Inferring the Effects of Permanent Policies from Temporary Interventions *

Ariel Goldszmidt

Santiago Lacouture
(Job Market Paper)

This Version: November 17, 2026

Click here for the latest version

Abstract

Many public policies are designed to be sustained or permanent, but the experiments used to evaluate them are only temporary. Responses to incentives take time to build when past choices shape future ones, so a temporary intervention may end before it reaches the effect of a permanent intervention. We develop and apply a framework that uses temporary experiments to predict responses to permanent policy interventions. The framework is built around a general dynamic choice model in which choices persist both because past choices shape future ones and because preferences and costs are serially correlated. We show that the path of treatment effects in a temporary experiment identify the model's key features, which we then use to evaluate the long-run effects of a permanent intervention. We apply the framework to the Seattle–Denver Income Maintenance Experiment, which was intended to inform a permanent negative income tax but paid most participants for an announced three-year period. We find that taking the experimental findings at face value understates by at least 25% the effect of a permanent policy on hours worked. We apply the framework to experiments in two other different settings, one discouraging social media use and one incentivizing exercise. We find that, made permanent, the social media intervention would reach two thirds of its long-run effect in the first week, but the exercise intervention only after about eight months.

*Goldszmidt: University of Chicago. Lacouture: University of Chicago. We are deeply grateful to Magne Mogstad, Alexander Torgovitsky, Manasi Deshpande, and Winnie van Dijk for their advising, guidance, and support. We also thank Stéphane Bonhomme, Santiago Franco, Lancelot Henry de Frahan, John Eric Humphries, Pauline Mourot, and Oscar Volpe for helpful comments and discussions. All errors are our own.

1 Introduction

Economists have long studied how and why choices persist over time. Those who work this year are likely to work next year (Heckman and Willis, 1977; Hyslop, 1999), smokers tend to keep smoking (Becker et al., 1994), and consumers tend to purchase brands they have purchased before (Keane, 1997; Dubé et al., 2010).

Observed persistence has two primary explanations (Heckman and Borjas, 1980): past choices shape future ones (*state dependence*), and individuals differ in preferences or face persistent circumstances (*unobserved heterogeneity*). The two explanations are hard to distinguish (Heckman, 1981; Torgovitsky, 2019), but have different implications for how individuals will respond to policies that change incentives to make these choices (Moffitt, 1992; Cawley and Ruhm, 2011). In particular, the source of persistence determines how the response to a temporary intervention relates to the response to the same intervention made permanent. If there is state dependence, the effect of an intervention grows while it lasts. Therefore, a temporary intervention, such as an experiment, may be a poor guide to the effects of the same program made permanent (Keane, 2015).[1]

Consider a subsidy that lowers the cost of working, either temporarily or permanently. If persistence in labor supply reflects only unobserved heterogeneity, then a temporary subsidy will increase hours worked only while it is in place, and making the subsidy permanent would extend that same increase in perpetuity. If instead there is state dependence in labor supply—for example, through human capital accumulation that raises future wages—then the extra hours the temporary subsidy causes today lead to more hours tomorrow, and its effects may compound and grow over time while it lasts and persist after it ends. Made permanent, the same subsidy keeps raising hours beyond the level the temporary one reached (Robins and West, 1980; Keane, 2015). The effects of the permanent and temporary subsidies on labor supply also diverge when workers are forward-looking, since a forward-looking worker values a subsidy she expects to be permanent differently from one she expects to be temporary (Metcalf, 1973). Inferring the effects of permanent policy changes using evidence from temporary interventions or experiments requires knowing how and why choices persist, and how forward-looking individuals are.

In this paper, we develop and apply a framework that uses temporary experiments to quantify the causes of persistence and to predict long-term policy effects.[2] We develop the framework around a dynamic choice model that allows for state dependence, unobserved heterogeneity, and forward-looking behavior. We show that the time series of treatment effects in temporary experiments reveals the model's features needed to predict long-term policy effects. We apply our

[1] Furthermore, experiments are almost always temporary, often out of budgetary or ethical reasons.

[2] By temporary experiments we mean designs in which participants are randomly assigned an incentive of fixed size for a fixed and known period, and their choices are measured both while the incentive lasts and after it ends. Such designs are standard in field experiments across economics (see the reviews in Harrison and List, 2004; Duflo et al., 2007; List, 2011; Gneezy et al., 2011).

framework in three empirical settings: labor supply, social media use, and physical exercise. In each setting, we use a temporary experiment to quantify the causes of choice persistence and to answer counterfactual questions about permanent policy changes.

Our framework starts with a general dynamic choice model in Section 2. The model features the explanations of choice persistence described above: past choices shape future ones through an accumulated 'stock' of past choices, generating state dependence, and individuals differ in unobserved preferences and face possibly serially correlated and predictable costs of choices. We study a class of interventions that change costs, either temporarily or permanently, and derive closed-form expressions for the difference between the short-run response to a temporary intervention and the long-run response to a permanent intervention. We show that the difference depends on three features of behavior. The first, *reinforcement*, is how strongly the accumulated history of choices affects the current choice. The second, *survival*, is how slowly the accumulated history depreciates. Together, reinforcement and survival determine *propagation*: how long past choices shape future ones. The third, *patience*, is how forward-looking individuals are.

We study identification of these features from data on choices in Section 3. We start with propagation. The identification challenge is that choices may be persistent either because propagation is high or because preferences and costs are serially correlated. Overcoming this challenge requires an exogenous and excludable instrument for past choices. Our approach exploits assignment in a temporary experiment as an (external) instrument. This instrument satisfies exogeneity if assignment to the intervention is random, and satisfies exclusion if the intervention changes participants' costs and beliefs about future costs only temporarily. Under these two conditions, propagation is identified by how fast the intervention's effects on choices fade out after the intervention ends. A classic alternative approach differences the choice equation and uses further lagged choices as (internal) instruments (Anderson and Hsiao, 1982; Arellano and Bond, 1991; Wooldridge, 2010). Under our model, these instruments satisfy exogeneity if individuals cannot anticipate future changes to their costs, a strong assumption on information and beliefs.

We next consider how to separate reinforcement from survival and how to identify patience. The decomposition of the two channels of propagation matters because the effects of a permanent policy will be greater when propagation comes mostly from reinforcement than when it comes mostly from survival. Patience matters because the effects of a permanent policy will be greater the more patient individuals are. We show that the same temporary experiment used to identify propagation also identifies survival, reinforcement, and patience. The decomposition is revealed by how much the intervention's effects 'drop' when its end arrives: the more propagation comes from survival, the less reinforcement sustains the effects, and the larger the drop in treatment effects. Patience is revealed by how much the intervention's effects 'bend' as its end approaches: the more patient individuals are, the more they anticipate the intervention's coming end, and the larger the bend in treatment effects.

We apply our framework to learn about the effects of permanent policies in three empirical

settings with temporary experiments. First, in Section 4, we apply the framework to the three-year arm of the Seattle–Denver Income Maintenance Experiment (SIME/DIME), the largest and most studied of the negative income tax (NIT) experiments of the 1970s (Christophersen, 1983). These experiments were designed to measure how low-income households would respond to a permanent NIT (Burtless and Haveman, 1984), but participants knew the experiment would only last a few years. From early on, researchers pointed out that the effects of an NIT would be different if it were implemented permanently, both because participants may respond differently to a program they know will end (Metcalf, 1973; Burtless and Greenberg, 1982) and because effects may build gradually when past work shapes future work (Robins and West, 1980); our framework allows for both possibilities.

Taking our framework to the data, we find that past work accounts for 62 percent of the quarter-to-quarter persistence in hours, mostly through work prior to the last quarter. Because the effects of past work accumulate, a permanent NIT would reduce hours worked in the long run by 27 to 58 percent more than the reduction observed at the end of the experiment. We also validate our framework by showing that the model estimated on the three-year arm of the experiment correctly predicts the build-up of responses over time in the held-out five-year arm.

Second, in Section 5, we apply the framework to the 'Bonus' experiment in Allcott et al. (2022), which paid adults to reduce their social media use for three weeks.[3] We find that persistent social media use mostly reflects persistent preferences and costs rather than habit formation: past use explains only 34 percent of observed week-to-week persistence. We estimate that a permanent incentive to reduce use would reach its long-run effect quickly: two thirds of that effect would be reached in the first week, and 89 percent within two months.

Third, in Section 6, we apply the framework to 'Study 1' in Charness and Gneezy (2009), which paid students to go to the gym for a month. The payments led to lasting effects on gym attendance, a finding widely read as evidence that temporary incentives can build strong exercise habits (Acland and Levy, 2015; Royer et al., 2015). Consistent with this interpretation, we estimate that past attendance causes 77 percent of the observed week-to-week persistence in attendance. We also find that the accumulated history of exercise depreciates slowly and reinforces future exercise nearly as fast as it depreciates, a combination which makes responses to permanent incentives continue to grow over time. We estimate that a permanent incentive scheme would raise weekly gym visits by 4.2 after 25 weeks, an effect 2.5 times as large as that observed in the experiment's last payment week.

Our results underscore the difficulty of extrapolating short-term experimental findings to the long-run effects of permanent policies. Nevertheless, we show that progress can be made by exploiting variation in when incentives are announced, how long they last, and when they begin.

[3]Allcott et al. (2022) combine the persistence of the bonus effects with evidence from screen-time limits to estimate a model of habit formation and self-control. We use the same data to measure how much of the persistence in use is caused by past use and how fast a permanent incentive would reach its long-run effect.

To this end, we provide suggestions for designing more informative experiments in Section 7. We propose three variations on the single-arm designs studied above: announcing the incentive to some participants early; varying the length of the incentive across participants; and staggering the start of the incentive across participants. These variations deliver additional and more informative ways to measure patience—the parameter estimated least precisely in our empirical applications—and provide opportunities to test the model's assumptions and validate its performance out of sample.

This paper contributes to three areas of literature. The first area of literature studies the identification and estimation of models with state dependence. One strand specifies and estimates dynamic choice models in which past choices affect future ones through specific mechanisms (e.g., experience, habits, addiction, or switching costs) (Keane and Wolpin, 1997; Gruber and Kőszegi, 2001; Dubé et al., 2010; Handel, 2013). Another strand develops (linear) panel data methods to estimate causal effects of lagged outcomes on current outcomes (Anderson and Hsiao, 1982; Arellano and Bond, 1991; Blundell and Bond, 1998; Arellano and Honoré, 2001; Bonhomme, 2025). We contribute to this literature by studying identification in a dynamic choice model using temporary experiments, and by interpreting the assumption behind the instruments of the classic panel data methods under that model.

The second area of literature to which we contribute studies whether temporary interventions have lasting effects (Allcott and Rogers, 2014; Charness and Gneezy, 2009; Royer et al., 2015; Schaner, 2018; Celhay et al., 2019; Cefala et al., 2024). We establish that results from such experiments not only provide tests of whether effects last but also identify flexible models of dynamic choice, allowing researchers to study why observed choices persist and to answer a variety of counterfactual questions. We show that the rate of decay of treatment effects after an experiment's end identifies propagation, while the path of treatment effects leading up to and immediately after the experiment's end decomposes propagation into reinforcement and survival, and identifies patience.

Finally, our paper contributes to a literature on the gap between short-run and long-run responses to incentives. Early studies of the NIT experiments argued that responses to the temporary experiments would differ from responses to a permanent NIT, both because households knew the experiments would end (Metcalf, 1974; Burtless and Greenberg, 1982) and because responses build slowly (Robins and West, 1980). More recent studies have considered various other mechanisms to seek to reconcile short-run (micro) and long-run (macro) estimates of labor supply elasticities (Chetty et al., 2011b; Keane and Rogerson, 2015), including optimization frictions (Chetty et al., 2011a; Chetty, 2012), human capital accumulation (Imai and Keane, 2004), and returns to effort that arrive only at job switches (Kleven et al., 2025). Our framework recovers the long-run effect of a permanent policy only from the path of treatment effects in a temporary experiment. The model allows both for anticipation of the experiment's end and for a response that builds as past choices accumulate, while remaining agnostic about the exact mechanism behind the gap.

The rest of the paper is organized as follows. Section 2 develops the model and studies responses to temporary and permanent interventions. Section 3 discusses identification of propa-

gation, reinforcement, survival, and patience. Section 4 applies our framework to labor supply in the SIME/DIME experiment. Section 5 applies it to the bonus experiment in Allcott et al. (2022). Section 6 applies it to the experiment in Charness and Gneezy (2009). Section 7 proposes experimental designs that better identify patience and test the model's assumptions. Section 8 concludes.

2 Model

We develop a model to study dynamic responses to temporary and permanent interventions.

2.1 A General Dynamic Choice Problem

In each period t, an individual chooses y_t, the choice of interest, and consumption of a numeraire good c_t. The history of past choices of y accumulates in a 'choice stock' S_t.

The individual has quasilinear flow preferences over the choice y_t, the stock S_t, and the numeraire c_t. Flow utility in period t is $u(y_t, S_t) + c_t$. The choice has a stock-independent per-unit opportunity cost p_t and generates a stock-dependent return $R(y_t, S_t)$, both measured in units of the numeraire. The per-period budget constraint is $c_t = m_t + R(y_t, S_t) - p_t y_t$, where m_t denotes resources available for numeraire consumption. The choice stock depreciates at rate $\delta \in (0, 1]$ (a share $1 - \delta$ of the choice stock *survives* into the next period): $S_{t+1} = (1 - \delta)S_t + y_t$.

The individual knows the current stock S_t and its law of motion, their preferences u and returns R, and current and past costs $\{p_s\}_{s \leq t}$. They form expectations over the unknown path of future costs $\{p_s\}_{s>t}$.[4] Let $\mathcal{F}_t$ denote the individual's information at the start of period t, which includes current and past costs (and possibly other signals about future costs) and accumulates over time. Let $\mathcal{E}_t$ denote expectation under the individual's subjective beliefs given $\mathcal{F}_t$. These beliefs need not coincide with the true data-generating process of costs.

The individual maximizes expected utility discounted by the factor $\beta \in [0, 1)$, which we call *patience*. By substituting the budget constraint into utility and defining $U(y, S) \equiv u(y, S) + R(y, S)$, we write the individual's problem as:

$$\max_{\{y_t\}_t} \sum_{t=0}^{\infty} \beta^t \mathcal{E}_0 \left[U(y_t, S_t) - p_t y_t + m_t \right] \quad \text{s.t.} \quad S_{t+1} = (1 - \delta)S_t + y_t, \quad S_0 \text{ given.} \tag{1}$$

Problem (1) is a discrete-time version of the model in Becker and Murphy (1988), extended to allow the stock to affect the return to the current choice and to allow for uncertainty over future costs. We assume that U is strictly concave in (y, S).

A key feature of the model is that the stock may affect the marginal payoff to the current choice ($U_{yS} \equiv \partial^2 U / \partial y \partial S \neq 0$), and thus past choices shape future ones. To see this, consider the

[4] Because utility is linear in the numeraire, resources m_t do not affect the choice y_t. We therefore leave m_t unrestricted, both in its level and in what the individual knows about it.

individual's optimal choice: let $y_t^\star = g_t\big(S_t; \mathcal{F}_t\big)$ denote the optimal policy function defined by the solution to (1), which may be nonlinear and time dependent.[5] If the individual's optimal choice today depends on S_t, then their past choices affect their current behavior. Let $\phi_t \equiv \frac{\partial g_t(S_t;\mathcal{F}_t)}{\partial S_t}$ denote the slope of g_t. Past choices may *reinforce* current choices ($\phi_t \neq 0$) for two reasons. First, the level of the stock affects the marginal payoff to the choice if $U_{yS} \neq 0$. For an individual with no patience, this channel is the only source of reinforcement: $\phi_t = -U_{yS}/U_{yy}$ when $\beta = 0$. Second, the choice affects future stocks and thus future payoff. A patient individual accounts for this effect when making their choice.

2.2 A Concrete Example: Labor Supply

One concrete model that fits the structure above is a standard model of labor supply with search costs that fall with work history (Hyslop, 1999; Eckstein and Lifshitz, 2011).[6] In this model, the individual is a worker, the choice y_t is hours worked in week t, and the stock S_t is accumulated work history up to that week.

The worker has flow utility $c_t - v(y_t)$, where v is the disutility of work, and earns a wage w per hour. For each hour worked, the worker incurs a history-independent *cost of working* p_t (e.g., commuting costs) and a history-dependent search cost $\gamma(S_t)$. This search cost is the money value of the time and effort it takes to find and arrange hours of work: contacting employers, arranging shifts, and rearranging other activities, such as child care, around work. The cost weakly falls with work history ($\gamma'(S) \leq 0$) because a worker with more work history has more of these arrangements already in place, so fewer must be made anew. The worker's budget constraint is therefore $c_t = m_t + \big(w - p_t - \gamma(S_t)\big)y_t$. The worker knows S_t, v, γ, w, and current and past costs of working, and forms expectations about future ones given information $\mathcal{F}_t$, which may include, for example, planned roadwork that will lengthen the commute.

The worker thus solves (1) with $u(y, S) = -v(y)$ and $R(y, S) = \big(w - \gamma(S)\big)\, y$, so that $U_{yS} = -\gamma'(S) \geq 0$: work history weakly raises the payoff to an extra hour by lowering its search cost. The worker's optimal hours $y_t^\star = g_t(S_t; \mathcal{F}_t)$ depend on work history for the two reasons above: work history can make an extra hour less costly, and a patient worker accounts for how this week's hours affect future search costs.[7]

More generally, work history may change the marginal payoff to an hour of work through preferences, u_{yS} (e.g., work habits), or through returns, R_{yS} (e.g., experience premia or search

[5] A unique optimal policy function exists if, for example, the choice lies in a bounded interval, costs are bounded and follow a Markov process under the individual's beliefs, and U is continuous and strictly concave (Stokey and Lucas, 1989, Theorems 9.6 and 9.8). It is differentiable in S if, in addition, U is twice continuously differentiable and strongly concave and the bounds on the choice do not bind (Blume et al., 1982; Santos, 1991).

[6] In the models of Hyslop (1999) and Eckstein and Lifshitz (2011), returning to work is harder after time out of work. In Hyslop (1999), a woman who did not work in the previous year must pay a search cost to receive a job offer; in Eckstein and Lifshitz (2011), the chance of receiving a job offer rises with employment in the previous year and with the number of years worked.

[7] The payoff is strictly concave in (y, S) if $v''(y)\gamma''(S)\, y > \gamma'(S)^2$.

costs); Appendix E.1.1 considers a life-cycle labor-supply model with human capital accumulation.

2.3 Other Examples

The general model above captures a feature common to many other dynamic problems in economics: choices made today leave behind lasting consequences that affect future payoffs and, in turn, future choices. The following examples illustrate this common structure in two additional settings.

Social media use. Research on addictive goods, from cigarettes (Becker et al., 1994) to social media (Allcott et al., 2022), studies how past consumption reinforces current consumption. Consider an individual choosing how much time to spend on social media platforms. Here y_t is hours in week t and S_t is the accumulated use. This stock changes the marginal payoff to current use through $u(\cdot)$ (e.g., habits or network formation making platform use more valuable) or $R(\cdot)$ (e.g., accumulated followers increasing earnings from platform use). p_t is the opportunity cost of an hour spent using the platform.

Physical exercise. Dynamic responses to incentives have also been studied in health behaviors (Celhay et al., 2019; Hussam et al., 2022), including exercise (Charness and Gneezy, 2009; Royer et al., 2015). Consider an individual choosing weekly exercise hours. Here y_t is hours spent exercising in week t and S_t is the accumulated 'fitness'. Fitness changes the marginal payoff to exercise through $u(\cdot)$ (e.g., fitness making exercise more enjoyable) or $R(\cdot)$ (e.g., employer or insurer incentives for fitness). p_t captures pecuniary costs or incentives, such as travel costs or incentives to attend the gym.

2.4 Policy Interventions and Propagation of Past Choices

We now study how past choices shape future ones in the model. To do so, we introduce a temporary policy intervention that changes today's cost, and with it today's choice. The change in today's choice affects, or *propagates to*, all future choices. We show that propagation combines two features of behavior, reinforcement and survival.

Policy interventions. We study a class of interventions that change costs. For concreteness, we focus on the labor supply example of the model, and use a transportation subsidy as the example of such an intervention in this section and the next. Let $\{p_t\}$ denote the path of costs absent any intervention. The intervention changes costs by b in a window $\mathcal{W} \equiv \{\underline{t}, \ldots, \bar{t}\}$ (where $\bar{t}$ may be infinite), to

$$\tilde{p}_t \equiv p_t - b\mathbf{1}[t \in \mathcal{W}]. \tag{2}$$

The intervention starts at the beginning of week $\underline{t}$, after work history $S_{\underline{t}}$ is inherited and before hours that week are chosen. By week $\underline{t}$, the worker knows b and $\mathcal{W}$: their information becomes

$\tilde{\mathcal{F}}_t = \mathcal{F}_t \cup \{b, \mathcal{W}\}$. The worker then solves (1) with $\tilde{p}_t$ in place of p_t and expectations $\tilde{\mathcal{E}}_t$ given $\tilde{\mathcal{F}}_t$ in place of $\mathcal{E}_t$. We assume that the intervention does not change the worker's beliefs about baseline costs $\{p_t\}$ or about future interventions.

Propagation. Suppose a one-week subsidy in week $\underline{t}$ ($\mathcal{W} = \{\underline{t}\}$) raises hours worked in that week. Because the subsidy tells the worker nothing about later costs, the additional hours affect later hours only through their effect on work history. Let $dS_{\underline{t}+k}$ denote the change in work history at week $\underline{t}+k$, $k > 0$, induced by the subsidy. The additional work history increases optimal hours in week $\underline{t}+k$ by $\phi_{\underline{t}+k} dS_{\underline{t}+k}$. Hence, the increase in work history in week $\underline{t}+k+1$ is

$$dS_{\underline{t}+k+1} = (1-\delta) dS_{\underline{t}+k} + \phi_{\underline{t}+k} dS_{\underline{t}+k} = r_{\underline{t}+k} dS_{\underline{t}+k}, \qquad r_{\underline{t}+k} \equiv (1-\delta) + \phi_{\underline{t}+k}.$$

Changes in the stock thus propagate at a rate $r_{\underline{t}+k}$.

This *propagation* rate combines two channels. The first, $1-\delta$, is the *survival* channel: part of the additional work history survives mechanically into the next week. The second, $\phi_{\underline{t}+k}$, is the *reinforcement* channel: the additional work history changes hours in week $\underline{t}+k$, and those induced hours become part of work history in week $\underline{t}+k+1$. Iterating this argument, we obtain the (impulse) response of hours k weeks later to the initial increase in hours:

$$\theta_{\underline{t}}^{(k)} \equiv \left. \frac{dy_{\underline{t}+k}^\star / db}{dy_{\underline{t}}^\star / db} \right|_{S_{\underline{t}},\, b=0,\, \mathcal{W}=\{\underline{t}\}} = \phi_{\underline{t}+k} \prod_{j=1}^{k-1} r_{\underline{t}+j}. \tag{3}$$

2.5 Dynamic Responses to Temporary and Permanent Interventions

Two intervention counterfactuals motivate our analysis. First, how many additional hours worked does a temporary intervention induce after it ends? Second, how much larger would the effect of an intervention be if it were implemented permanently rather than temporarily?

Temporary interventions. Suppose first that the intervention is a subsidy in place for only week $\underline{t}$, so $\mathcal{W} = \mathcal{W}^{\text{TEMP}} \equiv \{\underline{t}\}$. In that week, hours respond to the lower cost alone, since neither work history nor future expectations have changed.[8] We define the *cumulative lifetime effect* as the number of additional hours the subsidy generates after it ends, per additional hour worked in that week:

$$\text{CLE}_{\underline{t}} \equiv \frac{\sum_{k=1}^{\infty} \left. \frac{dy_{\underline{t}+k}^\star}{db} \right|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{TEMP}}}}{\left. \frac{dy_{\underline{t}}^\star}{db} \right|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{TEMP}}}} = \sum_{k=1}^{\infty} \theta_{\underline{t}}^{(k)}.$$

[8] Appendix E.1.2 extends the analysis in this section to temporary interventions of other lengths, normalizing instead by the response in the last week of the intervention; our applications use this extension.

The cumulative lifetime effect $\text{CLE}_{\underline{t}}$ answers the following question: given the hours response observed while a temporary subsidy is in place, how many additional hours does the subsidy generate after it ends?[9]

Permanent interventions. Suppose instead that the intervention is a permanent subsidy, so that $\mathcal{W} = \mathcal{W}^{\text{PERM}} \equiv \{\underline{t}, \underline{t}+1, \ldots\}$. The worker knows from week $\underline{t}$ onward that the cost of working will be lower by b than the worker would otherwise expect, in every future week. We define the *long-run multiplier* as the long-run increase in hours relative to the contemporaneous increase generated by a one-week subsidy:

$$\text{LRM}_{\underline{t}} \equiv \frac{\lim_{k\to\infty} \frac{dy^{\star}_{\underline{t}+k}}{db}\Big|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{PERM}}}}{\frac{dy^{\star}_{\underline{t}}}{db}\Big|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{TEMP}}}}.$$

The long-run multiplier $\text{LRM}_{\underline{t}}$ answers the following question: how much larger is the long-run effect of a permanent subsidy than the response to the same subsidy for one week?[10]

The two target parameters answer different questions about the same policy at two lengths. The first, $\text{CLE}_{\underline{t}}$, concerns the lasting effects of a temporary subsidy. The second, $\text{LRM}_{\underline{t}}$, concerns the difference in effects between temporary and permanent subsidies.

2.6 A Tractable Linear–Quadratic Specialization

To obtain a tractable characterization of the target parameters above, it is useful to impose more structure on preferences. Following Becker and Murphy (1988), we consider quadratic utility (which can be interpreted as a second order approximation to any smooth U):

$$U(y_t, S_t) = \alpha y_t + U_{yy}\frac{y_t^2}{2} + U_{yS} y_t S_t + \psi S_t + U_{SS}\frac{S_t^2}{2}, \tag{4}$$

where α is a *preference shifter*: a permanent difference across workers, such as ability or taste for leisure.[11] The benefit of quadratic utility is that it makes the policy function solving problem (1) affine in the stock and additively separable in expected costs.[12]

[9] A large experimental literature studies whether temporary incentives have lasting effects (e.g., Allcott and Rogers, 2014; Charness and Gneezy, 2009; Royer et al., 2015). Our model allows us to study not just whether effects last, but also how and why they propagate over time.

[10] Earlier analyses of the negative income tax experiments studied this question (Metcalf, 1973, 1974; Burtless and Greenberg, 1982). Our model allows for two reasons why the effects of temporary and permanent programs may differ: anticipation of the program's end and responses that build over time. We recover the long-run multiplier from the experimental effects alone (Section 3).

[11] The results below can be interpreted as approximations around a steady state in a non-quadratic model (see Appendix E.1.3 for formal details).

[12] In a model with a finite horizon, the optimal policy remains affine, but its coefficients vary with the date: the slope converges to the constant ϕ of Lemma 1 below exponentially as the end of the horizon recedes, so the results also describe long finite horizons (Appendix E.1.4).

