greater adoption when technological change occurs. The increase in adoption in later periods translates into an increase in subsidy spending in later periods. However, a regulator with a convex cost of public funds prefers to smooth spending over time. The regulator accomplishes this smoothing by decreasing the subsidy in later periods relative to earlier periods. Technological progress therefore generally favors a subsidy that decreases over time. And endogenizing technological progress strengthens this effect: if the regulator believes that early adoption makes costs decline faster, then the regulator has a stronger incentive to use a relatively high subsidy early so as to stimulate early adoption.⁶

³Boomhower and Davis (2014) estimate the fraction of inframarginal adopters in an energy efficiency program in Mexico. They find that more than 65% of households are inframarginal and that about half of adopters would have adopted in the absence of any subsidy. Chandra et al. (2010) find that the majority of hybrid electric vehicles that received a Canadian rebate were inframarginal. Gowrisankaran and Rysman (2012) find that consumers who buy digital camcorders in later periods have lower values for the good than did consumers who purchased in earlier periods.

⁴In the theory and simulations, our consumers have rational expectations about the evolution of technology costs and subsidies and attempt to time their adoption of the technology. However, we remain agnostic about whether consumers accurately internalize the present discounted value of the stream of benefits from the technology. This potential undervaluation of a stream of future technology benefits has been discussed at length in the literature on the “energy efficiency paradox” (e.g., Allcott and Greenstone, 2012; Busse et al., 2013; De Groote and Verboven, 2016).

⁵Relatedly, Conlon (2010) emphasizes that firms want to lower the price of a durable good over time in order to intertemporally price discriminate, but consumers’ willingness to wait for lower prices limits the rate at which firms can reduce prices.

⁶We show that endogenizing consumers’ willingness-to-pay for the technology (whether interpreted as peer effects or network effects) is formally equivalent to endogenizing the cost of the technology.

To assess the relative magnitude of these and other competing forces that together determine the shape of the efficient subsidy schedule, we combine our theoretical analysis with a dynamic discrete choice estimation of households' preferences for residential solar systems in California. Beyond learning whether the efficient subsidy increases or decreases over time, we use our theoretical model to decompose the efficient subsidy trajectory into its component drivers in order to understand precisely *why* the efficient subsidy has its shape. The California Solar Initiative (CSI) included a substantial subsidy for residential photovoltaic (rooftop solar) adoption between 2007 and 2014. This program spent nearly \$2.2 billion to obtain 1,940 MW of residential solar capacity. The residential solar subsidy declined step-wise over time from \$2.50/Watt to zero, with pre-subsidy installation costs over this period declining from around \$9/Watt to around \$4/Watt. We combine data on household-level installations with a dynamic discrete choice model to estimate the distribution of households' benefit of installing solar conditional on household demographics.⁷ The estimation assumes that households know the full time-path of subsidies and electricity prices but allows solar system prices to evolve stochastically. The estimated preferences show substantial heterogeneity in the private benefit of residential solar systems.

We find that the efficient subsidy increases over time, even though this type of subsidy is rarely enacted or discussed. An increasing subsidy allows the regulator both to delay spending and to price discriminate by offering a low initial subsidy to consumers with a high valuation of the technology and then raising the subsidy over time to encourage additional consumers to adopt. The benefit of increasing the subsidy over time is tempered by technological change that can make it more cost-effective for the regulator to lower the subsidy in later periods when consumers face a lower private cost of the technology. The combination of our calibrated regulator objective, estimated consumer preferences, and observed technological change cannot generate an efficient subsidy that resembles the sharply declining subsidy used by California and many other states to speed adoption of solar power.

Further, whereas previous literature on efficient subsidies has modeled consumers as myopic (e.g., Kalish and Lilien, 1983), we show that understanding consumer foresight is critical for designing the efficient subsidy and that facing forward-looking consumers can

⁷Because consumers are forward-looking, we should not use a static approach to estimate the implications of a change in the subsidy for adoption as, for instance, in Hughes and Podolefsky (2015). Instead, we must recover consumers' structural preferences that will remain stable even when the regulator changes expectations of future subsidies (Lucas, 1976).

increase public spending substantially. Consumers' rational expectations of future subsidies limit the regulator's ability to intertemporally price discriminate by offering a low subsidy in early periods and a high subsidy in later periods. In a world without technological progress, this effect increases total spending on solar subsidies by 4% and slows the rate at which the efficient subsidy for solar increases over time. Further, the regulator must offer forward-looking consumers a high subsidy in order to compensate them for forsaking the option to adopt solar at some later time. This effect becomes especially important when households anticipate that technology will improve over time. As a result, the regulator's total spending is over 600% greater when households have rational expectations of technological progress in the market for solar as opposed to a world in which technological progress happens but households are myopic.

Households' rational expectations also reduce the degree to which the regulator can take advantage of technological progress to reduce the total cost of the policy. When households do not anticipate that technology might improve over time, the regulator can take advantage of its understanding of technological progress to reduce spending by 97%. Myopic households actually obtain slightly more surplus in the absence of technological change because the regulator does not delay their adoption as strongly. However, when households are aware of the possibility of technological progress, the regulator can reduce its spending by only 79% as technological change increases households' surplus by \$700 million (230%). Expectations of technological progress increase households' incentives to wait until later periods to adopt the technology. Technological progress therefore increases the opportunity cost of adopting the technology today, which requires the regulator to offer households a larger subsidy than would be necessary in a world in which households were ignorant of technological progress.

Finally, whether or not we allow for technological change or household foresight, our benchmark model never generates an efficient subsidy that declines as strongly as did the one enacted in California. Either making the regulator much more patient or giving the regulator a much more convex cost of funds does generate a subsidy that declines over time, but in neither case does it decline anywhere near as sharply as did the actual subsidy. We can, however, generate the type of sharply declining subsidy seen in California in two special cases. The first case substantially increases the regulator's marginal value for solar electricity. In particular, our policymaker would need to value solar electricity at a level that is more than an order of magnitude greater than the estimates of the social value of solar electricity in

Baker et al. (2013). The regulator then designs a subsidy schedule that prioritizes obtaining adoption quickly because she does not want to defer the benefits of solar electricity. The second case studies extreme assumptions about the potential endogeneity of technology and preferences. In particular, it assumes that all realized declines in the cost of solar were induced by adoption in California (and therefore that there were no cost declines induced by adoption anywhere else in the world) and that the entire time trend in the preference for solar was due to peer effects arising from adoption in California (and therefore that panel quality was unchanged over our time frame and that consumer preferences changed only through peer effects). The regulator then uses a high early subsidy so as to lower costs in later periods. However, neither extreme assumption can generate a sharply declining subsidy on its own, and the combination of these two extreme assumptions about technology and preferences is quite unrealistic.

Our primary contribution is to ground the design of dynamic subsidy instruments in economic principles. Despite the prevalence of subsidies for durable investments, there has been little formal analysis of these instruments. Kalish and Lilien (1983) study the efficient subsidy trajectory in the presence of learning and of word-of-mouth diffusion. They argue that both channels call for a subsidy that declines over time.⁸ In their conclusion, they mention that a desire to avoid subsidizing high-value consumers could argue for an increasing subsidy schedule. Meyer et al. (1993) discuss how to design investment tax credits in order to obtain the “biggest bang for the buck.” They note that the investment incentive is determined by the credit offered to the marginal investor, whereas the regulator’s spending depends on the average credit offered to investors. Policymakers should aim to combine a high marginal credit with a low average credit. These papers’ informal observations illustrate the logic underpinning our intertemporal price discrimination channel. We formally demonstrate this channel, show how it depends on private actors’ expectations, and introduce new channels.

Three recent papers have discussed other dimensions of technology subsidies. First, Kremer and Willis (2016) study the efficient subsidy trajectory in the presence of positive spillovers, which act like our endogenously declining costs. They assume homogeneous private values for the technology and a regulator that differs from consumers only in internal-

⁸We obtain the same result in our analysis of endogenous technology/preferences, though tempered by competition with several forces that push in the opposite direction. Further, note that Kalish and Lilien (1983) assume myopic consumers. Our analysis suggests that both the level and the trajectory of the efficient subsidy will be sensitive to households’ expectations about their future preferences for solar.

izing spillovers, whereas we emphasize the implications of heterogeneous private values and a regulator that may also care about public spending and about the flow of social benefits from adoption. In particular, a price discrimination motive can arise only in a setting with heterogeneous private values and a regulator that dislikes spending public funds. Second, Newell et al. (2017) informally discuss how to structure subsidy payments to a given project. They argue that upfront payments make sense when the government can borrow more cheaply than the private sector. We formally analyze how to structure the upfront subsidies offered to different possible projects, allowing the regulator's discount rate to differ from private sector discount rates. Third, Dupas (2014) considers how learning and reference-dependence interact with subsidy policies when consumers who have already adopted the technology must choose whether to adopt it again. In her application, insecticide-treated bed nets last only a few years and a single household can use multiple bed nets. In our application, solar panels can last more than twenty years and a single household has only a single roof. We therefore study one-time adoption decisions, where preferences may depend on cumulative adoption but where we can ignore the effect of an individual's adoption decisions on the same individual's later preferences.

Our analysis links the dynamic public finance and industrial organization literatures. The dynamic public finance literature has long studied a government that commits to a schedule of future taxes in order to satisfy an exogenous revenue requirement (e.g., Judd, 1985; Chamley, 1986; Chari and Kehoe, 1999; Kocherlakota, 2010). Instead of considering households who save for the future, we here study households who time their adoption of a new technology, and instead of considering a government that commits to future taxes in order to fund a given level of services, we here study a government that commits to future subsidies in order to attain a given level of technology diffusion. Our focus on one-time adoption decisions relates to the industrial organization literature on monopolists' strategies for pricing durable goods. In particular, several papers have explored the conditions under which a monopolist finds intertemporal price discrimination to be optimal. When production is costless, a monopolist should commit to offering a constant price over time as long as all customers use the same discount rate, their valuations are constant over time, and the monopolist is at least as impatient as its customers (Stokey, 1979; Landsberger and Meilijson, 1985). In that case, all sales happen in the first instant. However, intertemporal price discrimination can be optimal when production costs are convex (Salant, 1989) or declining over time (Stokey,

1979). In the present paper, the dynamics of the efficient subsidy are more complicated because our regulator has a concave benefit function and a convex distaste for spending in each instant, two generalizations that are more compelling for a regulator than for a firm. We reserve the label of “price discrimination” for forces that arise only because of adopters’ equilibrium decisions, so that we disentangle other dynamic forces from intertemporal price discrimination motives.⁹

The next section describes the model. Section 3 theoretically analyzes the efficient subsidy trajectory. Section 4 introduces the empirical application, including the setting and the data. Section 5 describes the dynamic structural model for estimating the distribution of household values for solar photovoltaics. Section 6 combines the empirically estimated distribution of private values with the theoretical analysis in order to explore the determinants of the efficient subsidy trajectory for rooftop solar. Section 7 explores what assumptions might lead the efficient subsidy trajectory to resemble the actual CSI subsidy trajectory. The final section concludes. The appendix contains formal derivations, provides evidence that California households were forward-looking, outlines data details, reports sensitivity tests, describes the numerical calibration, and theoretically analyzes the case of a fixed regulator’s budget rather than a fixed adoption target.

2 Model

Our model of technology adoption includes households who are deciding whether to adopt a new durable technology and a regulator who encourages technology adoption via subsidies. Each household that has not yet adopted the technology faces a choice in each period whether to adopt the technology and does not face any further choices after adopting the technology.

⁹Many authors have also explored how a monopolist should price durable goods when it cannot commit to later periods’ prices (e.g., Coase, 1972; Stokey, 1981). In considering this literature’s implications for actual markets, Waldman (2003) criticizes the assumption that commitments are not possible. He notes that firms often do appear to commit to policies in practice. Similarly, it is easy to provide examples in which policymakers appear to successfully commit to a subsidy schedule. Our theoretical analysis focuses on this environment with commitment (as does the dynamic public finance literature). In our empirical application, we consider policies over relatively short timescales (less than 10 years) that were authorized by legislation, not just by executive action. Regulators have followed through on legislated subsidy schedules for residential solar in many states. Because these schedules are functions of time or of adoption rather than functions of underlying uncertain state variables (such as system cost), this follow-through is more suggestive of commitment than of equilibrium play.

The regulator commits in period 0 to a subsidy schedule, according to which it will offer subsidy s_t to any household that adopts the technology in period t .

Household i values the technology at v_{it} in time t . The household's value $v_{it} = h_{it} + \varepsilon_{i1t}$ depends on potentially time-varying characteristics of the household and of the technology (both captured in h_{it} and including factors like household demographics and changes in the technology's quality) and on shocks to the household's preference for the technology (captured in ε_{i1t}). Every household that adopts the technology at time t receives a subsidy s_t but must pay the technology's cost $C(t, Q_t, \omega_t) \geq 0$, where Q_t is cumulative technology adoption prior to period t and ω_t is a random variable that might, for instance, account for stochastic input costs in the technology's production. Thus, household i 's net benefit of adopting the technology in period t is $v_{it} - C(t, Q_t, \omega_t) + s_t$. If household i does not adopt the technology in a period, then it receives a benefit ε_{i0t} and has the choice of adopting the technology in the next period. We jointly define the stochastic preference shocks as $\vec{\varepsilon}_{it} = \{\varepsilon_{i1t}, \varepsilon_{i0t}\}$.

The technology's cost may change over time for several different reasons. First, cost may decline exogenously over time: $C_1(t, Q_t, \omega_t) \leq 0$, where the subscript indicates a partial derivative. Second, the technology's cost may also decline as a result of cumulative adoption ($C_2(t, Q_t, \omega_t) \leq 0$) for either of two reasons: cumulative adoption may lower costs through induced technological change or learning-by-doing, or cumulative adoption may increase households' value for solar through peer effects or network effects.¹⁰ Finally, cost also depends on the stochastic shocks captured by ω_t .

Forward-looking households form rational expectations over the evolution of technology costs $C(t, Q_t, \omega_t)$, subsidies s_t , and household and technology characteristics h_{it} . We denote the household's time t information set as Ω_t .¹¹ The household's value of choosing whether to adopt the technology at time t is:

$$V(\Omega_t, \vec{\varepsilon}_{it}) = \max \{ h_{it} - C(t, Q_t, \omega_t) + s_t + \varepsilon_{i1t}, \beta \mathbb{E}[V(\Omega_{t+1}, \vec{\varepsilon}_{i(t+1)}) | \Omega_t] + \varepsilon_{i0t} \}, \quad (1)$$

¹⁰Recall that household i 's net benefit for solar is $v_{it} - C(t, Q_t, \omega_t) + s_t$. We can interpret $C_2(t, Q_t, \omega_t) \leq 0$ as indicating peer effects or network effects because a reduction in costs is equivalent to a population-wide increase in the private value of solar.

¹¹Note that since technology costs $C(t, Q_t, \omega_t)$ are a function of time, cumulative adoption, and the random variable ω_t , each of these variables is a potential state variable and enters into Ω_t .

where $\beta \in [0, 1)$ is the per-period discount factor and \mathbb{E} is the expectation operator.¹² Forward-looking households have $\beta > 0$ and myopic households have $\beta = 0$.

The regulator commits to a subsidy schedule that will, in expectation, achieve a pre-determined level of adoption \hat{Q} after some given time $T > 0$.¹³ She knows the true distribution of potential adopters' values but does not know any particular household's value. Because the regulator commits to the subsidy schedule and cannot offer different subsidies to different households in the same period, households do not face any strategic incentives to obscure their technology valuations. The regulator dislikes spending money. Her distaste for spending money is $G(s_t [Q_{t+1} - Q_t]) > 0$, with $G(\cdot)$ strictly increasing and strictly convex. The convexity of $G(\cdot)$ could reflect political constraints or could reflect that the deadweight loss of taxation increases nonlinearly in revenue requirements. The regulator also receives instantaneous benefit $B(Q_t) > 0$ from cumulative adoption, with $B(\cdot)$ strictly increasing and strictly concave. In our application to adoption of solar photovoltaics, the benefit function will capture the regulator's value for production of solar electricity. As $B'(\cdot)$ and $G''(\cdot)$ become small, the regulator's problem becomes one of minimizing the present cost of subsidy spending.

When selecting the subsidy trajectory, the regulator has rational expectations about how the technology's cost will evolve and how households will respond to the offered subsidy,

¹²In order to focus on other effects, we ignore heterogeneity in potential adopters' discount rates. The implications of such heterogeneity depend on whether actors with high discount rates tend to have high or low private values for the technology. The case with a positive correlation between discount rates and private values corresponds to Stokey (1979). The case with a negative correlation arises when adopting the technology provides a stream of benefits that potential adopters discount to a present value. The assumption of a common discount rate corresponds to a well-known aspect of the empirical methodology, in which the econometrician must assume a common discount rate because the discount rate is generally not well identified by the data (Magnac and Thesmar (2002) discuss the conditions under which the discount rate can be identified, and De Groote and Verboven (2016) do estimate the discount rate in an empirical model of solar installation in Europe).

¹³The assumptions of a fixed terminal time T and of a fixed adoption target \hat{Q} are not critical to the theoretical analysis. These assumptions affect the transversality conditions for the regulator's problem, but they do not affect the necessary conditions that are the focus of the theoretical analysis. Further, the appendix shows that the theoretical analysis is robust to giving the regulator a fixed budget instead of a fixed adoption target. Intuitively, if the budget constraint does not bind, then that setting is equivalent to altering the present setting to allow \hat{Q} free, which would affect the transversality condition but not the other necessary conditions. If the budget constraint does bind, then the problems are effectively identical if we here fix \hat{Q} at the value that results from solving the problem with a budget constraint. The only adjustment to the channels analyzed below is to amplify the marginal cost of public funds by the shadow cost of the budget constraint.

though the regulator does not know which precise sequence of shocks to technology and preferences will be realized. At time 0, the regulator chooses the subsidy trajectory $\{s_t\}_{t=0}^T$ to maximize

$$\sum_{t=0}^T (1+r)^{-t} \mathbb{E}_0 [B(Q_t) - G(s_t [Q_{t+1} - Q_t])],$$

for given discount rate $r > 0$, for given initial adoption $Q_0 \geq 0$, and subject to the constraint that expected terminal adoption $\mathbb{E}_0[Q_{T+1}]$ equal $\hat{Q} > Q_0$. Potential adopters' decisions determine how the announced subsidy trajectory affects Q_t . Households expect the subsidy to drop to 0 after time T .

3 Theoretical Analysis

We now theoretically analyze the subsidy trajectory that efficiently incentivizes actors to adopt a new technology. The theoretical analysis relies on two specializations of the full setting described above, which together eliminate the need for expectation operators. First, we assume that technological progress is deterministic. We therefore drop ω_t from households' cost function, writing $C(t, Q_t)$. Second, we assume that households' preferences are fixed over time, so that we write v_i instead of v_{it} . Normalizing the measure of potential adopters to 1, the twice-differentiable cumulative distribution function $F(v_i) \in [0, 1]$ gives the number of households who are willing to pay no more than v_i for the technology. Define $f(v_i) \geq 0$ as the density function $F'(v_i)$.

