Dynamic Technology Subsidies*

Ashley Langer and Derek Lemoine

Department of Economics, University of Arizona
McClelland Hall 401, 1130 E Helen St, Tucson, AZ, 85721-0108, USA
alanger@email.arizona.edu
dlemoine@email.arizona.edu

January 2, 2016

Preliminary:
Not for distribution

We analyze how a policymaker should structure subsidies over time so as to induce agents to adopt a new durable good technology. We show that the gains from intertemporal price discrimination favor a subsidy that increases over time when agents are myopic, but the efficient subsidy tends to decline over time when agents anticipate future costs and policies. We apply our theoretical results to understanding the determinants of the efficient subsidy for rooftop solar photovoltaics. We empirically estimate the distribution of willingness to pay for solar among California households and simulate the key drivers of the efficient subsidy.

JEL: H21, H23, Q42, Q48, Q52

Keywords: subsidy, dynamics, technology, durable goods, solar, photovoltaics

*We are grateful to the University of Arizona Renewable Energy Network for supporting this work. Kyle Wilson and Max Rosenthal provided valuable research assistance.

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. And many U.S. states pay homeowners to install solar panels. These adoption subsidies achieve different social objectives, but all of these policymakers presumably prefer to obtain more adoption by spending less money. Yet the design of an efficient subsidy schedule has thus far remained an open policy question.

We formally analyze the dynamically efficient subsidy schedule. In each period, a regulator receives benefits from cumulative and contemporary adoption and dislikes spending money. In the initial period, the regulator commits to a subsidy schedule. Potential adopters have their own private values for the technology. When potential adopters are myopic, they choose whether to adopt the technology without considering how the subsidy and the cost of the technology will change in the future. When potential adopters are forward-looking, they time their adoption so as to maximize net private benefits. We show that the regulator's efficient subsidy trajectory is sensitive to whether potential adopters are forward-looking and to the distribution of their private values.

When potential adopters are myopic, whether the efficient subsidy increases or decreases over time depends primarily on a few forces. First, an intertemporal arbitrage condition must hold so that the regulator is indifferent to small deviations from the announced subsidy schedule. If a regulator is to be indifferent between raising the subsidy today or tomorrow, then the shadow benefit of the subsidy should grow as the regulator's discount rate. This effect is familiar from the Hotelling (1931) analysis of exhaustible resources. This increasing shadow benefit favors a subsidy that increases over time. The more impatient the regulator is, the more inclined she is to defer subsidy spending to later periods. However, early adoption also creates direct benefits in later instants. These direct benefits reduce the optimal growth rate of the subsidy's shadow benefit, similar to the Heal (1976) analysis of stock-dependent extraction costs in the case of nonrenewable resources. A greater marginal benefit of cumulative adoption favors a subsidy that declines over time, because the regulator finds earlier adoption more valuable.

Second, the regulator could more efficiently use its funds if it could price discriminate, but it does not know an individual adopter's value for the new technology. However, the regulator does know how the distribution of adopter private values changes over time as those with greater willingness to pay adopt the technology. The regulator can therefore structure the subsidy schedule so as to intertemporally price discriminate: by offering a low subsidy at first, the regulator allows actors with high private values to adopt the technology without receiving a big subsidy, and by raising the subsidy over time, the regulator induces actors with lower private values to also adopt the technology. This intertemporal price discrimination motive therefore favors an increasing subsidy schedule. It is strong when a concentrated distribution of willingness to pay or rapid improvements in technology create a lot of inframarginal adopters at a given subsidy.

When potential adopters are forward-looking, the regulator loses its complete freedom to time the subsidy so as to intertemporally price discriminate. Each individual actor chooses its time of adoption so as to maximize its own net benefits. The price discrimination channel disappears from the efficient subsidy trajectory. Instead, the regulator makes promises about later subsidies so as to induce earlier adoption. In particular, the regulator promises to keep the subsidy low in the near future, but when potential adopters are not perfectly patient, the regulator does not need to promise such low subsidies in the more distant future. The intertemporal price discrimination channel favored an increasing subsidy schedule at times with a lot of adoption, but this promise-keeping channel favors an increasing subsidy schedule when a lot of adoption happened in the past. By replacing the price discrimination channel with the promise-keeping channel, recognizing adopters' foresight favors a subsidy schedule that initially declines but eventually flattens out.

In order to understand how these considerations affect the effectiveness of subsidies in a real-world setting, we estimate the distribution of households' private willingness-to-pay for residential solar photovoltaic systems in California. Using data from California's largest residential solar subsidy program, the California Solar Initiative (CSI), we empirically show that households do appear to take changing future costs into consideration when choosing whether to install solar systems and therefore do appear to be forward-looking in their investment decision. We then use a dynamic discrete choice model to estimate the distribution of households' willingness-to-pay conditional on household demographics. The estimation assumes that households know the full time-path of subsidies but allows solar system prices to evolve stochastically conditional on the subsidies. The model suggests that there is substantial heterogeneity in the willingness-to-pay for residential solar systems.

We will combine this estimated distribution of willingness-to-pay with the theoretical results to determine the primary drivers of the efficient subsidy for rooftop solar. We will decompose the efficient subsidy into the analytic channels and will demonstrate how household foresight changes the dominant channels. We will also assess these channels' sensitivity to parameters such as the social discount rate, households' discount rates, the marginal cost of public funds, and the marginal social benefit of rooftop solar. Finally, we will compare installations and government spending under the actual subsidy program with the theoretically-derived efficient subsidy timepaths.

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. 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 revenue loss 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.

A larger literature has analyzed how monopolists should set prices over time. In particular, several 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 subsidy over time as long as all customers use the same discount rate 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). Our setting follows these in assuming that the regulator can commit to a subsidy schedule and in analyzing a case in which potential adopters have rational expectations. We avoid a corner solution (i.e., a constant subsidy) because we assume that the regulator has a concave benefit function. Further, we reserve the label of "price discrimination" for forces that arise only because of adopters' equilibrium decisions, so that we disentangle standard dynamic forces from intertemporal price discrimination motives.[1]

The next section contains the theoretical analysis. Section 2 introduces the data to be used in the empirical estimation. Section 3 reports evidence that California households anticipate future costs and subsidies. Sections 4 and 5 describe the dynamic structural model for estimating the distribution of household willingness to pay for solar photovoltaics and report the estimation results. 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. The final section concludes.

1 Theoretical Analysis

We begin by theoretically analyzing the subsidy trajectory that efficiently incentivizes actors to adopt a new technology. The next subsection describes the setting. Subsequent subsections analyze the cases in which households are myopic and are forward-looking.

[1]Many 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; Gul et al., 1986; Kahn, 1986; Besanko and Winston, 1990; Sobel, 1991). 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. The regulator indeed followed through on its commitments in our empirical application.

1.1 Setting

Actor i values a new technology at v_i and decides on a single time at which to adopt the technology. We normalize the measure of potential adopters to 1. The twice-differentiable cumulative distribution function $F(v_i) \in [0, 1]$ gives the number of actors who are willing to pay no more than v_i for the technology. We define $f(v_i) \geq 0$ as the density function (i.e., as $F'(v_i)$). We will often be interested in the case where $f'(v_i) < 0$, so that relatively few potential adopters are willing to pay a lot for the new technology.