Lemma 1. *Suppose (i) the cost process $\{p_t\}$ is uniformly bounded and (ii) $U_{yS} - (1-\delta)U_{yy} \neq 0$. Then, in the model of Section 2.1 with quadratic utility* (4)*, the optimal policy function for problem* (1) *is*

$$y_t^\star = g\big(S_t; \mathcal{F}_t\big) = \phi S_t + K\alpha + \pi_t,$$

where ϕ and K depend ***only*** *on the model parameters $(U_{yy}, U_{yS}, U_{SS}, \beta, \delta)$, and π_t is a function of the current cost p_t and expected future costs $\{\mathcal{E}_t[p_{t+k}]\}_{k\geq 1}$ at time t.*

Proof. For brevity, all proofs are collected in Appendix B. □

The affine form of the optimal policy has two useful implications. First, the marginal effect of inherited work history on current hours is constant: an additional unit of work history changes hours by ϕ, regardless of the worker's accumulated history or current information. Second, current and expected future costs of working shift labor supply through the *cost outlook* π_t, but do not change the sensitivity of hours to work history. Together, these features allow us to study propagation and answer counterfactual questions while leaving information in $\mathcal{F}_t$ and belief formation largely unrestricted.

Propagation and policy intervention targets. The form of the optimal policy function leads to simple, closed-form expressions for impulse responses and our two policy target parameters.

Proposition 1. *Suppose the conditions of Lemma 1 hold and let $r \equiv \phi + 1 - \delta$. Then the impulse response and the target parameters do not depend on $\underline{t}$ (we write them without the subscript). Namely,*

(i) The impulse response is $\theta^{(k)} = \phi r^{k-1}$ for $k \geq 1$.

(ii) If $|r| < 1$, the cumulative lifetime effect is

$$\text{CLE} = \frac{\phi}{\delta - \phi} = \frac{\phi}{1-r}.$$

(iii) If $|r| < 1$, the long-run multiplier is

$$\text{LRM} = \frac{\delta}{\delta - \phi}\frac{1-\beta(1-\delta)}{1-\beta r} = \frac{\delta}{1-r}\frac{1-\beta(1-\delta)}{1-\beta r}.$$

Part (i) of Proposition 1 shows that increases in work history propagate at a constant, time-independent rate r.

However, r does not by itself determine the effects of interventions with different durations. Parts (ii) and (iii) show that the target parameters depend also on how r splits into reinforcement, ϕ, and survival, $1-\delta$. For a given r, a larger ϕ raises the cumulative lifetime effect (part ii):

survival only carries the extra work history forward, while reinforcement turns that work history into hours.

Survival $1-\delta$ also matters for the long-run multiplier (part iii). In the long run, a permanent subsidy increases both hours and work history. A fraction δ of the work history decays each week, so hours must stay high to keep the stock stable. When survival is weaker, more work history must be rebuilt each week to sustain the long-run increase in hours. For a given propagation rate r, this makes the long-run increase in hours larger relative to the contemporaneous response, raising the long-run multiplier.

The long-run multiplier depends, in addition, on patience. A patient worker responds more to a permanent subsidy than to a temporary one when work history raises future hours. The factor

$$\frac{1-\beta(1-\delta)}{1-\beta r}$$

captures this difference between the initial responses to a permanent and a temporary subsidy. The factor equals one with no patience ($\beta = 0$).

Survival and reinforcement also determine how quickly a permanent subsidy approaches its long-run effect (See Appendix Proposition B.1). The ratio ϕ/δ determines the relative size of the first jump: in its first week, a permanent subsidy delivers a share $1-\phi/\delta$ of its long-run effect. The propagation rate then determines how fast the remaining gap closes: each week, a share $1-r$ of the gap closes.

Even in the convenient linear–quadratic model, answering counterfactual policy questions requires knowledge of how choices propagate, how that propagation decomposes into reinforcement and survival, and how patient individuals are. The next section studies identification of propagation, reinforcement, survival, and patience from data on choices.

3 Identification

We now consider what observed choice data reveal about the model's dynamic parameters. We first show that a regression of choices on lagged choices does not generally recover the propagation rate, but that an exogenous and excludable instrument for past choices does. We consider two sources of instruments and use the model to explain the economic content of the assumptions justifying each. We then study identification of reinforcement, survival, and patience using choice data in a temporary experiment.

3.1 From Model to Data

The researcher observes a panel of choices $\{Y_{it}\}_{i,t}$ for a population of individuals i over periods $t = 0, \ldots, T$. In labor supply, Y_{it} is worker i's hours in week t.[13] Individuals make choices as in Section 2 (Lemma 1):

$$Y_{it} \;=\; \phi_i S_{it} + K_i \alpha_i + \pi_{it}. \tag{5}$$

The coefficients ϕ_i and K_i depend on the parameter vector $\omega_i \equiv (U_{yy,i}, U_{yS,i}, U_{SS,i}, \beta_i, \delta_i)$, which we call worker i's *type*; types may differ across individuals but not over time. The individual's cost outlook in week t, π_{it}, is a weighted sum of current costs p_{it} and expected future costs $\{\mathcal{E}_{it}[p_{i,t+k}]\}_{k\geq 1}$ (see Appendix Equation (21)), where $\mathcal{E}_{it}$ denotes individual i's (subjective) expectation given their information set $\mathcal{F}_{it}$. This cost outlook is more favorable when costs are lower: a lower current cost raises π_{it}, and with it the choice. We write $\mathbb{E}[\cdot]$ and $\mathbb{E}[\cdot \mid \cdot]$ for expectations under the population distribution of individuals and their costs (which need not coincide with $\mathcal{E}_{it}$).[14]

The researcher does not observe the stocks S_{it} (accumulated work histories).[15] We use the law of motion $S_{it} = (1-\delta_i)S_{i,t-1} + Y_{i,t-1}$ to eliminate the stock in Equation (5), leaving

$$Y_{it} \;=\; r_i Y_{i,t-1} \;+\; \underbrace{\delta_i K_i \alpha_i}_{\equiv a_i} \;+\; \underbrace{\pi_{it} - (1-\delta_i)\pi_{i,t-1}}_{\equiv e_{it}}. \tag{6}$$

Equation (6) expresses the reduced form relationship between the current and lagged choices. Note that the coefficient $r_i = \phi_i + (1-\delta_i)$ on $Y_{i,t-1}$ does *not* equal the one-period impulse response $\theta_{it}^{(1)} = \phi_i$. The latter holds the stock fixed (Equation (3)), while Equation (6) does not. Heterogeneity in stable preferences enters (6) through the individual effect $a_i \equiv \delta_i K_i \alpha_i$. Current and expected future costs enter (6) through the cost outlooks π_{it} in the composite errors $e_{it} \equiv \pi_{it} - (1-\delta_i)\pi_{i,t-1}$. In labor supply, $a_i = \delta_i K_i \alpha_i$ captures workers' stable preferences, such as ability or taste for leisure, while e_{it} captures current and expected costs of working, such as childcare and commuting costs.

[13] We index periods from the first one observed, so the panel's $t = 0$ need not coincide with the first choice period $t = 0$ of Section 2. In particular, observing period 0 does not mean observing the initial stock S_{i0}. None of our identification results depend on when the panel starts.

[14] The researcher may also observe covariates, though we leave these implicit. The assumptions below can be read as conditional on observed covariates.

[15] Stocks cannot be constructed from observed choices because depreciation rates δ_i and initial stocks S_{i0} are unknown. Even if stocks S_{it} were observed, their realized variation would not generally identify reinforcement ϕ_i (even if $\phi_i = \phi$). The stock is generated by past choices. Since those choices themselves depend on α_i and on past π_{it}, individuals do not arrive at different stocks holding fixed the other determinants of their choices. A regression of Y_{it} on realized S_{it} therefore compares individuals or periods with systematically different stable preferences and cost outlooks. It need not recover the derivative $\partial g/\partial S$ that defines ϕ_i.

3.2 Anatomy of the Selection Problem

An ordinary least squares (OLS) regression of Y_{it} on $Y_{i,t-1}$ does not recover a weighted average of propagation rates r_i:

$$r^{\text{OLS}} \equiv \frac{\text{Cov}(Y_{i,t-1}, Y_{it})}{\text{Var}(Y_{i,t-1})} = \underbrace{\frac{\text{Cov}(Y_{i,t-1}, r_i Y_{i,t-1})}{\text{Var}(Y_{i,t-1})}}_{\substack{\text{weighted average of} \\ \text{propagation rates}}} + \underbrace{\frac{\text{Cov}(Y_{i,t-1}, a_i)}{\text{Var}(Y_{i,t-1})}}_{\substack{\text{stable heterogeneity} \\ \text{bias}}} + \underbrace{\frac{\text{Cov}\left(Y_{i,t-1}, e_{it}\right)}{\text{Var}(Y_{i,t-1})}}_{\substack{\text{cost-outlook} \\ \text{selection}}}. \tag{7}$$

The last two terms are nonzero because lagged choices depend on stable preferences and lagged cost outlooks.

Stable heterogeneity bias. The first bias arises because individuals differ in their stable preferences (a_i). An OLS regression attributes persistence from stable preferences to propagation from past choices. This is the classical identification challenge described by Heckman (1981). A larger a_i shifts the individual's entire choice path up, biasing r^{OLS} upward (unless preferences are strongly negatively correlated with cost outlooks). In labor supply, workers with a lower taste for leisure supply more hours in consecutive weeks, biasing OLS upward.

Cost-outlook selection. The second bias arises through the selection induced by current and lagged cost outlooks. This occurs in two ways. The first is direct selection: individuals with a more favorable lagged cost outlook (a higher $\pi_{i,t-1}$) have a higher lagged choice (Equation (5)) but, for a given lagged choice, a lower current choice (Equation (6)). The second is indirect selection: when costs are correlated over time or anticipated, individuals with a more favorable lagged cost outlook (and thus a higher lagged choice) also tend to have a more favorable current cost outlook, π_{it}, and hence a higher current choice.[16] An OLS regression attributes this dependence to propagation. In labor supply, the lagged work hours $Y_{i,t-1}$ depend on the lagged commuting cost outlook $\pi_{i,t-1}$. Because this outlook depends on expected future commuting costs, it will generally be correlated with the current cost outlook π_{it}.

The sign of this bias depends on the autocorrelation structure of $\{\pi_{it}\}$ and the depreciation rate δ_i, and is ambiguous in general. When costs are neither correlated nor anticipated, $\pi_{i,t-1}$ is unrelated to π_{it}. The only source of bias in this case is the direct selection, which generates a downward bias: the lagged choice depends positively on $\pi_{i,t-1}$ while e_{it} depends on it negatively. When costs are correlated over time or anticipated, the indirect selection may push the bias in either direction.[17]

[16] Anticipation matters even when costs are unrelated over time, because π_{it} depends on expected future costs as well as on the current one. A worker who learns this week that her commuting cost will rise next month adjusts her hours in every week until then, so $\pi_{i,t-1}$ and π_{it} respond to the same news.

[17] To illustrate, let propagation and depreciation be common across individuals and suppose each individual's costs follow an AR(1) process with known autocorrelation λ ($|\lambda| < 1$). For example, a persistent commuting cost means that a high cost this week predicts a high cost next week, and hence low hours in both weeks. When this persistence exceeds the survival of work history, $\lambda > 1 - \delta$, the bias is upward: hours look more persistent than

Heterogeneous propagation. The first term in (7) raises a different issue, one of aggregation rather than selection. If r_i varies across workers, then even if the two bias terms in (7) were 0, the OLS regression would recover a non-simple weighted average of the r_i.

Propagation is common ($r_i = r$ for all i) when all individuals have the same type ($\omega_i = \omega$ for all i). The two selection problems described above do not depend on whether propagation is common. We therefore study each instrumental-variables strategy below first under general heterogeneity and then under common propagation.

3.3 Identification of Propagation Using Instruments

The selection problem in Equation (7) suggests a familiar solution: find a source of variation in the lagged choice that is unrelated to the other determinants of current choices. Let h denote a known linear transformation of the choice equation and Z_{it} denote a candidate instrument. A valid instrument satisfies three conditions. First, it must be correlated with the lagged (transformed) choice (*relevance*). Second, it must be statistically unrelated to the (transformed) unobserved determinants of current choices, tastes and costs (*exogeneity*). Third, it must affect current choices only through its effect on the lagged choice (*exclusion*).

We consider two sources of instruments. *Internal* instruments use variation from the individual's further lagged choices, while *external* instruments use variation from (quasi-)experiments outside the individual's choice history. We use the model to explain the economic content of the exogeneity and exclusion assumptions in each case. Internal instruments are valid if changes in costs are neither persistent nor anticipatable, a strong assumption on costs, information, and beliefs. External instruments are valid if the variation is randomly assigned and changes costs and beliefs only temporarily. Many policy experiments satisfy both conditions by design.

3.3.1 Internal Instruments

A long literature studies identification of linear autoregressive models, such as Equation (6), from panel data (Anderson and Hsiao, 1982; Arellano and Bond, 1991; Arellano and Honoré, 2001; Wooldridge, 2010). Differencing Equation (6) removes the individual effect a_i, $\Delta Y_{it} = r_i \Delta Y_{i,t-1} + \Delta e_{it}$, but the lagged change $\Delta Y_{i,t-1}$ is correlated with Δe_{it}. The classic approach, first proposed by Anderson and Hsiao (1982), uses further lagged hours, $Y_{i,t-\ell}$ for some $\ell \geq 2$, as instruments for $\Delta Y_{i,t-1}$. In the notation above, h is the first difference operator ($h(Y_{i,t}) \equiv Y_{i,t} - Y_{i,t-1}$) and $Z_{it} \equiv Y_{i,t-\ell}$ for some $\ell \geq 2$. This subsection studies this classic approach, which we will refer to as the *internal instrument*, through the lens of the model in Section 2.

Exclusion. The internal instrument $Y_{i,t-\ell}$ satisfies exclusion as long as further lagged hours affect current hours only through work history, as in Section 2. Exclusion fails if further lagged hours

propagation alone implies. When commuting costs mean-revert faster than work history depreciates, $\lambda < 1 - \delta$, the bias is downward (Appendix Lemma E.2).

affect current hours through any other channel. For example, suppose working reveals information about future costs of working, so that $\mathcal{F}_{it}$ depends on $Y_{i,t-\ell}$. A change in $Y_{i,t-\ell}$ may then change Y_{it} by changing the expected future costs in the cost outlook π_{it}, violating exclusion.[18]

Exogeneity. The more substantive assumption for internal instruments is exogeneity. Exogeneity requires variation in further lagged hours to be statistically unrelated to the unobserved determinants of current work hours (after differencing out the individual effect a_i). Under the model, exogeneity restricts the process for the cost of working. A sufficient condition is that changes in costs be unforeseeable.

Assumption I.1 (Unpredictable cost innovations)**.** *There exists an individual-specific, $\mathcal{F}_{i0}$-measurable level $\bar{p}_i$ such that costs satisfy*

$$p_{it} = \bar{p}_i + \eta_{it}, \qquad \mathbb{E}[\eta_{it} \mid \mathcal{F}_{i,t-1}] = \mathcal{E}_{i,t-1}[\eta_{it}] = 0 \quad \forall t.$$

Assumption I.1 restricts all unobservables that enter the cost of working. The assumption says that the deviation $\eta_{it} \equiv p_{it} - \bar{p}_i$ from the worker's average cost $\bar{p}_i$ cannot be predicted from the worker's information at time $t-1$, under both the objective expectation $\mathbb{E}[\cdot \mid \mathcal{F}_{i,t-1}]$ and the subjective expectation $\mathcal{E}_{i,t-1}$.[19] The assumption rules out, for example, a rise in commuting costs that lasts several weeks, or a fare increase announced a month in advance. Conditioning on covariates makes the assumption more credible only to the extent that covariates explain all persistent and predictable movements of costs.

Assumption I.1 fails when unobserved costs persist or are anticipated. In this case, the sign of the bias of estimates based on further lagged hours as instruments is ambiguous in general. The reason for the ambiguity is that further lagged hours $Y_{i,t-\ell}$ are correlated with the differenced error Δe_{it}, so they carry a bias similar to that induced by cost-outlook selection (Section 3.2). Appendix Lemma E.2 shows that when the cost deviations η_{it} follow an AR(1) process, the sign of this bias, like that of the cost-outlook selection in OLS, depends on the autocorrelation structure of costs and on the depreciation rate.

What internal instruments deliver. Under Assumption I.1, further lagged hours are valid instruments.

Proposition 2 (Internal instruments)**.** *Suppose Assumption I.1 holds, and let $t \geq 3$.*

[18] Exclusion also fails if there is a second stock that fades at a different rate. For example, suppose eligibility for subsidized childcare depends only on hours worked last week. Eligibility then acts as a second stock that raises the return to working and depreciates fully each period ($\check{S}_{i,t+1} = Y_{it}$). Current hours then depend on both stocks, and eliminating work history as in Equation (6) leaves last week's eligibility, and hence hours two weeks earlier, in the error e_{it}. Further lagged hours are therefore correlated with the error through a channel other than work history, so they are not excludable.

[19] Past research has noted that the assumption behind lagged instruments restricts what individuals know (Abbring and Heckman, 2007; Bonhomme, 2025). Relative to them, we give a sufficient condition in terms of the economic primitives of our model.

(i) $\mathrm{Cov}(Y_{i,t-\ell}, \Delta e_{it}) = 0$ *for every* $3 \leq \ell \leq t$*. If, in addition,* $\delta_i = 1$ *for all* i*, the same holds for* $\ell = 2$*.*

(ii) For every $\ell \geq 3$ *(every* $\ell \geq 2$ *if* $\delta_i = 1$ *for all* i*) such that* $\mathrm{Cov}(Y_{i,t-\ell}, \Delta Y_{i,t-1}) \neq 0$*, the instrumental-variables (IV) estimand using* $Y_{i,t-\ell}$ *as an instrument for* $\Delta Y_{i,t-1}$ *is a weighted average of the propagation rates:*

$$r^{\text{int}} \equiv \frac{\mathrm{Cov}(Y_{i,t-\ell}, \Delta Y_{it})}{\mathrm{Cov}(Y_{i,t-\ell}, \Delta Y_{i,t-1})} = \frac{\mathrm{Cov}(Y_{i,t-\ell}, r_i\, \Delta Y_{i,t-1})}{\mathrm{Cov}(Y_{i,t-\ell}, \Delta Y_{i,t-1})} = \mathbb{E}[r_i W_i^{\text{int},\ell}], \tag{8}$$

where $W_i^{\text{int},\ell} \equiv \frac{(Y_{i,t-\ell} - \mathbb{E}[Y_{i,t-\ell}])\, \Delta Y_{i,t-1}}{\mathbb{E}[(Y_{i,t-\ell} - \mathbb{E}[Y_{i,t-\ell}])\, \Delta Y_{i,t-1}]}$*. If* $r_i = r$ *for all* i*, the estimand equals* r*.*

Part (i) of Proposition 2 establishes the familiar Anderson and Hsiao (1982) and Arellano and Bond (1991) sequential exogeneity condition. Part (ii) shows that the IV estimand based on this moment condition recovers a weighted average of propagation rates r_i. Workers whose further lagged hours are farther from the mean and workers with larger week-on-week changes in hours receive weights of larger magnitude. If propagation is common and several lags are relevant, then IV estimands based on different lags should all recover the same r, yielding overidentifying restrictions.

3.3.2 External Instruments

A second strategy uses a source of variation in work hours that comes from outside the individual's own choice history. Such external instruments are used widely in applied macroeconomics (Mertens and Ravn, 2013; Stock and Watson, 2018) and in some microeconomic studies of persistence.[20]

To illustrate the strategy, consider the transportation subsidy from Section 2.4. Let $D_i \in \{0, 1\}$ denote worker i's *treatment status*, which indicates whether the worker faces the subsidy. Control workers ($D_i = 0$) face the cost of working p_{it}, while treated workers ($D_i = 1$) face, through period $\bar{t}$,

$$\tilde{p}_{it} = p_{it} - b_i \mathbf{1}[t \in \mathcal{W}], \qquad \mathcal{W} = \{\underline{t}, \underline{t}+1, \ldots, \bar{t}\}, \tag{9}$$

where b_i is the reduction in worker i's cost of working while the subsidy is in place. The sign and magnitude of b_i may vary across workers i. Throughout, we assume that the panel covers the subsidy and at least two periods after it, $0 \leq \underline{t} \leq \bar{t} \leq T-2$. As in Section 2.4, $(\mathcal{F}_{it}, \mathcal{E}_{it})$ and $(\tilde{\mathcal{F}}_{it}, \tilde{\mathcal{E}}_{it})$ denote the information sets and expectations without and with the subsidy, respectively. Let TE_t denote the difference in mean hours worked in period t between treated and control workers:

$$\text{TE}_t \equiv \mathbb{E}[Y_{it} \mid D_i = 1] - \mathbb{E}[Y_{it} \mid D_i = 0]. \tag{10}$$

We consider using treatment status D_i as an instrument for lagged hours $Y_{i,t-1}$ in Equation

[20] Becker et al. (1994) use past cigarette prices and taxes as external instruments for past cigarette consumption.

(6) after the subsidy has ended; that is, $Z_{it} \equiv D_i$ and $t-1 > \bar{t}$. (For $t-1 \leq \bar{t}$, D_i is not excluded because it directly affects $\pi_{i,t-1}$, which enters e_{it}.)

Exclusion. Exclusion requires that the subsidy affect hours after period $\bar{t}$ only through the additional hours worked while the subsidy was in place. The subsidy must not change current or expected future costs after period $\bar{t}$.

Assumption X.1 (Exclusion)**.** *For every $t > \bar{t}$ and $k \geq 1$,*

$$\tilde{p}_{it} = p_{it} \qquad \textit{and} \qquad \tilde{\mathcal{E}}_{it}[\tilde{p}_{i,t+k}] = \mathcal{E}_{it}[p_{i,t+k}]\,.$$

The first requirement of Assumption X.1 may fail if, for example, the subsidy encouraged workers to move farther from work, creating a lasting change in commuting costs. The second requirement may fail if, for example, the subsidy leads treated workers to believe that they will receive more subsidies in the future.

Assumption X.1 restricts only how the transportation subsidy changes costs and expectations after the subsidy ends. Unlike Assumption I.1, Assumption X.1 allows workers to face persistent and predictable costs absent the subsidy. Assumption X.1 also leaves unrestricted what workers know or expect about the subsidy before it begins and while it is in place.

Exogeneity. Exogeneity requires treatment status to be unrelated to the unobserved determinants of hours. Exogeneity fails if treatment status depends on workers' preferences or costs (e.g., if the subsidy is only given to workers in neighborhoods with long commutes). A sufficient condition for exogeneity is random assignment, as in an experiment.[21]

Assumption X.2 (Random assignment)**.**

$$D_i \perp \left(S_{i0}, \alpha_i, \omega_i, \{p_{it}, \mathcal{F}_{it}, \{\mathcal{E}_{it}[p_{i,t+k}]\}_{k\geq 1}\}_{t=0}^{T}\right).$$

Random assignment ensures that treatment status is not systematically related to preferences, initial stocks, or the cost process, information, and beliefs that workers would have without treatment.[22] Unlike Assumption I.1, it places no restriction on whether the path of the cost of working is persistent or predictable. Conditioning on covariates makes Assumption X.2 more credible only to the extent that covariates explain all systematic differences in who is treated; in an experiment whose assignment probabilities vary with covariates, conditioning on those covariates makes the assumption hold by design.

What external instruments deliver. Define type-specific treatment effects $\text{TE}_t(\omega) \equiv \mathbb{E}[Y_{it} \mid D_i = 1, \omega_i = \omega] - \mathbb{E}[Y_{it} \mid D_i = 0, \omega_i = \omega]$. Under random assignment, $\text{TE}_t = \mathbb{E}[\text{TE}_t(\omega_i)]$. Let $r(\omega)$

[21] Appendix E.2.2 considers an alternative condition based on selection on permanent differences.

[22] The assumption allows for common shocks (such as recessions) that affect the labor supply of all workers, as long as these shocks do not differentially affect treated and control workers.

denote the propagation rate for workers with $\omega_i = \omega$. Under Assumptions X.1 and X.2, after the subsidy ends, each type's treatment effect propagates at that type's rate:

$$\text{TE}_t(\omega) = r(\omega)\text{TE}_{t-1}(\omega) \qquad \text{for every } t-1 > \bar{t}. \tag{11}$$

This observation leads to the following result.

Proposition 3 (External instruments)**.** *Suppose Assumptions X.1 and X.2 hold. Then for every* $k \geq 0$*,*

$$\text{TE}_{\bar{t}+1+k} = \mathbb{E}\left[r_i^k \, \text{TE}_{\bar{t}+1}(\omega_i)\right]. \tag{12}$$

For every $k \geq 1$ *with* $\text{TE}_{\bar{t}+k} \neq 0$*, the IV estimand using* D_i *as an instrument for* $Y_{i,\bar{t}+k}$ *is*

$$r^{\text{ext}}_{\bar{t}+1+k} \equiv \frac{\text{TE}_{\bar{t}+1+k}}{\text{TE}_{\bar{t}+k}} = \mathbb{E}[r_i W_i^{\text{ext},k}] \tag{13}$$

where $W_i^{\text{ext},k} = \frac{\text{TE}_{\bar{t}+k}(\omega_i)}{\mathbb{E}\left[\text{TE}_{\bar{t}+k}(\omega_i)\right]}$*. If* $r_i = r$ *for all* i*, then* $r^{\text{ext}}_{\bar{t}+1+k} = r$*.*

Proposition 3 shows that each ratio of successive post-subsidy treatment effects is a weighted average of propagation rates, in which each type's weight is its treatment effect in the previous period. These weights are all nonnegative when propagation rates are nonnegative and the treatment effect in the first post-subsidy week has the same sign for every type, i.e., when the subsidy raises hours for all types or lowers hours for all types.[23] Equivalently, the ratio of the treatment effect in any future period to the first post-subsidy effect, $\text{TE}_{\bar{t}+1+k}/\text{TE}_{\bar{t}+1}$, is the k-th moment of the distribution of propagation rates, weighted by the first post-subsidy effects.

Two implications follow from (12) (see Appendix E.2.3 for further discussion). First, the path of post-subsidy treatment effects reveals whether propagation is nonzero. If $r_i = 0$ for every worker, then every term with $k \geq 1$ in (12) vanishes, so all treatment effects from period $\bar{t}+2$ on must be zero. A nonzero effect two or more periods after the subsidy ends rules out this possibility. (A nonzero effect at $\bar{t}+1$ alone does *not* rule out $r_i = 0$, since the stock built while the subsidy was in place can affect period $\bar{t}+1$ choices even when $r_i = 0$.) Second, the shape of the path reveals whether propagation is common. When the weights $W_i^{\text{ext},k}$ are nonnegative, they shift toward types with higher propagation rates as k increases. Thus, the ratios of adjacent effects rise with time unless propagation is common among all types the subsidy moves.

Common propagation. Under common propagation, the weights no longer matter, and the post-subsidy treatment effect path is geometric:

$$\text{TE}_{t+1} = r \, \text{TE}_t \qquad \text{for every } t \geq \bar{t}+1. \tag{14}$$

[23] This condition is similar in spirit to the monotonicity assumption of Imbens and Angrist (1994): different types may respond by different amounts, as long as all respond in the same direction.