For analytic tractability, we conduct the theoretical analysis in continuous time, with households discounting at rate $\delta > 0$ and the regulator discounting at rate $r > 0$. All other definitions and notation extend in the natural way. We now have $Q_0 \in [0, 1]$ and $\hat{Q} \in (Q_0, 1]$. Note that $\dot{Q}(t)$ gives adoption at time t , where a dot indicates a derivative with respect to time.

Each actor i chooses the optimal time Ψ_i to adopt the technology, for given subsidy and cost trajectories:

$$\max_{\Psi_i} e^{-\delta\Psi_i} [v_i - C(\Psi_i, Q(\Psi_i)) + s(\Psi_i)].$$

The first-order necessary condition is¹⁴

$$\delta [v_i - C(\Psi_i, Q(\Psi_i)) + s(\Psi_i)] = \dot{s}(\Psi_i) - \dot{C}(\Psi_i, Q(\Psi_i)), \quad (2)$$

where $\dot{C}(t, Q(t))$ indicates the total derivative with respect to time. The left-hand side is the cost of waiting until the next instant: the actor delays receiving the instantaneous payoff $v_i - C(t, Q(t)) + s(t)$. The right-hand side is the benefit of waiting: when costs net of the subsidy are decreasing (i.e., when $\dot{C}(t, Q(t)) - \dot{s}(t) < 0$), then the actor can save money by adopting the technology later. The optimal time of adoption balances these costs and benefits. As potential adopters becomes perfectly patient ($\delta \rightarrow 0$), the cost of waiting disappears and they delay adoption until net costs reach their minimum. As potential adopters become perfectly impatient ($\delta \rightarrow \infty$), they adopt the technology as soon as their net benefit of adoption is positive.

Potential adopters' stopping problems generate the equilibrium conditions that constrain the regulator's choice of subsidy trajectory. Define $Y(t)$ as the value at which actors are just indifferent to adopting or not. The number of actors who have adopted the technology by t is $Q(t) = 1 - F(Y(t))$, which implies $\dot{Q}(t) = -f(Y(t))\dot{Y}(t)$. Note that $\dot{Y}(t) \leq 0$ along any efficient path: there is no reason for the regulator to adopt a subsidy that makes net costs strictly increase. Instead of selecting the subsidy at each instant, imagine that the regulator selects the quantity of adoption via $Y(t)$, with the subsidy determined by this choice and by actors' equilibrium conditions. Rearranging equation (2), the subsidy must evolve as

$$\dot{s}(t) = \delta [Y(t) - C(t, Q(t)) + s(t)] + \dot{C}(t, Q(t)). \quad (3)$$

Equilibrium adoption constrains both the level of the subsidy and also, for $\delta < \infty$, the change in the subsidy. The regulator's time 0 choice of subsidy schedule will be dynamically

¹⁴Define net costs as $z(t, Q(t)) \triangleq C(t, Q(t)) - s(t)$. These necessary conditions are sufficient if $\ddot{z}(\Psi_i, Q(\Psi_i)) > \delta^2 [v_i - z(\Psi_i, Q(\Psi_i))] + 2\delta\dot{z}(\Psi_i, Q(\Psi_i))$, where dots indicate total derivatives with respect to time. Substituting both equation (3) (obtained below) and its derivative with respect to time, we find that the necessary conditions are sufficient if $-\dot{Y}(\Psi_i) > \delta [v_i - Y(\Psi_i)]$. The left-hand side is positive, and it is strictly positive whenever current adoption is strictly positive. The inequality is satisfied for any v_i that is not too much greater than $Y(\Psi_i)$. But by the definition of $Y(\Psi_i)$ as the private value of the marginal adopter at time Ψ_i , it must be true that $v_i \leq Y(\Psi_i - dt)$ for any actor i who had not adopted prior to time Ψ_i . Using $Y(\Psi_i - dt) \approx Y(\Psi_i) - \dot{Y}(\Psi_i)dt$, we have $v_i - Y(\Psi_i) \leq -\dot{Y}(\Psi_i)dt$. Using that dt is small, the sufficient condition therefore holds at all times with strictly positive adoption.

inconsistent because the regulator commits to offering a given subsidy at time t in part to affect potential adopters at times $w < t$, but once time t arrives, those adoption decisions are in the past and thus irrelevant to a decision problem that starts from time t . However, we assume that the regulator is able to commit at time 0 to not revise its announced subsidy schedule. There is a long tradition in the dynamic public finance literature of assuming commitment (e.g., Judd, 1985; Chamley, 1986; Chari and Kehoe, 1999; Kocherlakota, 2010), and this assumption is particularly applicable in the case of subsidies for new technologies.¹⁵

Writing $y(t)$ for $\dot{Y}(t)$, the regulator solves

$$\begin{aligned} & \max_{y(t), s_0} \int_0^T e^{-rt} \left[B \left(1 - F(Y(t)) \right) - G \left(-s(t) f(Y(t)) y(t) \right) \right] dt \\ & \text{s.t. } \dot{Y}(t) = y(t) \\ & \quad \dot{s}(t) = \delta [Y(t) - C(t, 1 - F(Y(t))) + s(t)] + \dot{C}(t, 1 - F(Y(t))) \\ & \quad Y(0) = F^{-1}(1 - Q_0), \quad Y(T) = F^{-1}(1 - \hat{Q}) \\ & \quad s(0) = s_0, \quad s(T) = C(T, \hat{Q}) - Y(T) + J(Y(T), C(T, \hat{Q})). \end{aligned}$$

$J(v_i, C(T, \hat{Q}))$ is the present value to actor i of having the option to adopt the technology at time T , once the subsidy disappears for good. This actor adopts the technology at time T if and only if she has not adopted the technology at an earlier date and $v_i - C(T, \hat{Q}) + s(T) \geq J(v_i, C(T, \hat{Q}))$.

The appendix shows that the efficient subsidy evolves as:

$$\begin{aligned} \dot{s}(t) = & \left[-r \lambda(t) - B' f(Y(t)) - G' y(t) f(Y(t)) - \dot{\mu}(t) + r \mu(t) \right. \\ & \left. - [s(t)]^2 G'' f(Y(t)) \ddot{Q}(t) + [r \mu(t) - f(Y(t)) y(t) G'] C_2(t, Q(t)) f(Y(t)) \right] \\ & \left[G' f(Y(t)) - s(t) f(Y(t)) G'' y(t) f(Y(t)) \right]^{-1}, \end{aligned} \quad (4)$$

where $\lambda(t) \leq 0$ is the shadow value of the state variable $Y(t)$ (the negative sign means that the shadow benefit of adoption is positive) and the costate variable $\mu(t) \geq 0$ measures the

¹⁵For instance, our empirical application will consider California's subsidies for rooftop photovoltaic (solar) systems. Observers seem to have taken for granted that the regulator would follow its announced subsidy schedule.

degree to which the regulator is constrained at each instant by private actors' equilibrium behavior and rational expectations (i.e., it measures the cost of keeping promises made to those who adopted the technology in past instants). $\mu(0) = 0$ because the regulator is not constrained by past promises in the first instant.¹⁶ The denominator in (4) is positive (recall $y(t) \leq 0$), so whether the subsidy increases or decreases over time is determined by the terms in the numerator.

The first term in the numerator, $-r\lambda(t) \geq 0$, reflects the regulator's impatience. For now, ignore complications introduced by other channels. The analysis in the appendix shows that the shadow benefit of adoption must equal the social cost of each moment's spending on marginal adopters. In order for the regulator to be indifferent to small deviations in her policy trajectory, the shadow benefit of adoption ($-\lambda(t)$) must grow at the discount rate r , which keeps its present value constant over time.¹⁷ The efficient subsidy schedule therefore tends to increase because the impatient regulator will tolerate a greater social cost of spending in later instants. We call this first force for an increasing subsidy a *Hotelling channel*, due to its similarity to the Hotelling (1931) analysis of exhaustible resource extraction.

Second, the $-B' f(Y(t)) < 0$ reflects that raising today's subsidy in exchange for lowering tomorrow's subsidy not only shifts the shadow benefit of adoption forward in time but also provides benefits tomorrow by raising cumulative adoption.¹⁸ This effect of valuing the total stock of adoption is familiar from Heal (1976) models of resource extraction, in which extraction costs increase in the cumulative quantity extracted. This *adoption benefit channel* favors a decreasing subsidy schedule because it captures how waiting to spend money on the subsidy forgoes benefits in the interim.

Third, the $-G' y(t) f(Y(t)) \geq 0$ recognizes that the regulator cannot price discriminate within an instant. Recall that $y(t) \triangleq \dot{Y}(t) \leq 0$ measures the maximum gap between the private values of households adopting at time t . If the regulator offers a marginally greater

¹⁶Formally, $\mu(0) = 0$ is the transversality condition corresponding to the choice of s_0 .

¹⁷Imagine that the regulator deviates by reducing $Y(t)$ by ϵ and increasing $Y(t + \Delta t)$ by ϵ . And assume for the moment that the regulator's marginal cost of funds is unity (the convex cost of funds enters through other channels) and $C_2 = 0$. The regulator's savings today are $\epsilon s(t) f(Y(t))$, which by equation (A-1) equals $-\epsilon \lambda(t)$. The regulator invests this money and earns interest at rate r before spending $-\lambda(t + \Delta t)$ to obtain the later adoption. For the regulator to be indifferent to this deviation for Δt small, it must be true that $-\lambda(t + \Delta t) + \lambda(t) = -r\lambda(t)$. Letting Δt go to zero and using equation (A-2) in the derivative of equation (A-1) with respect to time, we have $\dot{s}(t) = -r\lambda(t)/f$.

¹⁸In footnote 17, the cost of delaying adoption should include $B' f(Y(t)) \Delta t$. The logic of the footnote would then imply that $\dot{s}(t) = [-r\lambda(t) - B' f]/f$.

subsidy to some adopter at time t , then it must offer that marginally greater subsidy to all adopters, including those who would have adopted at a lower subsidy. But if the regulator waits to offer the marginally greater subsidy in the next instant, then it avoids paying the extra money to the $-y(t) f(Y(t))$ inframarginal adopters at time t .¹⁹ The more inframarginal adopters there are at time t , the stronger the incentive to wait to offer the higher subsidy. This *price discrimination channel* thus favors an increasing subsidy, with its strength depending on the marginal cost of public funds and on the distribution of households' private values for the technology.

The next two terms in the numerator are $-\dot{\mu}(t) \leq 0$ and $r\mu(t) \geq 0$. When $\dot{\mu}(t) > 0$ (as must be true in early instants), the regulator's promises are accruing over time. This term then favors a decreasing subsidy, because the regulator promises low future subsidies to time t adopters. The $r\mu(t)$ term reflects how an impatient regulator is willing to promise a low time t subsidy in order to spur adoption in the instants just before t . These promises of a low time t subsidy favor a subsidy that increases from time t towards the ex post preferred level. The costate equation for $\mu(t)$ derived in the appendix is

$$-\dot{\mu}(t) + r\mu(t) = G' y(t) f(Y(t)) + \delta\mu(t).$$

This relationship has two implications. First, as agents become myopic ($\delta \rightarrow \infty$), $\mu(t) \rightarrow 0$ and $\dot{\mu}(t) \rightarrow 0$. These promise-keeping terms arise only from agents' attention to the value of waiting to adopt the technology. Second, for $\delta < \infty$, these promise-keeping terms offset the price discrimination channel, leaving only $\delta\mu(t) \geq 0$. We call this net effect of forward-looking agents a *constrained price discrimination channel*. The price discrimination channel reflected the regulator's ability to intertemporally price discriminate by raising the subsidy once adopters with greater private values have already claimed their subsidy. But when adopters are forward-looking, they might wait for the higher subsidies, which limits the regulator's ability to price discriminate. The constrained price discrimination channel is weakly positive, still favoring an increasing subsidy schedule. When adopters are perfectly patient ($\delta = 0$), the regulator must offer low subsidies forever, but when adopters are impatient, the regulator can offer higher subsidies in later instants without strongly disincentivizing adoption in early instants. Near the initial time, the constrained price discrimination channel is

¹⁹Recall that $\dot{Q}(t) = -y(t) f(Y(t))$.

approximately zero. Adopters' anticipation of future subsidies therefore eliminates the price discrimination channel in early instants, which works to tilt the efficient subsidy trajectory downward over those early instants.

The second line of equation (4) contains terms that are critical to the analysis of technical change. The $-[s(t)]^2 G'' f(Y(t))\ddot{Q}(t)$ captures the regulator's preference for smooth spending over time, driven by the convexity of the cost of public funds. $\ddot{Q}(t)$ describes how instantaneous adoption changes as time advances. When instantaneous adoption is increasing over time ($\dot{Q}(t) > 0$), this term favors a decreasing subsidy. Note that $\ddot{Q}(t) = -f(Y(t))\dot{y}(t) - f'(Y(t))[y(t)]^2$. First, if $\dot{y}(t) < 0$, then the measure of private values for which adoption is newly optimal is increasing over time. This case is especially plausible when technological progress, peer effects, or network effects make net installation costs fall over time. This greater adoption works to increase subsidy spending, which favors using a declining subsidy in order to smooth spending over time. Second, if $f'(Y(t)) < 0$, then the distribution of private values becomes thicker as more people find adoption to be optimal. Early adopters are in the tail of the distribution, so more adopters are on the margin in later instants. In this case, subsidy spending again tends to increase over time, which favors using a declining subsidy schedule. Putting these pieces together, this *smooth spending channel* favors a decreasing subsidy schedule in the plausible case with $\dot{y}(t) \leq 0$ and $f'(Y(t)) \leq 0$ and vanishes as the marginal cost of public funds becomes constant.

Now consider how anticipated declines in cost affect the efficient subsidy schedule. These cost declines could arise because of technological improvements or could reflect a market-wide increase in the private value of the technology, potentially because of peer effects. First, more strongly declining costs make $y(t)$ more negative for a given subsidy and cost. The more rapidly that costs are declining, the greater the number of inframarginal adopters at a given subsidy and the greater the incentive to price discriminate by using an increasing subsidy schedule. Second, the smooth spending channel depends on both the first and second derivatives of the cost function, via $y(t)$ and $\dot{y}(t)$. Rapid cost declines exacerbate the effect of moving to a thicker part of the distribution of private values (assuming $f'(Y(t)) < 0$), which favors a decreasing subsidy schedule. And if costs are declining at an accelerating rate, then $\dot{y}(t)$ tends to be negative, again favoring a decreasing subsidy schedule. Third, declining costs tend to reduce the shadow benefit $\lambda(t)$ of adoption by making it easier to obtain adoption in later instants. This effect weakens the Hotelling channel and thus favors a decreasing

subsidy schedule. Combining these pieces, we see that declining costs can strengthen the price discrimination channel that favors an increasing subsidy schedule but otherwise work to make the efficient subsidy decrease over time. The net effect on the efficient subsidy schedule is an empirical question that depends on the relative intensity of these channels in any particular application.

The final term on the second line of equation (4) adjusts the efficient subsidy for the potential endogeneity of net costs, whether due to induced technical change, peer effects, or network effects. When technical change and preferences are purely exogenous, this term vanishes because $C_2(t, Q(t)) = 0$. However, this term is negative when increasing adoption reduces the private cost borne by later adopters or increases later adopters' preference for the technology, because then $C_2(t, Q(t)) < 0$. This *endogenous cost channel* favors a declining subsidy because using a higher subsidy in earlier instants now carries the additional benefit of reducing the private net cost of adoption (and thus reducing the required subsidy) in later instants. Endogenizing either technological change or preferences therefore favors stimulating adoption through a larger early subsidy and taking advantage of lowered costs through a smaller later subsidy.²⁰

4 Empirical Setting and Data

The theoretical analysis shows that the efficient subsidy schedule for a durable technology can be sensitive to whether consumers are forward-looking and to the distribution of private values in the population. In order to quantitatively evaluate the determinants of the efficient subsidy schedule in a high-stakes setting, we focus on households' decisions about whether to install solar systems under the California Solar Initiative (CSI). We will pair our preference estimates with a calibrated regulator's objective to evaluate the efficient subsidy policy.

The CSI offered a state subsidy for residential solar installations from 2007 to 2014, administered by the California Public Utilities Commission. This program spent \$2.2 billion to obtain 1,940 MW of solar installations. In each of the three major California electric utilities (Pacific Gas and Electric (PG&E), San Diego Gas and Electric (SDG&E), and

²⁰The $r\mu(t)$ amplifies the endogenous cost channel in later instants. Substituting for $\lambda(t)$ in the Hotelling channel from the analysis in the appendix shows that one piece of the Hotelling channel cancels the $r\mu(t)$ in the endogenous cost channel, reducing the Hotelling channel to only those pieces that would arise under purely exogenous changes in technology and preferences.

Southern California Edison (SCE)), the subsidy started at \$2.50/Watt installed and declined over time to \$0. The CSI includes subsidies for both residential and commercial properties. We here focus on the residential component of the CSI.

In order to estimate households' preferences for solar during the CSI subsidy policy, we use three sources of data: (1) CSI data on installations and system costs, (2) electricity prices from each of the major utilities, and (3) local demographic and solar radiation data. We detail each of these sources in turn.

The CSI maintained data on all applications for residential solar subsidies under the program. The data include the application date, the household's zip code and utility, the subsidy received, and an extensive set of solar system characteristics, including system size, manufacturer, installer, and cost. Figure 1 shows the evolution of subsidies and pre-subsidy average system costs from July 2007 through May 2012 (when our estimation window ends).²¹ The general patterns are similar across the utilities. The cost of an average system does not change much in the initial periods when silicon costs are increasing and then declines over the majority of our estimation time-frame as technology advances and silicon costs fall. CSI subsidies decrease in a step-wise pattern, with the steps occurring at different times in each utility.²²

Electricity price data comes from each utility's rate statements. Electricity prices in California are based on a household's usage relative to a baseline. The marginal price of electricity is small up to the baseline and then increases steeply in monthly usage. Because solar is most cost-effective for households with high monthly electricity usage, we use the average of the marginal price per kilowatt-hour for usage over 200% of baseline.²³ This is in keeping with the results in Hughes and Podolefsky (2015), who find that electricity prices at the highest tiers are most predictive of solar panel adoption in California. Figure 2 shows the evolution of electricity prices in each utility over our estimation window. In general, prices for usage over 200% of baseline are substantially higher in PG&E than in the other two utilities, although this difference decreases somewhat over time.