The private cost of the new technology is $C(t)$. The technology's cost declines exogenously: $\dot{C}(t) \leq 0$, where a dot indicates a derivative with respect to time.[2] An actor who adopts the technology at time t receives a subsidy of $s(t)$. Actor i's net benefit of adopting is $v_i - C(t) + s(t)$. Let $Q(t)$ give the number of actors who adopted the technology prior to time t, so that $\dot{Q}(t)$ is adoption at time t. When potential adopters are myopic, they do not anticipate future changes in the technology's cost or in the subsidy. In that case, they compare the net benefit of adopting at time t to zero. When potential adopters are forward-looking, they anticipate all future changes in costs and in the subsidy, so that they compare the net benefit of adopting at time t to the net benefit of adopting at a later time, discounted at rate $\delta > 0$.[3]

The regulator commits at time 0 to offer a subsidy $s(t)$ to all $\dot{Q}(t)$ actors who adopt the technology at time t. The regulator aims to achieve total adoption Q_T by some given time $T > 0$. She knows the distribution of potential adopters' values but does not know any particular actor's value. The regulator dislikes spending money, with marginal cost of public funds $\gamma > 0$, and receives instantaneous benefit $B(Q(t), \dot{Q}(t))$ from adoption. This function is strictly increasing in each argument and strictly concave. We use subscripts to represent partial derivatives with respect to each argument. In our application to adoption of solar photovoltaics, the first argument of the benefit function captures the regulator's value from production of solar electricity (i.e., from cumulative adoption) and the second argument captures the regulator's value for installing new panels (i.e., from contemporary adoption), whether due to "green jobs" motivations or to voter perceptions.

When selecting the subsidy trajectory, the regulator correctly anticipates how the tech-

[2] We assume that the regulator is small relative to the world market, as in our empirical application. If we were to include standard learning models in which costs fall because of adoption, then we would gain two additional effects: first, the endogeneity of learning would favor a subsidy that starts high and declines, and second, the greater reduction in costs would act like making $\dot{C}(t)$ more negative. See Goulder and Mathai (2000) for similar insights.

[3] In 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 not identified by the data.

nology's cost will evolve and how actors will respond to the offered subsidy. At time 0, the regulator chooses the subsidy trajectory to maximize

$$\int_0^T e^{-rt} \left[B(Q(t), \dot{Q}(t)) - \gamma \, s(t) \, \dot{Q}(t) \right] \, \mathrm{d}t,$$

for given discount rate $r > 0$, given initial adoption $Q_0 \in [0, 1)$, and given terminal adoption $Q_T \in [Q_0, 1]$. Potential adopters' decisions determine how the announced subsidy trajectory affects $\dot{Q}(t)$ and thus $Q(t)$. We now specialize to the settings with myopic and forward-looking adopters.

1.2 Myopic Adopters

When potential adopters are myopic, they adopt the technology as soon as their net benefit of adoption is positive. Therefore, at time t, all actors with $v_i \geq C(t) - s(t)$ should adopt (or should already have adopted) the technology. Define $V(t)$ as the value at which actors are just indifferent to adopting or not: $V(t) \triangleq C(t) - s(t)$. The number of actors who have adopted the technology by t is $Q(t) = 1 - F(V(t))$, which implies $\dot{Q}(t) = -f(V(t)) \, \dot{V}(t)$. Note that $\dot{V}(t) \leq 0$ along any optimal path: there is no reason for the regulator to adopt a subsidy that makes net costs strictly increase.

Rather than imagining the regulator selecting the subsidy at each instant, imagine the regulator selecting the quantity of adoption, with the subsidy determined by this decision and by actors' equilibrium conditions. Differentiating the definition of $V(t)$ with respect to time, we see that the subsidy must evolve as:

$$\dot{s}(t) = \dot{C}(t) - \dot{V}(t).$$

Relabel the regulator's control $\dot{V}(t)$ as $y(t)$. The regulator's problem becomes:

$$\begin{aligned}
\max_{y(t)} \int_0^T e^{-rt} & \left[B\Big(1 - F(V(t)), -f(V(t)) \, y(t)\Big) + \gamma \, s(t) \, f(V(t)) \, y(t) \right] \, \mathrm{d}t \\
\text{s.t. } \dot{V}(t) =& y(t) \\
\dot{s}(t) =& \dot{C}(t) - y(t) \\
V(0) =& F^{-1}(1 - Q_0), \; V(T) = F^{-1}(1 - Q_T) \\
s(0) =& V(0) - C(0), \; s(T) = V(T) - C(T).
\end{aligned}$$

The Hamiltonian is:

$$\begin{aligned}
H(t, y(t), V(t), s(t), \lambda(t), \mu(t)) =& e^{-rt} \left[B\Big(1 - F(V(t)), -f(V(t)) \, y(t)\Big) + \gamma \, s(t) \, f(V(t)) \, y(t) \right] \\
&+ e^{-rt} \lambda(t) y(t) + e^{-rt} \mu(t) [\dot{C}(t) - y(t)].
\end{aligned}$$

$\lambda(t)$ gives the shadow value of $V(t)$ and so is negative: the shadow benefit of adoption is positive, and greater adoption corresponds to lower $V(t)$. $\mu(t)$ gives the shadow benefit of a higher subsidy (conditional on $V(t)$). It is negative because, for any given level of adoption, a higher subsidy means that the regulator is spending more money. The necessary conditions for a maximum are:

$$\begin{aligned}\lambda(t)-\mu(t)=&B_2\Big(1-F(V(t)),-f(V(t))\,y(t)\Big)\,f(V(t))-\gamma s(t)\,f(V(t)), \qquad (1)\\ -\dot{\lambda}(t)+r\lambda(t)=&-B_1\Big(1-F(V(t)),-f(V(t))\,y(t)\Big)\,f(V(t))\\ &-B_2\Big(1-F(V(t)),-f(V(t))\,y(t)\Big)\,f'(V(t))\,y(t)+\gamma s(t)\,f'(V(t))\,y(t),\\ -\dot{\mu}(t)+r\mu(t)=&\gamma f(V(t))\,y(t),\end{aligned}$$

in addition to the transition equations and the initial and terminal conditions. The first equation follows from the Maximum Principle and the other two are the costate (or adjoint) equations.