In this case, Proposition 3 establishes that post-subsidy treatment effects fade out at the propagation rate r. Every period after the subsidy ends identifies the same r, so the ratios in (13) must coincide across periods, which yields overidentifying restrictions. A violation indicates that exclusion or exogeneity fails or that propagation is heterogeneous.

3.3.3 Discussion

The two IV strategies leverage different sources of variation in past choices. Internal instruments use variation from the individual's own choice history. External instruments use variation from shocks to costs, such as the transportation subsidy in Equation (9). The two strategies recover different weighted averages of propagation rates.[24]

The instruments also differ in the assumptions that justify their validity. For internal instruments, validity is a matter of judgment: even after conditioning on the costs the researcher observes, all remaining changes in costs must neither persist nor be anticipated. For external instruments, validity is a matter of design: the researcher who implements the (quasi-)experiment chooses how the intervention is assigned and how long it lasts. To illustrate the difference, consider commuting costs. If commuting costs the researcher does not control for are persistent or can be anticipated by workers, internal instruments are invalid. In contrast, a randomly assigned transportation subsidy is a valid instrument as long as it changes the cost of working, and beliefs about it, only while it lasts.

A general external instrument identifies propagation under Assumptions X.1 and X.2. We show next that a specific type of external instrument also identifies reinforcement, survival, and patience.

3.4 Identification of Propagation Channels and Patience with Temporary Experiments

Our target parameters depend not only on the propagation rate, but also on its decomposition into reinforcement and survival and on patience (Proposition 1). Because each value of r is consistent with multiple values of ϕ, δ, and β, identifying these parameters requires additional variation or additional assumptions.

In this subsection, we consider identification of these parameters using the additional variation provided by *temporary (randomized) experiments*, a type of external instrument.[25] Section 3.3.2 used two features of a temporary experiment to identify propagation from the effects after it ends:

[24]Both strategies also furnish a test of full depreciation ($\delta_i = 1$ for all i). Under full depreciation, the stock is just last period's choice, so the error $e_{it} = \pi_{it}$ no longer contains the lagged cost outlook $\pi_{i,t-1}$. In that case, one more instrument becomes valid: the second lag for internal instruments (Proposition 2(i)), and D_i for the last treated period $\bar{t}$ for external ones (Appendix Proposition E.6).

[25]See Appendix Lemma B.5 for a more general treatment of identification. The basic logic does not depend on whether the variation is internal or external; Appendix Proposition E.5 gives an example of how survival and reinforcement can be identified using internal instruments under additional stationarity assumptions.

random assignment, and a treatment that changes the cost of working, and beliefs about it, only while it lasts. But such an experiment typically has two more features: the treatment is the same in every period it lasts, and treated workers know its size and when it ends. Together, these four features let the effects during the experiment identify survival, reinforcement, and patience. We develop the argument in the homogeneous case, $\omega_i = \omega$ for all i; Appendix E.2.5 studies the heterogeneous case.

3.4.1 Temporary Experiments

We denote by *temporary experiments* designs with five features: (i) individuals are randomly assigned to a treatment; (ii) the treatment changes costs, and individuals' expectations of future costs, only while it lasts; (iii) the treatment changes costs by the same amount in every period it lasts; (iv) treated individuals know this amount and when it ends, and learn nothing else about costs from treatment; and (v) choices are measured during and after the treatment. Field and laboratory experiments following this design are common in economics (List, 2011). Features (i) and (ii) are Assumptions X.2 and X.1, which we continue to assume. Features (iii) and (iv) motivate one additional assumption on what treated workers expect while the subsidy lasts:

Assumption X.3 (Known, constant intervention path)**.** *For every $t \in \mathcal{W}$ and every $k \geq 1$,*

$$\tilde{\mathcal{E}}_{it}[\tilde{p}_{i,t+k}] = \mathcal{E}_{it}[p_{i,t+k}] - b_i \mathbf{1}[t + k \in \mathcal{W}].$$

Assumption X.3 mirrors our assumption about the policy intervention in Section 2.4: treated workers take the subsidy at face value. It fails if, for example, they expect the subsidy to be extended, or read it as a sign that their commuting costs will change.

To be a temporary experiment, the transportation subsidy of Equation (9) must have the five features: (i) workers are randomly assigned to it; (ii)–(iii) it lowers their cost of working by b_i in every week of $\mathcal{W}$ and in no other week; (iv) treated workers are told its size and end date when it starts; and (v) hours are recorded during and after it. Such a design makes Assumptions X.1–X.3 credible.

3.4.2 Identifying Reinforcement, Survival, and Patience

Under Assumptions X.1 and X.2, the rate of decay of treatment effects *after* the experiment ends identifies the (common) propagation rate r (Proposition 3). With the addition of Assumption X.3, the path of effects *during* the experiment identifies δ, ϕ, and β.

Proposition 4 (Decomposing propagation with temporary experiments)**.** *Suppose $\omega_i = \omega$ for all i, Assumptions X.1, X.2, and X.3 hold, $\delta \in (0, 1)$, $\phi \neq 0$, and $\mathbb{E}[b_i \mid D_i = 1] \neq 0$.*

(i) If β is known and $\bar{t} - \underline{t} \geq 1$, then δ, and hence ϕ, is identified.

(ii) If $\bar{t} - \underline{t} \geq 2$, then β, δ, and ϕ are jointly identified.

Proposition 4(ii) shows that if the experiment lasts at least three periods ($\bar{t} - \underline{t} \geq 2$), then patience, survival, and reinforcement are all identified, and with them the cumulative lifetime effect and the long-run multiplier of Proposition 1. If patience β is known or assumed (as is common in empirical work estimating dynamic models), then a two-period experiment suffices to identify δ and ϕ (part (i)).[26]

To see why the result holds, take the difference between treated and control workers in Equation (6). Write $\text{TE}_t^{\pi} \equiv \mathbb{E}[\pi_{it} \mid D_i = 1] - \mathbb{E}[\pi_{it} \mid D_i = 0]$ for the *direct effect* of the subsidy on hours worked in period t, holding work history fixed. For every t,

$$\underbrace{\text{TE}_t - r\,\text{TE}_{t-1}}_{\text{propagated } \text{TE}_{t-1} \text{ residual}} = \underbrace{\text{TE}_t^{\pi}}_{\text{period's direct effect}} - \underbrace{(1-\delta)\,\text{TE}_{t-1}^{\pi}}_{\substack{\text{past direct effect that} \\ \text{does not carry over}}} . \tag{15}$$

Equation (15) is the effect of the subsidy on the error of Equation (6). Once r is identified, the left-hand side is observed. It is the part of today's treatment effect not explained by propagating yesterday's effect at rate r. The right-hand side is not observed, but for given δ and β, its path is known up to scale. The direct effects change while the subsidy is in place only with the number of subsidized periods that remain, because the subsidy is the same in every period and workers know when it will end. For an impatient worker ($\beta = 0$), the direct effect is the same in every subsidized period. For a patient worker ($\beta > 0$), an hour worked now also raises her hours in later subsidized periods, which are subject to the same subsidy. Her direct effect therefore changes as the end of the subsidy approaches, at a pace set by her patience β and reinforcement $\phi = r - (1 - \delta)$. The level of the direct effect also depends on the size of the subsidy and on how strongly a change in cost moves hours, but these two quantities form the scale, which is the same in every period. Comparing the left-hand side across periods removes this scale, and what remains reveals δ and β.

3.4.3 Intuition

Figure 1 illustrates how survival and patience shape the path of treatment effects. Survival determines the relative size of the two gaps between the observed effects and the dashed lines around the end of the subsidy (panel (a)); patience determines how the path departs, as the end approaches, from that of an impatient worker (panel (b)).

Identifying survival. In panel (a), the dashed lines propagate the effects at $\bar{t} - 1$ and $\bar{t}$ forward at rate r. While the subsidy is in place, the observed effects do not match the propagated effects, because of the two terms on the right-hand side of Equation (15). In the last subsidized period the

[26] Observing more treated periods overidentifies the model.

subsidy renews the direct effect, while the share $1-\delta$ of the previous direct effect does not carry over. We call the resulting departure from the dashed line the *interior gap*. For given patience, the interior gap grows with δ: propagation carries over a smaller share of the past direct effect, so less of the renewed direct effect is offset. Immediately after the subsidy ends, only the second term remains, so the effect falls below the dashed line; we call this difference the *boundary gap*. For given patience, the boundary gap shrinks with δ: propagation carries over a smaller share of the last direct effect, which in fact disappears. Thus, the relative size of the two gaps reveals δ.

Identifying patience. Panel (b) compares the observed path with that of an impatient worker (dashed). Her direct effect is the same in every subsidized period, so her effects keep converging toward the response to a permanent subsidy (dotted). A patient worker instead adjusts her responses as the end of the subsidy approaches. Because an hour she works while the subsidy is in place also raises her later subsidized hours ($\phi > 0$), her direct effect is larger early in the window and shrinks as the end approaches, until in the last subsidized period it equals that of an impatient worker. How fast it shrinks depends on her patience. As her direct effect shrinks, the treatment effects bend away from the dashed path before the subsidy ends: the more she values the future, the faster her direct effect shrinks and the more her path bends.

Figure 1: Identification of the propagation channels and patience with a temporary experiment

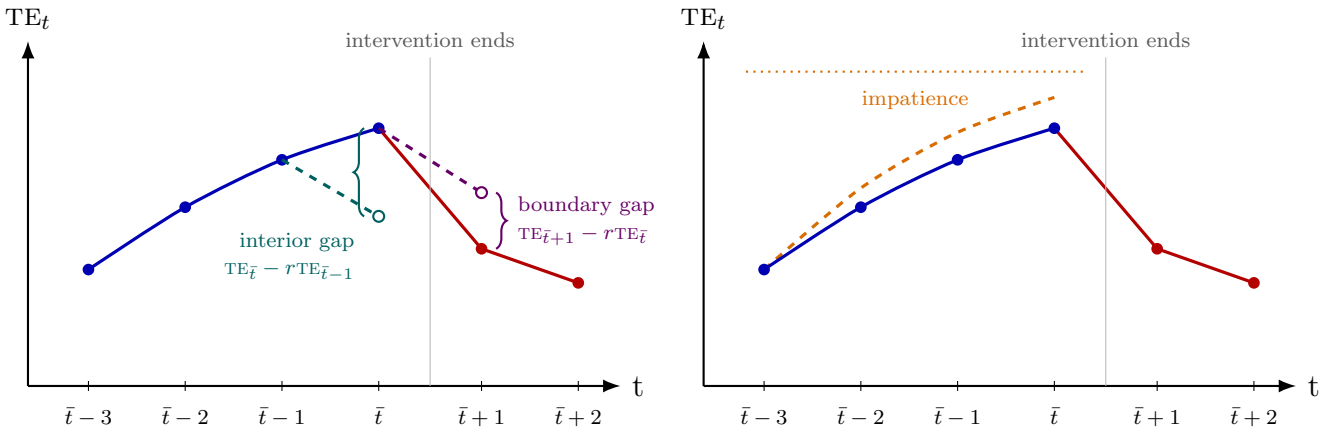

(a) Survival **(b)** Patience

Notes: This figure illustrates the identification of δ and β from an observed series of treatment effects. Blue denotes periods in which the intervention remains active and red denotes the post-treatment response. Panel (a) projects the effects at $\bar{t}-1$ and $\bar{t}$ one period forward at the propagation rate r (dashed lines, open circles). The interior and boundary gaps are the differences between the observed effects at $\bar{t}$ and $\bar{t}+1$ and these projections. Panel (b) compares the observed treatment-effect path with the response of an impatient worker (dashed orange); the dotted orange line denotes the long-run response to a permanent subsidy under impatience, toward which that response converges.

4 Labor Supply in the Negative Income Tax Experiments

We now take the framework to three temporary experiments and ask two questions in each: how much of the persistence in choice is caused by past choices, and what a permanent version of the temporary incentive would do.

We begin in this section by studying the Seattle–Denver Income Maintenance Experiment (SIME/DIME), the largest of the negative income tax (NIT) experiments of the 1970s (Christophersen, 1983). The experiments were designed to "represent the conditions of a fully implemented, permanent national program" (U.S. Department of Health and Human Services, Office of Income Security Policy, 1983, p.6). However, the experiments themselves were temporary: most participants received the NIT for up to three years and knew when it would end. Early studies gave two reasons why the response to the temporary program may differ from that to a permanent one: participants may respond differently to a program they know will end (Metcalf, 1973; Burtless and Greenberg, 1982), and responses may build gradually when past work affects future work (Robins and West, 1980). Our framework incorporates both reasons (Section 2) and identifies the structural parameters governing them from the path of treatment effects alone (Section 3).[27]

4.1 The Seattle–Denver Income Maintenance Experiment (SIME/DIME)

4.1.1 Experimental Design, Data, and Labor-Supply Effects

Experimental design. Between 1968 and 1982, the U.S. federal government sponsored four large-scale NIT experiments to study how cash transfers affect work. SIME/DIME was the largest, longest, and most generous of these experiments (Robins, 1985). The experiment began in Seattle in 1970 and in Denver in 1971, under contracts between the states of Washington and Colorado and the U.S. Department of Health, Education, and Welfare (Burtless and Haveman, 1984). The experiment recruited low-income single- and dual-headed households (Murarka and Spiegelman, 1978). Households were stratified by income, race, household composition, and location, then randomly assigned to a control group or to one of eleven NIT plans (Conlisk and Kurz, 1972).

Each plan combined an income guarantee and a tax on earnings. A family with no income received the full guarantee. Each additional dollar of income reduced the payment by the tax rate, until the payment reached zero at the plan's break-even level of income. Guarantees were \$3,800, \$4,800, or \$5,600 per year (ranging from slightly below to about 40% above the poverty line), and tax rates were 50%, 70%, or 80%.[28]

[27] Metcalf (1974) corrected the experimental effects using households' saving during the experiment to infer how they discount the future, and their spending to infer the share of leisure in future consumption. Burtless and Greenberg (1982) measured the bias by comparing the three- and five-year arms within a static model of income and substitution effects; however, the two arms differed in sample and in the mix of plans assigned. Robins and West (1980) imposed parametric assumptions on hours in the absence of the program to recover the speed of adjustment from the covariance of hours over time. Their model is the case $\delta = 1$ and $\beta = 0$ of our framework, in which only last quarter's work matters, which we reject in the data (see Section 4.2).

[28] In some plans, the tax rate also declined with income. The experiment reimbursed regular income and payroll

Families received their assigned plan for three, five, or twenty years. Participants knew in advance when the program would end. We study the three-year arm, the largest of the three.[29]

Data. The experimenters collected retrospective labor-supply histories through interviews conducted approximately quarterly during the experiment and for four quarters after it ended. Our outcome is quarterly hours worked, which matches the frequency of the interviews.[30] We restrict the sample to single heads and both heads of couples that stay intact through the follow-up. We also restrict to heads with complete labor-supply information over the treatment and follow-up window.

We make one last sample restriction to accommodate two features of the experiment: treatment probabilities varied across randomization strata, and attrition differed by treatment status. To account for both, we estimate the probability of treatment conditional on the assignment strata and baseline characteristics, and use the resulting propensity scores for augmented inverse-probability weighting (AIPW) in our estimation throughout.[31] We trim observations outside the common support of the estimated propensity scores.

Together, these restrictions yield an analysis sample of 2,728 individuals, consisting of 891 husbands, 891 wives, and 946 single heads. Appendix C.1.1 describes the sample construction in more detail.

Descriptive statistics. Labor supply is highly variable and persistent in the sample. Before enrollment, control-group individuals work an average of 274 hours per quarter (SD = 258). As a descriptive measure of persistence, a pooled regression of current quarterly hours on hours in the previous quarter yields a coefficient of $\hat{r}^{\text{OLS}} = 0.884$ (SE = 0.005).

Treatment effects. Figure 2 plots the AIPW estimates of the effect of the NIT on quarterly hours before, during, and after the three-year experiment. Treated individuals work about 21 fewer hours (SE = 7.0) per quarter over the twelve treatment quarters (7.2% of the control mean in these quarters). The reduction grows over the program and reaches 29 hours (SE = 9.2) in the final treatment quarter (9.9% of the control mean in that quarter). After the program ends, the reduction fades, from 16 hours (SE = 9.1) in the first post-treatment quarter to a statistically

taxes for families receiving payments, and fully taxed transfers from other welfare programs. Furthermore, the experimenters adjusted the plans for changes in the cost of living (Kurz et al., 1973).

[29]The longer-duration arms offered the same plans, although the distribution of assignments across plans differed. Below, we use the five-year arm to validate the dynamics estimated from the three-year arm. The twenty-year arm, a later addition to the experiment among a small subset of households already enrolled in the experiment, is excluded from our analysis.

[30]We use the public SIME/DIME files studied by Riddell and Riddell (2024) and Dutz et al. (WIP). The public files report monthly hours, but these are constructed from approximately quarterly interviews and contain little independent within-interview variation.

[31]See Appendix C.1.4 for details on the estimation of the propensity scores. Using the propensity scores to reweight the analysis sample is a common approach for accounting both for assignment based on observables and differential attrition under selection on observables; see, e.g., Wooldridge (2007). We augment the propensity-score weights with an outcome model using the same covariates, following the AIPW approach of Robins et al. (1994).

insignificant 6 hours (SE = 9.4) in the last observed quarter. These effects line up with the Final Report (see Appendix C.1.1).

4.1.2 Verifying Identifying Assumptions

In the rest of this section, we use the experimental variation to recover the parameters of the model, following the identification argument of Section 3. Before doing so, we check that the experiment meets the three assumptions that argument requires (Appendix C.1.2 gives further details).

Random assignment. Assumption X.2 holds by design because assignment to the NIT was random within strata, among invited families. Our sample, however, is a selected subsample: we keep only individuals that stayed in the experiment through the follow-up and couples that stayed together. If treatment changed who is part of our sample in a way related to how much they would work without it, assignment would no longer be random within the analysis sample (Ghanem et al., 2026). Studies of SIME/DIME find this concern to be small: the estimated reductions in labor supply barely change under a parametric correction for selective attrition (Robins and West, 1986), and the conclusions broadly hold under arbitrary unobservable selection (Dutz et al., WIP). We nonetheless account for selection on observable characteristics, including hours worked before enrollment, as in Riddell and Riddell (2024) through AIPW using the propensity scores. After this adjustment, the analysis sample is balanced on observable characteristics not included in the estimation of the propensity score (Appendix Table C.2).

Known, constant intervention path. Assumption X.3 holds because participants learned their plan and its duration at enrollment, and the change in the cost of working was constant over time.[32] A possible threat to constancy is that the NIT taxed earnings, so its increase in the cost of working is proportional to the wage. As treated workers lose work history, the increase in the cost of working would then shrink over time, since (pre-tax) wages may rise with work history. We argue that this threat is minor, for two reasons. First, even with large wage returns to work history, the treatment's increase in the cost of working falls by less than 5 percent by the end of the program (Appendix C.1.2). Second, it leaves propagation unaffected, because we identify propagation from the decay of the effects after the program, when the tax no longer applies.

Exclusion. Assumption X.1 holds because after the program ended, treated households returned to the control tax and transfer system. A possible threat to exclusion is income effects: by raising households' resources, the payments could lower their marginal utility of wealth (MUW), which would permanently change their incentives to work. We argue that income effects are not a concern in our setting, for two reasons. First, the Final Report found little evidence that the NIT

[32] We exclude the households later moved to the twenty-year program, whose program length changed while it was in place. The plans were also indexed to the cost of living, so the increase in the cost of working was constant in real terms throughout the program (Kurz et al., 1973).

affected household saving. Second, we estimate in Appendix C.1.2 a model relaxing quasilinearity and allowing the NIT to change the MUW. The resulting parameter estimates are similar to those reported below for the quasilinear model and the income terms are jointly insignificant. We therefore keep the base model, which is simpler and gives more precise estimates.

Taken together, these features make the NIT an intervention of the form in Equation (9).

Figure 2: Experimental effects of a negative income tax program on labor supply

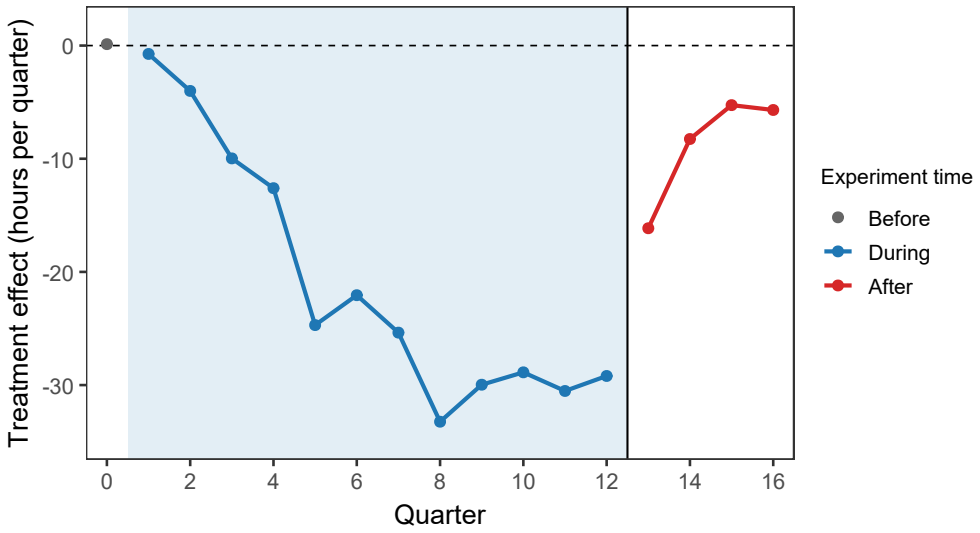

Notes: This figure reports treatment effects on hours worked per quarter for individuals in the three-year SIME/DIME negative income tax relative to the control group: pre-treatment (grey), during treatment (blue), post-treatment (red). Treatment effects are estimated by augmented inverse-probability weighting. The vertical line marks the end of the 12-quarter treatment period.

4.2 What the Experiment Reveals About Labor-Supply Dynamics

We estimate (r, δ, ϕ, β) jointly by the generalized method of moments (GMM) from the full path of treatment effects.[33] Appendices D and C.1.4 provide estimation details. This subsection reports the estimates, validates them against the five-year arm, and uses them to measure how much of the persistence in hours worked is causal.

4.2.1 Estimates of Propagation, Reinforcement, Survival, and Patience

Propagation rate. Table 1 reports estimates of the four model parameters. Columns (1) and (2) report the two conventional estimates of propagation from observed hours, and column (3) the estimate from the experimental path. The OLS regression of hours on lagged hours in column (1) gives $\hat{r}^{\text{OLS}} = 0.884$. This estimate mixes propagation with the two biases described in Section 3.2,

[33]Treated households had an incentive to underreport earnings (Greenberg and Halsey, 1983). Appendix C.1.3 shows that adjusting for plausible systematic reporting differences does not materially affect our estimates and conclusions below.

stable differences across workers and cost-outlook selection. The classic internal-instrument estimator in column (2) gives $\hat{r}^{\text{int}} = 0.519$. This estimate removes the bias from stable differences but not from cost-outlook selection; its overidentifying restrictions are rejected ($p < 0.001$).[34]

The path of treatment effects implies substantial propagation, $\hat{r}^{\text{exp}} = 0.862$. We fail to reject that the post-experiment decay of treatment effects is geometric with a constant rate ($p = 0.52$). We also fail to reject the overidentifying restrictions ($p = 0.57$).

The pattern across the three columns points to cost-outlook selection that downward biases the OLS and internal-instrument estimates. The OLS estimate is close to the experimental one because the upward bias from stable differences across workers offsets this downward bias. The internal-instrument estimator removes the upward bias but not the downward one, and therefore yields an estimate below the experimental estimate. This downward bias is consistent with costs of working that persist but fade within a few quarters (Appendix Lemma E.2), such as temporary road works that change commuting costs for only a few months.

Reinforcement. Reinforcement, ϕ, measures how inherited work history changes current hours. We estimate $\hat{\phi} = 0.235$: ten additional hours of work in one quarter raise hours in the next quarter by 2.35 hours. Because these induced hours propagate further into the future, the same ten hours generate 7.6 additional hours over the following four quarters. Past work thus modestly reinforces future work over the following year.

Survival. Survival, $1 - \delta$, measures how much of work history survives from one quarter to the next. We estimate $1 - \hat{\delta} = 0.627$ (SE $= 0.225$): 62.7% of work history survives quarter-to-quarter. A quarter's contribution to work history has a half-life of about 1.5 quarters, and about 15% of it survives after a year. We reject the null that only the last quarter of work history matters ($\delta = 1$) at the 10% level ($p = 0.055$).

Survival $1 - \hat{\delta}$ accounts for about 73% of propagation $\hat{r} = 1 - \hat{\delta} + \hat{\phi}$, while reinforcement $\hat{\phi}$ accounts for the other 27%. Propagation is thus mostly survival. Because most of each quarter's work history survives into the next, under a permanent incentive each paid quarter adds to a work history that keeps growing, so the long-run response takes time to build.

Patience. β measures workers' patience: how much they weigh the future. We estimate $\hat{\beta} = 0$: workers fully discount the future in choosing how many hours to work. However, the estimate is imprecise (SE $= 0.44$), so we cannot rule out moderate patience. Our estimates of ϕ and $1 - \delta$ are stable across a range of values of β (Appendix Figure A.1).

[34] See Table 1 notes for estimation details. Adding the available covariates lowers the OLS estimate slightly and leaves the internal-instrument estimate essentially unchanged (Appendix Table A.1).

4.2.2 Out-of-Sample Validation

We use treatment effects from the five-year treatment arm to conduct an out-of-sample validation of the estimated model. We fix the parameters (r, ϕ, δ, β) at the three-year estimates reported above, use five-year outcomes in quarters 1–6 to estimate the scale of that arm's initial response, and predict outcomes in quarters 7–18 (see Appendix C.1.5 for details and additional validation exercises).

Appendix Figure A.2 compares the predicted and observed paths. If the estimated dynamics were wrong, the predicted path would miss how the five-year response keeps building after quarter 12. We find instead that predictions closely match the continued build-up of the labor-supply response toward a plateau; the mean absolute prediction error is only 6.4 hours per quarter (8% of the average treatment effect). A joint test of the held-out prediction errors does not reject the model ($\chi^2(11) = 4.57$, $p = 0.95$).

4.2.3 How Much Labor-Supply Persistence is Causal?

Next, we ask how much of the observed persistence in hours worked ($\hat{r}^{\text{OLS}} = 0.88$) is causal. We use the estimated model to decompose observed persistence into the part caused by past work and the part due to other unobserved determinants.

Decomposition. By Lemma 1, observed hours are $Y_{it} = \phi_i S_{it} + K_i\alpha_i + \pi_{it}$, where $S_{it} = \sum_{k\geq 1}(1-\delta_i)^{k-1} Y_{i,t-k}$ is work history, $K_i\alpha_i$ stable preferences, and π_{it} the cost outlook. The covariance between current and lagged hours thus decomposes into a part that runs through work history and a part that does not:

$$\text{Cov}(Y_{it}, Y_{i,t-1}) = \underbrace{\text{Cov}(\phi_i\, S_{it}, Y_{i,t-1})}_{\text{causal effects of past work}} + \underbrace{\text{Cov}(K_i\alpha_i + \pi_{it}, Y_{i,t-1})}_{\text{other determinants of hours}}.$$

The first term is persistence caused by past work. The second is persistence in other determinants of hours, due both to stable differences across workers (like taste for leisure α_i) and to persistent cost outlooks (π_{it}).