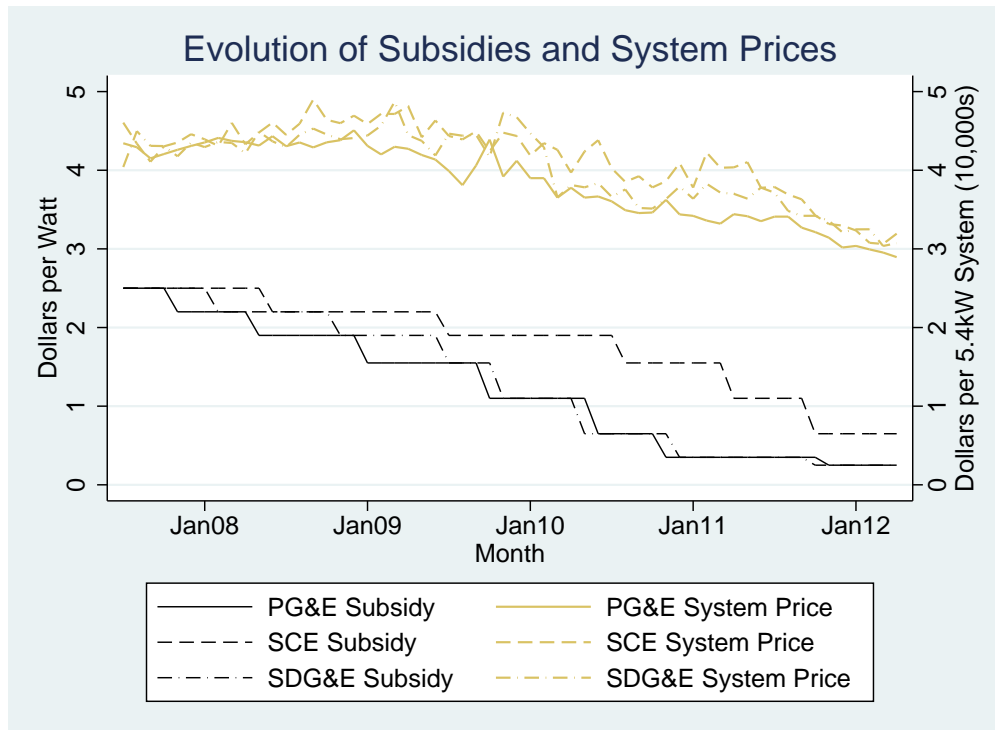
Finally, we use demographic data at the block group level from the American Community

²¹The average system costs are based on the average cost per Watt in each month for each utility multiplied by the average system size over the full period for all utilities of 5.4 kilowatts. There was also a 30% federal tax subsidy available for residential solar installation during this period that is not reflected in Figure 1.

²²Appendix Figure B1 shows how installation counts in each utility varied over time.

²³Appendix B.2 provides details on exactly how electricity rates are constructed.

Figure 1: Monthly CSI Subsidy and Average System Cost by Utility



Survey to allow preferences for solar to vary with consumer demographics. We focus on the demographics of owner-occupied households in each zip code under the assumption that all residential households that install solar systems are owner-occupied. Given the short panel of solar installation data, we do not allow demographics to vary over time. We supplement the demographic data with information from the California Secretary of State's office on Barack Obama's share of votes in the 2012 Presidential election at the precinct level. In order to account for geographic variation in solar generation potential across California, we construct the median solar direct normal irradiation at the zip code level from National Renewable Energy Laboratory data.

Table 1 summarizes the demographic data. The first column presents the average demographics for owner-occupied housing in the three California utilities in our sample. The second column presents average demographics for zip codes with households that install solar, weighted by the number of installations in each zip code. We see that households that install solar live, on average, in zip codes with slightly higher income and more expensive

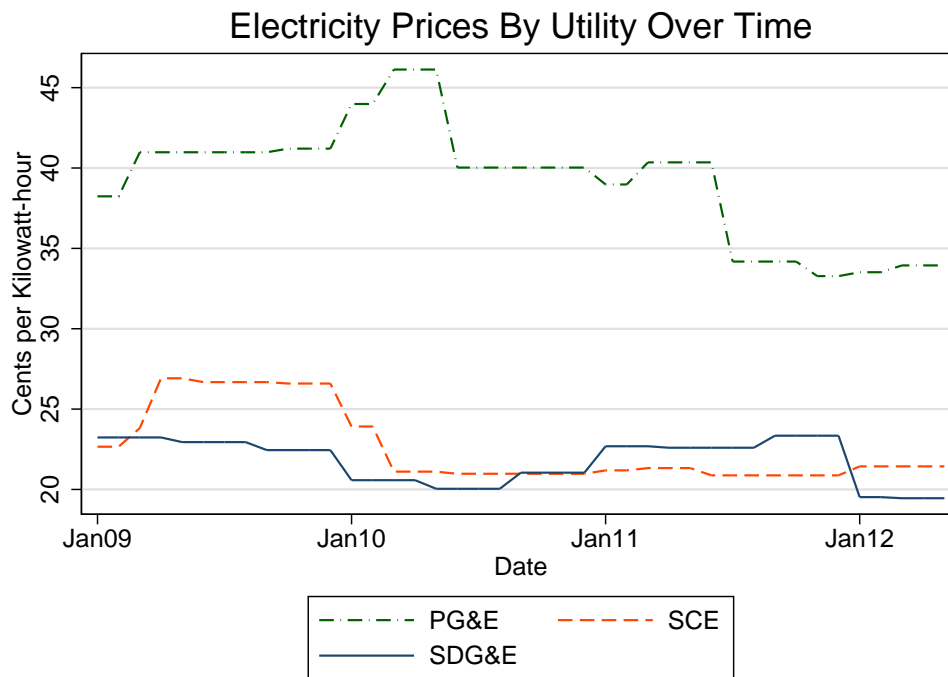


Figure 2: Evolution of electricity prices during the estimation window. Average marginal electricity price per kilowatt-hour for usage over 200% of baseline over the estimation window. Data details are in Appendix B2.

homes than owner-occupied households overall. Households that installed solar systems are in zip codes with greater median solar direct normal irradiance than households overall, which means that they are generally in areas with greater solar electricity generation potential. Households that install solar systems are in precincts that voted for Barack Obama at slightly lower rates than owner-occupied households overall, perhaps reflecting their higher income and home values or perhaps reflecting how political preferences are correlated with solar radiation in California. Households that install solar live in zip codes with approximately the same education, household size, and number of mortgages as owner-occupied households overall.²⁴

We limit our analysis to households in California zip codes that fall within one (and only one) of the three major utilities in order to have accurate information on the solar costs

²⁴Appendix B.2 provides additional details about how we handle potential complications in the data.

Table 1: Demographic Summary Statistics

	Average Owner-Occupied Household Demographics	
	Overall	Install Solar
Household Income (\$)	88,664	95,929
Home Value (\$)	476,483	512,723
Median Solar Radiation (kWh/m ² /day)	5.86	6.04
Democratic Vote Share	0.59	0.55
Years of Schooling	13.6	13.8
Number of Mortgages (0/1/2+)	0.94	0.96
Number of Household Members	2.6	2.6
Count	4,104,377	49,765

Data are at the block group level except for installations and solar radiation, which are at the zip code level, and Democratic vote share, which is at the precinct level. Owner-occupied households are assumed to have the average demographics of their zip code, weighted across block groups. Solar radiation is Direct Normal Irradiance.

and subsidies faced by households in each period. We begin our estimation window when the cap on the federal subsidy was removed in January 2009 because it was unclear whether households would have anticipated this federal policy change ahead of time. Additionally, we end our estimation window before the solar subsidy expires in the first utility (PG&E) because we do not have information on solar system prices after subsidy funds expire. These restrictions leave us with an estimation window of January 2009 through March 2012.

5 Econometrics

We estimate a dynamic model of residential solar system demand in order to understand how California households value residential solar and how different subsidy trajectories may change their installation decisions. In Appendix B.1, we provide novel reduced-form evidence that households are indeed forward-looking when deciding whether to install solar systems. It is therefore important to use a dynamic model of residential solar adoption because a static model would misestimate the benefits to households of adopting: some households that value solar above the current system cost will choose not to install for now as they wait for technology to advance and for costs to drop (see Aguirregabiria and Nevo, 2013).

Correctly accounting for the value of waiting for lower prices is critical to understanding the trade-offs regulators face in structuring solar subsidies. Our demand estimation approach is consistent with previous literature that has modeled residential solar installation decisions as a dynamic decision (e.g., Burr, 2014; Reddix II, 2014; De Groote and Verboven, 2016).

5.1 Estimation framework

Our empirical estimation parameterizes the household technology demand model presented in Section 2, without the simplifications made in the theoretical analysis. Households are assumed to be forward-looking and to value solar adoption in period t at $v_{it} = h_{it} + \varepsilon_{it}$, which should be thought of as the expected present discounted value of the stream of benefits from installing residential solar plus the upfront net benefit from installing solar. The upfront net benefit from installing solar is the current social, aesthetic, or reputational benefits to the household net of any nonmonetary fixed costs of researching solar systems or providers and the disruption of the solar installation process. For many households, the upfront nonmonetary cost of installing solar is likely to be large, which would make the upfront net benefit negative.

We parameterize the model in a few ways. First, we assume that $\vec{\varepsilon}_{it}$ is distributed i.i.d. extreme value type I and that $h_{it} = X'_{it}\gamma$, where X_{it} includes observable characteristics of the household and the technology. Allowing for heterogeneity in preferences based on observable differences across households differentiates our approach from Burr (2014) and is critical for understanding the shape of the full distribution of solar system valuations, which we have shown is an important input to the regulator's efficient subsidy.

Second, we introduce a parameter α_i that measures the disutility of spending money, which may vary by household, and we include the federal subsidy on solar that reduces the technology's cost by a fixed fraction ϕ . The system cost (net of subsidies) in a given period changes utility by $\alpha_i[(1 - \phi)C(t, Q_t, \omega_{it}) - s_{it}]$, where α_i should be negative and costs and subsidies are allowed to vary by utility.²⁵

Third, we specify how households form beliefs about the evolution of the states in Ω . We assume that technology costs evolve according to a first-order Markov process and that households understand that process. We also assume that technology costs evolve exoge-

²⁵This formulation assumes that the pass-through for the CSI subsidy is 100%. This assumption is consistent with recent empirical evidence in Pless and van Benthem (2017).

nously rather than depending upon the installed base of the technology. As a result, we have $C_{i(t+1)} = \gamma_0 + \gamma_1 C_{it} + \omega_{it}$, where ω_{it} is normally distributed and where we collapse the remaining cost function arguments into the subscript on C . We assume that households know the full time-path of subsidies and household preference shifters (such as electricity prices) from the beginning of the estimation window.

Finally, many owner-occupied households do not actually face the choice of installing residential solar. This may be because they live in a condominium or other multi-family dwelling and therefore do not have the right to install solar on their roof, or it may be because their roof's slope, orientation, and shading are not conducive to solar. We therefore also include a variable ζ_i that is the probability that a household is able to consider residential solar. This variable is fixed over time and can be thought of as a permanent, random shock to the household's preference for residential solar.²⁶

We estimate the model via maximum likelihood. The likelihood function in each period is

$$L_t = \prod_{i=1}^{N_t} \left\{ \left[\zeta_i \left(\frac{\exp(X'_{it}\gamma + \alpha_i \tilde{C}_{it})}{\exp(X'_{it}\gamma + \alpha_i \tilde{C}_{it}) + \exp(\beta \mathbb{E}[V(\Omega_{t+1}|\Omega_t)])} \right) \right]^{1\{i \text{ adopts in } t\}} * \left[(1 - \zeta_i) + \zeta_i \left(\frac{\exp(\beta \mathbb{E}[V(\Omega_{t+1}|\Omega_t)])}{\exp(X'_{it}\gamma + \alpha_i \tilde{C}_{it}) + \exp(\beta \mathbb{E}[V(\Omega_{t+1}|\Omega_t)])} \right) \right]^{1\{i \text{ does not adopt in } t\}} \right\}, \quad (5)$$

where $\tilde{C}_{it} \triangleq (1 - \phi)C_{it} - s_{it}$ is the net-of-subsidy cost of installing solar and N_t is the number of households that have not yet installed solar at the start of period t . In this formulation, the probability of adopting solar is equal to the probability ζ_i of considering solar times the probability of adopting solar conditional on considering it. The probability of not adopting solar is equal to the probability of not considering solar plus the product of the probability of considering solar and the probability of not adopting solar conditional on considering it. The variation in the data is not sufficient to precisely estimate the percentage of households who consider installing solar. We assume that consideration is independent of other household characteristics. Because the point estimate from a model that estimates $\zeta \triangleq \sum_i \zeta_i / N_0$ is 0.0660, we assume that 6.6% of the households who have not yet installed solar at the start

²⁶We discuss the role of the serial correlation that ζ_i introduces into the household's preference shocks in Appendix B.3.

of our time-frame are able to consider installing solar over our time-frame.²⁷

At each step of the likelihood maximization, we solve for the fixed point of the value function when all CSI subsidies equal zero and are expected to equal zero in the future. We then solve recursively for the value function in all previous periods. Households make an installation decision every month, using a monthly discount factor of 0.99 (for an annual discount rate of approximately 12%, which is consistent with Busse et al. (2013) and De Groot and Verboven (2016)). Standard errors are calculated as the square root of the inverse of the outer product of the Jacobian.

5.2 Identification

The CSI provides data on the applications to install solar, which includes information on cost of the solar system and the household's zip code but not the household's demographics. In order to estimate our model, we organize our data so that we have a count of the number of adopters and non-adopters in each period for each demographic bin in each utility.²⁸ To do this, we assume that solar adoptions within a zip code are randomly distributed across demographic bins within each block group and then aggregate the total number of households that adopt or do not adopt each period within each demographic bin in each utility.²⁹ This approach is similar to Nevo (2001) and Berry et al. (2004). There are over 2,500 zip codes in California, so identification of the demographic preference coefficients comes from the fact that, for instance, zip codes with high home prices have higher rates of solar installation conditional on solar costs than are zip codes with low home prices.

Identification of the cost parameters and of the demographic differences in the preference for solar systems and for spending should be thought of somewhat differently. Cost changes

²⁷Appendix B.3 analyzes sensitivity to assumptions about ζ . Neither the estimation results nor the conclusions of our simulations are sensitive to the assumed percentage of households that consider adopting solar.

²⁸We divide solar radiation into quintile bins, home values into four bins (less than \$200k, \$200-500k, \$500k-1 million and greater than \$1 million), and education into 3 bins (high school or less, more than high school to college completion, more than college).

²⁹Unfortunately, the joint distribution of demographics is not available at the zip code level in the ACS. To capture some of the correlation between demographic characteristics, we assume that adopters are drawn from the unconditional distribution of demographic characteristics in each block group within a zip code. Thus, if one block group has, on average, higher income and education while another block group has, on average, lower income and education, this tendency for income and education to be correlated will be captured in our simulations even though we do not know the actual covariance between income and education within a block group.

at the utility level are coming largely from changes in subsidies and changes in panel costs via technological advancement, input costs, and exchange rates. This means that unobserved shocks to local adoption are likely to be uncorrelated with average utility-level solar system costs.³⁰ We therefore argue that we are estimating the causal effect of changes in system costs on adoption.³¹ However, we do not claim that our estimates of the relationship between household demographics and preferences are causal. For example, while we are estimating how changes in the cost of a solar system will affect adoption and how this effect will vary over households with different home values, our estimates should not be interpreted as suggesting how a change in housing values will affect solar installation rates. This is because households in zip codes with high home values will differ from households in zip codes with low home values in unobservable ways that we are not capturing. We need causal estimates only for the effect of system cost on adoption because we will study policies that change net-of-subsidy costs, not demographics. We include demographic characteristics in order to better understand the density of consumer preferences, regardless of whether the shape of that density is being determined by observed or unobserved demographics.

Finally, we include a quadratic time-trend in our estimation to allow for the fact that solar system quality was probably increasing over this time period. It is possible that this time trend also captures changing preferences for solar, such as would come from peer effects (Bollinger and Gillingham, 2012).

5.3 Preference Estimates

The dynamic empirical model generates estimates of households' valuation of residential solar systems that we will use to quantitatively evaluate the efficient subsidy trajectory.

³⁰If there is imperfect competition among local solar installers, then local system costs could be correlated with local demand over the full period. We mitigate this concern by using monthly average costs by utility.

³¹One might be concerned about the identification of the system cost coefficient if upcoming declines in the subsidy increase current demand for solar systems and thereby increase current installation costs. There are two reasons why this concern may not be severe. First, what mattered for whether a household received a subsidy was the date that the household applied for the subsidy rather than the installation date. The household (or the installer) could then delay installing solar until the backlog of pending installations had subsided. In our data, we observe average delays between subsidy application and completed installation of approximately four months, with some delays lasting well over a year. These delays suggest an ability of installers to smooth demand spikes. Second, Pless and van Benthem (2017) find nearly 100% pass-through of subsidies for systems purchased by homeowners. If substantial system cost increases were occurring before subsidy changes, then their estimate of pass-through should have been substantially lower than 100%.

Table 2 presents the estimated coefficients, including the preference for installing a solar system of the median size and the preference for system costs. Both the preference for solar and the disutility of spending can vary by demographic group. The average household that considers installing solar has a negative valuation of solar, even after controlling for the fact that most households do not consider installing solar systems. This result is intuitive because the decision to adopt solar comes with substantial nonmonetary costs of researching whether solar is a good option for the household, of finding an installer to evaluate the home, of understanding whether financing is available to help with the upfront cost of solar, and of going through the installation process. Many households will likely perceive this cost of investigating solar to be substantial enough that they would require a considerable upfront payment to even evaluate whether solar is a reasonable option for them.

The benefit of installing solar is strongly increasing in the solar radiation in the household's zip code, which makes sense as the quantity of radiation directly determines the quantity of electricity generated and thus the electricity savings from installing solar. Preferences for solar are also strongly increasing in the electricity price faced by households, which makes sense as solar generation would directly offset electricity expenses under net metering policies (see Borenstein, 2017). Households with higher home values have a higher value for solar, likely reflecting that these homes have more potential roof space for solar panels and have better access to financing for solar panels. More educated households also have a higher value for solar, perhaps reflecting differences in the cost of collecting information about the benefits of solar. The magnitudes of the effect of radiation and electricity prices on preferences are substantially larger than those of home values and education: the difference in the preference for solar between the highest and lowest radiation bins³² is of a similar magnitude to the difference in the preference for solar between the highest and lowest electricity prices in the data,³³ whereas the difference in preference between the highest and lowest home value and education bins are approximately half that size.³⁴ We take this result

³²The highest radiation zip codes receive 0.7240 tenths of a kilowatt-hour per square meter per day of radiation whereas the lowest radiation zip codes only receive 0.5012. Therefore, the difference in preference for solar between these zip codes for otherwise identical households is $4.7961 \cdot (0.7240 - 0.5012) = 1.0686$.

³³The highest electricity price in the data is 0.4613 \$/kWh and the lowest is 0.1927 \$/kWh. Therefore, the difference in preference for solar between these zip codes for otherwise identical households is $4.0510 \cdot (0.4613 - 0.1927) = 1.0881$.

³⁴The difference in preference for solar between otherwise identical households with the highest and lowest home values and education in our data is $0.7711 \cdot (\log(1) - \log(0.2)) = 0.5390$ and $0.0986 \cdot 6 = 0.5916$, respectively.