Differentiate equation (1) with respect to time and suppress the arguments of $B(\cdot,\cdot)$:

$$\begin{aligned}\dot{\lambda}(t)-\dot{\mu}(t)=&B_2\,f'(V(t))\,y(t)-B_{12}\,[f(V(t))]^2y(t)-B_{22}\,f(V(t))\,f'(V(t))\,[y(t)]^2-B_{22}\,[f(V(t))]^2\dot{y}(t)\\ &-\gamma\dot{s}(t)\,f(V(t))-\gamma s(t)\,f'(V(t))\,y(t).\end{aligned}$$

Substitute in for $\dot{\lambda}(t)$ and $\dot{\mu}(t)$ from the costate equations, substitute in for $\lambda(t)-\mu(t)$ from equation (1), and rearrange to obtain:

$$\begin{aligned}\dot{s}(t)=&\frac{1}{\gamma}\Bigg\{r\overbrace{[\gamma s(t)-B_2]}^{\mu(t)-\lambda(t)}-B_1-\gamma y(t)\\ &-B_{12}\,f(V(t))y(t)-B_{22}\,\Big[f(V(t))\dot{y}(t)+f'(V(t))\,[y(t)]^2\Big]\Bigg\}. \qquad (2)\end{aligned}$$

This differential equation tells us how the subsidy changes along an efficient trajectory. The subsidy's trajectory is driven by four effects. First, the $r\,[\gamma s(t)-B_2]$ on the top line ensures that the regulator is indifferent to small deviations in the subsidy trajectory. At any point in time, the shadow benefit of raising the subsidy is the sum of the shadow benefit of starting the next instant with a higher subsidy (captured by $\mu(t)<0$) and the shadow benefit of obtaining greater adoption (captured by $-\lambda(t)>0$). This combined shadow benefit will typically be positive.[4] In order to keep the regulator indifferent to small deviations (and

[4] If $\gamma s(t)-B_2$ were negative, then the regulator would value adoption so highly that it would want to obtain adoption even in the absence of the dynamic benefits from cumulative adoption. We are interested in a setting in which the regulator is motivated by benefits from durable investments. We therefore assume that $\gamma s(t)-B_2>0$.

temporarily ignoring other complications), this shadow benefit should grow at the discount rate r, as is familiar from Hotelling (1931) models of exhaustible resource extraction. The efficient subsidy schedule therefore tends to increase because the total shadow benefit of the subsidy should be increasing along an efficient trajectory. As the regulator becomes more impatient (i.e., as r increases), the regulator dislikes spending money now to obtain later benefits and so uses a more sharply increasing subsidy schedule.

Second, the $-B_1$ on the top line 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. 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. The additional benefit of an earlier subsidy favors a decreasing subsidy schedule.

Third, the $-\gamma\, y(t)$ term recognizes that the regulator cannot price discriminate within a period. Recalling that $y(t) \triangleq \dot{V}(t) \leq 0$, this term favors an increasing subsidy schedule. If the regulator offers a marginally greater 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 $-y(t)\, f(V(t))$ inframarginal adopters at time t. This price discrimination motive favors an increasing subsidy. It becomes stronger as the marginal cost of public funds γ increases.

The bottom line in equation (2) tells us how the regulator's marginal benefit of adoption changes as a result of time t adoption. If that marginal benefit falls over time, then this channel favors a decreasing subsidy schedule. The first term has the same sign as B_{12}, which reflects how greater cumulative adoption changes the instantaneous benefit of contemporary adoption. In many cases, $B_{12} \approx 0$. The second term reflects how the instantaneous marginal benefit of contemporary adoption changes as the level of contemporary adoption changes. Contemporary adoption can increase either because the regulator increases the measure of private values for which adoption is optimal ($\dot{y}(t) < 0$) or because the distribution of private values becomes thicker as we move to lower values ($f'(V(t)) < 0$). Either reason reduces the marginal benefit of contemporary adoption, reflected in a negative second term and favoring a decreasing subsidy schedule.

In sum, we have seen that high adoption at a given subsidy, a high marginal cost of public funds, and a high degree of regulator impatience favor an increasing subsidy schedule, whereas a high marginal benefit from cumulative adoption and private values being concentrated below the highest levels favor a decreasing subsidy schedule. Now consider how anticipated improvements in technology affect the efficient subsidy schedule. First, on the top line, more strongly declining costs (i.e., more negative $\dot{C}(t)$) make $y(t)$ more negative for a given subsidy. 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 bottom line 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 $f'(V(t)) < 0$), which favors a decreasing subsidy schedule. If costs are declining at an accelerating rate ($\ddot{C}(t) < 0$), then $\dot{y}(t)$ tends to be negative, again favoring a decreasing subsidy schedule. Combining these pieces, we see that declining costs matter by increasing contemporary adoption and by affecting the change in adoption over time. The first effect favors an increasing subsidy schedule by increasing the gains from price discrimination, and the second favors a decreasing subsidy schedule when falling costs make adoption increase at a faster rate.

1.3 Forward-Looking Adopters

We have thus far considered the efficient subsidy schedule when potential adopters are completely myopic. However, when potential adopters anticipate future subsidies, the regulator can no longer induce adoption just by offering a large subsidy today; instead, the regulator must offer both a high subsidy today and a sufficiently small subsidy in the future.

Instead of adopting the technology as soon as instantaneous net benefits are greater than zero, actors now choose the optimal time to adopt the technology, for given subsidy and cost trajectories:

$$\max_T e^{-\delta T} \left[v_i - C(T) + s(T) \right].$$

The first-order necessary condition is[5]

$$\delta \left[v_i - C(T) + s(T) \right] = \dot{s}(T) - \dot{C}(T). \tag{3}$$

The left-hand side is the cost of waiting until the next instant: the actor delays receiving the instantaneous payoff $v_i - C(t) + s(t)$. The right-hand side is the benefit of waiting: when costs net of the subsidy are decreasing (i.e., when $\dot{C}(t) - \dot{s}(t) < 0$), then the actor can save money by adopting the technology later. At the optimal time to adopt, these costs and benefits cancel. As the potential adopter becomes perfectly patient ($\delta \to 0$), the cost of waiting disappears, and every actor waits until net costs hit their minimum.

Potential adopters' stopping problems generate the equilibrium conditions that constrain the regulator's choice of subsidy trajectory. As before, let the regulator control be which actors are marginal in each period, with $V(t)$ denoting the marginal actors' private value for the technology. Then, rearranging equation (3), the subsidy must evolve as

$$\dot{s}(t) = \delta \left[V(t) - C(t) + s(t) \right] + \dot{C}(t). \tag{4}$$

This problem in which the regulator commits to a subsidy trajectory that is constrained by private expectations is known as a dynamic Ramsey problem.[6] The regulator's time 0

[5] Define net costs as $z(t) \triangleq C(t) - s(t)$. For these necessary conditions to be sufficient, we need $\ddot{z}(T) \geq \delta^2 \left[v_i - z(T) \right] + 2\delta \dot{z}(T)$.