Each term can be decomposed further. Using the law of motion for work history, the causal part further decomposes into the effect of last quarter's hours and that of earlier work history:

$$\text{Cov}(\phi_i S_{it}, Y_{i,t-1}) = \underbrace{\text{Cov}(\phi_i Y_{i,t-1}, Y_{i,t-1})}_{\text{last quarter}} + \underbrace{\text{Cov}(\phi_i(1-\delta_i) S_{i,t-1}, Y_{i,t-1})}_{\text{earlier work history}}.$$

The non-causal part splits into stable differences across workers and changes in costs. We measure the stable part of the cost outlook by its average over quarter t and every later quarter observed,

$\bar{\pi}_{it} = \frac{1}{T-t+1} \sum_{s=t}^{T} \pi_{is}$:

$$\text{Cov}(K_i\alpha_i + \pi_{it}, Y_{i,t-1}) = \underbrace{\text{Cov}(K_i\alpha_i + \bar{\pi}_{it}, Y_{i,t-1})}_{\text{stable differences}} + \underbrace{\text{Cov}(\pi_{it} - \bar{\pi}_{it}, Y_{i,t-1})}_{\text{changes in costs}}.$$

The first term contains stable preferences $K_i\alpha_i$ in full, and also costs that persist over the remaining quarters; it therefore measures differences across workers that are stable over the rest of the panel, not stable preferences alone.

We conduct these covariance decompositions for control group workers using the estimates of $\phi_i = \phi$ and $\delta_i = \delta$ reported above (see Appendix C.1.6 for details).

Results. We estimate that the causal component accounts for 62.5% of observed quarter-to-quarter persistence, with 26.8% through last quarter's hours and 35.7% through earlier work history. Stable differences across workers account for 27.5% and changes in costs for 10.0%.

These shares speak to a long-standing question (Heckman, 1981; Moffitt, 1992). Estimates of state dependence in labor supply range from small to substantial (Ellwood, 1982; Hyslop, 1999; Arulampalam et al., 2000), and most come from models in which only the last period matters ($\delta = 1$), so that the effect of last period's work and the effect of past work as a whole coincide. We estimate that these two effects differ substantially: last quarter's hours cause 27% of persistence, while past work as a whole causes 62%. Imposing $\delta = 1$, which our data reject at the 10% level, would attribute 89% of persistence to last quarter's hours, more than past work as a whole causes in the unrestricted model (Appendix C.1.6). Earlier work history is also what makes the response to a permanent program keep building, as we study next.

4.3 From a Temporary Experiment to a Permanent Program

Finally, we turn to the question the experiments were designed to answer: how would hours worked respond to a permanent NIT? A permanent NIT would reduce hours by more than the experiment shows at its end,both because the response builds up as the effects of past work accumulate and because patient workers who know the program will last would respond more. Under our main estimates, a permanent NIT would reduce work by 37 hours per quarter in the long run, 27% more than the 29-hour reduction the experiment shows in its last quarter. With patient workers ($\beta = 0.95$), the long-run reduction is 46 hours, 58% more.

Constructing the permanent response. Figure 3 compares the experiment with a permanent NIT. It shows three dates: quarter 5, quarter 12 (the last quarter of the experiment), and the long run. The gray bars are the experiment's observed effects. The blue bars are the permanent NIT at our estimates, under which workers have no patience ($\hat{\beta} = 0$); the orange bars are the permanent NIT when workers are patient ($\beta = 0.95$), with the other parameters re-estimated from the same path of effects. The model pins down responses to the permanent version as a

Table 1: Estimated dynamics of labor supply

	(1) Observed persistence	(2) Internal IV	(3) Experiment
Panel A: Propagation			
Propagation rate, r	0.884	0.519	0.862
	(0.005)	(0.050)	(0.114)
Panel B: Reinforcement, survival, and patience			
Reinforcement, ϕ			0.235
			(0.155)
Survival, $1-\delta$			0.627
			(0.225)
Patience, β			0.000
			(0.444)
Individuals	1,010	1,010	2,728

Notes: This table presents estimates of the dynamic parameters of quarterly hours worked. Panel A reports estimates of the quarterly propagation rate r. Panel B reports reinforcement ϕ, survival $1-\delta$, and patience β. Columns (1) and (2) use the control group. Column (1) reports the pooled OLS coefficient from a regression of quarterly hours on their one-quarter lag over Q1–Q16, with standard errors clustered by individual. Column (2) reports two-step difference GMM estimates, instrumenting lagged hours with hours three to six quarters earlier, with the standard errors of Windmeijer (2005). Its first-stage F is 139.5, and the Hansen J test rejects its overidentifying restrictions ($p < 0.001$). Column (3) reports GMM estimates using the experimental treatment-effect path (first-stage $F = 10.2$); Hansen J tests do not reject the restrictions on the full path ($p = 0.57$) or a common propagation rate after the program ($p = 0.52$), and a one-sided score test rejects $\delta = 1$ at the 10% level ($p = 0.055$). Standard errors are reported in parentheses.

multiple of the response to a single-quarter temporary program (Proposition 1(iii) and Appendix Proposition B.1). To express responses in hours, we fit this single-quarter response, separately for each value of β, so that the model's three-year program matches the observed effects in quarters 5–8 (Appendix C.1.4). The right axis expresses each bar as a percentage of the effect observed in quarter 12 (dashed line).

The response builds up. Without patience, workers respond to a permanent NIT as to the temporary one while it runs, so anticipation of the end plays no role.[35] The blue bars therefore track the gray ones within the experiment. In quarter 5 (inside the fitting window), they are close by construction (24 hours against 25); in quarter 12 (outside the fitting window), the model implies 32 hours against 29 observed. The permanent response keeps building, because each quarter of lower hours lowers the work history the next quarter inherits. By quarter 12 the permanent program reaches 88% of its long-run effect, and it reaches 95% of this effect only in quarter 19, seven quarters after the experiment ended (Appendix Proposition B.1). In the long run, the

[35]Metcalf (1974) and Burtless and Greenberg (1982) also found that knowing the program would end changes the labor-supply response little. In their models, the reason is not that workers are impatient, but that a temporary program understates income effects and overstates substitution effects, and the two biases offset each other (Metcalf, 1974, p. 547).

permanent program reduces hours by 37 per quarter, 2.70 times the first-quarter effect of 13.7 hours (Proposition 1) and 27% above the experiment's last-quarter effect.

Patience. Our estimate of patience is imprecise, and patience changes the permanent response. A patient worker values an hour of work today partly for the hours it adds later through her work history; under a permanent NIT those later hours are taxed too, so she cuts her hours by more from the start (Proposition 1(iii)). The orange bars show permanent program effects under patience. In quarter 5, inside the fitting window, the permanent program already reduces hours by 29, against 25 observed; by quarter 12 it reduces hours by 39, 21% more than without patience and 34% more than the experiment showed. The larger response also builds into work history, and the long-run reduction rises to 46 hours, against 37 without patience (25% more) and 29 in the last quarter of the experiment (58% more).

Figure 3: Labor-supply responses to temporary and permanent negative income tax programs

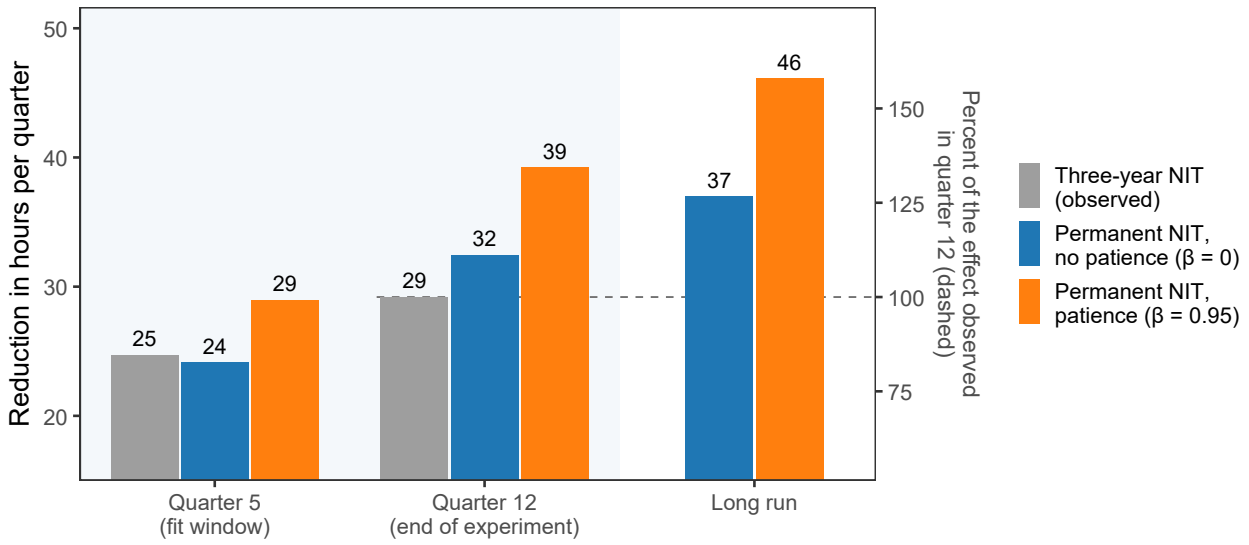

Notes: This figure compares the effects of the three-year NIT with the model-implied effects of a permanent NIT, in hours per quarter (left axis) and as a percentage of the effect observed in quarter 12 (right axis; dashed line). The gray bars depict the observed responses (estimates of Figure 2). The blue bars depict model-implied effects at the baseline estimates (at which $\hat{\beta} = 0$). The orange bars depict model-implied effects at the estimates obtained by imposing $\beta = 0.95$ and re-estimating the other parameters ($\hat{\phi} = 0.20$, $1 - \hat{\delta} = 0.67$, $\hat{r} = 0.88$). For each, the response to a single quarter of the program is fitted to the observed effects in quarters 5–8 using the model-implied path of the three-year program (Appendix C.1.4).

4.4 Takeaways

The NIT experiments were designed to measure how low-income households would respond to a permanent NIT, but participants knew the program would last only a few years. Our framework allows for both reasons why response to the temporary program may differ from that to a permanent one (that participants knew it would end and that past work affects future work), and it recovers the features that govern them from the path of treatment effects alone. Taking the

framework to the data, we find that past work accounts for 62 percent of the quarter-to-quarter persistence in hours, while stable differences across workers account for 28 percent and changes in costs for 10 percent. Most of the causal persistence runs through work before the last quarter, because most of work history survives from one quarter to the next. These dynamics, estimated from the three-year arm alone, closely predict the continued build-up of responses in the separate five-year arm.

Because the effects of past work accumulate, a permanent NIT would reduce hours worked in the long run by 27 to 58 percent more than the reduction observed at the end of the experiment, depending on how patient workers are. Taking the experiment's effects at face value would therefore understate the response the experiments were designed to measure.

5 Social Media Use and Incentives to Reduce It

Concern that young people overuse social media, or are addicted to it (Office of the Surgeon General, 2023), has led to calls for regulation, including school phone bans (UNESCO, 2023), age limits (Parliament of Australia, 2024), default time limits, and taxes or regulation of engagement (Romer, 2019; New York State Legislature, 2024). In this section, we use our framework to assess how much persistent social media use is driven by habit formation and evaluate the long-run effects of permanent policy interventions.

To do so, we use the "Bonus" experiment of Allcott et al. (2022), which paid a group of participants to reduce their use for three weeks. Allcott et al. (2022) used the same experiment to estimate a richer model with behavioral biases, in order to study how much social media use is due to self-control problems. We instead study how much of the persistence in use is due to past use, and what effects a permanent incentive would have.

5.1 The Digital Addiction Experiment

5.1.1 Experimental Design, Data, and Effects on Social Media Use

Experimental design. Allcott et al. (2022) ran a field experiment among U.S. adults with Android phones, using an app that recorded the time spent in six social media and browsing apps (Facebook, Instagram, Twitter, Snapchat, web browsers, and YouTube, "FITSBY"). After three weeks of baseline measurement, a (stratified) random quarter of the participants were assigned to a "Bonus" arm, offering \$50 for each one-hour reduction in their average daily FITSBY use below their baseline level (up to \$150) during a three-week period announced three weeks in advance. The bonus thus paid about \$2.50 per hour of use forgone. See Appendix C.2.1 and Allcott et al. (2022) for more details.

Data. We use the replication data provided by Allcott et al. (2022) (see Appendix C.2.1 for more details). The data consist of weekly FITSBY use (measured in minutes per day) for participants

over fifteen weeks: baseline (weeks 1–3), announcement (4–6), bonus (7–9), and follow-up (10–15). We exclude participants offered the separate time limit treatment and any participants not observed in all fifteen weeks.[36] The resulting sample contains 769 participants, 191 of whom are in the Bonus group. Throughout, we account for stratification using AIPW as in Section 4.

Descriptive statistics. Social media use is high, variable, and persistent among participants. At baseline (weeks 1–3), control-group participants use FITSBY apps for an average of 155 minutes per day (SD = 104). A pooled regression of weekly use on use in the previous week in the control group yields $\hat{r}^{\text{OLS}} = 0.90$ (SE = 0.009).

Treatment effects. Figure 4 plots the treatment-control difference in weekly FITSBY use before, during, and after the bonus. Before the announcement, use is similar in the two groups. In the three weeks between the announcement and the start of the bonus, the treatment reduces use by a statistically insignificant 5 minutes per day (SE = 5.4). Over the three paid weeks, the treatment reduces use by 61 minutes per day (SE = 5.6), 43% of the control mean. When the bonus ends, most of the reduction disappears at once: the effect falls to 21 minutes per day (SE = 6.0) in the first week after the bonus and to 11 minutes per day (SE = 5.9) in week 15.

5.1.2 Verifying Identifying Assumptions

Random assignment (Assumption X.2) holds by design, since assignment was random within strata. The **known, constant intervention path** (Assumption X.3) holds because participants knew the bonus's size and duration; kinks in the incentive schedule could make the shift in the cost of use vary across the paid weeks, but the payments depended on average use over the whole bonus period, and the effects are flat across the paid weeks. **Exclusion** (Assumption X.1) requires that the bonus left the cost of use and beliefs about it unchanged once it ended; income effects could violate it, but the payments were small and made once. Appendix C.2.2 gives the details.

5.2 What the Experiment Reveals About Social Media Habits

We estimate (r, δ, ϕ) by GMM from the path of treatment effects in weeks 7–15. Following Allcott et al. (2022), we fix $\beta = 0$.[37] Appendices D and C.2.3 provide estimation details. This subsection reports the estimates and uses them to measure how much of the persistence in use is causal.

5.2.1 Estimates of Propagation, Reinforcement, and Survival

Propagation rate. Table 2 reports the estimates. Columns (1), (2), and (3) report the OLS, internal-instrument, and experimental estimates of propagation. The three estimates are $\hat{r}^{\text{OLS}} =$

[36] We use participants offered the separate time limit to validate our estimates (Section 5.2.2).

[37] Allcott et al. (2022) argue that patient participants would have started cutting use as soon as the bonus was announced. Instead, use fell by only 5 minutes per day in the three announcement weeks. Appendix Figure A.3 shows the sensitivity of other parameter estimates to the choice of β.

Figure 4: Effects of the bonus experiment on social media use

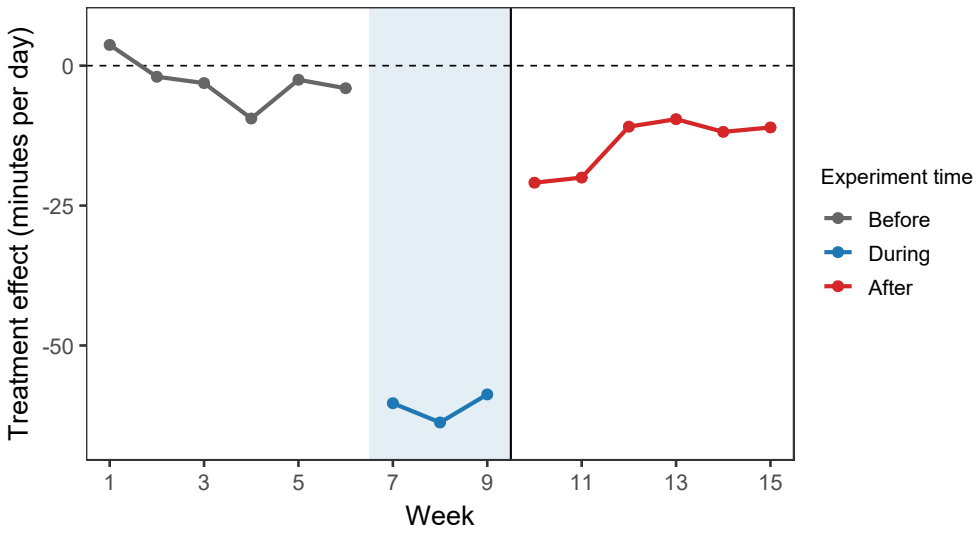

Notes: This figure reports weekly treatment effects of the bonus on FITSBY use, in minutes per day, before the bonus (gray; weeks 4–6 follow its announcement), while it was in place (blue; weeks 7–9), and after it ended (red). Effects are estimated by augmented inverse-probability weighting, on the replication data of Allcott et al. (2022).The vertical line marks the end of the bonus.

0.90, $\hat{r}^{\text{int}} = 0.28$, and $\hat{r}^{\text{exp}} = 0.85$ (SE $= 0.08$).[38] As in Section 4, the internal-instrument estimate falls far below the experimental one, pointing to a downward bias from cost-outlook selection, for example exam periods that change the cost of time on the platform for only a week or two.

Reinforcement. We estimate $\hat{\phi} = 0.075$ (SE $= 0.059$): ten additional minutes of daily use in one week raise daily use next week by 0.75 minutes. This additional use propagates further, but adds only about 35 total minutes of use over all future weeks.

Survival. We estimate $1 - \hat{\delta} = 0.78$ (SE $= 0.126$): about four fifths of accumulated use survives from one week to the next. A week's contribution to accumulated use has a half-life of 2.8 weeks, and 37% of it survives after four weeks. Accumulated use therefore extends well beyond the last week and we reject the restriction that only the last week matters ($\delta = 1$; $p < 0.001$).

Survival and reinforcement account for 91% and 9% of propagation, respectively. Propagation is thus mostly survival, and reinforcement is small relative to depreciation ($\hat{\phi}/\hat{\delta} = 0.34$): what a week of use builds lasts a few weeks but adds little to later use, so a permanent incentive builds little beyond its initial response and settles within weeks (Section 5.3).

5.2.2 Out-of-Sample Validation

The participants offered the time limit treatment, whom we do not use in estimation, provide an out-of-sample check on the estimates. The model implies that, while the bonus is in place, its

[38]The internal instrument's overidentifying restrictions are rejected at the 10% level ($p = 0.07$). We reject neither the null of a constant decay rate after the bonus ($p = 0.25$) nor the overidentifying restrictions ($p = 0.22$).

effect grows from week to week at a pace set by reinforcement and propagation, and that, once the bonus ends, its effect decays at rate r, whatever the values of ϕ and β (Proposition 3); so the effects in this group should build and fade as our estimates imply. Appendix Figure A.4 plots its effects against the paths our estimates imply from its week-7 and week-10 effects: the model predicts reductions of 50 and 53 minutes per day in weeks 8 and 9, against 54 and 53 observed. After the bonus, the decay line misses the weekly effects by less than a minute per day on average, and a joint test over the seven held-out weeks does not reject the model ($p = 0.23$).

5.2.3 How Much Persistence in Social Media Use is Causal?

We repeat the covariance decomposition of Section 4.2 to measure what share of observed persistence ($\hat{r}^{\text{OLS}} = 0.90$) is causal (see Appendix C.2.3 for details). Past use causes 34% of week-to-week persistence, with 8% coming through last week's use and 26% through earlier use. Stable differences across participants account for 56%, and changes in the cost of use account for 9%. Heavy users in one week are thus heavy users in the next mostly because of stable differences in preferences, not because of causal effects of past use.

Table 2: Estimated dynamics of social media use

	(1) Observed persistence	(2) Internal IV	(3) Experiment
Panel A: Propagation			
Propagation rate, r	0.898	0.278	0.853
	(0.009)	(0.148)	(0.078)
Panel B: Reinforcement and survival			
Reinforcement, ϕ			0.075
			(0.059)
Survival, $1-\delta$			0.778
			(0.126)
Individuals	577	577	769

Notes: This table presents estimates of the dynamic parameters of weekly FITSBY use. Panel A reports estimates of the weekly propagation rate r. Panel B reports reinforcement ϕ and survival $1-\delta$, estimated with $\beta = 0$. Column (1) reports the pooled OLS coefficient from a regression of weekly use on its one-week lag. Column (2) reports two-step difference GMM estimates, instrumenting lagged use with use three to six weeks earlier. Its first-stage F is 16.4, and the Hansen J test of its overidentifying restrictions has $p = 0.07$. Column (3) reports GMM estimates using the experimental treatment-effect path (first-stage $F = 95.8$); Hansen J tests do not reject the restrictions on the full path ($p = 0.22$) or a common propagation rate after the bonus ($p = 0.25$), and a one-sided score test rejects $\delta = 1$ ($p < 0.001$). Columns (1) and (2) use the 577 control participants observed in all fifteen weeks. See the notes to Table 1 for further estimation details. Standard errors are reported in parentheses.

5.3 From a Temporary Bonus to a Permanent Incentive

Figure 5 shows how fast a permanent incentive reaches its long-run response. The bars depict the share of the long-run effect reached in the first, third, and eighth weeks of the incentive, under our main point estimates (blue) and under constrained estimates in the limit $\beta \to 1$ (orange).[39] The long-run response is given by $\widehat{\text{LRM}} = 1.51$ times the one-week response (Proposition 1(iii)). We estimate a long-run response of 80 minutes per day, 56% of control-group use in the bonus weeks and 36% more than the effect observed in the last bonus week.

The response arrives fast. The blue bars show that a permanent incentive gets most of the long-run effect at once: 66% of the long-run response arrives in the first week. By Appendix Proposition B.1, that first-week share is $1 - \hat{\phi}/\hat{\delta}$. We estimate that reinforcement is small relative to depreciation, so little of the response runs through accumulated use. The remaining gap closes by $1 - \hat{r} = 15\%$ a week (22% of accumulated use that depreciates, minus 7% that reinforcement rebuilds), so the incentive reaches 76% of its long run by its third week (the duration of the bonus in the experiment) and 89% by its eighth week.[40]

Patience. The orange bars show that our conclusions are largely unchanged if we assume participants are patient. In the limit $\beta \to 1$, we estimate a similar propagation rate ($\hat{r} = 0.80$), higher reinforcement ($\hat{\phi} = 0.18$) and lower survival ($1 - \hat{\delta} = 0.62$). With these parameter values, we estimate that a smaller share of the long-run response is reached in the first week (52%), but still find that nearly all of the long-run response is reached by the eighth week (90%). Under both sets of estimates, a permanent incentive reaches most of its long-run effect within three weeks, the length of the bonus.

5.4 Takeaways

The Bonus experiment of Allcott et al. (2022) paid participants to reduce their social media use for three weeks, and we use it to measure how much of the persistence in use is due to past use, and what a permanent incentive would do. We find that persistent use mostly reflects stable differences across participants rather than habit formation: past use causes 34 percent of the week-to-week persistence in use, while stable differences account for 56 percent and changes in the cost of use for 9 percent. Accumulated use survives for a few weeks, but it reinforces later use only weakly. These dynamics are validated by how the bonus's effect builds and fades among the participants offered the time limit, whom we do not use in estimation.

Because accumulated use reinforces later use only weakly, a permanent incentive to reduce use

[39] As in Section 2.5, the permanent incentive raises the cost of use by the bonus's amount every week, and participants know it will last. As in Section 4.3, we express responses in minutes by fitting the one-week response to the effects in the three bonus weeks (Appendix C.2.3).

[40] The same rate also governs how the bonus's own effect fades once it ends. We estimate that the effect halves every 4.3 weeks, and 81% of the bonus's cumulative effect occurs within the fifteen weeks observed (Appendix C.2.4).

Figure 5: How fast a permanent incentive to reduce social media use takes full effect

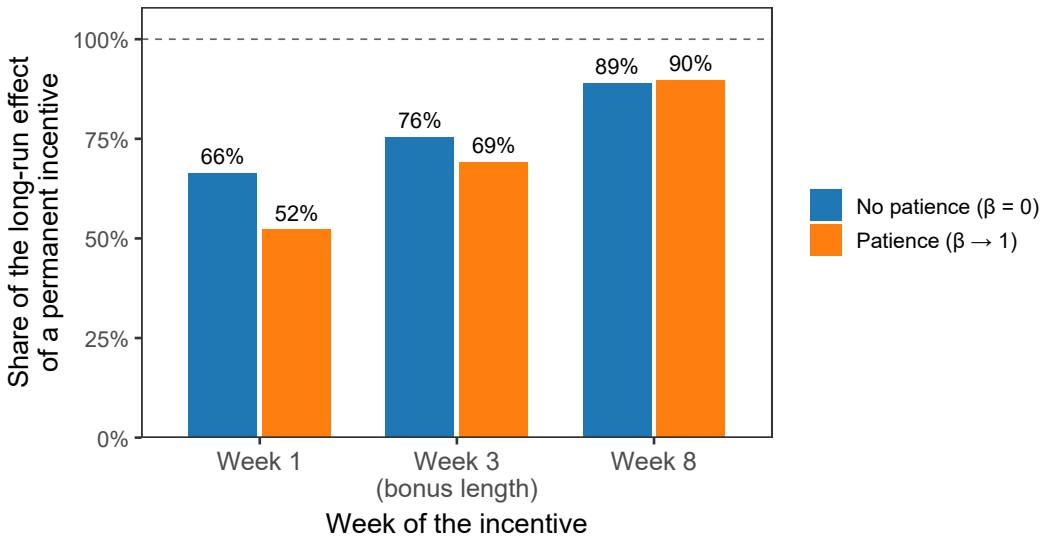

Notes: This figure reports the share of the long-run effect of a permanent incentive with the bonus's terms reached in its first, third, and eighth weeks (Appendix Proposition B.1). Blue bars use the estimates in Table 2 ($\beta = 0$); orange bars re-estimate ϕ and δ in the limit $\beta \to 1$ (Appendix Figure A.3).

would reach its long-run effect quickly: two thirds of that effect would arrive in the first week, and 89 percent within two months. At our estimates, a three-week experiment reveals most of what a permanent incentive would do.

6 Gym Attendance and Incentives to Exercise

Habit formation in physical exercise has been widely studied, with several experiments documenting that temporary incentives have lasting effects (Charness and Gneezy, 2009; Acland and Levy, 2015; Royer et al., 2015). In this section, we apply our framework to the experiment of Charness and Gneezy (2009) to quantify the dynamics of habit formation in exercise. We use the resulting estimates to answer counterfactual questions about a permanent policy to encourage exercising. Such policies have been discussed and implemented elsewhere: health insurers pay members every month in which they reach a visit target (Homonoff et al., 2020), and governments have subsidized exercise on an ongoing basis for very different populations and through very different instruments.[41]

[41] Canada, for instance, offered a tax credit for children's enrollment in physical-activity programs for close to a decade (Department of Finance Canada, 2017). U.S. Congress has considered subsidizing exercise through health savings accounts (U.S. House of Representatives, 2025).