Table 2: Dynamic Demand Estimates

Preference for Solar:	
Constant	-10.6419*** (0.5305)
Median Radiation (kWh/m ² /day/10)	4.7961*** (0.1166)
Electricity Price (\$ per kWh)	4.0510*** (0.2631)
log(Home Value (\$millions))	0.7711*** (0.1564)
Years of Schooling	0.0986*** (0.0331)
SCE	0.2885*** (0.0476)
SDG&E	1.2944*** (0.0606)
Time Trend	0.1247*** (0.0027)
Time Trend ²	-0.0021*** (0.0001)
Preference for Solar Costs:	
Cost (\$10,000s)	-0.5175*** (0.2416)
Cost (\$10,000s)*log(Home Value (\$millions))	-0.1105** (0.0732)
Cost (\$10,000s)*Years of Schooling	-0.0300* (0.0154)
-Log-likelihood	268,399
Months	41
Percent Who Consider Solar	6.6

Standard errors in parentheses. Electricity prices are the average rate in dollars per kilowatt-hour for usage over 200% of baseline. In Southern California Edison (SCE) and Pacific Gas and Electric (PG&E, the omitted utility for the utility fixed effects) this is the average of one rate for 201-300% of baseline and a second rate for 300+% of baseline. In San Diego Gas and Electric (SDG&E) this is an average of the summer and winter rates for usage over 200% of baseline.

as strong evidence that the economic viability of solar in terms of generating electricity for offsetting expensive electric bills is a more important determinant of preferences for solar adoption than are other differences between demographic groups.

We do also find that preferences for solar are significantly different across utilities and, importantly, that preferences for solar are increasing over time at a decreasing rate. In our baseline simulation results, we interpret this time trend as capturing unobservable changes in system quality, as would be expected of a new technology. We eliminate this trend in counterfactual simulations that hold technology constant. We also quantitatively assess the implications of attributing the entire time trend to peer effects.

We find that households get disutility from the net-of-subsidy cost of installing solar and that a household's cost sensitivity increases with both home value and schooling. This somewhat counterintuitive result likely arises because households with high home values and more education are better informed about the evolution of solar system prices and subsidies and are therefore better able to time their solar adoption to pay the lowest out-of-pocket solar price. The effect of education on price sensitivity is substantially larger than the effect of home value: the difference in price sensitivity between the highest and lowest education groups is more than twice as large as the difference in price sensitivity between the highest and lowest home value groups.³⁵

We tested models with other combinations of demographics. Log home value has a larger and more statistically significant effect on preferences than log income, and heterogeneity in preferences based on income is very small and statistically insignificant when home values are already included. Solar radiation and electricity prices do not have a large or statistically significant effect on price sensitivity, which is reasonable because radiation and electricity prices both increase the long-run returns to a solar system but have no impact on the upfront cost of the system. Whether a household voted Democratic, the number of mortgages on the home, and the number of household members do not statistically significantly impact either households' preferences for solar or their sensitivity to system cost once we control for the other demographics. Appendix B.3 discusses additional sensitivity checks of the model specification.

Before moving on to our simulations of the efficient subsidy for rooftop solar, it is helpful

³⁵For otherwise identical households, the difference in price sensitivity between the highest and lowest education groups is $-0.0300 \times 6 = -0.1800$, whereas the difference for the highest and lowest home value groups is $-0.1105 \times (\log(1) - \log(0.2))$.

to compare our estimates to the literature on the California Solar Initiative. Hughes and Podolefsky (2015) use a reduced-form, static analysis on a slightly longer estimation window to estimate the impact of CSI subsidies on solar system adoption. They find that solar installation in California would have been 53% lower in the absence of subsidies. Using our estimates we find that solar installation would have been 60% lower in the absence of subsidies. The fact that Hughes and Podolefsky (2015) reach a very similar conclusion despite using a different methodology with a different sample suggests that our model is in line with the literature's understanding of the aggregate effects of the CSI subsidy.³⁶

6 Results: The Efficient Subsidy for Rooftop Solar

We now use our structural estimates of household values for solar to simulate the efficient subsidy trajectory. The analysis in Section 3 disentangled the forces that could lead the subsidy to increase or decrease over time. We investigate which forces dominate in the setting of the California Solar Initiative.

We calibrate the regulator's benefit to the social value of solar energy (including emission displacement) from Baker et al. (2013), with the concavity of the regulator's benefit reflecting how the intermittent nature of solar energy reduces its marginal value once there is a lot of solar on the electric grid (from Gowrisankaran et al., 2016). We require the regulator to achieve a target of 0.68% adoption in 41 months, which is consistent with the actual policy. The appendix details the calibration and solution method. It also shows that results are qualitatively unchanged for an alternative target of 1.5% adoption.

In order to understand why a regulator might choose a particular subsidy trajectory, we focus on two key drivers of the efficient subsidy schedule: whether consumers are forward-looking, and whether technology is improving. Figure 3 plots the efficient subsidy (left) and expected monthly subsidy spending (right) for the different combinations of assumptions about household foresight and technical change. Table 3 reports the present value of subsidy spending and the initial and terminal subsidy. It also reports the consumer surplus obtained

³⁶Hughes and Podolefsky (2015) remove all observations close to subsidy changes from their sample in order to abstract from the dynamics of the solar adoption decision. They also begin their estimation window in 2007 at the start of the CSI program, whereas our estimation window begins in 2009 when the cap on the federal tax credit was removed.

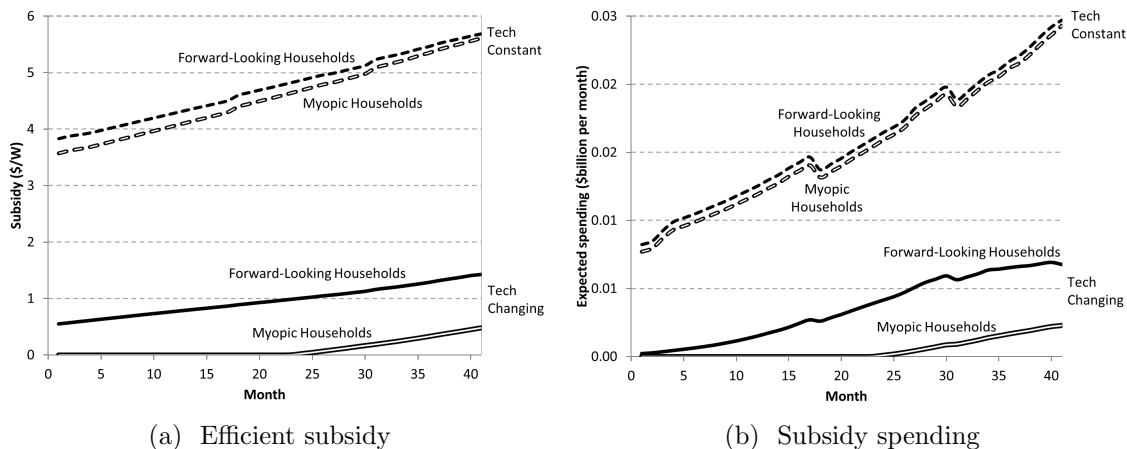


Figure 3: The efficient subsidy (left) and expected monthly spending (right) when households are forward-looking (solid) and myopic (hollow). Connected lines allow for exogenous technical change and dashed lines hold technology constant over time. The jumps in spending correspond to the jumps in PG&E electricity prices seen in Figure 2.

by households in the absence of any subsidy and under the efficient subsidy policy.³⁷ Figure 4 depicts monthly adoption along the efficient subsidy trajectory for forward-looking (solid) and myopic (dotted) households as well as adoption by forward-looking households when they are offered the subsidy that would be efficient for myopic households (dashed).

The first main result is that both the subsidy level and total expected spending are substantially greater when technology is fixed rather than improving over time. If technology is fixed, then there are only two things that could convince initially reluctant consumers to adopt the technology: either the regulator offers them a larger subsidy, or they happen to receive a set of stochastic draws that makes adoption especially attractive in some future period.³⁸ The second effect alone is not strong enough to achieve the targeted level of adoption. When technology is progressing, two forces work to reduce the regulator's spending. First, more and more consumers would adopt the technology even if the subsidy and preferences were fixed over time. The regulator can therefore use a smaller subsidy to achieve a given level of adoption in each period. Second, the regulator will change the timing of adoption along the efficient trajectory so as to substitute technological progress for subsidy spend-

³⁷Appendix C.2 explains how we calculate consumer surplus for forward-looking and myopic households.

³⁸For instance, an installer approaches a household in a given period and provides an unsolicited quote for a solar system.

Table 3: The present value of spending, initial and terminal subsidies, and consumer surplus along the efficient subsidy trajectory, along with the consumer surplus obtained in the absence of subsidies.

	Myopic		Forward-Looking	
	Technology Constant	Technology Improving	Technology Constant	Technology Improving
Present value of spending (\$million)	495	15	514	110
Present value of spending (\$/W)	3.7	0.1	3.8	0.8
Initial subsidy (\$/W)	3.6	0.0	3.8	0.6
Terminal subsidy (\$/W)	5.6	0.5	5.7	1.4
Consumer surplus without any subsidy (\$million)	26	183	84	918
Consumer surplus with efficient subsidy (\$million)	202	194	304	1003

ing. Comparing the left and right panels of Figure 4, we see that introducing technological progress leads the regulator to delay adoption until later periods, when she does not need as high a subsidy to obtain adoption.

The second main result is somewhat less intuitive: it is more expensive to obtain adoption from forward-looking households than it is to obtain adoption from myopic households, especially when technology is improving. This is because forward-looking households account for future changes in the subsidy and in technology and for the possibility of future preference shocks. Today’s subsidy must compensate forward-looking households not just for losses from installing solar but also for forsaking the option to adopt solar in some later period. When technology is fixed, households’ foresight increases the regulator’s spending by \$19 million (4%). The effect of foresight is even stronger in the presence of technological change because households’ option to adopt in a later period becomes especially valuable: when technology is improving, households’ foresight increases the regulator’s spending by \$96 million (637%).

Households’ foresight is good for the households, however: the regulator’s spending increases because households capture more surplus when they are forward-looking. As we will see below, households’ foresight constrains the regulator’s ability to intertemporally price discriminate and thereby avoid “over-subsidizing” high-value households in early periods. As a result, the subsidy program increases myopic households’ surplus by only \$175 million (\$12 million) when technology is constant (improving) but increases forward-looking households’ surplus by \$219 million (\$86 million). In the absence of a subsidy, technological change

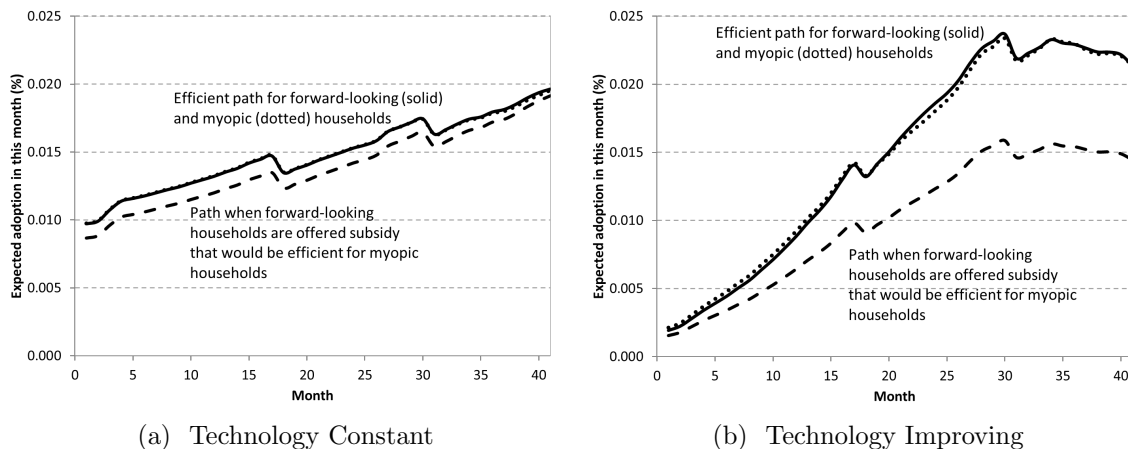


Figure 4: Monthly expected adoption along the efficient subsidy schedule for forward-looking households (solid), along the efficient subsidy for myopic households (dotted), and for forward-looking households who are offered the subsidy that would be efficient for myopic households (dashed). The jumps in adoption correspond to the jumps in PG&E electricity prices seen in Figure 2.

increases forward-looking households' surplus by \$833 million, but technological change increases myopic households' surplus by only \$156 million because they do not optimize the timing of their adoption. With the efficient subsidy, the contrast is even starker: myopic households actually *lose* \$7 million in surplus under technological change, whereas forward-looking households gain \$700 million.³⁹ This difference arises because, when households are myopic, the forward-looking regulator can optimize the subsidy schedule to increase adoption in later periods when costs have fallen (see Figure 4) and can thereby capture more of the benefits of technological change. The delay in adoption reduces the present value of consumer surplus for myopic households below what it was with constant technology. In contrast, forward-looking households' awareness of future subsidies and technological improvements forces the regulator to share the cost reductions enabled by technological change.

We now consider the slopes of the efficient subsidy trajectories in Figure 3 in more detail. Critically, we do not need to speculate about why some subsidy trajectories increase strongly

³⁹Comparisons of consumer surplus between myopic and forward-looking households must be undertaken with caution since the calculation is somewhat different for the two consumer types. Clearly, forward-looking consumers' ability to time adoption should weakly increase their welfare relative to a situation where they were forced to adopt solar in the first moment that their current benefits exceed their current costs. That is close to, but not exactly, the calculation that is presented here. See Appendix C.2 for details.

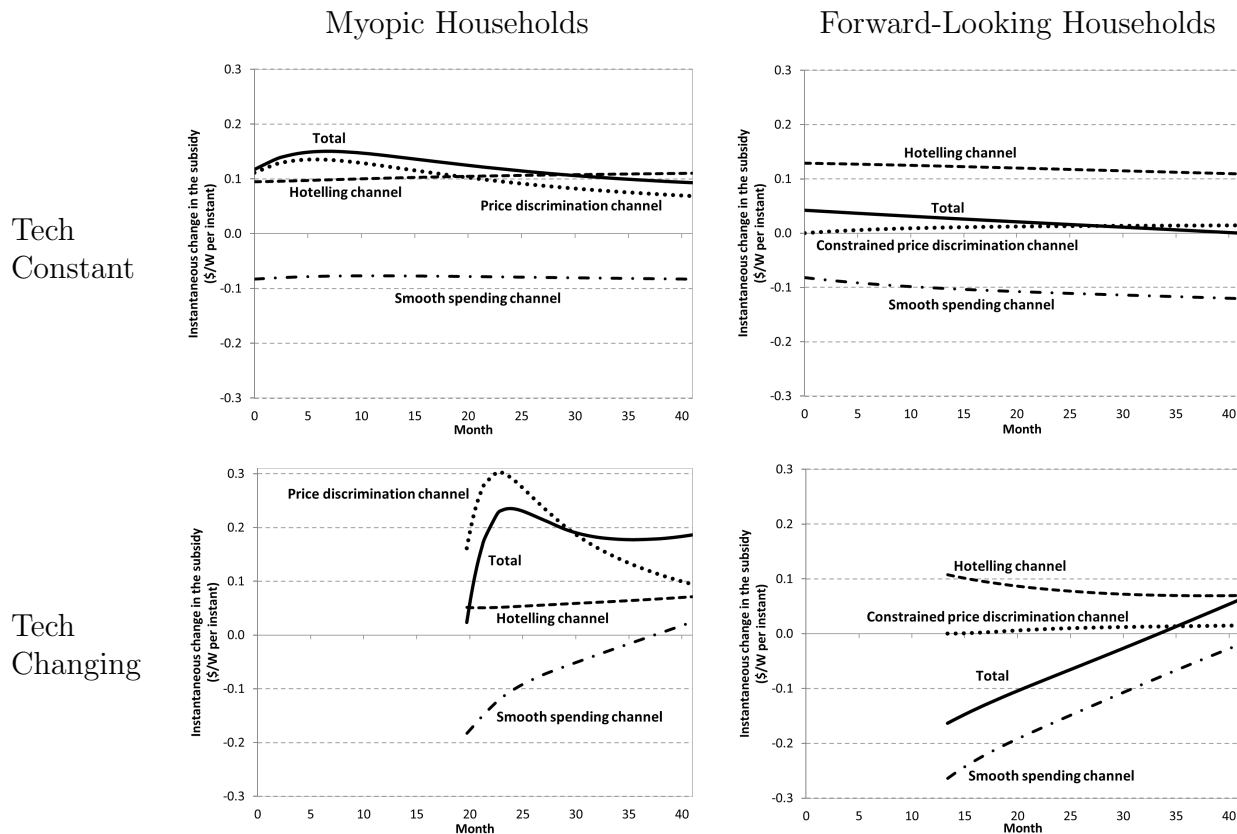


Figure 5: The change in the efficient subsidy at each instant ($\dot{s}(t)$, labeled “Total”), as well as each component from equation (4). The adoption benefit component (not plotted) is negative but very small in magnitude. When technology is changing, the regulator chooses to start the subsidy after the initial period to take advantage of technological progress.

and some are flatter. Instead, we use the theoretical analysis from Section 3 to disentangle the multiple forces that determine whether the efficient subsidy increases or decreases over time.⁴⁰

In Figure 5, the thick bold line gives the instantaneous change in the subsidy ($\dot{s}(t)$). The

⁴⁰The only approximations are that we impose the theoretical setting’s restrictions that preferences are fixed over time and that technology evolves deterministically. We also fix electricity prices at their initial levels. The appendix plots the efficient subsidy and adoption trajectories under these restrictions. The efficient subsidy trajectories are qualitatively similar to the cases in Figure 3. The main differences are that these restrictions increase the level of the subsidy and that they change the subsidy trajectory in the case with forward-looking households and technological progress so that the subsidy’s start is delayed and the subsidy declines over an initial interval once it does begin. The appendix explains these differences in more detail, emphasizing how they arise from eliminating stochasticity in households’ preferences.

other lines are the components of that instantaneous change identified in the theoretical analysis, so that their vertical sum also equals $\dot{s}(t)$. When households are myopic and technology is constant (top left panel), the efficient subsidy increases strongly over time because the price discrimination channel is large. By starting with a relatively small subsidy and raising it over time, the regulator avoids paying a large subsidy to households who would adopt even for a smaller subsidy. In the full model, the regulator reduces spending by 3.4% by using an increasing subsidy instead of the unique constant subsidy that exactly achieves the adoption target. However, when households are forward-looking (top right panel of Figure 5), their expectations and ability to time adoption constrain the regulator's ability to intertemporally price discriminate. The constrained price discrimination channel begins at zero and increases only slowly. The efficient subsidy therefore increases more slowly for forward-looking households.⁴¹ In the full model, the efficient subsidy now reduces spending by only 2.9% relative to a constant subsidy.

Figure 4 highlights forward-looking households' incentive to delay adopting the solar technology in the full model. The dashed lines give the adoption rate per month if the forward-looking households were offered the subsidy that would be efficient for myopic households. The left panel shows the case without technological change. Here, forward-looking households' incentives to delay adoption arise from the increasing subsidy schedule (seen in Figure 3). We see that the slightly more sharply increasing subsidy offered to myopic households would dampen adoption by forward-looking households in early periods. Their willingness to wait for the high later subsidies limits the regulator's ability to price discriminate through an increasing subsidy schedule.