[6] For more on solving dynamic Ramsey problems, see Ljungqvist and Sargent (2004).

choice of subsidy schedule will be dynamically inconsistent because the regulator commits to offering a given subsidy at time t in part to affect potential adopters at times $s < t$, but once time t arrives, those adoption decisions are in the past and thus irrelevant to a decision problem that starts from time t. There is a long tradition in the optimal taxation literature of analyzing this Ramsey problem with full commitment (e.g., Judd, 1985; Chamley, 1986), and it can be particularly applicable in the case of subsidies for new technologies.[7]

Again using $y(t)$ for $\dot{V}(t)$, the regulator solves

$$\begin{aligned}
\max_{y(t)} \int_0^T & e^{-rt} \left[B\Big(1 - F(V(t)), -f(V(t))\, y(t)\Big) + \gamma\, s(t)\, f(V(t))\, y(t) \right] \mathrm{d}t \\
\text{s.t. } \dot{V}(t) =& y(t) \\
\dot{s}(t) =& \delta\, [V(t) - C(t) + s(t)] + \dot{C}(t) \\
V(0) =& F^{-1}(1 - Q_0),\ V(T) = F^{-1}(1 - Q_T) \\
s(0) =& -\dot{C}(0)/\delta - V(0) + C(0).
\end{aligned}$$

We set the initial condition on $s(\cdot)$ to be consistent with Q_0 and $\dot{s}(0) = 0$. The Hamiltonian is

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

The costate variable $\lambda(t)$ is the same as in the myopic setting. The costate variable $\mu(t)$ now measures the degree to which the regulator is constrained at each instant by private actors' equilibrium behavior and foresight: it measures the cost of keeping promises made to those who adopted the technology in past periods. At time 0, this promise-keeping cost must be zero, because the regulator has not yet made promises that it is forced to maintain. The necessary conditions for a maximum are:

$$\lambda(t) = f(V(t))\, B_2\Big(1 - F(V(t)), -f(V(t))\, y(t)\Big) - \gamma s(t) f(V(t)), \tag{5}$$

$$\begin{aligned}
-\dot{\lambda}(t) + r\lambda(t) =& -f(V(t))\, B_1\Big(1 - F(V(t)), -f(V(t))\, y(t)\Big) \\
&- f'(V(t))\, y(t)\, B_2\Big(1 - F(V(t)), -f(V(t))\, y(t)\Big) + \gamma s(t) y(t)\, f'(V(t)) + \delta\mu(t),
\end{aligned}$$

$$-\dot{\mu}(t) + r\mu(t) = \gamma\, f(V(t))\, y(t) + \delta\mu(t), \tag{6}$$

$$\mu(0) = 0.$$

[7]For instance, our empirical application will consider California's subsidies for rooftop photovoltaic (solar) systems. Most observers take for granted that the regulator will follow the announced subsidy schedule.

in addition to 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 initial condition for the promise-keeping constraint.

Solving for $\mu(t)$ in equation (6), we find

$$\mu(t) = -\gamma \int_0^t e^{(r-\delta)(t-i)} f(V(i))\, y(i)\, \mathrm{d}i \geq 0.$$

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. The regulator's promises accrue as past adoption accrues: the regulator becomes more bound by the commitments it has made in order to obtain past adoption. 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 bigger $\mu(t)$ is: past regulators helped themselves by constraining the future.

Differentiate equation (5) with respect to time to obtain

$$\begin{aligned}\gamma \dot{s}(t)\, f(V(t)) + \gamma s(t)\, f'(V(t))\, y(t) + \dot{\lambda}(t) =& B_2\, f'(V(t))\, y(t) - B_{12}\, [f(V(t))]^2 y(t) \\ &- B_{22}\, f(V(t))\, f'(V(t))\, [y(t)]^2 - B_{22}\, [f(V(t))]^2 \dot{y}(t),\end{aligned}$$

where we suppress arguments on $B(\cdot,\cdot)$. Substitute for $\dot{\lambda}(t)$ from the costate equation on $V(t)$, substitute for $\lambda(t)$ from equation (5), and rearrange:

$$\begin{aligned}\dot{s}(t) =& \frac{1}{\gamma}\Bigg\{ r\, \overbrace{[\gamma s - B_2]}^{-\lambda(t)} - B_1 + \delta \mu(t)/f(V(t)) \\ & - B_{12}\, f(V(t)) y(t) - B_{22}\, \left[f(V(t)) \dot{y}(t) + f'(V(t))\, [y(t)]^2 \right] \Bigg\}. \end{aligned} \tag{7}$$

This equation defines the dynamics of the efficient subsidy when potential adopters correctly anticipate future subsidies and costs. The bottom line and the $r\,[\gamma s - B_2] - B_1$ on the top line also appeared in the myopic case. However, the remaining term differs. In the myopic case, we had a term that 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 price discrimination term has therefore disappeared.

The new term $\delta\mu(t)/f(V(t))$ reflects the time t cost of keeping past promises, as described above. This promise-keeping cost is weakly positive, favoring an increasing subsidy schedule.

All else equal, a regulator who obtained a lot of adoption prior to time t must have promised a low time t subsidy, which makes her want to raise the subsidy for later times. When adopters are perfectly patient ($\delta = 0$), the regulator must offer that low subsidy forever, but when adopters are impatient, the regulator can offer a high subsidy in late periods without strongly disincentivizing adoption in early periods.

Replacing the price discrimination channel from the myopic case with this promise-keeping channel changes the subsidy schedule in different ways as time passes. The price discrimination channel favored an increasing subsidy schedule at times with high adoption, because these are the times with a lot of inframarginal adopters. In contrast, the promise-keeping channel favors an increasing subsidy some time *after* adoption was high, because the regulator must have promised not to raise the subsidy for a while in order to induce that high adoption. Near the initial time, the promise-keeping cost is approximately zero while the price discrimination channel is large if costs are declining rapidly. Recognizing that adopters anticipate future subsidies therefore eliminates a force for an increasing subsidy in early periods but introduces a new force for an increasing subsidy in later periods, potentially once a lot of adoption would have already occurred in the myopic case. Recognizing adopters' foresight therefore tilts the subsidy schedule down at first (to induce adoption in early periods) but can eventually tilt the subsidy schedule back up (to induce adoption in later periods).

2 Data

Our theory shows that subsidy schedules may have different effects depending on consumers' foresight or myopia and on the distribution of willingness-to-pay in the population. In order to be more concrete about how the effects of an existing subsidy might change with different time-paths, we use data from the California Solar Initiative to estimate the distribution of willingness-to–pay for residential solar in the California population and then simulate counterfactual subsidy time paths.

The California Solar Initiative (CSI) was a state subsidy for distributed solar installations. We focus on the residential solar component of the policy, which subsidized residential solar installations on a per-watt-installed basis that was decreasing with the total amount of solar installed in each of the three major California utilities (Pacific Gas and Electric (PG&E), San Diego Gas and Electric (SDG&E), and Southern California Edison (SCE)). The CSI subsidy is described in detail in Burr (2014) and others, so we will focus on the data that we are using in our analysis.

The CSI maintains data on all applications for residential solar subsidies under the program. We assume that all households that install solar between the second quarter of 2007 and the third quarter of 2014 apply for the CSI subsidy and use the CSI data as the universe of households that install solar systems during this timeframe.[8] The data includes the application date, the household's zip code and utility, and an extensive set of characteristics of the solar system the household will install including the system size, manufacturer and installer, and total cost.

Our empirical model (described below) will only describe households' willingness to pay for a standard residential solar system. To this end, we assume that each household is choosing whether to install a system of the median size over the entire period, 5.29 kilowatts (kW). We also assume that the price the household will expect to pay is the average per-Watt price in the household's utility in the given quarter, multiplied by 5.29. Figure 1 shows the evolution of system costs (absent any subsidies) for each utility over the time-frame of our data. Costs rise in the initial periods when silicon costs are increasing and then decline over the majority of our sample as technology advances and silicon costs fall.