6.1 The Incentives-to-Exercise Experiment

6.1.1 Experimental Design, Data, and Effects on Attendance

Experimental design. In their 'Study 1', Charness and Gneezy (2009) recruited 120 students at the University of Chicago and obtained the records of their visits to the campus gym (to which all had access). Forty students, recruited in separate sessions, form the control group. The other 80 students were paid \$25 to visit the gym at least once in the following week. A week later, half of these 80 students were chosen at random and offered a further \$100 to visit the gym at least eight times over the next four weeks (the five-week group); the other half form the one-week group. Appendix C.3.1 provides further details on the experiment and the data.

Data. We use the public replication data of Charness and Gneezy (2009), which report weekly gym visits (from the gym's swipe-card records) for all 120 participants over twenty weeks: eight before the payments, five with payments (weeks 1–5: the week of the \$25 payment and the four weeks of the \$100 offer), and seven after them (weeks 6–12). We restrict the sample to the five-week and control groups and hold out the one-week group to validate the estimated dynamics.

Descriptive statistics. Gym attendance is low, variable, and persistent among participants. Before the experiment, control-group participants visit the gym 0.59 times per week on average (SD = 1.08). A pooled regression of weekly visits on visits in the previous week in the control group yields $\hat{r}^{\mathrm{OLS}} = 0.84$ (SE = 0.04): one more gym visit in a given week is associated with 0.84 more visits the following week.

Treatment effects. Figure 6 plots the treatment-control difference in weekly visits, before, during, and after the payments. Before the payments, the two groups attend equally often (which is reassuring given the separate recruitment of the two groups). Payments raise attendance by 0.9 visits per week (SE = 0.23) in the first two weeks and 1.7 (SE = 0.20) in the last three. After payments stop, the effect falls to 0.8 visits per week (SE = 0.27) and then stays flat, averaging 0.7 (SE = 0.25) over the seven weeks after payments (120% of the control group mean).

6.1.2 Verifying Identifying Assumptions

Random assignment (Assumption X.2) is not explicit in the experimental implementation, because the control group was recruited in separate sessions. Nonetheless, the treatment assignment looks as good as random on observable characteristics (Charness and Gneezy, 2009), and the treatment effects before the start of the experiment are zero (Figure 6). The **known, constant intervention path** (Assumption X.3) holds because participants knew the terms of each payment once it was offered. One minor concern is that the \$100 offer came as a surprise a week into the payments, so in the first week participants did not know the path that followed. To account for this, we drop the first week, before the \$100 offer was known; this does not affect our conclusions

(see Section 6.2.2).[42] **Exclusion** (Assumption X.1) requires that the payments left the cost of a visit, and beliefs about it, unchanged once they ended; income effects could violate it, but the payments were small, at most $125 in total.

Figure 6: Experimental effects of payments on gym attendance

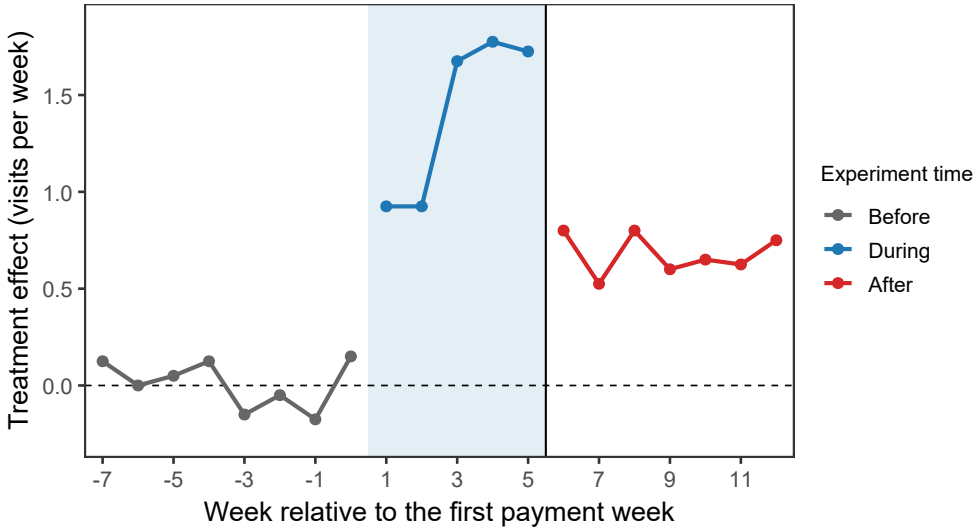

Notes: This figure reports weekly treatment effects of the payments on gym visits, the difference between the five-week and control groups of Study 1 of Charness and Gneezy (2009), in visits per week, before the payments (gray), while they were paid (blue, shaded), and after them (red). The vertical line marks the end of the payments.

6.2 What the Experiment Reveals About Exercise Habits

We estimate (r, δ, ϕ, β) jointly by GMM from the path of treatment effects in the five weeks with payments and the five following weeks. Appendices D and C.3.3 provide estimation details. This subsection reports the estimates, validates the estimates against the one-week group, and uses them to measure how much of the persistence in gym attendance is causal.

6.2.1 Estimates of Propagation, Reinforcement, Survival, and Patience

Propagation rate. Table 3 reports the estimates. Columns (1), (2), and (3) report the OLS, internal-instrument, and experimental estimates of propagation. The three estimates are $\hat{r}^{\text{OLS}} = 0.84$, $\hat{r}^{\text{int}} = 0.65$, and $\hat{r}^{\text{exp}} = 0.97$ (SE $= 0.07$).[43] As in Section 4, the internal-instrument estimate

[42] Strictly speaking, the $100 did not change the cost of each additional visit, so the change could vary over time: the schedule is nonlinear in the number of visits made. One way to interpret the incentive is as lowering the cost of each visit for a participant who plans to reach a target (whatever it is). Because the payment depends only on how many visits participants made, not on when, each visit counts equally toward the target, and thus the change in the marginal cost does not vary over time (see Appendix C.3.2).

[43] The internal instruments are weak (first-stage $F = 3.9$); their overidentifying restrictions are not rejected ($p = 0.79$). We reject neither a constant decay rate after the payments ($p = 0.20$) nor the overidentifying restrictions ($p = 0.10$).

falls below the experimental one, pointing to a downward bias from costs of a visit that persist for a few weeks but fade, such as an exam period or an illness.

Reinforcement. We estimate $\hat{\phi} = 0.19$ (SE $= 0.05$): one additional visit in a week raises attendance in the next week by 0.19 visits through fitness (our name for whatever past exercise builds, such as physical condition, know-how, and routines). Because visits propagate, the same induced visit generates about 0.7 additional visits over the following four weeks.

Survival. We estimate $1 - \hat{\delta} = 0.79$ (SE $= 0.11$): about four fifths of fitness survives from one week to the next. A week's contribution to it has a half-life of 2.9 weeks, and about 38% of it survives after four weeks. Fitness therefore carries more than the last week's attendance, and we reject the restriction that only the last week matters ($\delta = 1$; $p < 0.001$).

Survival and reinforcement account for 81% and 19% of propagation, respectively. Although propagation is mostly survival, reinforcement nearly replaces the fitness that wears off ($\hat{\phi}/\hat{\delta} = 0.87$), so a permanent incentive keeps building for months (Section 6.3).

Patience. We estimate $\hat{\beta} = 0$: participants choose visits as if they ignored the consequences of current visits for their future attendance. This point estimate reflects the fact that observed effects do not shrink as the end of the incentive approaches. However, the estimate is imprecise (SE $= 0.34$), so we cannot rule out some patience. Reassuringly, our estimates of ϕ and $1 - \delta$ are largely insensitive to β: imposing the limit $\beta \to 1$ gives $\hat{\phi} = 0.17$ and $1 - \hat{\delta} = 0.77$ (column (4) of Appendix Table A.2).

6.2.2 Out-of-Sample Validation and Robustness

Validation. The one-week group, which we do not use in estimation, provides an out-of-sample check on the estimated propagation rate.[44] The model implies that, once its single paid week ends, its effect on attendance decays at rate r, whatever the values of ϕ and β (Proposition 3), so its decay should match the rate estimated from the five-week group. Appendix Figure A.5 plots its effects against a decay at $\hat{r}^{\text{exp}}$; estimated on its own, the one-week group's effect decays at 0.89 (SE $= 0.14$), statistically indistinguishable from 0.97. The decay line also tracks the weekly effects closely: its root mean squared deviation from them is 0.07 visits per week. Thus, the propagation rate estimated from a month of payments also describes how the effect of a single paid week fades.

Robustness. Appendix Table A.2 addresses the concerns of Section 6.1.2 and the imprecision of $\hat{\beta}$. Column (2) drops the first week, before the \$100 offer was known. Column (3) compares the five-week group with the one-week group over the four weeks of the \$100 offer and the seven

[44] This comparison cannot speak to patience: in the first week both groups had the same information, a single paid visit, so their attendance is the same in that week whatever their degree of patience. Had they known the length of their incentive from the start, a comparison of their first-week responses would have revealed patience (Section 7).

weeks after them. (Recall that assignment to these groups was random.) Column (4) imposes the limit $\beta \to 1$. No alternative estimate of propagation, reinforcement, or survival differs from the headline estimate by more than its standard error, and all of them imply the same picture for a permanent incentive: reinforcement close to depreciation ($\hat{\phi}/\hat{\delta}$ between 0.74 and 0.79), and a long-run response several times the initial one.

6.2.3 How Much Persistence in Gym Attendance is Causal?

We repeat the covariance decomposition of Section 4.2 to measure what share of observed persistence ($\hat{r}^{\text{OLS}} = 0.84$) is causal (see Appendix C.3.3 for details). Applying the decomposition to control-group visits, we find that past attendance causes 77% of week-to-week persistence, with 22% coming through last week's attendance and 55% through earlier attendance.[45] Stable differences across participants account for 20%, and changes in the cost of a visit for the remaining 3%. Regular gym attendance thus largely reflects past attendance, especially visits more than one week earlier.

6.3 From a Temporary Experiment to a Permanent Incentive

Figure 7 shows how a permanent incentive with the experiment's terms would build. The blue bars give its effect in weeks 5, 9, 16, and 25 of the incentive, at our estimates. The purple bars give the same effect if only last week's attendance mattered ($\delta = 1$). The gray bar is the effect observed in week 5, the last week with payments. As in Section 4.3, we express responses in visits by fitting the one-week response to the effects in the first three weeks with payments (Appendix C.3.4). In week 5, outside the fitting window, the model (1.7 visits per week) matches the data (1.7). At our estimates, the long-run response of a permanent incentive is 7.4 visits per week, about four times the effect observed in the last week with payments.

The response builds slowly. The blue bars show that a permanent incentive takes months to approach its long run. Only 13% of the long-run response arrives in the first week. The remaining gap closes by $1 - \hat{r} = 2.9\%$ a week, so it takes 24 weeks to halve. The response grows to 2.3 visits per week in week 9, 3.3 in week 16, and 4.2 in week 25, 2.5 times the effect observed in the last week with payments but still only 57% of the long run. A month of payments thus shows under a quarter of what a permanent incentive would do, and the rest takes months.

The role of survival. To show where the slow build-up comes from, we repeat the exercise with $\delta = 1$, keeping reinforcement and the one-week response at the base estimates. The purple bars show the resulting effects in the same weeks. We find that the response would level off at 1.2 visits

[45] Because fitness survives for several weeks, the shares depend on the fitness participants had built before our data begin (week 1), which the twelve weeks of control-group data we use (weeks 1–12; Appendix C.3.1) cannot pin down. We report the midpoint between two extremes: With fitness starting at zero when the data begin, past attendance causes 58% of persistence; at the steady state the model implies, 96%.

Table 3: Estimated dynamics of gym attendance

	(1) Observed persistence	(2) Internal IV	(3) Experiment
Panel A: Propagation			
Propagation rate, r	0.836	0.648	0.971
	(0.043)	(0.159)	(0.075)
Panel B: Reinforcement, survival, and patience			
Reinforcement, ϕ			0.185
			(0.045)
Survival, $1-\delta$			0.786
			(0.110)
Patience, β			0.000
			(0.343)
Individuals	40	40	80

Notes: This table presents estimates of the dynamic parameters of weekly gym visits. Panel A reports estimates of the weekly propagation rate r. Panel B reports estimated reinforcement ϕ, survival $1-\delta$, and patience β. Columns (1) and (2) use the control group over weeks 1–12. Column (1) reports the pooled OLS coefficient from a regression of weekly visits on their one-week lag, with standard errors clustered by individual. Column (2) reports two-step difference GMM estimates, instrumenting lagged visits with visits three to six weeks earlier, collapsed. Its first-stage F is 3.9, and the Hansen J test does not reject its overidentifying restrictions ($p = 0.79$). Column (3) reports GMM estimates using the treatment-effect path of the five-week group relative to the control group over weeks 1–10 (first-stage $F = 50.7$); Hansen J tests do not reject the restrictions on the full path ($p = 0.10$) or a common propagation rate after the payments ($p = 0.20$), and a one-sided score test rejects $\delta = 1$ ($p < 0.001$). See the notes to Table 1 for further estimation details. Standard errors are reported in parentheses.

per week within two weeks, below the 1.7 observed in the last week with payments, and far below the blue bars in later weeks. Reinforcement alone would add only about a fifth to the one-week response. The data reject this case: the $\delta = 1$ model would have the effect of the payments vanish within two weeks of their end, whereas the observed effect stays near 0.7 visits per week for seven weeks (Section 6.2).

6.4 Takeaways

The experiment paid students to go to the gym for a month, and its lasting effects are widely read as evidence that temporary incentives can build exercise habits. Consistent with this reading, we find that past attendance causes 77 percent of the week-to-week persistence in attendance, while stable differences across participants account for 20 percent. Most of the causal persistence runs through attendance before the last week, because fitness survives from one week to the next. These dynamics are validated by how the effect of a single paid week fades in the one-week group, which we do not use in estimation.

Fitness also raises attendance almost as fast as it depreciates, so the response to a permanent incentive keeps building for months. In its twenty-fifth week, a permanent incentive would raise

Figure 7: Effects of temporary and permanent payments on gym attendance

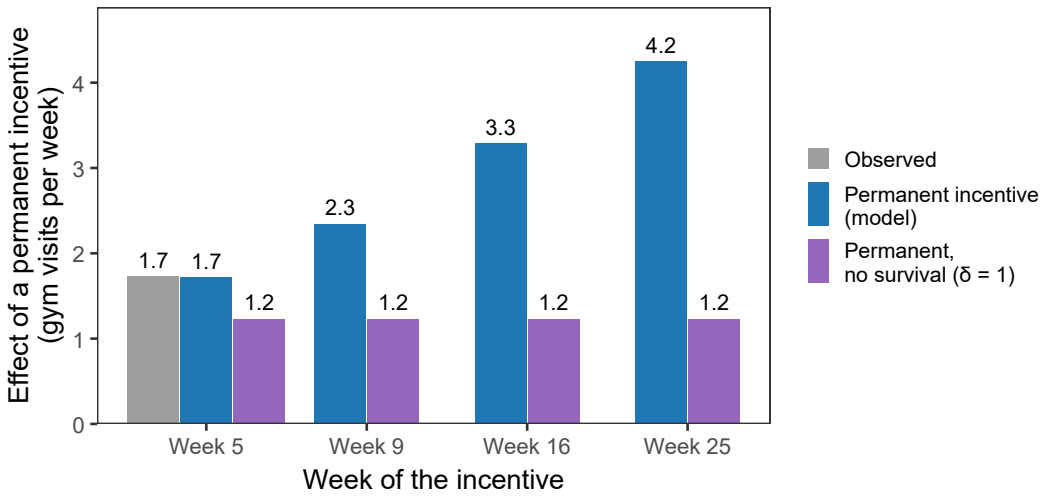

Notes: This figure reports the effect of a permanent incentive with the experiment's terms on weekly visits in weeks 5, 9, 16, and 25 of the incentive (Appendix C.3.4). Blue bars use the estimates in Table 3 ($\hat{\beta} = 0$), with the one-week response fitted to the effects observed in weeks 1–3. Purple bars set $\delta = 1$, holding $\hat{\phi}$ and the one-week response (1 visit per week) at their estimates. The gray bar is the effect observed in week 5, the last week of the intervention.

attendance by 4.2 visits per week, 2.5 times the effect observed in the last week with payments and still only 57 percent of its long run. Holding reinforcement at our estimate, a model in which only last week's attendance matters would miss this build-up: the response would level off at 1.2 visits per week within two weeks. A month of payments thus reveals only the beginning of what a permanent incentive would do.

7 Implications for Experimental Design

The empirical applications in Sections 4–6 highlight two practical challenges with implementing our framework. The first is learning about patience. In the NIT and gym applications, it is the least precisely estimated parameter, with point estimates close to no-patience and large standard errors. For social media, we follow Allcott et al. (2022) and fix $\beta = 0$. The second challenge is testing and validating the model. Our out-of-sample validations rely on arms the experiments happened to include: the five-year NIT arm, the social-media participants offered another incentive scheme, and the group offered gym attendance incentives for one-week only. The social-media experiment provides no such validation arm.

In this section, we use our framework to offer practical suggestions for experiment design that sharpen identification of patience and provide additional tests of assumptions. Beyond the size of the incentive, the researcher makes three choices in the design of a temporary experiment (Section 3.4.1): when participants learn about the incentive, how long the incentive lasts, and when the incentive begins. We study what a second treated arm adds when it changes one of these design choices: its participants are informed about the incentive earlier, receive it for longer,

or begin to receive it later. We show that additional arms of these types add more information about patience, provide additional overidentifying restrictions, or both (Table 4).

Throughout, we assume that Assumptions X.1–X.3 hold for each treated arm, and that participants learn the full path of the incentive when they learn of it. For brevity, we defer formal propositions and proofs to Appendix F.

Table 4: Alternative experimental designs: information about patience and overidentifying restrictions

Design	Second arm	Reveals patience from	Overidentifying restrictions
Single arm (Section 3.4)	none; the baseline arm is treated from $\underline{t}$ through $\bar{t}$	how the effects move as the end approaches (Section 3.4.3)	(1) After the end, the effects fade out at the same rate in every period. (2) If $\bar{t} - \underline{t} \geq 3$, survival and patience are overidentified by the effects while treated.
Lead announcement	when participants learn: told one period earlier, at $\underline{t} - 1$	its response in the announcement period, before any change in costs	(1) and (2) in each arm. (3) The propagation rate and the boundary gap are the same in both arms. (4) The initial difference between the arms fades out at the propagation rate, even while treated.
Variable duration	how long it lasts: treated through $\bar{t}_2 > \bar{t}$	the difference between the arms while both are treated	(1) and (2) in each arm, and (3). (5) The shorter arm's estimates predict the longer arm's path.
Staggered start	when it begins: treated from $\underline{t} + s$ through $\bar{t} + s$	none	(1) and (2) in each arm, and (3). (6) In event time, the two paths of treatment effects coincide.

Notes: This table compares the single-arm temporary experiment design studied in Section 3.4 with three designs that add a second treated arm. Each row is a design; in the last three, the second arm changes one design choice of the baseline arm and nothing else. The first column names the design. The second column describes the second arm: the design choice it changes and its treatment dates. The third column reports the evidence that reveals patience; all of it requires reinforcement, as in Proposition 4(ii), since without it future incentives leave today's choice unchanged. The fourth column lists the overidentifying restrictions each design implies. The boundary gap is the drop in the effect, beyond its fade-out at the propagation rate, in the first period after the incentive ends (Section 3.4.3). See Appendix F for formal propositions and proofs.

Lead announcement. Patience, the weight individuals put on the future, is revealed by how individuals respond to what they anticipate. In the single-arm design studied in Section 3, patience is identified from how individuals change their behavior as they approach the known end date of the incentive (Section 3.4.3). Such changes may be small and difficult to distinguish from noise in sample data.

An additional arm told about the incentive a period before it starts gives a second measure (Proposition F.1). In that period its costs have not yet changed, so its response is pure anticipation of the incentive, which is nonzero only if individuals are patient and past choices reinforce future ones.[46]

[46] Announcing treatment with further lead time adds more periods of pure anticipation, all further carrying information about patience (Proposition F.1).

From the first treated period on, both arms face the same incentive and the same future, so they differ only in the additional stock the announced arm built in anticipation. The lead announcement arm also adds two overidentifying restrictions (Appendix F.2). First, the initial difference between the arms must fade out at the propagation rate, even while the incentive is in place. Second, the boundary gap must be the same in both arms.

Variable (incentive) duration. Adding an arm that receives the same incentive for longer provides another way to measure patience. While both arms are treated, they face the same incentive each period but different incentives in the future, so their responses can only differ if individuals are patient. The difference between their responses reveals patience (Proposition F.2).

The longer arm also adds overidentifying restrictions that test the model. First, a model estimated on one arm must predict the other arm's entire path of treatment effects; equivalently, the parameters estimated from the two arms must coincide.[47] Second, after each arm's incentive ends, its effects must fade out at the same rate, and the two boundary gaps must coincide (Appendix F.3).

Staggered start. Our argument identifying reinforcement, survival and patience from the path of effects (Proposition 4) assumes that responses depend only on event time (i.e., time relative to the incentive's start and end), not on calendar time. With only one treatment arm, this assumption is untestable, because event time and calendar time are collinear.

Adding an arm that receives the same incentive for the same number of periods, but which starts s periods later, separates event time from calendar time. The staggered arm provides an overidentifying restriction: in event time, the two arms' treatment-effect paths must coincide (Proposition F.3).

Practical considerations. These extended experimental designs introduce some practical costs and challenges. A lead announcement requires early recruitment of participants, and comes with a risk of information about the incentive spilling over across arms. Variable duration requires lengthening the panel, and comes with a risk of (differential) attrition across arms. A staggered start also requires longer panels (potentially much longer, depending on the researcher's goals), and comes with a risk that participants enrolled in later waves may differ from those enrolled in earlier waves.

One benefit of our proposed designs is that dividing a fixed sample size between more treatment arms need not sacrifice much statistical power for estimating model parameters. If overidentifying restrictions are not rejected, researchers can combine moment restrictions across arms to obtain more efficient estimates.

[47] This observation underlies our out-of-sample validation exercises in the NIT and gym applications. Those applications required adjusted tests to account for other differences across treatment arms (see Sections 4.2 and 6.2).

8 Conclusions

To understand why choices persist and answer counterfactual questions about policies of different lengths, we develop a framework that combines a dynamic choice model with temporary (randomized) experiments. In the model, choices persist because past choices shape future ones, because individuals differ in stable preferences, and because individuals face serially correlated costs. Answering counterfactual questions about policies of different lengths requires knowledge of how choices propagate over time, how that propagation decomposes into reinforcement and survival, and how patient individuals are.

We show that these key features of behavior can all be identified from the path of treatment effects in a temporary experiment. How effects decay after the experiment ends reveals propagation; how effects drop immediately after the experiment ends distinguishes survival from reinforcement; and how effects bend as the end approaches reveals patience. Importantly, we show that a temporary experiment identifies these features without assuming that individuals cannot anticipate changes in their costs. The experiment therefore separates causal propagation from persistence caused by serially correlated or predictable costs.

We apply our framework to temporary experiments in three settings: the Seattle–Denver Income Maintenance Experiment (labor supply), the bonus experiment of Allcott et al. (2022) (social media use), and the experiment of Charness and Gneezy (2009) (exercise). In the Seattle–Denver experiment, we find that past work accounts for almost two thirds of the persistence in labor supply; that most propagation comes from survival; and that the effects observed at the end of the three-year NIT meaningfully understate the long-run effects of a permanent NIT. In the social-media experiment, we find that past use accounts for about a third of the persistence in use, and that a permanent incentive to reduce social media use would approach its full long-run effect within weeks because reinforcement from past use is small relative to the depreciation of accumulated use. In the exercise experiment, we find that past attendance accounts for more than three quarters of the persistence in gym visits, and that the response to a permanent incentive would continue to build over a long period of time because reinforcement nearly offsets depreciation in fitness.

Finally, we propose three practical extensions adding a second treated arm to the temporary experiment design. Each changes one feature: participants in the second arm learn of the incentive earlier, receive it for longer, or start it later. These extensions furnish additional strategies to identify patience and provide opportunities to test and validate the model.