The bottom row of Figure 5 explains why introducing technical change flattens the efficient subsidy trajectory. When households are myopic, the regulator delays the start of the subsidy for several months in order to take advantage of technological progress. Once the subsidy does begin, advancing technology amplifies the price discrimination channel by increasing the number of inframarginal households, but the smooth spending channel is now strongly negative. The net effect is to flatten the subsidy trajectory. Technological change has an even stronger effect in the case of forward-looking households: the constrained price

⁴¹Using an increasing subsidy does not eliminate early adoption when households are forward-looking and technology is constant because households who are not perfectly patient will adopt the technology as long as the subsidy is not increasing too fast. The more patient that households are, the less freedom the regulator has to use an increasing subsidy.

discrimination channel is unchanged in the first instants and only slightly altered in later instants even as the effect on the smooth spending channel is even greater. In the full model, the enhanced incentive to price discriminate when households are myopic and technology is changing allows the regulator to reduce spending by 9.6% relative to a case with a constant subsidy, but the regulator can reduce spending relative to a constant subsidy by only 2.0% when households are forward-looking.

The right panel of Figure 4 shows how households' foresight flattens the efficient subsidy trajectory in the full model with technological change. The efficient subsidy is initially fixed at zero for myopic households because the regulator waits to take advantage of technological improvements.⁴² Offering that constant subsidy to forward-looking households does not substantially dampen adoption in early periods. However, forward-looking households would substantially delay adoption in the middle periods if offered the subsidy designed for myopic households because they are willing to wait for the subsidy to become nonzero. The efficient subsidy offered to forward-looking households changes more smoothly in order to convince enough forward-looking households not to postpone adopting the technology.

7 Comparison to Existing Subsidy Paths

We have seen that the efficient subsidy increases over time, but the California regulator used a strongly declining subsidy. What types of assumptions can generate that type of subsidy trajectory? The left panel of Figure 6 varies parameters in the regulator's objective. We see that giving the regulator a much more convex cost of funds and increasing its marginal benefit of adoption tenfold can both make the subsidy decline over time.⁴³ A more convex cost of funds makes the regulator use a declining subsidy in order to smooth spending over time, and a regulator with a greater marginal benefit of adoption is less inclined to defer adoption to later periods. Additional experiments showed that reducing the regulator's discount rate by half has similar, albeit weaker, effects as increasing the regulator's marginal

⁴²When households are myopic and technology is changing, expected adoption in the absence of a subsidy would reach 0.64% by the end of the 41-month interval, just shy of the target of 0.68%. The regulator therefore does not find it costly to concentrate the subsidy in later periods. In contrast, when households are forward-looking, expected adoption would reach only 0.46% in the absence of a subsidy.

⁴³Specifically, the case with more convex funds reduces the spending level at which the cost of funds doubles to 0.1% of the original value, and the case with a greater marginal benefit of adoption increases the marginal benefit of the first unit of adoption by the stated amount.

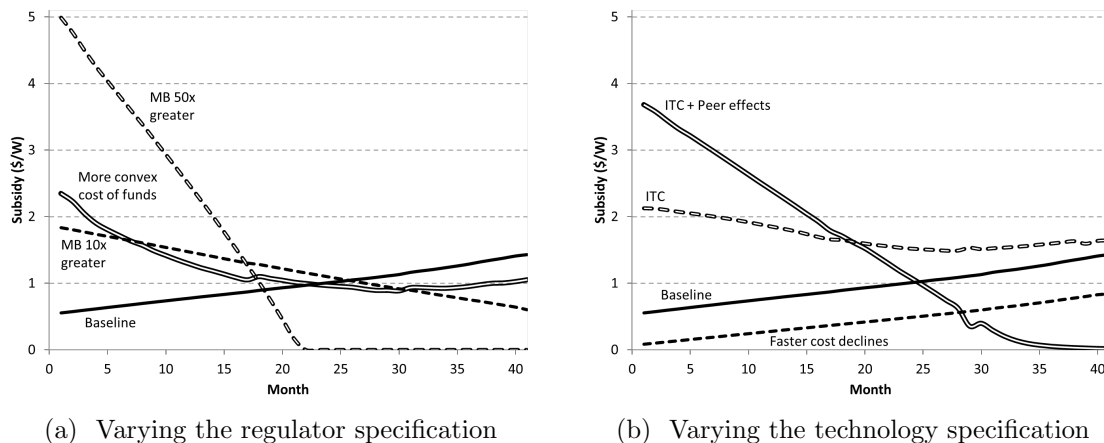


Figure 6: Left: The effect on the efficient subsidy of varying parameters in the regulator’s objective. Right: The effect on the efficient subsidy of different assumptions about the pace and origin of technical progress, as described in the text. All cases use forward-looking households and allow for technological change.

benefit tenfold. A more patient regulator is less inclined to defer spending to later periods. Yet even in these cases, the efficient subsidy still does not decline nearly as sharply as the actual subsidy.

However, a final case is different: increasing the regulator’s marginal benefit of adoption fiftyfold can generate the type of sharply declining subsidy trajectory seen in practice. In our baseline calibration, the adoption benefit channel was trivial, but the same theoretical decomposition here shows that a large adoption benefit channel drives the declining subsidy trajectory. The regulator values solar so much that it wants to speed up adoption in order to obtain the benefits of solar electricity sooner, even at the cost of having to offer a greater subsidy early on. These results are especially intriguing because it is plausible that California regulators did value solar installations to a much greater degree than recommended by the economic analyses to which we calibrated the regulator’s objective. This disagreement about the marginal social value of solar would simultaneously explain why California regulators chose to use a subsidy to spur adoption at all and justify why California regulators designed such a sharply declining subsidy.

The right panel of Figure 6 varies assumptions about the pace and origin of technical change. Assuming that costs fall twice as fast as they did in reality does not substantially

affect the shape of the subsidy trajectory. The remaining cases explore extreme assumptions about the potential endogeneity of technology and preferences. The theoretical analysis showed that allowing early adoption to lower costs (or increase value) among later adopters favors a declining subsidy. Here, the induced technical change (“ITC”) case assumes that *all* realized changes in cost were due to past adoption in California, as described in Appendix C.2. This extreme assumption rules out installations anywhere else in the world having an effect on solar costs, so it severely overstates—and therefore bounds—the potential size of the effect from ITC. Even in this extreme case, the efficient subsidy decreases only slowly before flattening out. The case with “ITC + Peer effects” additionally assumes that the *entire* time trend in the preference for solar was due to past adoption in California. We have heretofore interpreted this time trend as picking up improvements in the quality of solar technology, so the present interpretation of that trend as arising entirely from peer effects assumes that, absent installations in California, solar system quality would have been constant over this period. This combination of extreme assumptions about ITC and peer effects does generate a sharply declining efficient subsidy. In ongoing work, we are assessing whether the extreme assumption about peer effects is sufficient to generate a sharply declining subsidy on its own. Preliminary results suggest not.

8 Conclusions

We have demonstrated the forces that determine how to design a subsidy to induce adoption of a new technology over time. In particular, we have shown that a regulator can reduce its overall spending by using an increasing subsidy schedule as a means of intertemporally price discriminating. The regulator’s ability to price discriminate is constrained by consumers’ willingness to wait for later, higher subsidies. Quantitatively, rational expectations of future subsidies increase the regulator’s spending by 4% in the absence of technological change. Further, when consumers have rational expectations about technological progress, the regulator must offer them a relatively large subsidy in order to compensate them for forsaking their option to adopt the technology in a later period. Quantitatively, technological progress would reduce the regulator’s spending by 97% if households were myopic, but technological progress reduces the regulator’s spending by only 79% when households are forward-looking. Rational expectations increase the total cost of the policy program more than sixfold in a

world with technological progress.

Future work should explore the implications of rational expectations and technological dynamics in other policy environments. For instance, economists commonly recommend emission taxes that increase over time and subsidies for research that would improve future technology. Yet many economic models abstract from consumer and firm expectations of policy and of technology. In addition, our results suggest that a regulator could reduce spending if it could convince households to ignore the possibility of higher subsidies or better technology in the future. Future theoretical analysis could consider when regulators may have an incentive to mislead households, and future empirical work could test for such effects in actual policy environments.

References

- Aguirregabiria, Victor and Aviv Nevo (2013) “Recent developments in empirical IO: Dynamic demand and dynamic games,” in Daron Acemoglu, Manuel Arellano, and Eddie Dekel eds. *Advances in Economics and Econometrics. Tenth World Congress of the Econometric Society. Volume 3, Econometrics*, New York: Cambridge University Press.
- Allcott, Hunt and Michael Greenstone (2012) “Is there an energy efficiency gap?” *Journal of Economic Perspectives*, Vol. 26, No. 1, pp. 3–28.
- Baker, Erin, Meredith Fowlie, Derek Lemoine, and Stanley S. Reynolds (2013) “The economics of solar electricity,” *Annual Review of Resource Economics*, Vol. 5, No. 1, pp. 387–426.
- Berry, Steven, James Levinsohn, and Ariel Pakes (2004) “Differentiated products demand systems from a combination of micro and macro data: The new car market,” *Journal of Political Economy*, Vol. 112, No. 1, pp. 68–105.
- Bollinger, Bryan and Kenneth Gillingham (2012) “Peer effects in the diffusion of solar photovoltaic panels,” *Marketing Science*, Vol. 31, No. 6, pp. 900–912.
- Boomhower, Judson and Lucas W. Davis (2014) “A credible approach for measuring infra-marginal participation in energy efficiency programs,” *Journal of Public Economics*, Vol. 113, pp. 67–79.

- Borenstein, Severin (2017) “Private net benefits of residential solar PV: The role of electricity tariffs, tax incentives, and rebates,” *Journal of the Association of Environmental and Resource Economists*, Vol. 4, No. S1, pp. S85–S122.
- Burr, Chrystie (2014) “Subsidies, tariffs and investments in the solar power market,” working paper.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer (2013) “Are consumers myopic? Evidence from new and used car purchases,” *American Economic Review*, Vol. 103, No. 1, pp. 220–56.
- Chamley, Christophe (1986) “Optimal taxation of capital income in general equilibrium with infinite lives,” *Econometrica*, Vol. 54, No. 3, pp. 607–622.
- Chandra, Ambarish, Sumeet Gulati, and Milind Kandlikar (2010) “Green drivers or free riders? An analysis of tax rebates for hybrid vehicles,” *Journal of Environmental Economics and Management*, Vol. 60, No. 2, pp. 78–93.
- Chari, V.V. and Patrick J. Kehoe (1999) “Optimal fiscal and monetary policy,” in John B. Taylor and Michael Woodford eds. *Handbook of Macroeconomics*, Vol. Volume 1, Part C: Elsevier, pp. 1671–1745.
- Coase, R. H. (1972) “Durability and monopoly,” *Journal of Law and Economics*, Vol. 15, No. 1, pp. 143–149.
- Conlon, Christopher T. (2010) “A dynamic model of prices and margins in the LCD TV industry,” working paper.
- De Groote, Olivier and Frank Verboven (2016) “Subsidies and myopia in technology adoption: Evidence from solar photovoltaic systems,” CEPR Discussion Paper DP11438.
- Dupas, Pascaline (2014) “Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment,” *Econometrica*, Vol. 82, No. 1, pp. 197–228.
- Gowrisankaran, Gautam, Stanley S. Reynolds, and Mario Samano (2016) “Intermittency and the value of renewable energy,” *Journal of Political Economy*, Vol. 124, No. 4, pp. 1187–1234.

- Gowrisankaran, Gautam and Marc Rysman (2012) “Dynamics of consumer demand for new durable goods,” *Journal of Political Economy*, Vol. 120, No. 6, pp. 1173–1219.
- Heal, Geoffrey (1976) “The relationship between price and extraction cost for a resource with a backstop technology,” *The Bell Journal of Economics*, Vol. 7, No. 2, pp. 371–378.
- Hotelling, Harold (1931) “The economics of exhaustible resources,” *The Journal of Political Economy*, Vol. 39, No. 2, pp. 137–175.
- Hughes, Jonathan E. and Molly Podolefsky (2015) “Getting green with solar subsidies: Evidence from the California Solar Initiative,” *Journal of the Association of Environmental and Resource Economists*, Vol. 2, No. 2, pp. 235–275.
- Judd, Kenneth L. (1985) “Redistributive taxation in a simple perfect foresight model,” *Journal of Public Economics*, Vol. 28, No. 1, pp. 59–83.
- Kalish, Shlomo and Gary L. Lilien (1983) “Optimal price subsidy policy for accelerating the diffusion of innovation,” *Marketing Science*, Vol. 2, No. 4, pp. 407–420.
- Kocherlakota, Narayana (2010) *The New Dynamic Public Finance*, Princeton: Princeton University Press.
- Kremer, Michael and Jack Willis (2016) “Guns, latrines, and land reform: Dynamic Pigouvian taxation,” *American Economic Review: Papers and Proceedings*, Vol. 106, No. 5, pp. 83–88.
- Landsberger, Michael and Isaac Meilijson (1985) “Intertemporal price discrimination and sales strategy under incomplete information,” *The RAND Journal of Economics*, Vol. 16, No. 3, pp. 424–430.
- Lucas, Robert E. (1976) “Econometric policy evaluation: A critique,” *Carnegie-Rochester Conference Series on Public Policy*, Vol. 1, pp. 19–46.
- Magnac, Thierry and David Thesmar (2002) “Identifying dynamic discrete decision processes,” *Econometrica*, Vol. 70, No. 2, pp. 801–816.
- Meyer, Laurence H., Joel L. Prakken, and Chris P. Varvares (1993) “Designing an effective investment tax credit,” *Journal of Economic Perspectives*, Vol. 7, No. 2, pp. 189–196.

- Nevo, Aviv (2001) “Measuring market power in the ready-to-eat cereal industry,” *Econometrica*, Vol. 69, No. 2, pp. 307–342.
- Newell, Richard G., William A. Pizer, and Daniel Raimi (2017) “Federal government subsidies for clean energy,” working paper.
- Pless, Jacquelyn and Arthur A. van Benthem (2017) “The surprising pass-through of solar subsidies,” Working Paper 23260, National Bureau of Economic Research.
- Reddix II, Kenneth (2014) “Powering demand: Solar photovoltaic subsidies in California,” job market paper, University of North Carolina.
- Salant, Stephen W. (1989) “When is inducing self-selection suboptimal for a monopolist?” *The Quarterly Journal of Economics*, Vol. 104, No. 2, pp. 391–397.
- Stokey, Nancy L. (1979) “Intertemporal price discrimination,” *The Quarterly Journal of Economics*, Vol. 93, No. 3, pp. 355–371.
- (1981) “Rational expectations and durable goods pricing,” *The Bell Journal of Economics*, Vol. 12, No. 1, pp. 112–128.
- Waldman, Michael (2003) “Durable goods theory for real world markets,” *The Journal of Economic Perspectives*, Vol. 17, No. 1, pp. 131–154.

Appendix

The first section formally derives the efficient subsidy trajectory analyzed in the main text. The second section is an empirical appendix that presents evidence that households are forward-looking when considering whether to install solar, provides additional information on our data and empirical approach, and reports sensitivity tests for the dynamic estimation. The third section describes the numerical calibration of the policymaker’s problem. It also describes the solution techniques. The fourth section reports results for an alternate target of 1.5%. The final section analyzes the efficient subsidy trajectory for the case of a fixed budget. It shows that the results are effectively identical to the case with a fixed adoption target, with only one minor difference.

A Formal Analysis of the Theory Model

The Hamiltonian is

$$\begin{aligned}
 H(t, y(t), Y(t), s(t), \lambda(t), \mu(t)) = & e^{-rt} \left[B \left(1 - F(Y(t)) \right) - G \left(-s(t) f(Y(t)) y(t) \right) \right] \\
 & + e^{-rt} \lambda(t) y(t) \\
 & + e^{-rt} \mu(t) \delta [Y(t) - C(t, 1 - F(Y(t))) + s(t)] \\
 & + e^{-rt} \mu(t) [C_1(t, 1 - F(Y(t))) - C_2(t, 1 - F(Y(t))) f(Y(t)) y(t)].
 \end{aligned}$$

$\lambda(t)$ gives the (current) shadow value of $Y(t)$. The costate variable $\mu(t)$ measures the degree to which the regulator is constrained at each instant by private actors’ equilibrium behavior and rational expectations: it measures the cost of keeping promises made to those who adopted the technology in past periods.

The necessary conditions for a maximum are:

$$\begin{aligned} \lambda(t) = & -s(t) f(Y(t)) G' \left(-s(t) f(Y(t)) y(t) \right) \\ & + \mu(t) C_2(t, 1 - F(Y(t))) f(Y(t)), \end{aligned} \quad (\text{A-1})$$

$$\begin{aligned} -\dot{\lambda}(t) + r\lambda(t) = & -f(Y(t)) B' \left(1 - F(Y(t)) \right) \\ & + s(t) f'(Y(t)) y(t) G' \left(-s(t) f(Y(t)) y(t) \right) + \delta\mu(t) \\ & + \delta\mu(t) f(Y(t)) C_2(t, 1 - F(Y(t))) \\ & + \mu(t) \left[-C_{12}(t, 1 - F(Y(t))) f(Y(t)) + C_{22}(t, 1 - F(Y(t))) [f(Y(t))]^2 y(t) \right. \\ & \left. - C_2(t, 1 - F(Y(t))) f'(Y(t)) y(t) \right], \end{aligned} \quad (\text{A-2})$$

$$-\dot{\mu}(t) + r\mu(t) = f(Y(t)) y(t) G' \left(-s(t) f(Y(t)) y(t) \right) + \delta\mu(t), \quad (\text{A-3})$$

$$\mu(0) = 0,$$

along with the transition equations and the initial and terminal conditions. The first equation follows from the Maximum Principle, the next two equations are the costate (or adjoint) equations, and the final equation is the transversality condition corresponding to the choice of s_0 .

Solving for $\mu(t)$ in equation (A-3), we find

$$\mu(t) = - \int_0^t e^{-(\delta-r)(t-i)} f(Y(i)) y(i) G' \left(-s(i) f(Y(i)) y(i) \right) di \geq 0. \quad (\text{A-4})$$

Whereas costate variables in standard optimal control problems are forward-looking, here the costate variable $\mu(t)$ is backward-looking. The costate variable $\mu(t)$ is the discounted value of all past adoption. In the first instant, the regulator is not bound by past commitments, but over time the regulator becomes more bound by the commitments it has made in order to obtain past adoption. The regulator's promises accrue as past adoption accrues, and these promises decay at rate $\delta - r$. The more impatient that potential adopters are (the higher is δ), the faster that past commitments fade away and the smaller $\mu(t)$ is. The more impatient that the regulator is (the higher is r), the more commitments tend to accumulate and the larger $\mu(t)$ is: the regulator obtained early benefits by making promises about later dates' subsidies.