Figure 1: Quarterly system cost by utility

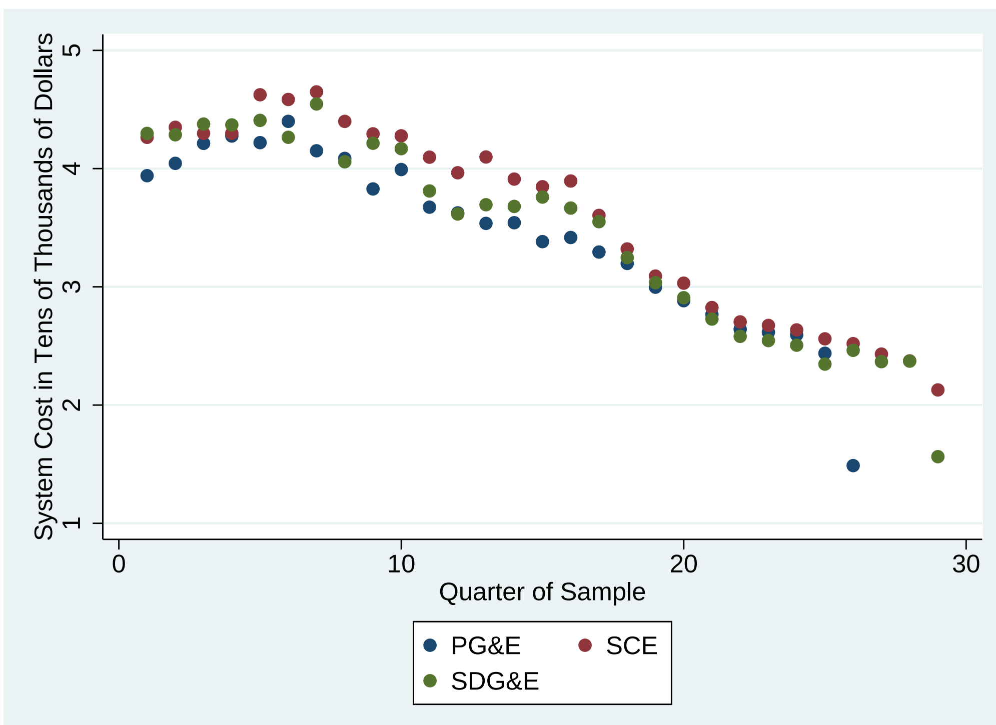

In addition to the information on solar system installations and costs, we use demographic data at the zip code level from the American Community Survey to allow preferences for

[8]There is some feeling that not all households applied for the CSI subsidy at the end of the timeframe. We can test removing the later periods from the estimation to see if it affects our results.

solar to vary with consumer demographics. Given the short panel of solar installation data, we do not allow demographics to vary over time. We also 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. Finally, in a future version of this analysis we will also use data from the National Renewable Energy Laboratory (NREL) on the average solar irradiation at the zip code level to account for variation in the solar generation potential in different areas of California.

3 Evidence that Consumers are Forward-Looking

Before we estimate consumers' willingness-to-pay for residential solar, it is important to understand whether consumers actually do think about future costs and benefits (or prices and subsidies) of solar when they make their investment decisions. There are a few reasons why we might think that the dynamic time-path 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 time-path of solar subsides 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 show that households in California are behaving as if they are forward-looking, we look at evidence from discontinuous changes in subsidies and the more frequent but smaller changes in input (silicon and panel) costs, holding constant current solar module costs.

To this end, we regress the weekly counts of residential solar subsidy applicants on measures of current and future 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. If the subsidy is about to decrease, then forward-looking households will choose to install solar now rather than delay.[9] Similarly, holding constant the current solar module costs, if silicon costs are high then this means that the future cost of solar modules will increase, so forward-looking households will want to install now. Finally, 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 these hypotheses, we combine the CSI application data with a solar module price index and a silicon price index from Bloomberg. We also include the 3 month change in the dollar-yuan exchange rate,[10] and a set of controls including the VIX (a measure

[9] This effect has been shown previously in Burr (2014).

[10] We will replace this with the 3 month exchange rate future minus the current exchange rate in a future

of the S&P 500 volatility), 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. We hope to capture overall macro-economic incentives to invest in solar with this set of controls.[11]

Table 1: Evidence that Consumers are Forward-Looking

Dependent Variable: Weekly CSI Applications	
Within 2 months of subsidy drop	0.299***
	(0.055)
Solar Module Price Index	-2.667***
	(0.066)
Silicon Price Index	0.012**
	(0.006)
3 month change in Dollar-Yuan Exchange Rate	-1.955***
	(0.694)
VIX (Instrumented)	-0.032**
	(0.015)
Number of weeks	
R^2	

Controls include the Dollar-Yuan exchange rate, the Euro-Yuan exchange rate, the 6 month T-bill interest rate and the DOW REIT index. Instruments are month-over-month changes in 10 year bond yields for Greece, Italy, Russia, and Spain. F=26.79

The regression results presented in Table 1 suggest that consumers are forward-looking in their decision to install solar (or at least that solar installers are forward-looking). The coefficient estimates suggest that more people submit applications to install solar systems when subsidies are about to decline, as we would expect. Higher solar module prices lead to lower installation, but conditional on solar module prices, higher silicon prices (which suggest that module prices will be higher in the future) lead to more installation now. 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 and there are fewer installations today. It is also interesting to note that economic uncertainty, as captured by the instrumented VIX, tends to decrease solar installations, as would be expected if households account for option value (which again implies that households are not fully myopic).

draft.

[11]We instrument for the VIX with month-over-month changes in 10 year bond yields for several economically-volatile developed countries (Greece, Italy, Russia, and Spain) in order to focus on the effect of market uncertainty rather than the underlying changes in the macro-economy that are associated with economic uncertainty in the 2008-2010 period.

4 Dynamic Model

4.1 Model

Given that consumers are forward-looking in their decision about whether to invest in solar energy, we need a dynamic choice model in order to understand the distribution of willingness-to-pay in California and the effect of any given subsidy schedule. To this end, we will estimate a dynamic choice model to that is similar to Burr (2014) except that we will explicitly estimate how consumer demographics are correlated with willingness-to-pay and we will allow the evolution of solar prices to be stochastic rather than fully known by consumers. Allowing willingness-to-pay to be correlated with demographics is critical for recovering a realistic distribution of willingness-to-pay to use in the subsidy simulations. Allowing the price path of solar systems to be stochastic makes for a more realistic situation where subsidy schedules must be set under price uncertainty.