References

Abbring, J. H. and Heckman, J. J. (2007). Econometric evaluation of social programs, part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium

policy evaluation. In Heckman, J. J. and Leamer, E. E., editors, *Handbook of Econometrics*, volume 6B, pages 5145–5303. Elsevier. 15

Acland, D. and Levy, M. R. (2015). Naiveté, projection bias, and habit formation in gym attendance. *Management Science*, 61(1):146–160. 3, 37

Allcott, H., Gentzkow, M., and Song, L. (2022). Digital Addiction. *American Economic Review*, 112(7):2424–2463. 3, 5, 7, 32, 33, 34, 36, 43, 46, 54, 55, 91, 92

Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–3037. 4, 9

Anderson, T. W. and Hsiao, C. (1982). Formulation and estimation of dynamic models using panel data. *Journal of Econometrics*, 18(1):47–82. 2, 4, 14, 16

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The review of economic studies*, 58(2):277–297. 2, 4, 14, 16

Arellano, M. and Honoré, B. (2001). Panel data models: some recent developments. In *Handbook of econometrics*, volume 5, pages 3229–3296. Elsevier. 4, 14

Arulampalam, W., Booth, A. L., and Taylor, M. P. (2000). Unemployment persistence. *Oxford Economic Papers*, 52(1):24–50. 29

Ashenfelter, O. (1978). Estimating the effect of training programs on earnings. *Review of Economics and Statistics*, 60(1):47–57. 111

Becker, G. S., Grossman, M., and Murphy, K. M. (1994). An empirical analysis of cigarette addiction. *American Economic Review*, 84(3):396–418. 1, 7, 16

Becker, G. S. and Murphy, K. M. (1988). A theory of rational addiction. *Journal of political Economy*, 96(4):675–700. 5, 9

Blume, L., Easley, D., and O'Hara, M. (1982). Characterization of optimal plans for stochastic dynamic programs. *Journal of Economic Theory*, 28(2):221–234. 6

Blundell, R. and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1):115–143. 4

Bonhomme, S. (2025). Back to feedback: Dynamics and heterogeneity in panel data. Working paper 2512.17576, arXiv. 4, 15

Burtless, G. and Greenberg, D. (1982). Inferences concerning labor supply behavior based on limited-duration experiments. *The American Economic Review*, 72(3):488–497. 3, 4, 9, 23, 30

Burtless, G. and Haveman, R. H. (1984). Policy lessons from three labor market experiments. Discussion Paper 746-84, Institute for Research on Poverty, University of Wisconsin–Madison. 3, 23

Cawley, J. and Ruhm, C. J. (2011). The economics of risky health behaviors. In Pauly, M. V., McGuire, T. G., and Barros, P. P., editors, *Handbook of Health Economics*, volume 2, pages 95–199. Elsevier. 1

Cefala, L., Kaur, S., Schofield, H., and Shamdasani, Y. (2024). Habit formation in labor supply. Technical report, Working paper. 4

Celhay, P. A., Gertler, P. J., Giovagnoli, P., and Vermeersch, C. (2019). Long-run effects of temporary incentives on medical care productivity. *American Economic Journal: Applied Economics*, 11(3):92–127. 4, 7

Charness, G. and Gneezy, U. (2009). Incentives to exercise. *Econometrica*, 77(3):909–931. 3, 4, 5, 7, 9, 37, 38, 39, 46, 55, 94

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, 80(3):969–1018. 4

Chetty, R., Friedman, J. N., Olsen, T., and Pistaferri, L. (2011a). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records. *Quarterly Journal of Economics*, 126(2):749–804. 4

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011b). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3):471–475. 4

Christophersen, G. (1983). Final Report of the Seattle-Denver Income Maintenance Experiment Volume 1: Design and Results. 3, 23

Conlisk, J. and Kurz, M. (1972). The Assignment Model of the Seattle and Denver Income Maintenance Experiments. 23

Department of Finance Canada (2017). Report on federal tax expenditures: Concepts, estimates and evaluations 2017. Technical report, Government of Canada, Ottawa. Part 9: evaluation of the Children's Fitness Tax Credit and the Children's Arts Tax Credit. 37

Dubé, J.-P., Hitsch, G. J., and Rossi, P. E. (2010). State dependence and alternative explanations for consumer inertia. *RAND Journal of Economics*, 41(3):417–445. 1, 4

Duflo, E., Glennerster, R., and Kremer, M. (2007). Using randomization in development economics research: A toolkit. In Schultz, T. P. and Strauss, J. A., editors, *Handbook of Development Economics*, volume 4, pages 3895–3962. Elsevier, Amsterdam. 1

Dutz, D., Lacouture, S., Mogstad, M., and Torgovitsky, A. (WIP). The Negative Income Tax experiments 50 years later: What can we learn by accounting for attrition? 24, 25

Eckstein, Z. and Lifshitz, O. (2011). Dynamic female labor supply. *Econometrica*, 79(6):1675–1726. 6, 83

Ellwood, D. T. (1982). Teenage unemployment: Permanent scars or temporary blemishes? In Freeman, R. B. and Wise, D. A., editors, *The Youth Labor Market Problem: Its Nature, Causes, and Consequences*, pages 349–390. University of Chicago Press. 29

Ghanem, D., Hirshleifer, S., and Ortiz-Becerra, K. (2026). Testing attrition bias in field experiments. *Journal of Human Resources*, 61(2):690–731. 25

Gladden, T. and Taber, C. (2000). Wage progression among less skilled workers. In Card, D. E. and Blank, R. M., editors, *Finding Jobs: Work and Welfare Reform*. Russell Sage Foundation, New York. 83

Gneezy, U., Meier, S., and Rey-Biel, P. (2011). When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives*, 25(4):191–210. 1

Greenberg, D. and Halsey, H. (1983). Systematic misreporting and effects of income maintenance experiments on work effort: Evidence from the Seattle-Denver experiment. *Journal of Labor Economics*, 1(4):380–407. 26, 86

Gruber, J. and Kőszegi, B. (2001). Is addiction "rational"? Theory and evidence. *Quarterly Journal of Economics*, 116(4):1261–1303. 4

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–2682. 4

Harrison, G. W. and List, J. A. (2004). Field experiments. *Journal of Economic Literature*, 42(4):1009–1055. 1

Heckman, J. J. (1981). Heterogeneity and state dependence. *Studies in labor markets*, 31:91–140. 1, 13, 29

Heckman, J. J. and Borjas, G. J. (1980). Does unemployment cause future unemployment? Definitions, questions and answers from a continuous time model of heterogeneity and state dependence. *Economica*, 47(187):247–283. 1

Heckman, J. J., Lochner, L. J., and Todd, P. E. (2006). Earnings functions, rates of return and treatment effects: The Mincer equation and beyond. In Hanushek, E. A. and Welch, F., editors, *Handbook of the Economics of Education*, volume 1, pages 307–458. Elsevier, Amsterdam. 83

Heckman, J. J. and Willis, R. J. (1977). A beta-logistic model for the analysis of sequential labor force participation by married women. *Journal of Political Economy*, 85(1):27–58. 1

Homonoff, T., Willage, B., and Willén, A. (2020). Rebates as incentives: The effects of a gym membership reimbursement program. *Journal of Health Economics*, 70:102285. 37

Hussam, R., Rabbani, A., Reggiani, G., and Rigol, N. (2022). Rational habit formation: Experimental evidence from handwashing in India. *American Economic Journal: Applied Economics*, 14(1):1–41. 7

Hyslop, D. R. (1999). State dependence, serial correlation and heterogeneity in intertemporal labor force participation of married women. *Econometrica*, 67(6):1255–1294. 1, 6, 29

Imai, S. and Keane, M. P. (2004). Intertemporal labor supply and human capital accumulation. *International Economic Review*, 45(2):601–641. 4

Imbens, G. W. and Angrist, J. D. (1994). Identification and Estimation of Local Average Treatment Effects. *Econometrica*, 62(2):467–475. 18

Keane, M. and Rogerson, R. (2015). Reconciling micro and macro labor supply elasticities: A structural perspective. *Annual Review of Economics*, 7(1):89–117. 4

Keane, M. P. (1997). Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics*, 15(3):310–327. 1

Keane, M. P. (2015). Effects of permanent and transitory tax changes in a life-cycle labor supply model with human capital. *International Economic Review*, 56(2):485–503. 1

Keane, M. P. and Wolpin, K. I. (1997). The career decisions of young men. *Journal of Political Economy*, 105(3):473–522. 4

Kleven, H., Kreiner, C., Larsen, K., and Søgaard, J. (2025). Micro versus macro labor supply elasticities: The role of dynamic returns to effort. *American Economic Review*, 115(9):2849–2890. 4

Kurz, M., Spiegelman, R. G., and Brewster, J. (1973). The payment system for the seattle and denver income maintenance experiments. *SRI International*. 24, 25

List, J. A. (2011). Why economists should conduct field experiments and 14 tips for pulling one off. *Journal of Economic Perspectives*, 25(3):3–16. 1, 20

Mertens, K. and Ravn, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the United States. *American Economic Review*, 103(4):1212–1247. 16

Metcalf, C. E. (1973). Making inferences from controlled income maintenance experiments. *The American Economic Review*, 63(3):478–483. 1, 3, 9, 23

Metcalf, C. E. (1974). Predicting the effects of permanent programs from a limited duration experiment. *The Journal of Human Resources*, 9(4):530–555. 4, 9, 23, 30

Moffitt, R. A. (1992). Incentive effects of the U.S. welfare system: A review. *Journal of Economic Literature*, 30(1):1–61. 1, 29

Murarka, B. A. and Spiegelman, R. G. (1978). Sample Selection in the Seattle and Denver Income Maintenance Experiments. *SRI International*. 23

New York State Legislature (2024). Stop Addictive Feeds Exploitation (SAFE) for Kids Act, S. 7694-A. Signed into law June 20, 2024. 32

Office of the Surgeon General (2023). Social media and youth mental health: The U.S. Surgeon General's Advisory. Technical report, U.S. Department of Health and Human Services, Washington, DC. 32

Parliament of Australia (2024). Online Safety Amendment (Social Media Minimum Age) Act 2024. 32

Riddell, C. and Riddell, W. C. (2024). Welfare versus work under a negative income tax: Evidence from the Gary, Seattle, Denver, and Manitoba income maintenance experiments. *Journal of Labor Economics*, 42(2). 24, 25, 78, 79

Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866. 24, 97

Robins, P. K. (1985). A comparison of the labor supply findings from the four negative income tax experiments. *Journal of human Resources*, pages 567–582. 23

Robins, P. K. and West, R. W. (1980). Labor supply response over time. *Journal of Human Resources*, pages 524–544. 1, 3, 4, 23, 79, 80

Robins, P. K. and West, R. W. (1986). Sample attrition and labor supply response in experimental panel data: A study of alternative correction procedures. *Journal of Business & Economic Statistics*, 4(3):329–338. 25

Romer, P. M. (2019). A tax that could fix Big Tech. The New York Times. May 6, 2019. 32

Royer, H., Stehr, M., and Sydnor, J. (2015). Incentives, Commitments, and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company. *American Economic Journal: Applied Economics*, 7(3):51–84. 3, 4, 7, 9, 37

Santos, M. S. (1991). Smoothness of the policy function in discrete time economic models. *Econometrica*, 59(5):1365–1382. 6

Schaner, S. (2018). The persistent power of behavioral change: Long-run impacts of temporary savings subsidies for the poor. *American Economic Journal: Applied Economics*, 10(3):67–100. 4

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *Economic Journal*, 128(610):917–948. 16

Stokey, N. L. and Lucas, Jr., R. E. (1989). *Recursive methods in economic dynamics*. Harvard University Press. 6, 61

Torgovitsky, A. (2019). Nonparametric inference on state dependence in unemployment. *Econometrica*, 87(5):1475–1505. 1

UNESCO (2023). *Global Education Monitoring Report 2023: Technology in Education: A Tool on Whose Terms?* UNESCO, Paris. 32

U.S. Department of Health and Human Services, Office of Income Security Policy (1983). *Overview of the Seattle–Denver Income Maintenance Experiment Final Report*. U.S. Government Printing Office, Washington, DC. 23, 79, 80

U.S. House of Representatives (2025). One big beautiful bill act, H.R. 1, 119th congress (house-passed version). Passed the House on May 22, 2025. Its provision allowing health savings account funds to pay for fitness expenses was removed before enactment. 37

West, R. W. (1980). Effects on wage rates: An interim analysis. *Journal of Human Resources*, 15(4):641–653. 83

Windmeijer, F. (2005). A Finite Sample Correction for the Variance of Linear Efficient Two-Step GMM Estimators. *Journal of Econometrics*, 126(1):25–51. 30

Wooldridge, J. M. (2007). Inverse probability weighted estimation for general missing data problems. *Journal of Econometrics*, 141(2):1281–1301. 24

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA, 2nd edition. 2, 14

A Additional Exhibits

Table A.1: Observed persistence and internal-instrument estimates with observable controls

	Labor supply	Social media use	Gym attendance
Panel A. Observed persistence (OLS, control group)			
$\hat{r}^{\mathrm{OLS}}$, no controls	0.884	0.898	0.836
	(0.005)	(0.009)	(0.043)
$\hat{r}^{\mathrm{OLS}}$, + observables	0.820	0.891	0.834
	(0.008)	(0.010)	(0.045)
Panel B. Internal-instrument propagation (difference GMM, lags $t-3$ to $t-6$)			
$\hat{r}^{\mathrm{int}}$, no controls	0.519	0.278	0.648
	(0.050)	(0.148)	(0.159)
$\hat{r}^{\mathrm{int}}$, + observables	0.512	—	—
	(0.049)		

Notes: Each column reports one application; standard errors in parentheses. Panel A reports the pooled OLS coefficient of the outcome on its own one-period lag in the control group, without fixed effects and with standard errors clustered by individual. "+ observables" adds all available time-invariant covariates (labor supply: randomization strata, age, education, family size, and baseline earnings, hours and employment; social media: age, gender, race, education and income; gym attendance: sex), estimated on the control-group subsample with complete covariates. Panel B reports two-step difference GMM (Arellano–Bond) estimates in the control group, instrumenting the lagged outcome with outcomes dated three to six periods earlier, with differenced time dummies and Windmeijer-corrected standard errors; for gym attendance, with 40 control participants, the instrument set is collapsed. In Panel B, "+ observables" adds strictly exogenous time-varying controls (other household income, welfare receipt and the number of children), available only in the labor-supply data; time-invariant covariates difference out of the estimating equation.

Figure A.1: Robustness to patience in labor supply

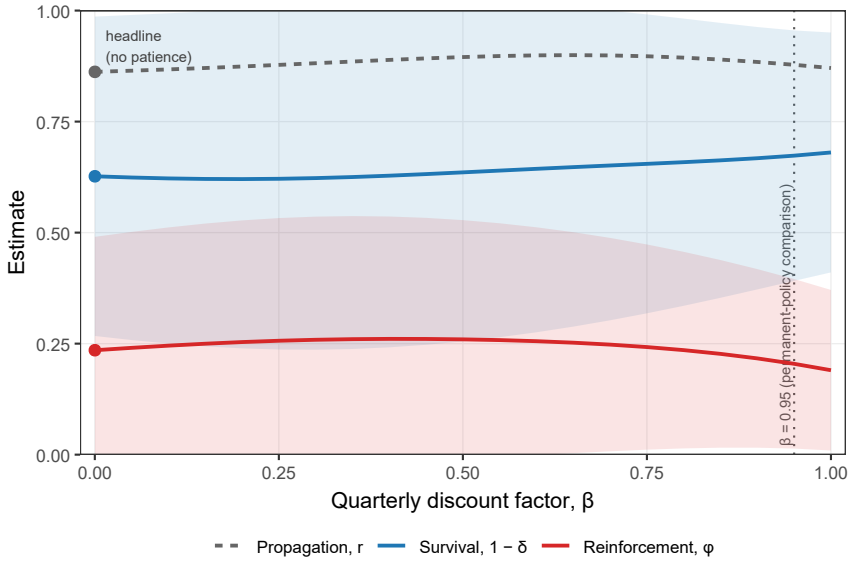

Notes: Each point on the horizontal axis re-estimates the model fixing the quarterly discount factor β and leaving all other parameters free, using the same doubly-robust specification, sample, and moments as the headline estimates. Shaded bands are 90% confidence intervals for ϕ and $1-\delta$. The headline estimates sit at $\beta = 0$, the case of no patience, and the dotted line marks $\beta = 0.95$, the degree of patience used for the permanent program in Figure 3.

Figure A.2: Out-of-sample validation: the five-year NIT arm

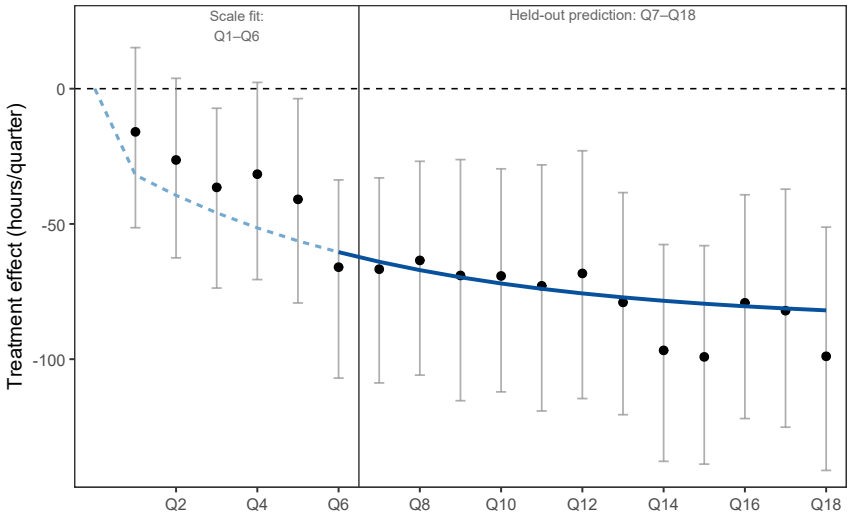

Notes:
This figure compares the treatment effects estimated in the five-year arm with the path predicted by the dynamic parameters estimated from the three-year arm. Points show treatment effects of the five-year SIME/DIME NIT on hours per quarter ($N = 558$), estimated by augmented inverse-probability weighting as in the three-year analysis and reweighted to the three-year arm's composition (Appendix C.1.5); bars are 90% family-cluster bootstrap confidence intervals ($B = 500$). The line is the model-implied path with all dynamic parameters fixed at the three-year estimates and the change in the cost of working starting in the first program quarter. Five-year outcomes are used only to estimate the scale of that change, by GLS on quarters 1–6 (a first-quarter effect of −31.9 hours, SE 10.8; dashed segment); quarters 7–18 are predicted out of sample (solid segment). The held-out prediction errors have an RMSE of 9.6 hours per quarter (MAE 6.4), and a joint test does not reject the model ($\chi^2(11) = 4.57$, $p = 0.95$).

Figure A.3: Robustness to patience in social media use

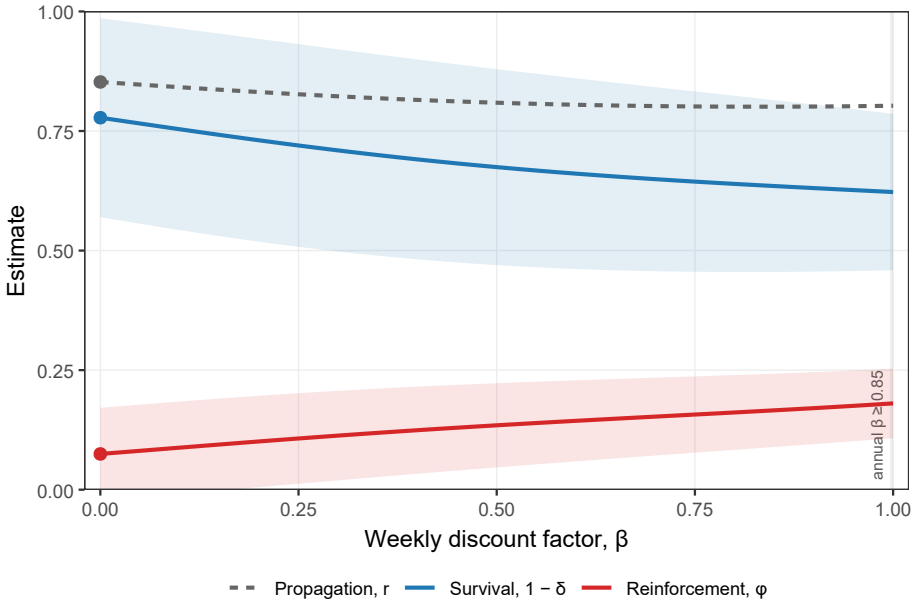

Notes: Each point on the horizontal axis re-estimates the model fixing the weekly discount factor β and leaving r, δ and the response to the bonus free, with the same sample, treatment effects and moments (weeks 7–15) as Table 2. Shaded bands are 90% confidence intervals for ϕ and $1-\delta$. The markers at $\beta = 0$ are the estimates in Table 2: $\hat{r} = 0.853$ (SE 0.078), $\hat{\phi} = 0.075$ (SE 0.059) and $1-\hat{\delta} = 0.778$ (SE 0.126). In the limit $\beta \to 1$, $\hat{r} = 0.803$ (SE 0.086), $\hat{\phi} = 0.180$ (SE 0.044) and $1-\hat{\delta} = 0.622$ (SE 0.099); these estimates give the orange bars in Figure 5. Reinforcement rises with β and is 0.18 for any annual discount factor of at least 0.85, marked at the right edge; propagation stays between 0.80 and 0.85. The estimation criterion varies by less than 2% across values of β, so the data do not determine β; Section 5.2 explains why we follow Allcott et al. (2022) in setting $\beta = 0$.

Figure A.4: Out-of-sample validation: the bonus among participants offered the time limit

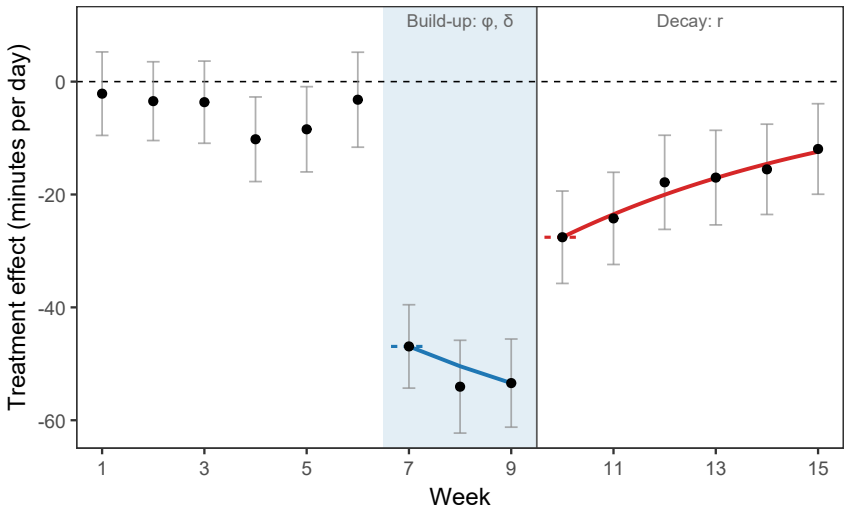

Notes: This figure reports weekly treatment effects of the bonus on FITSBY use, in minutes per day, among the 1,121 participants also offered the time limit treatment, with 90% confidence intervals; Allcott et al. (2022) randomized the bonus independently of the time limit, and these participants are not used in estimation. Effects are estimated by augmented inverse-probability weighting, as in Figure 4. The blue line is the build-up during the bonus implied by the reinforcement and survival estimated in Table 2, with $\beta = 0$ and its level set to the week-7 effect (dashed); the red line decays from the week-10 effect (dashed) at the estimated propagation rate, $\hat{r} = 0.853$. After the bonus, the held-out prediction errors have an RMSE of 1.1 minutes per day (MAE 0.9), and a joint test over the seven held-out weeks does not reject the model ($\chi^2(5) = 6.87$, $p = 0.23$).

Figure A.5: Out-of-sample validation: the decay of a one-week payment

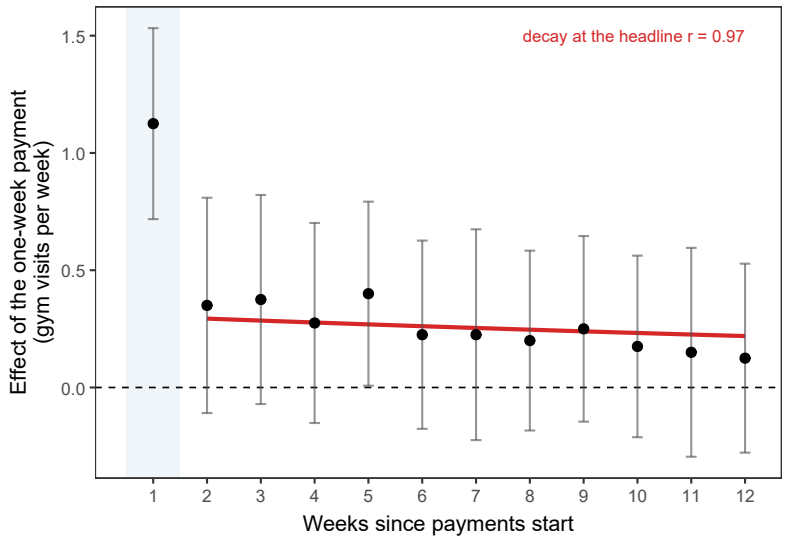

Notes: This figure reports week-specific differences in gym visits between the one-week group and the control group of Study 1 of Charness and Gneezy (2009), with 90% confidence intervals; the one-week group was paid $25 to visit the gym at least once in the first week (shaded), and is not used in estimation. The line decays from the second week at the propagation rate estimated from the five-week group, $\hat{r} = 0.97$ (Table 3); its level is fitted to the effects in weeks 2 to 12.

Table A.2: Estimated dynamics of gym attendance: robustness

	(1) Headline	(2) Without week 1	(3) One-week group as comparison	(4) Patient ($\beta \to 1$)
Comparison group	Control	Control	One week	Control
Payment weeks	1–5	2–5	2–5	1–5
Propagation rate, r	0.971	0.934	0.908	0.940
	(0.075)	(0.060)	(0.104)	(0.070)
Reinforcement, ϕ	0.185	0.243	0.268	0.174
	(0.045)	(0.061)	(0.084)	(0.045)
Survival, $1-\delta$	0.786	0.691	0.639	0.766
	(0.110)	(0.104)	(0.160)	(0.098)
Patience, β	0.000	0.000	0.000	$\to 1$
	(0.343)	(2.558)	(1.301)	(0.533)
ϕ/δ	0.865	0.786	0.744	0.743
Overidentification test	9.265	11.167	13.239	9.412
	[0.099]	[0.132]	[0.066]	[0.152]
Individuals	80	80	80	80

Notes: This table reports GMM estimates of the dynamic parameters of weekly gym visits from the experimental treatment-effect path of the five-week group. Column (1) reproduces column (3) of Table 3. Column (2) drops the first week, with its \$25 payment for a single visit, from the payment period, using the four weeks of the \$100 offer and the seven following weeks, with $\beta = 0$. Column (3) uses the one-week group as the comparison group over the same weeks, with $\beta = 0$; both groups were paid in the first week, and the \$100 offer was randomized among them at the return visit. Column (4) takes the limit $\beta \to 1$. In columns (2)–(4), the standard error of the imposed β is computed from the score at the restricted estimates, allowing the other parameters to adjust. Overidentification tests are Hansen J statistics. Standard errors are reported in parentheses, p-values in brackets.

B Proofs

B.1 Proofs for Section 2

In this section, we derive the Euler equation, several properties of the roots of the characteristic polynomial, and a sufficiency lemma. We then prove Lemma 1 and Proposition 1. We close by studying the path of the response to a permanent subsidy and its speed of convergence to the long run (Proposition B.1). Throughout this appendix, we also write the discount rate $\rho \equiv \beta^{-1} - 1$, so that $\beta = 1/(1+\rho)$. When $\beta = 0$ the problem is static in each period, and the results below hold with the limits of their expressions as $\rho \to \infty$.

B.1.1 Preliminaries

Euler equation. We characterize optimal choices in the model using a perturbation argument. Consider increasing y_t by a small amount dy and decreasing y_{t+1} by $(1-\delta)dy$, leaving y unchanged in all other periods. This perturbation raises S_{t+1} by dy but leaves S_{t+2} and all subsequent stocks unchanged. On the optimal path, this perturbation must not change expected utility. With U as in (4), this gives the Euler equation

$$\begin{aligned} &\alpha + U_{yy} y_t^\star + U_{yS} S_t - p_t \\ &\quad + \frac{1}{1+\rho}\, \mathcal{E}_t\big[\, \psi + U_{yS} y_{t+1}^\star + U_{SS} S_{t+1} - (1-\delta)\big(\alpha + U_{yy} y_{t+1}^\star + U_{yS} S_{t+1} - p_{t+1}\big)\big] = 0. \end{aligned} \tag{16}$$

The perturbation holds p_{t+1} fixed while raising S_{t+1}, which is valid because, as in Section 2.1, the cost does not depend on the stock:

$$p_t \text{ does not depend on } \{S_s\}_{s \le t}. \tag{17}$$

Everything the stock changes therefore runs through preferences $u(\cdot)$ and returns $R(\cdot)$. We assume (17) throughout. (Without it, the optimal policy remains affine when costs are linear in the stock, but ϕ inherits the coefficient on S_t in the cost process and is thus no longer a function of $(U_{yy}, U_{yS}, U_{SS}, \rho, \delta)$ alone.)

Roots of the characteristic polynomial. Define

$$\begin{aligned} A &\equiv U_{yy} + \frac{U_{SS} - (1-\delta)U_{yS}}{1+\rho}, \\ B &\equiv U_{yS} + \frac{(1-\delta)(U_{SS} - (1-\delta)U_{yS})}{1+\rho}, \\ C &\equiv \frac{U_{yS} - (1-\delta)U_{yy}}{1+\rho}. \end{aligned}$$

Further define $a_2 \equiv C$, $a_1 \equiv A - (1-\delta)C$, and $a_0 \equiv B - (1-\delta)A = (1+\rho)a_2$. Let P denote the characteristic polynomial

$$P(r) \equiv a_2 r^2 + a_1 r + a_0.$$

Condition (ii) of Lemma 1 ensures that P is a nondegenerate quadratic ($a_2 \neq 0$).[48] The roots of P play an important role in the identification argument to follow. The following lemma establishes several properties of these roots.