The result that $\mu(t) \geq 0$ implies from equation (A-1) that $\lambda(t) \leq 0$: because lower $Y(t)$ corresponds to greater adoption $Q(t)$, this negative sign means that the shadow benefit of

adoption is positive. Differentiate equation (A-1) with respect to time to obtain:

$$\begin{aligned} \dot{\lambda}(t) = & -\dot{s}(t) f(Y(t)) G' - s(t) f'(Y(t)) y(t) G' \\ & - s(t) f(Y(t)) G'' \left[-\dot{s}(t) f(Y(t)) y(t) - s(t) f'(Y(t)) [y(t)]^2 - s(t) f(Y(t)) \dot{y}(t) \right] \\ & + \dot{\mu}(t) C_2(t, 1 - F(Y(t))) f(Y(t)) \\ & + \mu(t) C_{12}(t, 1 - F(Y(t))) f(Y(t)) \\ & - \mu(t) C_{22}(t, 1 - F(Y(t))) [f(Y(t))]^2 y(t) \\ & + \mu(t) C_2(t, 1 - F(Y(t))) f'(Y(t)) y(t), \end{aligned}$$

where we suppress the argument of $G(\cdot)$. Substituting for $\dot{\lambda}(t)$ from the costate equation (A-2) and rearranging yields the solution for $\dot{s}(t)$ given in the main text.

B Empirical Appendix

B.1 Evidence that Consumers are Forward-Looking

Before we estimate consumers’ private values for residential solar, we must understand whether consumers actually do think about future solar costs and subsidies when they make their investment decisions. There are a few reasons why we might think that the dynamic trajectory of prices and subsidies would matter for households’ solar installation decisions. First, there are examples in the literature on goods other than residential solar where consumers make decisions in expectation of future policy changes (e.g., Mian and Sufi, 2012). Second, it is not necessary that households themselves are informed about the future trajectory of solar subsidies as long as solar installers are informed. If installers use the fact that subsidies will be changing to encourage households to install solar now, then households will act as if they were directly informed about the subsidy schedule.

To understand whether households in California behave as if they are forward-looking, we begin by looking at the monthly installations in each utility along with the dates on which subsidies change (recall that the subsidy levels are shown in Figure 1 in the text). Figure B1 shows that there are large increases in the number of permits filed to install solar in the last full month before a utility’s subsidy declines.¹ This suggests that households are anticipating the subsidy decline.

While Figure B1 provides some basic evidence that consumers are forward-looking, it does not control for other factors that might affect households’ decision to adopt. In order to better understand whether households are reacting to dynamic incentives, we regress the weekly counts of residential solar subsidy applications on measures of current and future

¹This effect has also been shown in Burr (2014) and Hughes and Podolefsky (2015).

Figure B1: Monthly Installations and Dates of Subsidy Declines by Utility



solar system costs and controls. In particular, we focus on future drops in the CSI subsidy, changes in the prices of solar modules and of silicon, and changes in the US-China exchange rate. As we saw above, when the subsidy is about to decrease, more households will choose to install solar now if households are forward-looking. Similarly, for a given current cost of solar modules, high input (silicon) costs might suggest that the cost of solar modules is about to increase, so forward-looking households will want to install now. Since China produces a large number of solar panels, if the US-China exchange rate is increasing, then panels will be less expensive in the future and forward-looking households should wait to invest in solar.

In order to test whether household behavior responds to these measures of future costs, we combine the CSI application data with a solar module price index and a silicon price index from Bloomberg. We also include the realized 3-month change in the dollar-yuan exchange rate and a set of controls that includes the VIX (a measure of the expected volatility of U.S. equities), the current dollar-yuan exchange rate, the current Euro-yuan exchange rate, the 6-month Treasury bill interest rate, and the DOW real estate investment trust index. These controls aim to capture macroeconomic conditions that affect the incentive to invest in solar.²

²We instrument for the VIX with month-over-month changes in 10-year bond yields for several econom-

Table B1: Evidence That Consumers are Forward-Looking

Dependent Variable: Log Weekly CSI Applications	
Within 2 months of subsidy drop	0.300*** (0.055)
Solar Module Price Index	-2.701*** (0.619)
Silicon Price Index	0.012** (0.006)
3 month change in Dollar-Yuan Exchange Rate	-2.045*** (0.694)
VIX (Instrumented)	-0.032** (0.014)
Number of utility-weeks	650
R^2	0.6779

Standard errors clustered by week. Controls include the Dollar-Yuan exchange rate, the Euro-Yuan exchange rate, the 6 month T-bill interest rate, the DOW REIT index, utility fixed effects and utility-specific time trends. Instruments are month-over-month changes in 10 year bond yields for Greece, Italy, Russia, and Spain. $F=482.55$ Results are not substantially different if the VIX is not instrumented.

The regression results presented in Table B1 suggest that consumers are forward-looking in their decision to install solar (or at least that solar installers are forward-looking and convey this information to households). The coefficient estimates suggest that more households submit applications to install solar systems when subsidies are about to decline.³ Higher solar module prices reduce installation, but conditional on solar module prices, higher silicon prices (which suggest that module costs may be higher in the future) increase installation. Similarly, if the US-China exchange rate is increasing (conditional on the current level of the exchange rate) then solar panels will likely be less expensive in the future. Indeed, our estimates suggest that an increasing US-China exchange rate reduces installations today. Finally, it is interesting that economic uncertainty, as captured by the instrumented VIX, tends to decrease solar installations, as would be expected if households account for option

ically volatile developed countries (Greece, Italy, Russia, and Spain) in order to isolate sources of market uncertainty that should not directly affect California households through channels such as employment or income.

³Since the CSI subsidies were tied to cumulative installation in each utility, it is possible that higher adoption is causing subsidy declines. To test for this, we conducted a separate analysis restricted to media markets served by two utilities. If local conditions were leading households to install solar and thereby trigger declines in the subsidy, then we would see increased adoption in households in one utility when their neighboring utility was about to have a subsidy decline. We found no evidence that this was true, leading us to conclude that reverse-causality is not a major problem for this descriptive analysis.

value when timing their installation decision (which again implies that households are not fully myopic).

Given that households behave as if they are forward-looking, it is important to use a dynamic model of residential solar adoption in order to estimate the benefits of installing a solar system. If we used a static model of solar adoption, we would underestimate the benefits to households of adopting: some households that value solar above the current system cost will choose not to install for now as they wait for technology to advance and for costs to drop. As we show in the main test, correctly accounting for the value of waiting for lower prices is critical to understanding the trade-offs regulators face in structuring solar subsidies.

B.2 Data Details

Sample

When we estimate our dynamic solar adoption model, we limit our estimation sample in a few important ways to make sure that our model fits what actually happened in California during the CSI program. First, we limit the timeframe of our analysis. In the early months of the CSI, there was a 30% federal tax credit available for residential solar installation, represented by ϕ . However, the federal tax credit was initially capped at \$2,000 (which was strictly binding for most systems in California), and this cap was only lifted on January 1, 2009. Because we do not know whether households anticipated this federal policy change, we limit our analysis to installations that occurred after the cap was lifted.

Additionally, because the CSI data only includes information on installations that applied for CSI subsidy funding, we need to end our analysis when the CSI data ends. There is some evidence that some installers were not submitting applications for CSI subsidies at the very end of the program when the subsidy was very low, so we end our estimation period in May of 2012 when all utilities still had a subsidy (our results are fundamentally unchanged if we end the estimation window 3 months earlier). These restrictions leave us with a 41-month time-frame for estimating households’ preferences. We solve recursively for the value function by working backwards from the value function’s fixed point in the first period in which all utilities’ subsidies were zero, but we limit our maximum likelihood estimation window to only those periods where we were sure that all three utilities had positive subsidy levels and there were enough subsidy applications for the reported system costs to be reliable.

We geographically limit our estimation sample by only estimating preferences for households in those zip codes that are wholly part of one of the three major California utilities. We therefore exclude households in zip codes that are not serviced by PG&E, SCE, or SDG&E and households in zip codes that are serviced by more than one of those utilities. We found that households did not react to subsidy changes by utilities with customers in the household’s media market but who were not actually serving that particular zip code, so we are not concerned about spillover effects from subsidy changes in nearby zip codes.

Finally, some households in our data actually “lease” residential solar systems rather than purchasing them outright, meaning that an outside firm pays the upfront cost of the system and then shares the benefits of the solar generation over time. Since the reported costs for these systems may differ from purchased systems (Podolefsky, 2013), we might be concerned about treating these leases as purchases if the rate of leasing varies systematically with subsidy changes. Third-party ownership is reported in our data, and while we did find some reduced-form evidence that leased systems report lower costs than purchased systems, we did not find a substantial systematic change in the leasing rate near subsidy changes and therefore treat both leased and purchased systems identically. We assess sensitivity to this assumption below.

Electricity prices

We use the average marginal electricity price for usage above 200% of baseline as the electricity price in our analysis. We collect price data from each utility’s rate reports, which list rates for each tier of electricity usage as well as the dates for which the given rate is effective. We assume that the rate in each month is the rate that is effective for the majority of days in that month.

Each utility’s pricing structure is slightly different. San Diego Gas and Electric (SDG&E) has a single marginal price for usage over 200% of baseline in any given month, but this price varies based on “summer” and “winter” months. We average the current month summer and winter prices in this case to get a sense of the average annual cost of electricity. Southern California Edison (SCE) and Pacific Gas and Electric (PG&E) have identical announced prices for summer and winter months, but do charge different marginal prices for usage between 201 and 300% of baseline and for usage over 300% of baseline (although PG&E charged the same marginal price for both of these tiers starting in June of 2010). Here again we average the two prices to get the average marginal price of electricity over 200% of baseline.

Cost evolution

We use the average per-Watt installed system price in each of the three utilities in each month to estimate the parameters of the AR1 model for system costs, assuming that all utilities have the same transition function. When expressed in tens of thousands of dollars for the average 5.4kW system, we estimated costs to evolve as $C_{i(t+1)} = 6.468 \times 10^{-4} + 0.9925C_{it} + \omega_{it}$, where ω_{it} is mean-zero and has a standard deviation of 0.1611.

B.3 Sensitivity of Dynamic Estimation

Variables and Timing

In addition to the empirical specification presented, we estimated models that included log income, an indicator for Democrat, the number of home mortgages, the number of people in the household, and, instead of median radiation in each zip code, mean radiation in each zip code. Income has a similar effect to home prices but has a larger standard error. Democrat, mortgages, and number of people in the household were not statistically significant in predicting installation decisions. Replacing median radiation with mean radiation produced very similar results.

In reality, the step-wise declines in the CSI subsidy were triggered once a utility achieved a predefined level of total solar adoption. Since it was broadly announced when the quantity targets were approaching, we make the assumption that the subsidies declined on set dates rather than at set installation quantities. In the theoretical setting, the regulator understands how the subsidy will determine expected adoption over time, so she could just as easily announce the subsidy as a function of time or of cumulative adoption. However, we could be concerned that households who were unsure when the subsidy would drop might adopt solar earlier to ensure that they receive the higher subsidy. Empirically, running the model with subsidy changes moved one period forward resulted in coefficients that were nearly identical to and statistically indistinguishable from the baseline coefficients, so we use the actual timing of the subsidy schedule in our baseline results.

Finally, there is some uncertainty as to when the subsidies officially ended in each utility’s service area. The lowest subsidy level was low enough that it appears that some installers found that the cost of submitting the paperwork for the subsidy did not justify the benefit to be received. In order to test whether the uncertainty surrounding the end date of the subsidy is affecting our results, we estimated a model where we assumed that the subsidies ended three periods earlier than they do in our baseline specification. Recall that this does not change our estimation window since we end our estimation before the actual end of the CSI. Our results are fundamentally unchanged.

Share of Households that Consider Installing Solar

As explained in the text, we estimate a model where the percent of households who consider installing solar is a free parameter and find that the likelihood is maximized when 6.6% of households consider installing solar. In that model, other coefficients are quite similar to our baseline model, although standard errors are larger. In our primary specifications, we therefore assume that 6.6% of households consider installing solar. We limit the share of households that consider installing solar primarily because we are concerned that if we were to estimate a model with i.i.d. draws and every household considering adopting solar, we would be obscuring a substantial amount of serial correlation in the residuals that actual

households receive each month. Some homes are just not suited for solar systems, whether that is because of the age of the building or the orientation of the roof (if the majority of suitable roof space is north-facing, then a standard solar system will not generate much electricity and will not pay for itself in reduced electricity bills). By including a permanent, random shock to preferences in the form of the “considered” variable, we allow for a degree of serial correlation that is otherwise missing from the model.

Columns 2-4 of Table B2 shows how both the coefficient estimates and the efficient initial and terminal subsidy with technological change and forward-looking households would change under alternative assumptions about the percent of households that consider installing solar (column 1 of Table B2 shows the baseline results from the main text). While there are statistically significant differences across specifications (particularly in the constant benefit of installing solar, which has to adjust in order for total adoption to remain constant with many more households considering solar), the overall takeaways remain unchanged. With forward-looking households and technological change, the efficient subsidy is always clearly increasing. The one difference is that the level of the subsidy is smaller when more households consider solar, because individual households’ decisions are more influenced by single-period stochastic draws, which makes households behave more myopically and therefore require a lower subsidy.

Understanding how the assumption of how many households consider adopting solar affects the results of our paper also provides insight into the role of the stochastic errors in the design of the efficient subsidy schedule (discussed in Section C.3). If we had assumed that every household considered installing solar, then there would be nearly 20x more households deciding whether to adopt solar in the initial period (and similar increases in other periods). Given that total adoption is unchanged, adoption now occurs only if a household gets a very uncommon draw of residuals in a period. Since there is no serial correlation in a household’s draws (other than whether the household considers solar), forward-looking households will not expect to receive another draw of this magnitude within a reasonable time-frame. This model without serial correlation in the residuals leads the forward-looking households who are responding to the subsidy to behave very similarly to myopic households and adopt as soon as they receive a set of draws that makes adoption attractive, even if technology or subsidies are set to improve in the future. The efficient subsidy for forward-looking households therefore converges towards the efficient subsidy for myopic households. In fact, if households were myopic, the efficient subsidy in column 4 of Table B2 would start at 0 and would increase to 0.6 \$/W, which is very similar to the efficient subsidy with foresight presented in the table. Given that we believe (both from introspection and estimation) that the percent of households who consider installing solar is low and that there is serial correlation in the stochastic errors, we base our simulations on the model where 6.6% of households consider adopting solar.

Table B2: Dynamic Demand Estimate Sensitivity

	Baseline	Percent Who Consider Solar			No
		5%	50%	100%	Third Party
Benefit of Solar:					
Constant	-10.6419*** (0.53)	-10.42*** (0.54)	-12.63*** (0.53)	-12.86*** (0.53)	-10.59*** (0.57)
Median Radiation (kWh/m ² /day/10)	4.80*** (0.12)	4.98*** (0.12)	4.29*** (0.10)	4.26*** (0.10)	3.67*** (0.12)
Electricity Price (\$ per kWh)	4.05*** (0.26)	4.12*** (0.26)	3.98*** (0.26)	3.81*** (0.26)	2.88*** (0.32)
log(Home Value (\$millions))	0.77*** (0.16)	0.78*** (0.16)	0.76*** (0.16)	0.75*** (0.16)	0.84*** (0.17)
Years of Schooling	0.10*** (0.03)	0.10*** (0.03)	0.09** (0.03)	0.07** (0.03)	0.11*** (0.04)
SCE	0.29*** (0.05)	0.29*** (0.05)	0.32*** (0.05)	0.29*** (0.05)	0.25*** (0.06)
SDG&E	1.29*** (0.06)	1.35*** (0.06)	1.16*** (0.05)	1.12*** (0.05)	1.03*** (0.06)
Time Trend	0.12*** (0.003)	0.12*** (0.003)	0.13*** (0.003)	0.13*** (0.003)	0.15*** (0.003)
Time Trend ²	-0.0021*** (1e-4)	-0.0021*** (1e-4)	-0.0021*** (1e-4)	-0.0021*** (1e-4)	-0.0031*** (1e-4)
Cost of Solar:					
Cost (\$10,000s)	-0.52** (0.24)	-0.54** (0.24)	-0.53** (0.25)	-0.62** (0.25)	-0.16 (0.26)
Cost (\$10,000s)*log(Home Value (\$millions))	-0.11** (0.07)	-0.10 (0.07)	-0.13* (0.07)	-0.13* (0.07)	-0.17** (0.08)
Cost (\$10,000s)*Years of Schooling	-0.03* (0.02)	-0.03* (0.02)	-0.03 (0.02)	-0.02 (0.02)	-0.04** (0.02)
-Log-likelihood	268,399	268,396	268,411	268,413	197,073
Months	41	41	41	41	41
Percent Who Consider Solar	6.6	5	50	100	6.6
Efficient Subsidy:					
(w/ foresight & tech change)					
Initial subsidy (\$/W)	0.6	0.7	0	0	1.1
Terminal subsidy (\$/W)	1.4	1.6	0.8	0.7	2.9

Standard errors in parentheses. Electricity prices are the average rate in dollars per kilowatt-hour for usage over 200% of baseline. In Southern California Edison (SCE) and Pacific Gas and Electric (PG&E, the omitted utility for the utility fixed effects) this is the average of one rate for 201-300% of baseline and a second rate for 300+% of baseline. In San Diego Gas and Electric (SDG&E) this is an average of the summer and winter rates for usage over 200% of baseline.

Removing Leased Systems

Column 5 of Table B2 shows how our results would change if we maintained the assumption that 6.6% of households consider installing solar but limit adoption to only household-owned systems. This means that all installations by “third parties,” where the household is technically leasing the system from a solar company that pays the upfront costs of the system, are now considered to be part of the “outside good”.⁴ Interestingly, the results are fairly similar under this specification, except for a change to the quadratic time-trend to reflect the increasing appeal of third-party installations over our estimation window. For household-owned systems, this increasing appeal of third-party installations looks like a decrease in the unobserved quality of solar relative to the outside good that includes leasing solar. But again, the efficient subsidy with consumer foresight and technological change is clearly increasing. The efficient subsidy is also now substantially higher than the base case given that the outside good of not purchasing a solar system outright is now more attractive to households.

C Numerical Calibration and Solution Method

We begin by describing the calibration. We then describe how we solve the stochastic and deterministic models.