In our empirical model, households face the choice in each period of whether to invest in a solar system of a given size. If the household decides to purchase the solar system, then they receive the benefits of that system in each period in the future but do not make an additional investment decision. In this way, the model is very similar to Rust (1987). In each period, t, each household i who has not yet installed solar has a choice of investing in solar and receiving the benefit:

$$U_{i1t} = \alpha_{iX} - \alpha_{iP} P_t + \varepsilon_{i1t} + \sum_{\tau=1}^{\infty} \beta^{t+\tau}(\alpha_{iX} + \varepsilon_{i1\tau}),$$

where $\alpha_{iX} - \alpha_{iP} P_t + \varepsilon_{i1t}$ is the value of solar in the current period, including the up-front cost of solar installation P_t and a logit error ε_{i1t}. The household also receives a stream of future benefits from installing solar which are assumed to continue indefinitely into the future, $\sum_{\tau=1}^{\infty} \beta^{t+\tau}(\alpha_{iX} + \varepsilon_{i1\tau})$. If the household chooses not to install solar, then it receives:

$$U_{i0t} = \varepsilon_{i0t} + \beta \mathbb{E} V(\Omega' | \Omega),$$

where ε_{i0t} is a logit draw and $\beta \mathbb{E} V(\Omega'|\Omega)$ is the discounted continuation value. Here, $\Omega = \{P_t, W_t\}$ is the set of state variables in the current period, which includes the current cost of a solar system (P_t) and the change in the subsidy between periods t and $t+1$ (W_t).

Consistent with policymakers committing to a subsidy schedule, we assume that the entire subsidy schedule is known to households. However, the time-path of solar system prices is not known to households, who only know that solar costs (inclusive of any subsidy) evolve according to an AR1 process:

$$P_{t+1} = \gamma_0 + \gamma_1 P_t + \gamma_2 W_t + \mu_t, \tag{8}$$

where μ_t is a normal residual.

4.2 Implementation

We estimate our model using simulated maximum likelihood. The CSI provides data on the applications to install solar, which includes information on the price paid for the solar system and the household's zip code but not the household's demographics. In order to estimate how preferences vary with demographics, we simulate over the distribution of demographics within each zip code as given by the American Community Survey as in Nevo (2001) and Berry, Levinsohn, and Pakes (2004).[12] There are over 2,500 zip codes in California, so identification of the demographic preference coefficient coefficients is coming from the fact that, for instance, particularly high-income zip codes are more likely to have higher rates of solar installation conditional on prices than low-income zip codes. We assume that only owner-occupied households consider installing solar.

At each step of the likelihood maximization, we solve the value function recursively from the final period for each demographic group. At the final period the CSI subsidy is zero and we assume that the subsidy will be zero in all periods after the end of our data. We assume that households make a decision on whether or not to install solar each quarter, and we assume a quarterly discount factor of 0.99. Standard errors are calculated as the square root of the inverse of the outer product of the Jacobian.

5 Empirical Results

5.1 Preference Parameters

Table 2 presents the estimated coefficients from our dynamic demand model. The first set of estimates are for the benefit of installing a solar system of the median size. On average, households have a negative value of installing solar, which makes sense given the low overall installation rate for solar. Households with higher income have a higher value of solar, but this effect is both economically and statistically insignificant. Democratic voters have a lower value of solar, although this could be picking up the fact that Democratic voters are clustered in the North and along the coast in California where there is less sun. Having a bachelor's degree and living in a small house (less than 3 bedrooms) both have a statistically insignificant effect on the value of installing solar, although the point estimates for both are negative.

The second set of estimates capture how price-responsive households are when it comes to installing solar systems. Overall, households are strongly price sensitive. Having a bachelor's degree makes households somewhat less price sensitive. While statistically insignificant, democratic voters and people in small houses are both less price sensitive, which makes sense if Democratic voters are more likely to have green preferences and small houses are

[12] Unfortunately, the joint distribution of demographics is not available at the zip code level in the ACS, so we draw from each demographic's unconditional distribution separately.

more likely to install smaller systems. Households with higher income are actually more price sensitive, but again this coefficient is economically and statistically quite small.

Table 2: Dynamic Demand Estimates

Benefit of Solar		-4.13***
		(0.11)
	* Household Income ($10,000)	0.0026
		(0.0089)
	* Democratic Voter	-0.308***
		(0.103)
	* Bachelor's Degree	-0.136
		(0.117)
	* Small House	-0.53
		(0.341)
Price ($10,000)		-0.815***
		(0.039)
	* Household Income ($10,000)	-0.0001
		(0.0029)
	* Democratic Voter	0.055
		(0.035)
	* Bachelor's Degree	0.087**
		(0.038)
	* Small House	0.11
		(0.107)
	Quarters	29

Notes

Our model simplifies actual decisions to install a solar system. Our simplifications are justified because we are interested in estimating the correct distribution of willingness-to-pay for a standard residential solar system, as opposed to the magnitude of different preference coefficients. With this in mind, it is worth understanding where some standard complications would enter our estimation. First, one of the main incentives for consumers in California to adopt solar is the steep tiered pricing structure of electricity. Households that use a lot of electricity pay a substantially higher marginal rate than households that use less. This tiered structure means that solar is most cost-effective for households with the greatest total electricity use. While we are not able to use actual electricity consumption data in our model, we hope to capture much of this incentive by allowing willingness-to-pay to vary with demographics such as house size and household income.

Similarly, we do not observe the solar generating potential of each individual home in California. Roof angle, size, and shading all affect the electricity generation from a solar system at a given house and higher costs systems will only partially offset these effects. The small house variable will capture some of this variation, but much of it must remain in the error term. At the moment, differences in solar potential due to weather and latitude differences are not included in the model, but they will be added in a future draft.

The other complication that is currently omitted from the empirical model is the relationship of installations in California during the CSI period to other state and federal solar incentive programs, including those that consumers might expect to occur after the CSI finishes. Our empirical model assumes that subsidies disappear by the last period of the data and are expected to remain at 0 into the future. To the extent that this is not true and households expect that there will be some subsidy in the future, households might be inclined to wait for solar costs to decline further and new incentive policies to be enacted, which would make them look less price-responsive than they actually are. This is largely a problem for our analysis if the expectations of future policies are somehow heterogeneous across the population in a way that changes the distribution of willingness-to-pay substantially.

Finally, federal subsidies for solar are available over this period. We will build these into the analysis in the next round of revisions, with the most important component being the lifting of the cap on federal subsidies for solar systems in 2008. Not including these subsidies at the moment is likely making households look less price sensitive than they actually are since households are not facing the full cost of solar systems that we are including in the model.

5.2 Willingness-to-pay

The estimated distribution of willingnesses-to-pay for solar systems is shown in Figure 2 (with 2000 draws of the error term). All of the willingnesses-to pay are substantially negative (meaning that consumers would need to be paid to adopt solar). While it seems reasonable that the average willingness-to-pay for solar might still be negative (as many homes are shaded or households do not want to spend the time and energy to think about installing solar), it seems that the heterogeneity in willingness-to-pay is likely larger than we have estimated. We anticipate that incorporating solar irradiation and additional demographic controls into our dynamic model will expand the support of the distribution substantially. To that end, rather than use this estimated distribution to conduct our policy simulations, we will use a simulated density until we have updated the dynamic estimation.