Lemma B.1. *(i) ϕ is a root of the polynomial $\phi \mapsto C\phi^2 + (A + (1-\delta)C)\phi + B$ if and only if $1 - \delta + \phi$ is a root of the characteristic polynomial P. If $1 - \delta + \phi$ is one root of P, the other root is*

$$\frac{1+\rho}{1-\delta+\phi} = -\frac{A + C\phi}{C}.$$

(ii) P has two distinct real roots.

(iii) P has one root satisfying $|r| < \sqrt{1+\rho}$ and one root satisfying $|r| > \sqrt{1+\rho}$. In particular, P has at most one stable root.

(iv) P has exactly one stable root if and only if

$$\left| \frac{A - (1-\delta)C}{C} \right| > 2 + \rho$$

or, equivalently,

$$(2+\rho)|U_{yS} - (1-\delta)U_{yy}| < -(1+\rho)U_{yy} - [(1-\delta)^2 U_{yy} - 2(1-\delta)U_{yS} + U_{SS}].$$

Strict concavity of U in (y, S) is not sufficient for the existence of a stable root.

(v) Both roots of P are positive if and only if

$$\frac{A - (1-\delta)C}{C} < 0;$$

otherwise, both roots are negative. If both roots are positive, then the stable root (if one exists) is the smaller of the two, and vice versa if both roots are negative.

Proof. (i) Let $r \equiv 1 - \delta + \phi$. If ϕ is a root of $\phi \mapsto C\phi^2 + (A + (1-\delta)C)\phi + B$, then r is a root of

[48] As $\rho \to \infty$, $a_2 \to 0$ and P becomes linear.

the quadratic

$$\begin{aligned}
&C(r-(1-\delta))^2+(A+(1-\delta)C)(r-(1-\delta))+B\\
&\quad=Cr^2+(A-(1-\delta)C)r+C(1-\delta)^2-(A+(1-\delta)C)(1-\delta)+B\\
&\quad=Cr^2+(A-(1-\delta)C)r+(B-(1-\delta)A)\\
&\quad=P(r).
\end{aligned}$$

By the same argument, if r is a root of P, then $\phi\equiv r-(1-\delta)$ is a root of $C\phi^2+(A+(1-\delta)C)\phi+B=0$.

Next, we use Vieta's formula to calculate that the product of the two roots is

$$\frac{B-(1-\delta)A}{C}=\frac{U_{yS}-(1-\delta)U_{yy}}{[U_{yS}-(1-\delta)U_{yy}]/(1+\rho)}=1+\rho.$$

Thus if $r=1-\delta+\phi$ is one root, the other root is $(1+\rho)/r$. Using Vieta's formula again, the sum of the two roots is

$$1-\delta+\phi+\frac{1+\rho}{1-\delta+\phi}=-\frac{A-(1-\delta)C}{C}=1-\delta-\frac{A}{C}.$$

Subtracting $1-\delta+\phi$ from both sides gives

$$\frac{1+\rho}{1-\delta+\phi}=-\frac{A+C\phi}{C}.$$

(ii) P is a nondegenerate quadratic polynomial ($a_2=C\neq 0$ by condition (ii) of Lemma 1), so it has two distinct real roots if and only if its discriminant

$$\Delta_P\equiv(A-(1-\delta)C)^2-4C(B-(1-\delta)A)=(A-(1-\delta)C)^2-4(1+\rho)C^2$$

is strictly positive. Plugging in the definitions of A and C, we compute

$$\begin{aligned}
(1+\rho)^2\Delta_P&=\big[(1+\rho)U_{yy}+(U_{SS}-2(1-\delta)U_{yS}+(1-\delta)^2U_{yy})\big]^2\\
&\quad-4(1+\rho)[U_{yS}-(1-\delta)U_{yy}]^2\\
&=\big[(1+\rho)U_{yy}-(U_{SS}-2(1-\delta)U_{yS}+(1-\delta)^2U_{yy})\big]^2\\
&\quad+4(1+\rho)U_{yy}(U_{SS}-2(1-\delta)U_{yS}+(1-\delta)^2U_{yy})\\
&\quad-4(1+\rho)[U_{yS}-(1-\delta)U_{yy}]^2\\
&=\big[(1+\rho)U_{yy}-(U_{SS}-2(1-\delta)U_{yS}+(1-\delta)^2U_{yy})\big]^2\\
&\quad+4(1+\rho)[U_{yy}U_{SS}-U_{yS}^2].
\end{aligned}$$

The first term is nonnegative because it is a square, while the second is strictly positive by strict concavity of U. Thus Δ_P is strictly positive.

(iii) By parts (i) and (ii), the two roots of P are real, distinct, and their product is $1+\rho$. Thus one must have magnitude strictly less than $\sqrt{1+\rho}$ and the other magnitude strictly greater than $\sqrt{1+\rho}$.

(iv) By (iii), since $\rho > 0$, P has at most one root in the interval $(-1, 1)$. If $P(-1)$ and $P(1)$ have different signs, then P has exactly one root in $(-1, 1)$. We calculate

$$\frac{P(-1)}{C}\frac{P(1)}{C} = [1 + (1+\rho) - (A - (1-\delta)C)/C][1 + (1+\rho) + (A - (1-\delta)C)/C],$$

which is negative (implying that $P(-1)$ and $P(1)$ have opposite signs) if and only if

$$\left|\frac{A - (1-\delta)C}{C}\right| > 2 + \rho.$$

By inserting the definitions of A and C, we rewrite this condition as

$$\left|\frac{(1+\rho)U_{yy} + [(1-\delta)^2 U_{yy} - 2(1-\delta)U_{yS} + U_{SS}]}{U_{yS} - (1-\delta)U_{yy}}\right| > 2 + \rho.$$

The numerator of the fraction on the left-hand side is strictly negative by strict concavity of U, so we can further rewrite the condition as

$$(2+\rho)|U_{yS} - (1-\delta)U_{yy}| < -(1+\rho)U_{yy} - [(1-\delta)^2 U_{yy} - 2(1-\delta)U_{yS} + U_{SS}].$$

To see that concavity is insufficient for this condition to hold, take $U_{yy} = -1$, $U_{SS} = -0.5$, and $U_{yS} = 0.69$. Then $U_{yy}U_{SS} - U_{yS}^2 = 0.0239 > 0$, so U is strictly concave. Set $\delta = \rho = 0.5$. Then $a_2 = 0.7933$, $a_1 = -1.96$, and $a_0 = 1.19$. The roots of the characteristic polynomial P are 1.074 and 1.396, both of which are outside the unit circle.

(v) The product of the two roots of P is $1 + \rho > 0$, so the two roots have the same sign. The sum of the two roots is

$$-\frac{A - (1-\delta)C}{C},$$

so both roots are positive if and only if

$$\frac{A - (1-\delta)C}{C} < 0.$$

In this case, the root satisfying the stability condition (if one does) must be the smaller of the two. □

Sufficiency of the Euler and transversality conditions. To prove Lemma 1, we need the following auxiliary lemma.

Lemma B.2. *Suppose $\beta > 0$, condition (i) of Lemma 1 holds, and $\{y_t\}$ is feasible for* (1)*, satisfies the Euler equation* (16) *and the transversality condition*

$$\lim_{t\to\infty} \frac{1}{(1+\rho)^t} \mathcal{E}_0[\mu_t S_t] = 0, \tag{18}$$

where

$$\mu_t \equiv \psi + U_{yS} y_t^\star + U_{SS} S_t - (1-\delta)[\alpha - p_t + U_{yy} y_t^\star + U_{yS} S_t]$$

is the marginal value of stock in period t, and achieves expected utility $> -\infty$. Then $\{y_t\}$ is the unique optimal solution to (1)*.*

Proof. Consider any alternative feasible path $\{\tilde{y}_t\}_{t=0}^\infty$. We will show that

$$\mathcal{D} \equiv \lim_{T\to\infty} \sum_{t=0}^{T} \frac{1}{(1+\rho)^t} \mathcal{E}_0[U(y_t, S_t) - p_t y_t - U(\tilde{y}_t, \tilde{S}_t) + p_t \tilde{y}_t] > 0,$$

so $\{\tilde{y}_t\}_{t=0}^\infty$ cannot be optimal.

First, note that it suffices to consider alternative feasible paths satisfying

$$\lim_{T\to\infty} \sum_{t=0}^{T} \frac{1}{(1+\rho)^t} \mathcal{E}_0[\tilde{y}_t^2 + \tilde{S}_t^2] < \infty.$$

To see this, note that because $(\tilde{y}_t, \tilde{S}_t) \mapsto U(\tilde{y}_t, \tilde{S}_t) - p_t \tilde{y}_t$ is a strictly concave linear-quadratic function and p_t is uniformly bounded under the individual's beliefs, there exist finite constants $M_0, M_1 > 0$ such that $U(\tilde{y}_t, \tilde{S}_t) - p_t \tilde{y}_t < M_0 - M_1(\tilde{y}_t^2 + \tilde{S}_t^2)$ for all t. Then

$$\sum_{t=0}^{T} \frac{1}{(1+\rho)^t} \mathcal{E}_0[U(\tilde{y}_t, \tilde{S}_t) - p_t \tilde{y}_t] \leq \frac{M_0}{1 - 1/(1+\rho)} - M_1 \sum_{t=0}^{T} \frac{1}{(1+\rho)^t} \mathcal{E}_0[\tilde{y}_t^2 + \tilde{S}_t^2]$$

for all T. Any path with $\lim_{T\to\infty} \sum_{t=0}^{T} (1+\rho)^{-t} \mathcal{E}_0[\tilde{y}_t^2 + \tilde{S}_t^2] = \infty$ therefore provides expected utility $-\infty$. Since $\{y_t\}$ provides expected utility $> -\infty$ by assumption, such a path cannot be optimal. Moreover, since $\{y_t\}$ provides expected utility $> -\infty$, it also follows that

$$\lim_{T\to\infty} \sum_{t=0}^{T} (1+\rho)^{-t} \mathcal{E}_0[y_t^2 + S_t^2] < \infty.$$

The next part of the proof follows the proof of Theorem 4.15 of Stokey and Lucas (1989).

Define

$$F_t(S, S') \equiv U(S' - (1-\delta)S, S) - p_t[S' - (1-\delta)S].$$

Rewriting in terms of F_t, the problem (1) becomes

$$\max_{\{S_{t+1}\}} \sum_{t=0}^{\infty} \frac{1}{(1+\rho)^t} \mathcal{E}_0[F_t(S_t, S_{t+1})] \qquad \text{s.t.} \quad S_0 \text{ given},$$

the Euler equation (16) becomes

$$F_{t,S'}(S_t, S_{t+1}) + \frac{1}{1+\rho} \mathcal{E}_t[F_{t+1,S}(S_{t+1}, S_{t+2})] = 0,$$

and the transversality condition (18) becomes

$$\lim_{t\to\infty} \frac{1}{(1+\rho)^t} \mathcal{E}_0[F_{t,S}(S_t, S_{t+1}) S_t] = 0.$$

We now calculate

$$\begin{aligned}
\mathcal{D} &\equiv \lim_{T\to\infty} \sum_{t=0}^{T} \frac{1}{(1+\rho)^t} \mathcal{E}_0[F_t(S_t, S_{t+1}) - F_t(\tilde{S}_t, \tilde{S}_{t+1})] \\
&> \lim_{T\to\infty} \sum_{t=0}^{T} \frac{1}{(1+\rho)^t} \mathcal{E}_0[F_{t,S}(S_t, S_{t+1})(S_t - \tilde{S}_t) + F_{t,S'}(S_t, S_{t+1})(S_{t+1} - \tilde{S}_{t+1})] \\
&= \lim_{T\to\infty} \sum_{t=0}^{T-1} \frac{1}{(1+\rho)^t} \mathcal{E}_0 \left[\left(\frac{1}{1+\rho} F_{t+1,S}(S_{t+1}, S_{t+2}) + F_{t,S'}(S_t, S_{t+1}) \right) (S_{t+1} - \tilde{S}_{t+1}) \right] \\
&\qquad + \frac{1}{(1+\rho)^T} \mathcal{E}_0[F_{T,S'}(S_T, S_{T+1})(S_{T+1} - \tilde{S}_{T+1})] \\
&= - \lim_{T\to\infty} \frac{1}{(1+\rho)^T} \mathcal{E}_0[F_{T,S'}(S_T, S_{T+1}) \tilde{S}_{T+1}].
\end{aligned}$$

The second line follows from strict concavity of F_t. The third line follows from rearranging the sum and noting that $S_0 = \tilde{S}_0$ by feasibility. The fourth line follows from the Euler equation and the transversality condition. To complete the proof, we use the Cauchy–Schwarz inequality to bound the limit above:

$$\begin{aligned}
&\left| \lim_{T\to\infty} \frac{1}{(1+\rho)^T} \mathcal{E}_0[F_{T,S'}(S_T, S_{T+1}) \tilde{S}_{T+1}] \right| \\
&\qquad \leq \lim_{T\to\infty} \left(\frac{1}{(1+\rho)^T} \mathcal{E}_0[F_{T,S'}(S_T, S_{T+1})^2] \right)^{1/2} \left(\frac{1}{(1+\rho)^T} \mathcal{E}_0[\tilde{S}_{T+1}^2] \right)^{1/2}.
\end{aligned}$$

Since both $\{(y_t, S_t)\}$ and $\{(\tilde{y}_t, \tilde{S}_t)\}$ are square-summable sequences, both terms on the right-hand side must converge to 0 as $T \to \infty$. We conclude that $\mathcal{D} > 0$, as was to be shown. □

B.1.2 Proof of Lemma 1

The proof proceeds as follows. We conjecture an optimal policy of the form $y_t^\star = \phi S_t + K\alpha + \pi_t$, where π_t is a function of current and expected future costs but not of S_t or α. We solve for the unique values of ϕ, K, and π_t for which this policy satisfies the Euler equation (16) and the transversality condition (18), then verify that it provides expected utility $> -\infty$.[49] By Lemma B.2, the conjecture is then the unique optimal policy.

The conjectured policy together with the stock law of motion implies

$$S_{t+1} = (1 - \delta + \phi)S_t + K\alpha + \pi_t$$

and

$$\mathcal{E}_t[y_{t+1}^\star] = \phi(1 - \delta + \phi)S_t + \phi(K\alpha + \pi_t) + K\alpha + \mathcal{E}_t[\pi_{t+1}].$$

Plugging these expressions into the Euler equation (16) gives

$$\begin{aligned}
&\alpha + U_{yy}(\phi S_t + K\alpha + \pi_t) + U_{yS}S_t - p_t \\
&+ \frac{1}{1+\rho}\Bigg[\psi + [U_{yS} - (1-\delta)U_{yy}](\phi(1-\delta+\phi)S_t + \phi[K\alpha + \pi_t] + \mathcal{E}_t[K\alpha + \pi_{t+1}]) \\
&\qquad + [U_{SS} - (1-\delta)U_{yS}]((1-\delta+\phi)S_t + K\alpha + \pi_t) - (1-\delta)(\alpha - \mathcal{E}_t[p_{t+1}])\Bigg] = 0.
\end{aligned}$$

Collecting terms, we have

$$\begin{aligned}
&\left[C\phi^2 + (A + (1-\delta)C)\phi + B\right] S_t + [A + C\phi]\,(K\alpha + \pi_t) + C\mathcal{E}_t[K\alpha + \pi_{t+1}] \\
&\qquad = -\left(1 - \frac{1-\delta}{1+\rho}\right)\alpha + p_t - \frac{1-\delta}{1+\rho}\mathcal{E}_t[p_{t+1}] - \frac{\psi}{1+\rho},
\end{aligned} \tag{19}$$

where A, B, C are as defined in Appendix B.1.1.

Solution for ϕ. Since costs do not depend on the stock (17), we can match coefficients on S_t in (19) to obtain a quadratic equation for ϕ:

$$C\phi^2 + (A + (1-\delta)C)\phi + B = 0. \tag{20}$$

Thus ϕ is a root of the left-hand side quadratic. (The coefficients of this quadratic depend only on $(U_{yy}, U_{yS}, U_{SS}, \beta, \delta)$, so ϕ depends only on these parameters.) By Lemma B.1(i), $1 - \delta + \phi$ must

[49] When $\beta = 0$, Lemma B.2 does not apply, and it is not needed: the problem is static in each period, so the optimal choice solves $\alpha + U_{yy}y_t + U_{yS}S_t - p_t = 0$. This solution is the limit of the policy derived below as $\rho \to \infty$: $\phi \to -U_{yS}/U_{yy}$, $K \to -1/U_{yy}$, and $\pi_t \to p_t/U_{yy}$.

be a root of the characteristic polynomial P. Two values of ϕ satisfy (20). By Lemma B.1(iii), one satisfies $|1-\delta+\phi| < \sqrt{1+\rho}$, while the other satisfies $|1-\delta+\phi| > \sqrt{1+\rho}$. Since $S_{t+1} = (1-\delta+\phi)S_t + K\alpha + \pi_t$, the latter choice of ϕ yields a policy under which $\mu_t S_t$ grows at a rate faster than $1+\rho$, violating the transversality condition (18). Thus ϕ must be the unique root of (20) satisfying $|1-\delta+\phi| < \sqrt{1+\rho}$.

As $\rho \to \infty$, we have $C \to 0$, $A \to U_{yy}$, and $B \to U_{yS}$, so (20) reduces to $U_{yy}\phi + U_{yS} = 0$ and the selected root tends to $-U_{yS}/U_{yy}$, as stated in Section 2. (The other root diverges.)

Solution for K. Matching coefficients on α in (19) yields an equation for K:

$$K = -\frac{1 - \frac{1-\delta}{1+\rho}}{A + C(1+\phi)}.$$

Since we already found a unique solution for ϕ, this equation uniquely determines K.

Solution for π_t. With ϕ chosen to be the unique root of (20) satisfying $|1-\delta+\phi| < \sqrt{1+\rho}$ and $K = -\frac{1-\frac{1-\delta}{1+\rho}}{A+C(1+\phi)}$, (19) reduces to a recursive equation for π_t:

$$(A + C\phi)\pi_t + C\mathcal{E}_t[\pi_{t+1}] = p_t - \frac{1-\delta}{1+\rho}\mathcal{E}_t[p_{t+1}] - \frac{\psi}{1+\rho}.$$

By Lemma B.1(i) and (iii), $\left|-\frac{C}{A+C\phi}\right| = \left|\frac{1-\delta+\phi}{1+\rho}\right| < 1$. Because information accumulates over time (Section 2.1), iterated expectations apply, and the recursion can be solved forward to obtain

$$\pi_t = \frac{1}{A+C\phi}\sum_{k=0}^{\infty}\left(\frac{1-\delta+\phi}{1+\rho}\right)^k\left(\mathcal{E}_t[p_{t+k}] - \frac{1-\delta}{1+\rho}\mathcal{E}_t[p_{t+k+1}] - \frac{\psi}{1+\rho}\right). \tag{21}$$

Since we already found a unique solution for ϕ, this equation determines the unique bounded solution for π_t for all t.

Verifying the transversality condition. We next verify that the conjectured policy $y_t^\star = \phi S_t + K\alpha + \pi_t$ satisfies the transversality condition (18). Write $r \equiv 1-\delta+\phi$.

Since costs p_t are uniformly bounded under the individual's beliefs (condition (i) of Lemma 1) and the geometric factor $r/(1+\rho)$ in (21) has modulus less than 1, π_t is uniformly bounded as well: $\sup_t |\pi_t| \leq \Pi < \infty$. Because $|r| < \sqrt{1+\rho}$ by Lemma B.1(iii) and $1 < \sqrt{1+\rho}$, we can choose ν such that

$$\max\{|r|, 1\} < \nu < \sqrt{1+\rho}.$$

Iterating the law of motion from S_0 yields the upper bound

$$\left(\mathcal{E}_0[S_t^2]\right)^{1/2} \leq |r|^t \left(\mathcal{E}_0[S_0^2]\right)^{1/2} + (|K\alpha| + \Pi) \sum_{j=0}^{t-1} |r|^j = O(\nu^t).$$

Because $y_t^\star$ and μ_t are affine in $(S_t, \alpha, \pi_t, p_t)$, $(\mathcal{E}_0[y_t^2])^{1/2}$ and $(\mathcal{E}_0[\mu_t^2])^{1/2}$ are also $O(\nu^t)$. The Cauchy–Schwarz inequality then gives

$$\frac{1}{(1+\rho)^t} \left|\mathcal{E}_0[\mu_t S_t]\right| \leq \frac{\left(\mathcal{E}_0[\mu_t^2]\right)^{1/2} \left(\mathcal{E}_0[S_t^2]\right)^{1/2}}{(1+\rho)^t} = O\left(\left(\frac{\nu^2}{1+\rho}\right)^t\right).$$

Since $\nu^2 < 1 + \rho$, we conclude that $\frac{1}{(1+\rho)^t} \left|\mathcal{E}_0[\mu_t S_t]\right| \to 0$ as $t \to \infty$.

Verifying finite expected utility. Finally, we verify finite expected utility. The same bound used above gives $\mathcal{E}_0[y_t^2 + S_t^2] = O(\nu^{2t})$, so

$$\sum_{t=0}^{\infty} \frac{1}{(1+\rho)^t} \mathcal{E}_0\left[y_t^2 + S_t^2\right] < \infty.$$

By an argument similar to the one employed in the proof of Lemma B.2, it follows that the path $\{(y_t, S_t)\}$ provides finite expected utility. □

B.1.3 Proof of Proposition 1

(i) Consider the one-week subsidy that defines $\theta_t^{(k)}$ in (3). It changes the date-t choice by dy_t, holding the stock S_t fixed. It leaves later costs unchanged and, because the intervention does not change beliefs about baseline costs (Section 2.4), expectations of them as well. Write dS_{t+k} and dy_{t+k} for the differences between the perturbed and unperturbed paths k periods after t. By construction, $dS_t = 0$, so

$$dS_{t+1} = (1-\delta)\, dS_t + dy_t = dy_t.$$

From $t+1$ onward the individual follows the policy of Lemma 1. Because later costs and expectations of them are unchanged, $d\pi_{t+k} = 0$ for every $k \geq 1$, and because α is unchanged, $d(K\alpha) = 0$. Hence for every $k \geq 1$,

$$dy_{t+k} = \phi\, dS_{t+k}, \qquad dS_{t+k+1} = (1-\delta)\, dS_{t+k} + dy_{t+k} = r\, dS_{t+k}.$$

Iterating from $dS_{t+1} = dy_t$ gives $dS_{t+k} = r^{k-1} dy_t$, and hence $dy_{t+k} = \phi r^{k-1} dy_t$ for every $k \geq 1$. Since both changes are proportional to db, their ratio is the impulse response:

$$\theta_t^{(k)} = \left. \frac{dy_{t+k}^\star / db}{dy_t^\star / db} \right|_{S_t,\, b=0,\, \mathcal{W}=\{t\}} = \phi\, r^{k-1},$$

which does not depend on t.

(ii) The temporary subsidy lowers the cost at t only and conveys no information about later costs. By (21), it therefore changes π_t, through the $k = 0$ term, by $d\pi_t = -db/(A + C\phi)$, and leaves π_{t+k} unchanged for every $k \geq 1$. Hence $dy_t^\star = d\pi_t$, and from $t+1$ on, the additional choice affects later choices only through the stock, so $dy_{t+k}^\star = \theta^{(k)} dy_t^\star$. Using the result of part (i), we have

$$\text{CLE} = \sum_{k=1}^{\infty} \theta^{(k)} = \frac{\phi}{1-r} = \frac{\phi}{\delta - \phi}.$$

The sum converges because $|r| < 1$.

(iii) A permanent unit decrease in all costs $\{p_{t+\ell}\}_{\ell \geq 0}$ lowers the cost and its expectation at every later date by one, since the intervention does not change beliefs about baseline costs (Section 2.4). By (21), it therefore changes π_{t+k} for every $k \geq 0$ by

$$\begin{aligned} d\pi_{t+k} &= -\frac{1}{A + C\phi} \sum_{j=0}^{\infty} \left(\frac{1 - \delta + \phi}{1 + \rho} \right)^j \left(1 - \frac{1-\delta}{1+\rho} \right) \\ &= -\frac{1}{A + C\phi} \frac{1}{1 - \frac{1-\delta+\phi}{1+\rho}} \left(1 - \frac{1-\delta}{1+\rho} \right) \\ &= -\frac{1}{A + C\phi} \frac{\rho + \delta}{\rho + \delta - \phi}. \end{aligned}$$

Note that $d\pi_{t+k}$ does not depend on k. Since

$$y_{t+k}^\star = \phi(1 - \delta + \phi)^k S_t + \sum_{\ell=1}^{k} \phi(1 - \delta + \phi)^{\ell-1} (K\alpha + \pi_{t+k-\ell}) + K\alpha + \pi_{t+k},$$

and $dS_t = 0$, we have

$$\begin{aligned} dy_{t+k}^\star &= -\frac{1}{A + C\phi} \frac{\rho + \delta}{\rho + \delta - \phi} \left[1 + \sum_{\ell=1}^{k} \phi(1 - \delta + \phi)^{\ell-1} \right] \\ &= -\frac{1}{A + C\phi} \frac{\rho + \delta}{\rho + \delta - \phi} \left[1 + \phi \frac{1 - (1 - \delta + \phi)^k}{\delta - \phi} \right]. \end{aligned}$$

As $k \to \infty$,

$$\begin{aligned} dy^{\star}_{t+k} &\to -\frac{1}{A + C\phi} \frac{\rho + \delta}{\rho + \delta - \phi} \left[1 + \frac{\phi}{\delta - \phi} \right] \\ &= -\frac{1}{A + C\phi} \frac{\rho + \delta}{\rho + \delta - \phi} \left[\frac{\delta}{\delta - \phi} \right] \\ &= -\frac{1}{A + C\phi} \frac{1 - \beta(1 - \delta)}{1 - \beta r} \left[\frac{\delta}{\delta - \phi} \right]. \end{aligned}$$

As noted in the proof of part (ii), the contemporaneous response to a temporary unit decrease in the cost is $-1/(A + C\phi)$. Dividing the long-run response by this contemporaneous response gives

$$\text{LRM} = \frac{\delta}{\delta - \phi} \frac{1 - \beta(1 - \delta)}{1 - \beta r}.$$

□

B.1.4 Convergence to the Long Run

Proposition 1(iii) characterizes the long-run response to a permanent intervention. The next result gives the path to it; the applications use this path to report the share of the long run reached after a given number of periods.