C.1 Calibration

We here describe the calibration of the regulator’s benefit from cumulative adoption and the cost of public funds. Assume that the regulator’s benefit function is quadratic:

$$B(Q_t) = \gamma_1 Q_t + \gamma_2 Q_t^2.$$

We calibrate γ_1 as the marginal social benefit of solar from Baker et al. (2013). They simulate a 5 kW array, whereas we use a 5.4 kW array. They report a south-facing array in San Francisco as generating 7,220 kWh (AC) per year. Assume that this energy production is evenly distributed across months. They report the value of solar to the electric grid (“weighted average λ ”) as \$0.055/kWh. They report the emission displacement rate as 1.11 pounds of carbon dioxide (CO₂) per kWh. Using the U.S. government’s year 2015 social cost of carbon (with a 3% discount rate) of \$36/tCO₂, we have the emission benefit of a 5 kW array as \$0.0181/kWh. Over the course of a month, the combined grid and emission benefit of a 5.4 kW array is

$$(5.4/5)(7220/12)(0.0181 + 0.055)/10^4 = 0.0048$$

⁴System cost in this specification are also only for household-owned systems and are allowed to evolve differently from system costs in the main estimation: $C_{i(t+1)} = \tilde{\gamma}_0 + \tilde{\gamma}_1 C_{it} + \tilde{\omega}_{it}$. We estimate that $\tilde{\gamma}_0 = 0.0168$, $\tilde{\gamma}_1 = 0.9880$, and the standard deviation of $\tilde{\omega}_{it}$ is 0.1787 for household-owned systems.

in tens of thousands of dollars, which we use as the marginal social benefit of the first array to be adopted (i.e., the first 5.4 kW array to be adopted generates \$48 per month of social value). In the simulations, Q_t measures the fraction of the population that has adopted solar. The parameter γ_1 must adjust for the size N of the population:

$$\gamma_1 = 0.0048N.$$

Assume that the marginal social value of an installed solar array falls by $x\%$ by the time we reach $z\%$ adoption (e.g., because of concerns about intermittency). Then

$$\gamma_1 + 2\gamma_2 z/100 = (1 - x/100)\gamma_1,$$

which implies

$$\gamma_2 = -\frac{1}{2} \frac{x}{z} \gamma_1 \leq 0.$$

Gowrisankaran et al. (2016) estimate that the intermittency of solar electricity would impose costs of \$46/MWh if solar photovoltaics provided 20% of electricity in Arizona.⁵ For our 5.4 kW array, this works out to a cost of

$$(5.4/5) * (7220/12) * (46/1e3)/1e4 = 0.0030$$

in tens of thousands of dollars, which means that our marginal benefit of solar would decline by 62.5% (yielding $x = 62.5$) if solar photovoltaics provided 20% of electricity in California. Using California’s 2014 electricity consumption of 296,843 GWh,⁶ providing 20% of electricity means providing 59,369 GWh. To generate this much electricity, we require $59369e6/[(5.4/5) * 7220]$ arrays (or nearly 8 million arrays). We therefore have $z = 100 * 59369e6/[(5.4/5) * 7220]/N$, or an adoption rate of 185%.

Assume that the regulator’s cost of funds is quadratic:

$$G(z) = g_1 z + g_2 z^2,$$

where z is subsidy spending in tens of thousands of dollars. The argument of $G(\cdot)$ in the main text is $s_t[Q_{t+1} - Q_t]$, which is total subsidy spending. When Q is normalized to be the fraction of the population (as in the simulations), we have $z = s_t[Q_{t+1} - Q_t]N$. Barrage (2016) collects estimates of the marginal cost of public funds from the literature. Averaging the estimates for the U.S. yields \$1.35 for the marginal cost of the first dollar of spending. We therefore have:

$$g_1 = 1.35.$$

⁵This scenario assumes that conventional sources of electricity are reoptimized around the 20% solar penetration rate. This number does not account for how the marginal value of electricity may decline in solar penetration.

⁶http://energyalmanac.ca.gov/electricity/total_system_power.html

Let the marginal cost of funds double when spending reaches x dollars. So the cost of funds doubles when we have $z = x/10^4$. Thus, we have

$$g_1 + 2g_2 \frac{x}{10^4} = 2g_1,$$

which implies

$$g_2 = \frac{1}{2} \frac{g_1}{x} 10^4.$$

There is not much literature on the curvature of the cost of public funds. The traditional marginal cost of public funds is likely to be linear over the sums of interest, but our cost of funds is meant to be a broad measure whose curvature captures the regulator’s distaste for disbursing a lot of money all at one time. This distaste could be driven by political constraints that are beyond traditional economic estimates of the marginal cost of public funds. We assume that this distaste is such that the marginal cost of public funds would double if the regulator were to allocate its entire actual cumulative spending to a single instant. This assumption yields x equal to \$148 million.

The regulator’s horizon is 41 months. As in the empirical model, we assume that households use a monthly discount rate of 1% and we measure time in months. We assume that the regulator uses the same discount rate as its households. We begin with 0.071% of households having adopted solar ($Q_0 = 0.00071$).

C.2 Solving the stochastic setting

We now describe how we solve the setting in which each household can receive a new draw of each ε_i in each period and in which installation costs may evolve stochastically. For any candidate subsidy trajectory, we simulate the empirical model over 100 draws of the stochastic component of the system cost evolution in order to obtain adoption in each period.⁷ This simulation gives expected adoption at the end of the horizon and gives the regulator’s expected value from committing to the candidate subsidy trajectory. We use the Knitro solver in Matlab to search for the 41-element trajectory of per-month subsidies that maximizes the regulator’s expected value while matching expected adoption at the end of the horizon to \hat{Q} .

For the simulations with peer effects and induced technological change (ITC), the continuation value depends upon households’ expectation of total adoption across all households in the period. In order to solve for the appropriate continuation values, we calculate the current period adoption given a discretized set of potential cumulative adoption levels next period. These represent households’ potential assumptions as to the total adoption this period and therefore the cumulative adoption next period, which determine next period’s preferences and solar costs via peer effects and ITC, respectively. We then solve for the fixed point in households’ assumptions about next period’s starting cumulative adoption that generates

⁷Experiments with more draws did not yield substantially different results.

the associated amount of total current period adoption. The intuitive argument for the existence of a fixed point relies on the monotonicity of current period adoption in next period’s cumulative adoption: if a household believes that many other households will adopt today, then, all else equal, that household will prefer to delay adoption in order to face a lower system cost (or higher preferences) tomorrow. Similarly, if a household believes that few households will adopt today, then the benefit of delaying adoption until tomorrow decreases and the household is more likely to adopt today. If we cannot find a fixed point adoption level, then we assume that the continuation value is the mean of the continuation value at the two adjacent adoption bins for the same system price.

Given the well-known difficulties with estimating ITC and peer effects, we instead make extreme assumptions about ITC and peer effects in order to bound their effects on the efficient subsidy. For ITC, we assume that all price declines observed over the CSI are due to induced technological change rather than exogenous factors such as technological improvements or input costs. For peer effects, we likewise assume that the entire time trend in household preferences for solar is coming from peer effects rather than from other factors such as improvements in system quality. These assumptions are much more extreme than we believe is realistic, but we feel that they create an upper bound on the potential effect of endogenous technology and preference changes.

In the cases with induced technical change, we calibrate the pace of technical change to the efficient adoption pathway in the presence of exogenous technical change.⁸ This calibration ensures that the efficient subsidies in the case with induced technical change differ due to the endogeneity of technology, not due to differences in eventual costs. Regressing log costs on cumulative adoption yields the following cost function:

$$\xi(t, Q(t)) = e^{\hat{\theta}_0 + \hat{\theta}_1 Q(t) + \hat{\theta}_2 Q(t)^2 + \hat{\theta}_3 Q(t)^3},$$

with $\hat{\theta}_0 = 1.55$, $\hat{\theta}_1 = -142$, $\hat{\theta}_2 = 25717$, and $\hat{\theta}_3 = -1894075$. The fit to the exogenous cost trajectory (along that scenario’s efficient adoption trajectory) is quite good, with an R^2 of 0.996.

In simulations with peer effects, we use a simple univariate OLS regression of the predicted time trend as estimated from our model on cumulative adoption up to that point in the CSI. We use the estimated coefficient from this regression and the cumulative adoption up to that point in place of the time trend in our calculation of households’ likelihood of installing solar in a given period. It is important to reiterate that this means that we are capturing peer effects at the state level rather than at the local level. One additional adoption changes all consumers’ preferences for solar in the same way, regardless of whether the adoption is across the street or across the state. This is consistent with a regulator who is unable to differentiate subsidies across households within a period and does not focus on targeting adoption by households who have large local effects on their neighbors.

⁸In ongoing work, we calibrate ITC to the realized adoption pathway rather than to the efficient adoption pathway.

Consumer Surplus Calculation

We estimate our model assuming that households are forward-looking, based on the evidence provided in Appendix B.1. In order to calculate the expected consumer surplus of having the choice to install solar for forward-looking households under the counterfactual of no subsidy or the efficient subsidy, we use the standard expected consumer surplus equation:

$$E[CS_{fore sight}] = \sum_{i=1}^{N_1} \zeta_i \log \left[\exp(X'_{i1} \gamma + \alpha_i \tilde{C}_1) + \exp(\beta E[V(\Omega_2|\Omega_1)]) \right] + S,$$

where \tilde{C}_1 is the average net-of-subsidy system cost in the initial period and S is a constant that will be differenced away when we calculate the change in consumer surplus across subsidy paths. We calculate the consumer surplus for myopic households by calculating expected utility in each period and discounting it back to the initial period:

$$E[CS_{myopic}] = \sum_{i=1}^{N_1} \zeta_i \sum_{t=1}^T \beta^{t-1} \left[Pr_{it}(X'_{it} \gamma + \alpha_i \tilde{C}_t - \log(Pr_{it})) + (1 - Pr_{it})(-\log(1 - Pr_{it})) \right] + S,$$

where Pr_{it} is the probability that household i adopts solar in period t . This formulation subtracts γ , Euler’s constant, from the standard formulation of a household’s expected utility of choosing to adopt and from the household’s utility of choosing to wait. Because households do not make a choice after they adopt, including the expectation of the EV1 draw would lead to higher expected consumer surplus estimates in situations where adoption by myopic households happens later in the policy window. As a result, expected consumer surplus is not directly comparable between myopic and forward-looking households, but the differences between the myopic households’ expected consumer surpluses under different subsidy policies should be comparable to the differences for forward-looking households.

C.3 Solving the deterministic setting

We now describe how we use the theoretical analysis to solve the setting in which each household receives only a single draw of each $\tilde{\varepsilon}_i$ for all time and in which the evolution of installation costs is deterministic.

To start, consider how the structural empirical model provides the desired distribution over private values v_i from the theoretical setting. The household’s private value from adopting solar is equal to the price at which the household would be indifferent between adopting solar and not adopting solar if adoption were a now-or-never decision. From equation (1) and the definition of h_{it} , household i is indifferent to adopting solar at time t when

$$X'_{it} \gamma + \alpha_i \tilde{C}_{it} + \varepsilon_{i1t} = \varepsilon_{i0t} + \beta E[V(\Omega_{t+1}|\Omega_t)],$$

which we can write as

$$b_i + \alpha_i \tilde{C}_{it} + \varepsilon_{i1t} - \varepsilon_{i0t} = \beta \mathbb{E}[V(\Omega'|\Omega)], \quad (\text{C-5})$$

with

$$b_i = X'_{i0} \gamma.$$

For these simulations, we fix demographics—in particular, electricity prices—at their initial values. When adoption is a now-or-never decision, the expectation on the right-hand side of equation (C-5) is zero. Household i 's private value for solar is then

$$v_i = -\frac{1}{\alpha_i} (b_i + \varepsilon_{i1t} - \varepsilon_{i0t}),$$

where each $\varepsilon_{i* t}$ is a draw from a type I generalized extreme value distribution with location parameter 0 and scale parameter 1 and where each estimated α_i is negative. Rewrite as

$$\varepsilon_{i1t} - \varepsilon_{i0t} = -\alpha_i v_i - b_i.$$

The difference of two type I generalized extreme value random variables is itself a random variable following a logistic distribution with location parameter 0 and scale parameter 1. The cumulative distribution function $P(\cdot)$ of v_i is:

$$P(v_i) = \frac{1}{1 + e^{\alpha_i v_i + b_i}}.$$

The number of households with the same demographic characteristics as household i is N_i . Aggregate across demographic groups to obtain the cumulative distribution function for private values v across all demographic groups:

$$F(v) = \frac{\sum_i N_i \frac{1}{1 + e^{\alpha_i v + b_i}}}{\sum_i N_i}.$$

Differentiating yields the density function of private values:

$$f(v) = \frac{-\sum_i N_i \frac{\alpha_i e^{\alpha_i v + b_i}}{(1 + e^{\alpha_i v + b_i})^2}}{\sum_i N_i}.$$

Differentiating a second time yields:

$$f'(v) = \frac{-\sum_i N_i \frac{\alpha_i^2 e^{\alpha_i v + b_i} (1 - e^{\alpha_i v + b_i})}{(1 + e^{\alpha_i v + b_i})^3}}{\sum_i N_i}.$$

Now consider the evolution of the private cost of installing solar. The empirical estimates account for changes in both the monetary cost of installing solar panels and the preference

for solar, which we interpret as changes in the quality of solar panels. We take $C(t, Q(t)) = \xi(t, Q(t)) - \chi(t)$ to measure the quality-adjusted cost, with $\xi(t, Q(t))$ the direct monetary cost and $\chi(t)$ a discount to reflect quality improvements since time 0. As described above for b_i , the trend in quality will be valued by demographic groups differentially via α_i . We abstract from the time trend’s demographic dependence by using a population-weighted average of α_i when calibrating $\chi(t)$ to the empirical estimates. The empirical estimates yield $\chi(t) = \chi_0 t + \chi_1 t^2$, with $\chi_0 > 0$ and $\chi_1 < 0$.

We regress the cost of installation (beginning with the elimination of the federal subsidy cap) against a constant and its lagged value (as explained in Appendix B.2):

$$\xi_{t+1} = \theta_0 + \theta_1 \xi_t,$$

where t is measured with 1 as the first month of the estimation window and costs are measured in tens of thousands of dollars per 5.4 kW system. The initial cost of installing is assumed to be $\xi(0) = 4.4$, with the initial period serving as the reference for quality ($\chi(0) = 0$). Subtracting ξ_t from each side and passing to the continuum, we have

$$\dot{\xi}(t) = \theta_0 + [\theta_1 - 1]\xi(t),$$

where we ignore the potential dependence of ξ on $Q(t)$ because we are here focusing on the case with exogenous technical change. Solve the differential equation:

$$\begin{aligned} \dot{\xi}(t) + [1 - \theta_1]\xi(t) &= \theta_0 \\ \Leftrightarrow e^{[1-\theta_1]t} \left\{ \dot{\xi}(t) + [1 - \theta_1]\xi(t) \right\} &= e^{[1-\theta_1]t} \theta_0 \\ \Leftrightarrow \int_0^t e^{[1-\theta_1]s} \left\{ \dot{\xi}(s) + [1 - \theta_1]\xi(s) \right\} ds &= \int_0^t e^{[1-\theta_1]s} \theta_0 ds \\ \Leftrightarrow \xi(t) &= \frac{\theta_0}{1 - \theta_1} [1 - e^{-[1-\theta_1]t}] + e^{-[1-\theta_1]t} \xi(0). \end{aligned} \tag{C-6}$$

Substituting into $\dot{\xi}(t)$, we have

$$\dot{\xi}(t) = \theta_0 e^{-[1-\theta_1]t} - [1 - \theta_1] e^{-[1-\theta_1]t} \xi(0).$$

Differentiating with respect to time, we have:

$$\ddot{\xi}(t) = -\theta_0(1 - \theta_1)e^{-[1-\theta_1]t} + [1 - \theta_1]^2 e^{-[1-\theta_1]t} \xi(0).$$

We also must solve for $J(Y(T), C(T, Q(T)))$ in the setting with forward-looking households. Recall that $J(v_i, C(T, \hat{Q}))$ is the present value to household i of having the option to adopt the technology at time T , once the subsidy disappears for good. At time T , a household that has yet to adopt the technology solves:

$$J(v_i, C(T, Q(T))) = \max_{\Psi_i} e^{-\delta(\Psi_i - T)} [v_i - C(\Psi_i, Q(\Psi_i))],$$

for $\Psi_i > T$. In our application, it is always true that $v_i \leq 0$ for households who have yet to adopt solar at time T . Such households will never adopt the technology at a later time Ψ_i , once the subsidy is gone. We can therefore fix $J(v_i, C(T, Q(T))) = 0$, which yields the terminal subsidy $s(T)$ as described in the main text.

To solve the setting with myopic households, note that $\mu(t) = 0$, $s(t) = C(t, 1 - F(Y(t))) - Y(t)$, and equation (A-1) implicitly defines $y(t)$ as a function of $\lambda(t)$ and $Y(t)$. We then have two differential equations ($\dot{Y}(t) = y(t)$ and the costate equation for $Y(t)$) in two variables. We know $Y(0)$ and $Y(T)$. For any guess for $\lambda(T)$, we solve the system backwards from time T to time 0 and compare the obtained $Y(0)$ to the desired $Y(0)$. We use the Matlab `fzero` root-finding function to search for the $\lambda(T)$ that yields trajectories that satisfy the initial conditions. We use Matlab’s `ode15s` solver with an analytic Jacobian to solve the system of differential equations for any guess of $\lambda(T)$.

To solve the setting with forward-looking households, note that equation (A-1) implicitly defines $y(t)$ as a function of $\lambda(t)$, $s(t)$, $\mu(t)$, and $Y(t)$. We then have four differential equations (the transition and costate equations) in the same four variables. We know $Y(0)$, $\mu(0)$, $Y(T)$, and $s(T)$. For any guess for $\lambda(T)$ and $\mu(T)$, we solve the system backwards from time T to time 0 and compare the obtained $Y(0)$ and $\mu(0)$ to the desired $Y(0)$ and $\mu(0)$. We use the Knitro solver in Matlab to search for the terminal conditions that yield trajectories that satisfy the initial conditions. We use Matlab’s `ode15s` solver with an analytic Jacobian to solve the system of differential equations for any guess of terminal conditions.

Figure C2 plots the resulting efficient subsidy and cumulative adoption. The subsidy trajectories are qualitatively similar to those in the full model analyzed in the main text, except for two differences that are related to the stochastic preference draws contained in the full model but not in the present model. First, the level of the subsidy is here greater because the regulator cannot expect some households to receive favorable draws in later periods (which makes the adoption target more difficult to achieve). Second, the efficient subsidy is here decreasing over an initial interval (following a delayed start) when households are forward-looking and technology is changing, whereas in the main text it increased throughout the 41-month horizon. In the full model, households that receive a favorable draw act somewhat myopically because they expect to have a lower value for solar in the future, which means that they may adopt the technology even without being promised a lower future subsidy. In contrast, here preferences are fixed over time, so many households will adopt the technology only if the subsidy declines sufficiently rapidly to offset the lure of waiting for technological change.

D Results for a Target of 1.5% Adopting Solar

We here replicate the main text’s numerical results for a higher target of 1.5% of the population adopting solar. Figure D3 shows that the efficient subsidy still increases over time.