6 Evaluating the Efficient Subsidy Schedule for Rooftop Solar (Examples for now)

We now combine the empirically estimated distribution of private values for rooftop solar technologies with the theoretical analysis of Section 1. The simulations are still in progress, as they await the new empirical results. For now, this section describes how we translate the

Figure 2: Distribution of Estimated Willingness-to-Pay

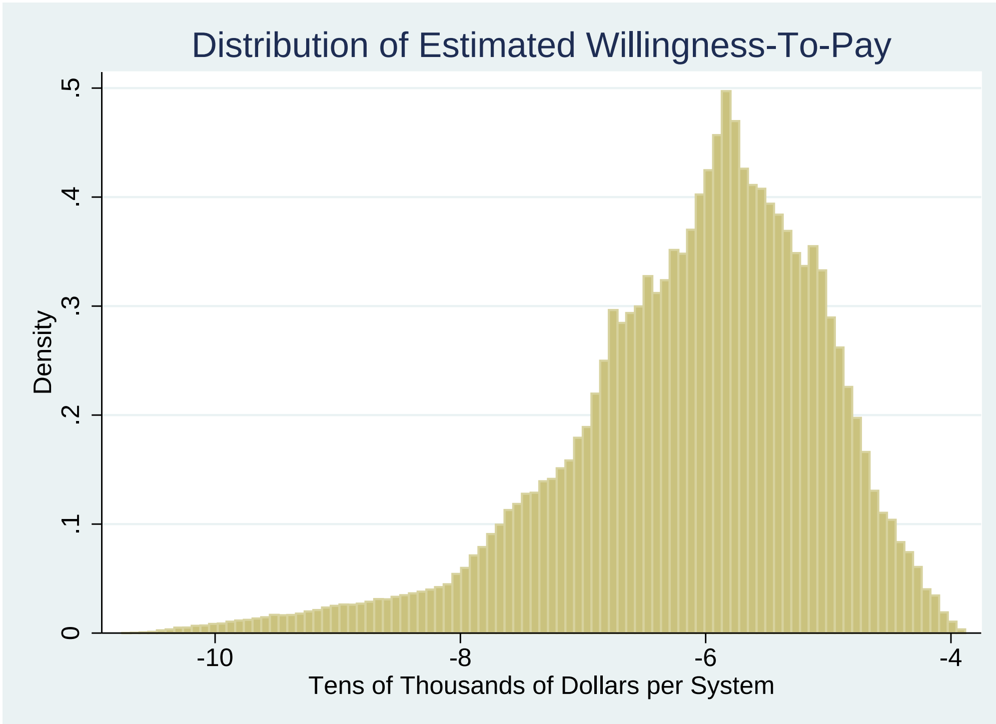

empirical setting into the theoretical setting, outlines the solution method, and illustrates results using a distribution of private values from an early round of empirical analysis.

The empirical setting estimated private values for solar of the form

$$v_j = \frac{1}{a_j}(b_j + \epsilon_j),$$

for each demographic group j. With $a_j < 0$, the cumulative distribution function for v_j is

$$F_j(v) = 1 - e^{-e^{-a_j v + b_j}}.$$

The number of households in demographic group j is N_j. The cumulative distribution function $F(v)$ for willingness to pay across the entire population is

$$F(v) = \frac{\sum_j N_j F_j(v)}{\sum_j N_j}.$$

Using the realized cost trajectory, the private cost of adoption evolves as

$$C(t) = 4.763 - 0.087\,t,$$

with time t measured in quarters and costs measured in tens of thousands of dollars.

All of the trajectory simulations assume that the policymaker achieves 50% penetration ($Q_T = 0.5$) by the end of the 29-quarter horizon ($T = 29$). The marginal cost of public funds is unity ($\nu = 1$), and both households and the regulator have discount rates of 0.4% per year. The regulator has benefit function:

$$B(Q(t), \dot{Q}(t)) = 4.763\, Q(t) - 2.3815\, Q(t)^2 + \dot{Q}(t) - 5000\, \dot{Q}(t)^2.$$

All of these parameters are chosen arbitrarily for now. The next step is to calibrated them so that the discount rates match the empirical setting and so that the regulator's benefit from cumulative adoption matches estimates of the social benefit of solar, as from Baker et al. (2013).

In the setting with myopic adopters, the terminal condition $Q(T) = Q_T$ gives both $V(T)$ and $s(T)$. The initial value $Q_0 = 0.0182$ is determined by the distribution of private values and the time 0 cost of adoption. Solving equation (2) for $\ddot{s}(t)$ (via $\dot{y}(t)$), we have a system of differential equations in $V(t)$, $s(t)$, and $\dot{s}(t)$, with known $V(T)$ and $s(T)$. We solve this system by guessing a value of $\dot{s}(T)$, solving the system backwards in time from T to 0, and checking whether $Q(0) = Q_0$. If $Q(0)$ and Q_0 differ, then we try a new guess for $\dot{s}(T)$.

In the setting with forward-looking adopters, we solve equation (7) for $\dot{y}(t)$ in order to obtain a system of equations in $s(t)$, $V(t)$, $y(t)$, and $\mu(t)$. We fix Q_0 (and thus $V(0)$) to be equal to the Q_0 used in the myopic setting, and we fix $s(0)$ so that $Q(0)$ solves equation (4) with $\dot{s}(0) = 0$ (i.e., we fix the subsidy's initial value so that Q_0 adopters would have found it optimal to adopt with the constant subsidy $s(0)$, assuming they did not anticipate the new subsidy policy). And recall that we know $\mu(0) = 0$. For any guess $\dot{y}(0)$, we solve the system forward in time from 0 to T and check whether $Q(T) = Q_T$.

Figure 3 illustrates results for the cases of myopic (left) and forward-looking (right) adopters. The top panels plot cumulative adoption $Q(t)$ and the instantaneous rate of adoption $\dot{Q}(t)$ along the efficient trajectories. We see that the efficient adoption trajectories are similar whether or not potential adopters are forward-looking. However, the efficient subsidy trajectory (middle row) depends strongly on whether potential adopters are forward-looking. We see that the efficient subsidy increases over time when adopters are myopic but decreases over time when adopters are forward-looking. The bottom row explores the determinants of these subsidy trajectories, following the analysis of equations (2) and (7). The change in the subsidy ($\dot{s}(t)$) is the sum of the plotted channels. In this initial example, we see that the common channels have similar magnitudes across the two settings. In the setting with myopic adopters, the price discrimination motive starts out positive and quickly declines, but in the setting with forward-looking adopters, the promise-keeping channel is always small. The initially positive price discrimination channel explains why the efficient subsidy initially increases strongly in the case of myopic adopters, and replacing this channel with the small promise-keeping channel explains why the efficient subsidy decreases in the case of forward-looking adopters.