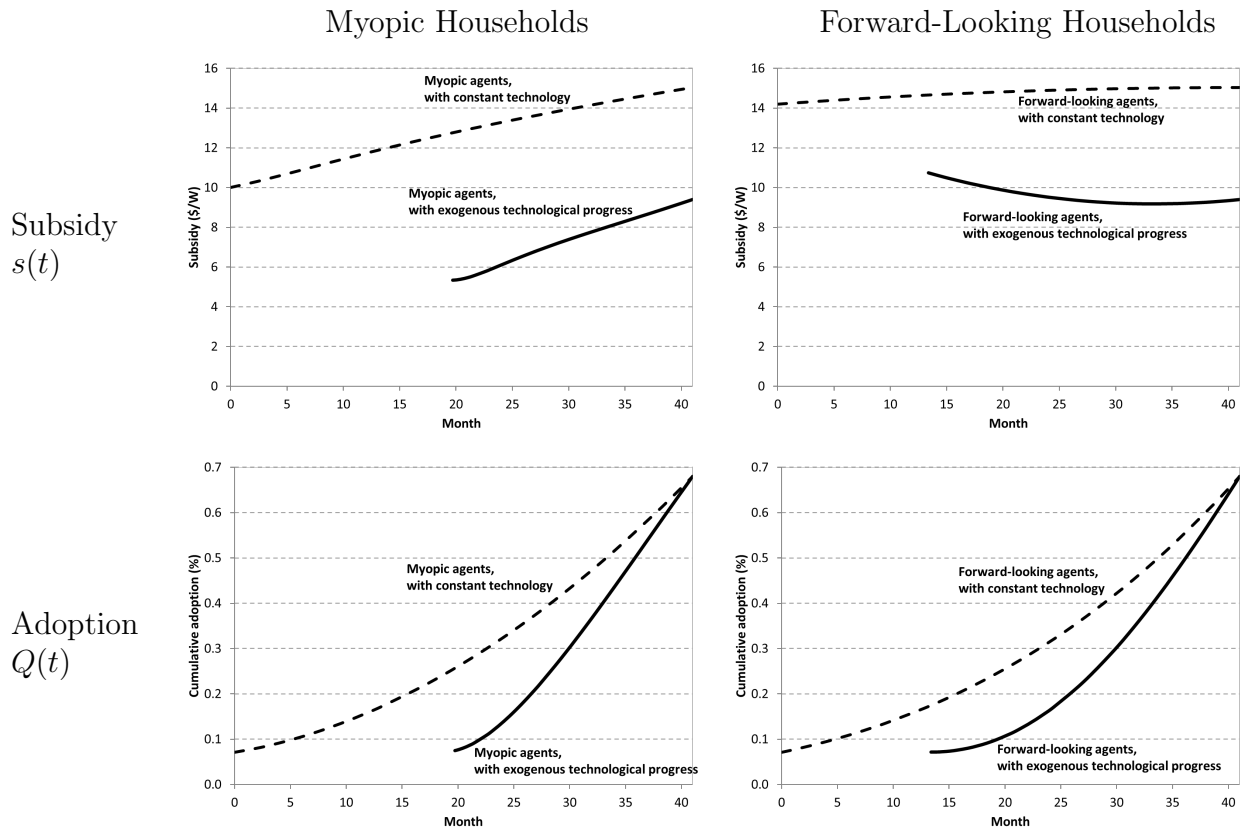


Figure C2: The efficient subsidy and corresponding adoption trajectories in the model from the theory section.

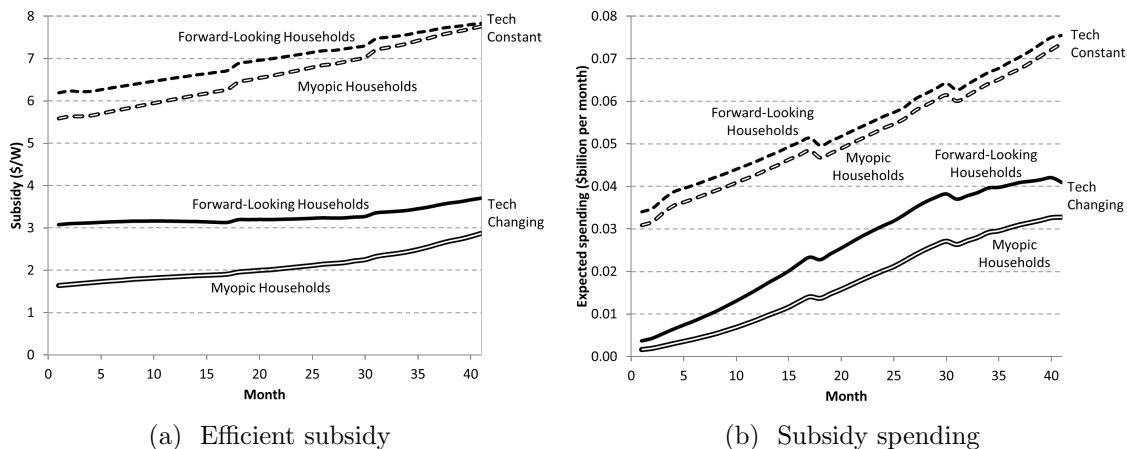


Figure D3: The efficient subsidy (left) and expected monthly spending (right) when households are forward-looking (solid) and myopic (hollow) and the target is 1.5% adoption. Connected lines allow for exogenous technical change and dashed lines hold technology constant over time.

The higher target requires a larger subsidy, but the dynamics of each subsidy trajectory are similar to the results in the main text. Table D3 also shows the same patterns analyzed in the main text. Figure D4 is also qualitatively similar to the main text, except that now the case with myopic households and no technological progress demonstrates a much more strongly increasing subsidy in the first instants, driven by a strong price discrimination channel. The price discrimination channel is especially important in early instants because the marginal adopter is far in the tail of the distribution of households and the larger subsidy required for the 1.5% target then generates many inframarginal adopters. The smooth spending channel now also favors a more strongly increasing subsidy in the first instants when households are myopic and there is no technical change. Spending tends to decline over the first instants as the weakening of the price discrimination channel slows the subsidy’s initially rapid rate of increase. As a result, the measure of values that find adoption newly optimal falls over these instants ($\dot{y}(t) > 0$). The regulator smooths spending by not flattening the subsidy trajectory too quickly.

E Theoretically Analyzing a Fixed Budget

We now consider a setting in which the regulator has a fixed budget but is free to choose cumulative adoption. Let $Z(t)$ denote cumulative spending, with $\hat{Z} > 0$ the fixed budget. Assume that the budget is small enough that the regulator wants to exhaust it. The regulator

Table D3: The present value of spending, initial and terminal subsidies, and consumer surplus along the efficient subsidy trajectory for a target of 1.5%, along with the consumer surplus obtained in the absence of subsidies.

	Myopic		Forward-Looking	
	Technology Constant	Technology Improving	Technology Constant	Technology Improving
Present value of spending (\$million)	1680	540	1776	809
Present value of spending (\$/W)	5.3	1.7	5.6	2.6
Initial subsidy (\$/W)	5.6	1.6	6.2	3.1
Terminal subsidy (\$/W)	7.8	2.9	7.8	3.7
Consumer surplus without any subsidy (\$million)	26	183	84	918
Consumer surplus with efficient subsidy (\$million)	474	454	695	1367

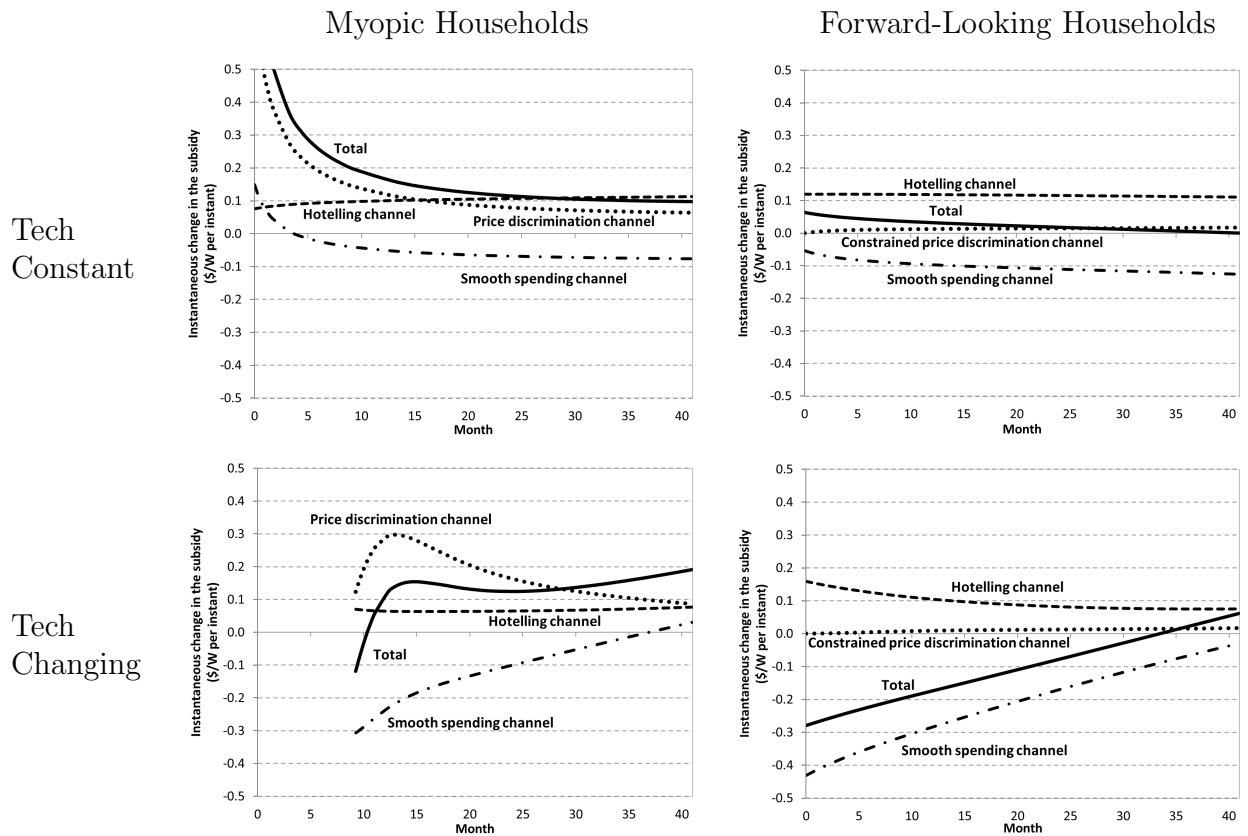


Figure D4: The change in the efficient subsidy at each instant ($\dot{s}(t)$, labeled “Total”) for a target of 1.5%, as well as each component identified in the theoretical analysis. The adoption benefit component (not plotted) is negative but very small in magnitude.

solves:

$$\begin{aligned}
& \max_{y(t), s_0, \hat{Q}} \int_0^T e^{-rt} \left[B \left(1 - F(Y(t)) \right) - G \left(-s(t) f(Y(t)) y(t) \right) \right] dt \\
& \text{s.t. } \dot{Y}(t) = y(t) \\
& \quad \dot{s}(t) = \delta [Y(t) - C(t, 1 - F(Y(t))) + s(t)] + \dot{C}(t, 1 - F(Y(t))) \\
& \quad \dot{Z}(t) = -s(t) f(Y(t)) y(t) \\
& \quad Y(0) = F^{-1}(1 - Q_0), \quad Y(T) = F^{-1}(1 - \hat{Q}) \\
& \quad s(0) = s_0, \quad s(T) = C(T, 1 - F(Y(T))) - Y(T) + J(Y(T), C(T, 1 - F(Y(T)))) \\
& \quad Z(0) = 0, \quad Z(T) = \hat{Z}.
\end{aligned}$$

The Hamiltonian is

$$\begin{aligned}
& H(t, y(t), Y(t), s(t), Z(t), \lambda_Y(t), \mu(t), \lambda_Z(t)) \\
& = e^{-rt} \left[B \left(1 - F(Y(t)) \right) - G \left(-s(t) f(Y(t)) y(t) \right) \right] \\
& \quad + e^{-rt} \lambda_Y(t) y(t) + e^{-rt} \mu(t) \delta [Y(t) - C(t, 1 - F(Y(t))) + s(t)] \\
& \quad + e^{-rt} \mu(t) \left[C_1(t, 1 - F(Y(t))) - C_2(t, 1 - F(Y(t))) f(Y(t)) y(t) \right] \\
& \quad - e^{-rt} \lambda_Z(t) s(t) f(Y(t)) y(t).
\end{aligned}$$

$\lambda_Z(t)$ gives the (current) shadow value of $Z(t)$.

The necessary conditions for a maximum are:

$$\begin{aligned}
\lambda_Y(t) &= -s(t) f(Y(t)) \left[-\lambda_Z(t) + G' \left(-s(t) f(Y(t)) y(t) \right) \right] \\
&\quad + \mu(t) C_2(t, 1 - F(Y(t))) f(Y(t)), \tag{E-7} \\
-\dot{\lambda}_Y(t) + r\lambda_Y(t) &= -f(Y(t)) B' \left(1 - F(Y(t)) \right) \\
&\quad + s(t) f'(Y(t)) y(t) \left[-\lambda_Z(t) + G' \left(-s(t) f(Y(t)) y(t) \right) \right] + \delta\mu(t) \\
&\quad + \delta\mu(t) f(Y(t)) C_2(t, 1 - F(Y(t))) \\
&\quad + \mu(t) \left[-C_{12}(t, 1 - F(Y(t))) f(Y(t)) + C_{22}(t, 1 - F(Y(t))) [f(Y(t))]^2 y(t) \right. \\
&\quad \quad \left. - C_2(t, 1 - F(Y(t))) f'(Y(t)) y(t) \right], \\
-\dot{\mu}(t) + r\mu(t) &= f(Y(t)) y(t) \left[-\lambda_Z(t) + G' \left(-s(t) f(Y(t)) y(t) \right) \right] + \delta\mu(t), \\
-\dot{\lambda}_Z(t) + r\lambda_Z(t) &= 0, \\
\lambda_Y(T) &= 0, \\
\mu(0) &= 0,
\end{aligned}$$

along with the transition equations and the initial and terminal conditions on $Z(\cdot)$. The first equation follows from the Maximum Principle, the next two equations are the costate (or adjoint) equations, the second-to-last equation is the transversality condition corresponding to the choice of \hat{Q} , and the final equation is the transversality condition corresponding to the choice of s_0 .

Differentiate equation (E-7) with respect to time and suppress the argument of $G(\cdot)$:

$$\begin{aligned}
\dot{\lambda}_Y(t) &= [-\dot{s}(t) f(Y(t)) - s(t) f'(Y(t)) y(t)] [-\lambda_Z(t) + G'] \\
&\quad + s(t) f(Y(t)) \dot{\lambda}_Z(t) \\
&\quad - s(t) f(Y(t)) G'' \left[-\dot{s}(t) f(Y(t)) y(t) - s(t) f'(Y(t)) [y(t)]^2 - s(t) f(Y(t)) \dot{y}(t) \right] \\
&\quad + \dot{\mu}(t) C_2(t, 1 - F(Y(t))) f(Y(t)) \\
&\quad + \mu(t) C_{12}(t, 1 - F(Y(t))) f(Y(t)) \\
&\quad - \mu(t) C_{22}(t, 1 - F(Y(t))) [f(Y(t))]^2 y(t) \\
&\quad + \mu(t) C_2(t, 1 - F(Y(t))) f'(Y(t)) y(t).
\end{aligned}$$

Substitute for $\dot{\lambda}_Y(t)$, $\dot{\lambda}_Z(t)$, and $\delta\mu(t)$ from the costate equations and rearrange to obtain:

$$\dot{s}(t) = \left[-r\lambda_Y(t) + \overbrace{r\lambda_Z(t)s(t)f(Y(t))}^{\text{new channel}} - B'f(Y(t)) - [G' - \lambda_Z(t)]y(t)f(Y(t)) - \dot{\mu}(t) + r\mu(t) \right. \\ \left. - [s(t)]^2 G''f(Y(t))\ddot{Q}(t) + (r\mu(t) - f(Y(t))y(t)[G' - \lambda_Z(t)])C_2(t, Q(t))f(Y(t)) \right] \\ \left[[G' - \lambda_Z(t)]f(Y(t)) - s(t)f(Y(t))G''y(t)f(Y(t)) \right]^{-1},$$

where we suppress the argument of $B(\cdot)$. We see the same effects as in the main text's setting with a fixed adoption target \hat{Q} , plus an additional effect. $\lambda_Z(t) < 0$ is the shadow value of additional spending, so that $-\lambda_Z(t)s(t)$ is the shadow cost of time t spending. This shadow cost is driven by the scarcity of funds. We see the price discrimination channel and the endogenous cost channel become amplified by this additional cost of funds. We also see a new channel that works to make the efficient subsidy decline over time. The regulator has a fixed budget, and all else equal, an impatient regulator chooses to consume more of this budget earlier.

Combine this new channel with the first (Hotelling) channel and substitute for $\lambda_Y(t)$ from equation (E-7) to obtain:

$$\dot{s}(t) = \left[r s(t) G' f(Y(t)) - r\mu(t)C_2(t, Q(t))f(Y(t)) - B'f(Y(t)) \right. \\ \left. - [G' - \lambda_Z(t)]y(t)f(Y(t)) - \dot{\mu}(t) + r\mu(t) \right. \\ \left. - [s(t)]^2 G''f(Y(t))\ddot{Q}(t) + (r\mu(t) - f(Y(t))y(t)[G' - \lambda_Z(t)])C_2(t, Q(t))f(Y(t)) \right] \\ \left[[G' - \lambda_Z(t)]f(Y(t)) - s(t)f(Y(t))G''y(t)f(Y(t)) \right]^{-1}.$$

We have replaced the Hotelling channel and the new channel with $r s(t) G' f(Y(t))$. This remaining term is the exact same term as the Hotelling channel in the main text, once we substitute for $\lambda(t)$ in equation (4) from equation (A-1). Thus, the main text's analysis of $\dot{s}(t)$ also applies to a setting with a fixed budget instead of a fixed adoption target, only replacing G' with $[G' - \lambda_Z(t)]$.

References from the Appendix

Baker, Erin, Meredith Fowlie, Derek Lemoine, and Stanley S. Reynolds (2013) “The economics of solar electricity,” *Annual Review of Resource Economics*, Vol. 5, No. 1, pp. 387–426.

- Barrage, Lint (2016) “Optimal dynamic carbon taxes in a climate-economy model with distortionary fiscal policy,” working paper.
- Burr, Chrystie (2014) “Subsidies, tariffs and investments in the solar power market,” working paper.
- Gowrisankaran, Gautam, Stanley S. Reynolds, and Mario Samano (2016) “Intermittency and the value of renewable energy,” *Journal of Political Economy*, Vol. 124, No. 4, pp. 1187–1234.
- Hughes, Jonathan E. and Molly Podolefsky (2015) “Getting green with solar subsidies: Evidence from the California Solar Initiative,” *Journal of the Association of Environmental and Resource Economists*, Vol. 2, No. 2, pp. 235–275.
- Mian, Atif and Amir Sufi (2012) “The effects of fiscal stimulus: Evidence from the 2009 Cash for Clunkers Program,” *Quarterly Journal of Economics*, Vol. 127, No. 3, pp. 1107–1142.
- Podolefsky, Molly (2013) “Tax evasion and subsidy pass-through under the solar investment tax credit,” job market paper, University of Colorado Boulder.