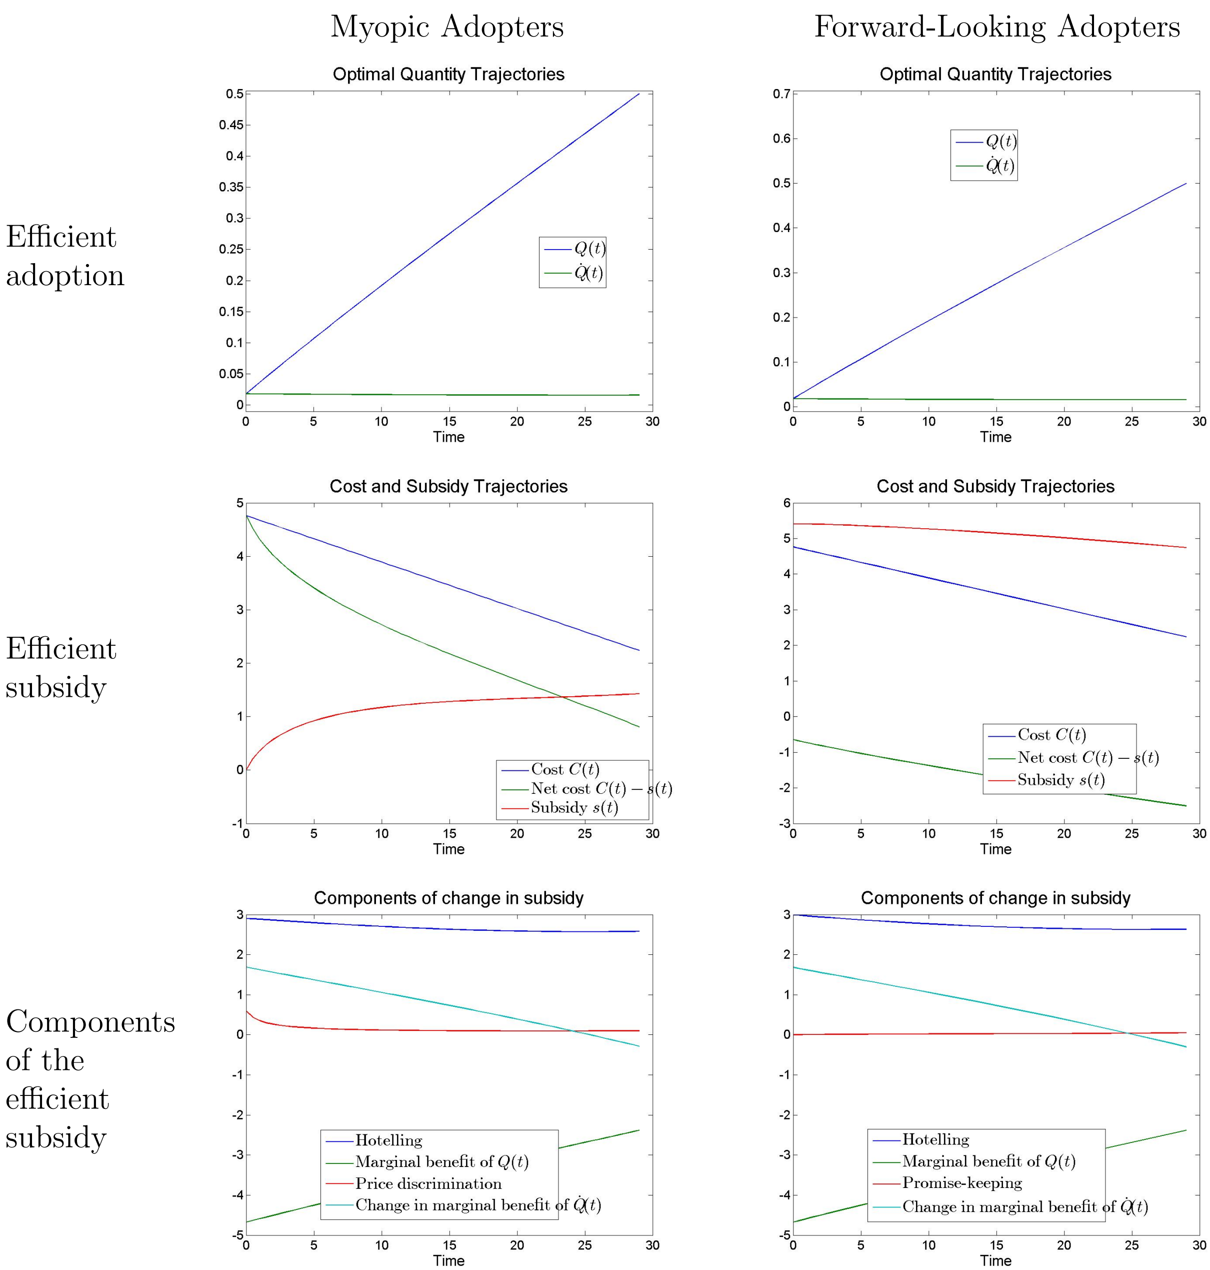

Figure 3: Examples of the efficient adoption (top) and subsidy (middle) schedules in the case of myopic (top) and forward-looking (bottom) households. The bottom row plots the components of the efficient subsidy trajectories, from equations (2) and (7) for myopic and forward-looking adopters, respectively. The costs and subsidies are measured in tens of thousands of dollars, and time is measured in quarters. These plots are placeholders until we have final estimation results.

Our next steps are to implement the final empirical estimates into these simulations, to more carefully calibrate the regulator's parameters, and to use the simulations to explore the importance of factors like discount rates and the marginal benefit of adoption.

7 Conclusions

TBA

References

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.

Besanko, David and Wayne L. Winston (1990) "Optimal price skimming by a monopolist facing rational consumers," *Management Science*, Vol. 36, No. 5, pp. 555–567.

Burr, Chrystie (2014) "Subsidies, tariffs and investments in the solar power market," working paper.

Chamley, Christophe (1986) "Optimal taxation of capital income in general equilibrium with infinite lives," *Econometrica*, Vol. 54, No. 3, pp. 607–622.

Coase, R. H. (1972) "Durability and monopoly," *Journal of Law and Economics*, Vol. 15, No. 1, pp. 143–149.

Goulder, Lawrence H. and Koshy Mathai (2000) "Optimal CO2 abatement in the presence of induced technological change," *Journal of Environmental Economics and Management*, Vol. 39, No. 1, pp. 1–38.

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson (1986) "Foundations of dynamic monopoly and the Coase conjecture," *Journal of Economic Theory*, Vol. 39, No. 1, pp. 155–190.

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.

Judd, Kenneth L. (1985) "Redistributive taxation in a simple perfect foresight model," *Journal of Public Economics*, Vol. 28, No. 1, pp. 59–83.

Kahn, Charles (1986) "The durable goods monopolist and consistency with increasing costs," *Econometrica*, Vol. 54, No. 2, pp. 275–294.

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.

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.

Ljungqvist, Lars and Thomas J. Sargent (2004) *Recursive Macroeconomic Theory*, Cambridge, Mass: The MIT Press, 2nd edition.

Meyer, Laurence H., Joel L. Prakken, and Chris P. Varvares (1993) "Policy watch: Designing an effective investment tax credit," *Journal of Economic Perspectives*, Vol. 7, No. 2, pp. 189–196.

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.

Rust, John (1987) "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, Vol. 55, No. 5, pp. 999–1033.

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.

Sobel, Joel (1991) "Durable goods monopoly with entry of new consumers," *Econometrica*, Vol. 59, No. 5, pp. 1455–1485.

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.