Proposition B.1 (The path under a permanent incentive)**.** *Suppose the conditions of Lemma 1 hold, let $r \equiv \phi + 1 - \delta$, and assume $|r| < 1$.*

(i) For every $k \geq 0$,

$$\left. \frac{dy^{\star}_{\underline{t}+k}}{db} \right|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{PERM}}} = \left. \frac{dy^{\star}_{\underline{t}}}{db} \right|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{TEMP}}} \times \frac{1 - \beta(1 - \delta)}{1 - \beta r} \left[1 + \phi \frac{1 - r^k}{1 - r} \right].$$

As $k \to \infty$, this expression converges to the contemporaneous response to a temporary subsidy multiplied by the long-run multiplier LRM*. Convergence is monotonic when $\phi > 0$ and $0 < r < 1$.*

(ii) The share of the long-run response reached k periods after the start is

$$\frac{\left. \dfrac{dy^{\star}_{\underline{t}+k}}{db} \right|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{PERM}}}}{\text{LRM} \times \left. \dfrac{dy^{\star}_{\underline{t}}}{db} \right|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{TEMP}}}} = 1 - \frac{\phi}{\delta} r^k,$$

so the remaining gap shrinks by a factor r every period. When $\phi > 0$ and $0 < r < 1$, the gap halves every $\ln 2/\ln(1/r)$ periods, and a share $s \in (1-\phi/\delta, 1)$ is reached after $\ln\big[(1-s)\delta/\phi\big]/\ln r$ periods.

(iii) The gap after k periods—the long-run response minus the response at $\underline{t}+k$—equals

$$\frac{1-\beta(1-\delta)}{1-\beta r}\sum_{j>k}\theta^{(j)}$$

times the contemporaneous response to the temporary subsidy.

The speed of convergence is governed by the propagation rate r. The share of the long-run response reached after a given number of periods also depends on ϕ/δ. Patience changes the level of the whole path through the factor $\frac{1-\beta(1-\delta)}{1-\beta r}$.

Proof. Write $dy^{\star}_{\underline{t}+k}$ for the response at $\underline{t}+k$ to the permanent subsidy and $d\pi_{\underline{t}+k}$ for the corresponding change in the cost outlook, as in the proof of Proposition 1(iii). As shown in that proof, $d\pi_{\underline{t}+k} = -\frac{1}{A+C\phi}\frac{1-\beta(1-\delta)}{1-\beta r}$ for every $k \geq 0$, and the contemporaneous response to the temporary subsidy equals $-1/(A+C\phi)$.

Starting from $dS_{\underline{t}} = 0$ and iterating $dS_{\underline{t}+k+1} = r\,dS_{\underline{t}+k} + d\pi_{\underline{t}+k}$ forward gives

$$dS_{\underline{t}+k} = d\pi_{\underline{t}}\frac{1-r^k}{1-r}.$$

Plugging this expression into $dy^{\star}_{\underline{t}+k} = \phi\,dS_{\underline{t}+k} + d\pi_{\underline{t}+k}$ gives part (i). Since $|r| < 1$, the path converges to the contemporaneous response times $\frac{1-\beta(1-\delta)}{1-\beta r}\big[1+\frac{\phi}{1-r}\big] = \frac{1-\beta(1-\delta)}{1-\beta r}\frac{\delta}{\delta-\phi} = \text{LRM}$. The convergence is monotone in k when $\phi > 0$ and $0 < r < 1$ because r^k is monotone in this case.

Dividing the path by its limit gives

$$\frac{1+\phi(1-r^k)/(1-r)}{\delta/(\delta-\phi)} = \frac{\delta-\phi r^k}{\delta},$$

which is part (ii). The expressions for the half-life and for the time to reach share s then follow by solving $r^k = 1/2$ and $(\phi/\delta)r^k = 1-s$ for k.

To derive part (iii), note that, by (i), the gap is the contemporaneous response times

$$\frac{1-\beta(1-\delta)}{1-\beta r}\,\frac{\phi\,r^k}{\delta-\phi},$$

and $\sum_{j>k}\theta^{(j)} = \sum_{j>k}\phi\,r^{\,j-1} = \phi\,r^k/(1-r)$. □

B.2 Proofs for Section 3.3

This subsection starts with two lemmas: a general one establishing that a valid instrument identifies r (Lemma B.3), and one giving the error of Equation (6) under Assumption I.1 (Lemma B.4). The proofs of Propositions 2 and 3 verify the conditions of Lemma B.3 for internal and external instruments and extend it to propagation rates that differ across individuals. We treat temporary experiments in Appendix B.3.

B.2.1 Preliminaries

Individual parameters. Section 3.1 lets the type $\omega_i \equiv (U_{yy,i}, U_{yS,i}, U_{SS,i}, \beta_i, \delta_i)$, the parameters that govern propagation, vary across individuals. We assume that the conditions of Lemma 1 hold for every type, so Lemmas 1 and B.1 apply individual by individual. The resulting slope ϕ_i, loading K_i, and propagation rate $r_i = 1 - \delta_i + \phi_i$ therefore depend only on ω_i, which is time-independent and known to the individual at date 0.

We assume throughout that the expectations in Section 3 are finite.

A statistical identification result. The following lemma isolates the statistical content of the identification argument for r.

Lemma B.3 (Identification of r from a valid instrument)**.** *Suppose observed choices satisfy* (6) *with common propagation, $r_i = r$. Let h be a known linear transformation and Z_{it} a candidate instrument. If* $\mathrm{Cov}(Z_{it}, h(a_i + e_{it})) = 0$ *and* $\mathrm{Cov}(Z_{it}, h(Y_{i,t-1})) \neq 0$*, then*

$$r = \frac{\mathrm{Cov}(Z_{it}, h(Y_{it}))}{\mathrm{Cov}(Z_{it}, h(Y_{i,t-1}))}, \tag{22}$$

so r is identified.

Proof. Applying h to (6) and using linearity gives $h(Y_{it}) = r\, h(Y_{i,t-1}) + h(a_i + e_{it})$. Taking the covariance of both sides with Z_{it} and using the first condition gives $\mathrm{Cov}(Z_{it}, h(Y_{it})) = r\ \mathrm{Cov}(Z_{it}, h(Y_{i,t-1}))$. Dividing by $\mathrm{Cov}(Z_{it}, h(Y_{i,t-1})) \neq 0$ gives (22). Both covariances are functions of the joint distribution of observables, so r is identified. □

The differenced error under unpredictable cost innovations. The next lemma gives the error of (6) under Assumption I.1; the proof of Proposition 2 uses it.

Lemma B.4 (The error under unpredictable costs)**.** *Suppose observed choices satisfy* (5) *and Assumption I.1 holds, so that $p_{it} = \bar{p}_i + \eta_{it}$ with $\mathbb{E}[\eta_{it} \mid \mathcal{F}_{i,t-1}] = \mathcal{E}_{i,t-1}[\eta_{it}] = 0$. Then*

$$\pi_{it} = \bar{\pi}_i + \xi_i\, \eta_{it}, \qquad \xi_i \equiv -\frac{r_i}{U_{yS,i} - (1-\delta_i) U_{yy,i}}, \tag{23}$$

with $\xi_i \neq 0$*, and*

$$e_{it} = \delta_i \bar{\pi}_i + \xi_i\big(\eta_{it} - (1-\delta_i)\eta_{i,t-1}\big), \qquad \Delta e_{it} = \xi_i\big(\eta_{it} - (2-\delta_i)\eta_{i,t-1} + (1-\delta_i)\eta_{i,t-2}\big). \tag{24}$$

Proof. By the subjective half of Assumption I.1, $\mathcal{E}_{i,t-1}[\eta_{it}] = 0$, applied at each horizon through iterated expectations under the individual's beliefs, $\mathcal{E}_{it}[p_{i,t+k}] = \bar{p}_i$ for every $k \geq 1$ and $\mathcal{E}_{it}[p_{it}] = p_{it}$. Substituting into (21), shows that every forward term is constant over time for each individual and that the only period-t variation enters through the $k = 0$ term. Collecting all constants into $\bar{\pi}_i$ gives (23), with $\xi_i = 1/(A_i + C_i\phi_i)$, which is nonzero; Lemma B.1(i) gives the expression for ξ_i in (23).

Substituting (23) into $e_{it} = \pi_{it} - (1-\delta_i)\pi_{i,t-1}$ from (6) and using $\bar{\pi}_i - (1-\delta_i)\bar{\pi}_i = \delta_i\bar{\pi}_i$ gives the first expression in (24). Differencing once more gives the second, since $\eta_{it} - \eta_{i,t-1} - (1-\delta_i)(\eta_{i,t-1} - \eta_{i,t-2}) = \eta_{it} - (2-\delta_i)\eta_{i,t-1} + (1-\delta_i)\eta_{i,t-2}$. □

Two features of (24) drive the next proof. First, the differenced error Δe_{it} loads on exactly three innovations, those in periods t, $t-1$, and $t-2$, with weights $\big(1,\ -(2-\delta_i),\ 1-\delta_i\big)$. Second, the weight on the period $t-2$ innovation vanishes if and only if $\delta_i = 1$.

B.2.2 Internal Instruments

Proof of Proposition 2. (i) Fix ℓ with $3 \leq \ell \leq t$. By Lemma B.4,

$$\Delta e_{it} = \xi_i \sum_{j=0}^{2} w_{ij}\, \eta_{i,t-j}, \qquad (w_{i0}, w_{i1}, w_{i2}) = \big(1,\ -(2-\delta_i),\ 1-\delta_i\big).$$

By (5) and (23), $Y_{i,t-\ell}$ is a function of $S_{i,t-\ell}$, α_i, ω_i, $\bar{\pi}_i$ and $\eta_{i,t-\ell}$, all known to the individual at $t-\ell$ (Section 2.1 and Assumption I.1), so it is $\mathcal{F}_{i,t-\ell}$-measurable. So are ξ_i and the weights w_{ij}, which depend only on ω_i. Each innovation in Δe_{it} is dated $t-2$ or later, hence after $t-\ell$, and the objective half of Assumption I.1 gives $\mathbb{E}[\eta_{i,t-j} \mid \mathcal{F}_{i,t-j-1}] = 0$. Since $t-\ell \leq t-j-1$ and information accumulates, $\mathcal{F}_{i,t-\ell} \subseteq \mathcal{F}_{i,t-j-1}$, so by iterated expectations, for $j \in \{0,1,2\}$,

$$\mathbb{E}[Y_{i,t-\ell}\, \xi_i w_{ij}\, \eta_{i,t-j}] = \mathbb{E}[Y_{i,t-\ell}\, \xi_i w_{ij}\ \mathbb{E}[\eta_{i,t-j} \mid \mathcal{F}_{i,t-j-1}]] = 0.$$

The same argument with $Y_{i,t-\ell}$ replaced by 1 gives $\mathbb{E}[\Delta e_{it}] = 0$. Hence $\mathrm{Cov}(Y_{i,t-\ell}, \Delta e_{it}) = 0$.

If $\delta_i = 1$ for all i, then $w_{i2} = 0$, so Δe_{it} loads only on η_{it} and $\eta_{i,t-1}$. In this case, the same argument as above applies with $\ell = 2$.

(ii) Differencing (6) gives $\Delta Y_{it} = r_i \Delta Y_{i,t-1} + \Delta e_{it}$. For every lag covered by part (i), $\mathrm{Cov}(Y_{i,t-\ell}, \Delta e_{it}) = 0$, so

$$\mathrm{Cov}(Y_{i,t-\ell}, \Delta Y_{it}) = \mathrm{Cov}(Y_{i,t-\ell}, r_i\, \Delta Y_{i,t-1}),$$

and dividing by $\mathrm{Cov}(Y_{i,t-\ell}, \Delta Y_{i,t-1}) \neq 0$ gives (8). Since $\mathrm{Cov}(Y_{i,t-\ell}, X) = \mathbb{E}[(Y_{i,t-\ell} - \mathbb{E}[Y_{i,t-\ell}])X]$ for any X, the ratio equals $\mathbb{E}[r_i W_i^{\mathrm{int},\ell}]$, with the weights $W_i^{\mathrm{int},\ell}$ of the statement, which average one. If $r_i = r$ for all i, the numerator equals r times the denominator, so the ratio equals r. This is Lemma B.3 with $h = \Delta$ and $Z_{it} = Y_{i,t-\ell}$. □

B.2.3 External Instruments

Proof of Proposition 3. Fix a type ω and a date t with $t - 1 > \bar{t}$.

Exclusion. By Assumption X.1, for every $s > \bar{t}$, $\tilde{p}_{is} = p_{is}$ and $\tilde{\mathcal{E}}_{is}[\tilde{p}_{i,s+k}] = \mathcal{E}_{is}[p_{i,s+k}]$ for every $k \geq 1$. By (21), the cost outlook π_{is} therefore does not depend on D_i for $s > \bar{t}$. Since $t - 1 > \bar{t}$, neither does $e_{it} = \pi_{it} - (1 - \delta_i)\pi_{i,t-1}$.

Exogeneity. By exclusion, e_{it} equals its value at $D_i = 0$, a function of ω_i and of $\{p_{is}, \{\mathcal{E}_{is}[p_{i,s+k}]\}_{k\geq 1}\}_{s\leq t}$, and the individual effect $a_i = \delta_i K_i \alpha_i$ is a function of (α_i, ω_i). By Assumption X.2, both are independent of D_i, so $\mathbb{E}[a_i + e_{it} \mid D_i = 1, \omega_i = \omega] = \mathbb{E}[a_i + e_{it} \mid D_i = 0, \omega_i = \omega]$.

The type-specific path. By (6),

$$\begin{aligned}\mathrm{TE}_t(\omega) &= \mathbb{E}[Y_{it} \mid D_i = 1, \omega_i = \omega] - \mathbb{E}[Y_{it} \mid D_i = 0, \omega_i = \omega] \\ &= r(\omega)\,\mathrm{TE}_{t-1}(\omega) + \Big(\mathbb{E}[a_i + e_{it} \mid D_i = 1, \omega_i = \omega] - \mathbb{E}[a_i + e_{it} \mid D_i = 0, \omega_i = \omega]\Big) = r(\omega)\,\mathrm{TE}_{t-1}(\omega),\end{aligned}$$

where the last equality uses exogeneity. This is (11). Applying (11) at $t = \bar{t} + 2, \ldots, \bar{t} + 1 + k$, we have

$$\mathrm{TE}_{\bar{t}+1+k}(\omega) = r(\omega)^k\,\mathrm{TE}_{\bar{t}+1}(\omega) \qquad \text{for every } k \geq 0. \tag{25}$$

Aggregation. By Assumption X.2, ω_i has the same distribution in both arms, so $\mathrm{TE}_t = \mathbb{E}[\mathrm{TE}_t(\omega_i)]$ for all t. Taking expectations of both sides of (25) gives (12).

The ratio. With $\Pr(D_i = 1) \in (0, 1)$, $\mathrm{Cov}(D_i, Y_{is}) = \mathrm{Var}(D_i)\,\mathrm{TE}_s$ for every s, so the instrumental-variables estimand using D_i as the instrument for $Y_{i,\bar{t}+k}$ in (6) is

$$\frac{\mathrm{Cov}(D_i, Y_{i,\bar{t}+1+k})}{\mathrm{Cov}(D_i, Y_{i,\bar{t}+k})} = \frac{\mathrm{TE}_{\bar{t}+1+k}}{\mathrm{TE}_{\bar{t}+k}},$$

for $k \geq 1$ with $\mathrm{TE}_{\bar{t}+k} \neq 0$. By (11) and aggregation, the numerator is $\mathrm{TE}_{\bar{t}+1+k} = \mathbb{E}[r_i\,\mathrm{TE}_{\bar{t}+k}(\omega_i)]$ and the denominator is $\mathrm{TE}_{\bar{t}+k} = \mathbb{E}[\mathrm{TE}_{\bar{t}+k}(\omega_i)]$, so their ratio is (13).

If $r_i = r$ for all i, the numerator of (13) is r times the denominator, so $r^{\mathrm{ext}}_{\bar{t}+1+k} = r$; this is Lemma B.3 with h the identity and $Z_{it} = D_i$. □

B.3 Proofs for Section 3.4

This subsection proves Proposition 4. It first states the identification argument in general form (Lemma B.5), which Appendix E.2.4 also applies to lagged choices. It then derives three properties of the treatment effects in a temporary experiment: the gap identity (Equation (15)), the difference equation they follow inside the window, and the closed form of the direct effect. The proof of Proposition 4 uses the difference equation to write the gaps as known functions of (δ, β) up to scale, and applies Lemma B.5. The subsection closes with the test of full depreciation, which uses the gap identity.

B.3.1 Identification from Structured Moments

The following lemma provides general conditions for identification of δ and β.

Lemma B.5 (Identification from structured moments)**.** *Suppose observed choices satisfy* (6) *with common propagation, $r_i = r$. Let h be a known linear transformation and Z_{it} a source of variation for the period-t transition. For r given, define*

$$m_t^h(r) \equiv \mathrm{Cov}\big(Z_{it},\, h(Y_{it}) - r\, h(Y_{i,t-1})\big) = \mathrm{Cov}\big(Z_{it},\, h(a_i + e_{it})\big),$$

the part of the transformed choice that propagating the lagged choice does not explain, projected onto Z_{it}. Suppose r is identified and that at periods $t_1, \ldots, t_J$,

$$m_{t_j}^h(r) = F_j(\delta, \beta; r), \qquad j = 1, \ldots, J,$$

for known functions F_j. If the map $(\delta, \beta) \mapsto \big(F_1, \ldots, F_J\big)(\delta, \beta; r)$ is injective on the parameter space, then (δ, β) is identified, and so is $\phi = r - (1 - \delta)$. If β is known, a single moment that is one-to-one in δ suffices. If the moments are known only up to a common nonzero factor, as when the size of the cost change is unknown, the same conclusions hold for their ratios.

Proof. Each $m_{t_j}^h(r)$ is a function of the joint distribution of $(Z_{it}, Y_{it}, Y_{i,t-1})$ at $t = t_j$ and of r, so once r is identified the vector $\big(m_{t_1}^h(r), \ldots, m_{t_J}^h(r)\big)$ is known. If two parameter values (δ, β) and (δ', β') generated this vector, injectivity would make them equal; hence (δ, β) is the unique value consistent with the data, and ϕ follows from $r = 1 - \delta + \phi$. With β known the same argument applies to δ alone, and the ratios of the moments are known when the moments are known up to a common factor. □

The proof of Proposition 4 applies the lemma with $Z_{it} = D_i$ and h the identity, so that the moments are proportional to the gaps G_s defined below; the proof of Proposition E.5 applies it with $h = \Delta$.

B.3.2 Decomposition with Temporary Experiments

Let TE_t be the treatment effect of Equation (10), $\text{TE}_t^S \equiv \mathbb{E}[S_{it} \mid D_i = 1] - \mathbb{E}[S_{it} \mid D_i = 0]$ the mean difference in stocks, and $\text{TE}_t^\pi \equiv \mathbb{E}[\pi_{it} \mid D_i = 1] - \mathbb{E}[\pi_{it} \mid D_i = 0]$ the mean difference in cost outlooks (the direct effect of Subsection 3.4.2). For a given r, the gap at date s is the part of the next treatment effect that propagating the current one does not explain,

$$G_s \equiv \text{TE}_{s+1} - r\,\text{TE}_s.$$

In the language of Section 3.4.3, $G_{\bar{t}}$ is the boundary gap (the transition out of the last subsidized period) and $G_{\bar{t}-1}$ is the interior gap (the transition into it). With common parameters ($\omega_i = \omega$, as in Section 3.4), the terms in α_i cancel by random assignment when (5) is differenced across arms, so $\text{TE}_s = \phi\,\text{TE}_s^S + \text{TE}_s^\pi$; differencing (6) the same way gives Equation (15), which in terms of gaps reads $G_s = \text{TE}_{s+1}^\pi - (1-\delta)\text{TE}_s^\pi$ (the gap identity).

The proof ofProposition 4 uses the treatment-effect path inside the intervention window. The following difference equation characterizes that path.

Evolution of treatment effects. Let $b_{t,k}$ denote the arm difference in the mean date-t forecast of the cost k periods ahead,

$$b_{t,k} \equiv \mathbb{E}\Big[\tilde{\mathcal{E}}_{it}[\tilde{p}_{i,t+k}] \mid D_i = 1\Big] - \mathbb{E}[\mathcal{E}_{it}[p_{i,t+k}] \mid D_i = 0]\,,$$

where for $k = 0$ the forecast is the current cost itself. Then $b_{t,k} = -b \cdot \mathbf{1}\{t+k \in \mathcal{W}\}$ for every $t \geq \underline{t}$ with $b \equiv \mathbb{E}[b_i \mid D_i = 1]$ the mean cost reduction. At treated dates, this follows from Equation (9) and Assumption X.2 for $k = 0$, and also uses Assumption X.3 for $k \geq 1$. After the window, it follows from Assumptions X.1 and X.2 for every $k \geq 0$.

Lemma B.6. *Suppose Assumptions X.1, X.2, and X.3 hold, and the parameters are common across individuals. Then for every $t \geq \underline{t}$, the treatment effects satisfy the second-order linear difference equation*

$$a_2\text{TE}_{t+2} + a_1\text{TE}_{t+1} + a_0\text{TE}_t = \mathcal{B}_t, \tag{26}$$

where (a_2, a_1, a_0) are as in Appendix B.1.1 and the forcing term is

$$\mathcal{B}_t = b_{t+1,0} - \frac{1-\delta}{1+\rho} b_{t+1,1} - (1-\delta)\left(b_{t,0} - \frac{1-\delta}{1+\rho} b_{t,1}\right). \tag{27}$$

$\mathcal{B}_t$ depends on t only through which of the dates t, $t+1$, $t+2$ lie in $\mathcal{W}$, and $\mathcal{B}_t = 0$ for $t > \bar{t}$.

Proof. Use the stock law of motion to substitute for $S_{i,t+1}$ in the Euler equation (16) written for

individual i, and rearrange:

$$AY_{it} + BS_{it} + C\mathcal{E}_{it}[Y_{i,t+1}] + \frac{\psi}{1+\rho} + \alpha_i - \frac{1-\delta}{1+\rho}\alpha_i = p_{it} - \frac{1-\delta}{1+\rho}\,\mathcal{E}_{it}[p_{i,t+1}]\,. \tag{28}$$

Write (28) one period ahead and subtract $(1-\delta)$ times (28):

$$\begin{aligned} &C\mathcal{E}_{i,t+1}[Y_{i,t+2}] + AY_{i,t+1} - (1-\delta)C\mathcal{E}_{it}[Y_{i,t+1}] - (1-\delta)AY_{it} + B[S_{i,t+1} - (1-\delta)S_{it}] \\ &\quad + \delta\left(\frac{\psi}{1+\rho} + \alpha_i - \frac{1-\delta}{1+\rho}\alpha_i\right) \\ &\quad = p_{i,t+1} - \frac{1-\delta}{1+\rho}\,\mathcal{E}_{i,t+1}[p_{i,t+2}] - (1-\delta)\left(p_{it} - \frac{1-\delta}{1+\rho}\,\mathcal{E}_{it}[p_{i,t+1}]\right). \end{aligned}$$

With $S_{i,t+1} - (1-\delta)S_{it} = Y_{it}$ and $B - (1-\delta)A = U_{yS} - (1-\delta)U_{yy}$, this equation becomes

$$\begin{aligned} &C\mathcal{E}_{i,t+1}[Y_{i,t+2}] + AY_{i,t+1} - (1-\delta)C\mathcal{E}_{it}[Y_{i,t+1}] + (U_{yS} - (1-\delta)U_{yy})\,Y_{it} \\ &\quad + \delta\left(\frac{\psi}{1+\rho} + \alpha_i - \frac{1-\delta}{1+\rho}\alpha_i\right) \\ &\quad = p_{i,t+1} - \frac{1-\delta}{1+\rho}\,\mathcal{E}_{i,t+1}[p_{i,t+2}] - (1-\delta)\left(p_{it} - \frac{1-\delta}{1+\rho}\,\mathcal{E}_{it}[p_{i,t+1}]\right). \end{aligned}$$

Take expectations conditional on $D_i = 1$ and $D_i = 0$ and difference.

The individual-specific terms enter only through $\delta\big(\psi/(1+\rho) + \alpha_i[1-(1-\delta)/(1+\rho)]\big)$, which has the same mean in both arms by random assignment (Assumption X.2) and differences out. On the right-hand side, the arm differences of the mean costs and mean cost forecasts are $b_{t+1,0}$, $b_{t+1,1}$, $b_{t,0}$, and $b_{t,1}$ (a current cost is its own forecast), so the right-hand side differences to $\mathcal{B}_t$ in (27). For the forecasts of choices, $S_{i,s+1}$ is known at s, so by (5)

$$\mathcal{E}_{is}[Y_{i,s+1}] = \phi\, S_{i,s+1} + K\alpha_i + \mathcal{E}_{is}[\pi_{i,s+1}],$$

and by (21) and iterated expectations, $\mathcal{E}_{is}[\pi_{i,s+1}]$ is (21) at $s+1$ with every forecast made at s. Differencing across arms replaces each forecast of $p_{i,s+1+k}$ by the arm difference $b_{s,1+k}$, whereas differencing $\pi_{i,s+1}$ itself replaces it by $b_{s+1,k}$; the two are equal because $b_{t,k}$ depends only on the date $t+k$. The mean forecast of the cost outlook therefore differs across arms by TE^{π}_{s+1}, and, by the decomposition above,

$$\mathbb{E}\left[\mathcal{E}_{is}[Y_{i,s+1}] \mid D_i = 1\right] - \mathbb{E}\left[\mathcal{E}_{is}[Y_{i,s+1}] \mid D_i = 0\right] = \phi\,\mathrm{TE}^{S}_{s+1} + \mathrm{TE}^{\pi}_{s+1} = \mathrm{TE}_{s+1}, \qquad s = t, t+1.$$

This leaves

$$\underbrace{C}_{=a_2}\,\mathrm{TE}_{t+2} + \underbrace{\big[A-(1-\delta)C\big]}_{=a_1}\,\mathrm{TE}_{t+1} + \underbrace{\big(U_{yS}-(1-\delta)U_{yy}\big)}_{=a_0}\,\mathrm{TE}_t = \mathcal{B}_t. \;\square$$

Evaluating (27) at the three dates used below, with the cost reduction b constant across the window and $\bar{t}-2$, $\bar{t}-1$, $\bar{t} \in \mathcal{W}$, gives

$$\mathcal{B}_{\bar{t}} = (1-\delta)\, b, \tag{29}$$

$$\mathcal{B}_{\bar{t}-1} = -\big[\delta + (1-\delta)^2\beta\big]\, b, \tag{30}$$

$$\mathcal{B}_{\bar{t}-2} = -\delta\,[1-\beta(1-\delta)]\, b. \tag{31}$$

The direct effect. Section 3.4.2 uses one property of the direct effect TE^{π}_t: given δ and β, its path inside the window is known up to scale. This paragraph derives it. Write $H_m \equiv \sum_{k=0}^{m-1}(\beta r)^k$ for $m \geq 0$ (so $H_0 = 0$, $H_{m+1} = 1+\beta r\, H_m$, and $H_m > 0$ for $m \geq 1$ since $|\beta r| < 1$). At every treated date, the treated individuals' mean forecast of the cost k periods ahead differs from the control individuals' by $b_{t,k}$, as shown above. Differencing (21) across arms and writing $\xi \equiv 1/(A+C\phi)$, gives, for $t \in \mathcal{W}$,

$$\mathrm{TE}^{\pi}_t = \xi \sum_{k\geq 0}(\beta r)^k \big[b_{t,k} - \beta(1-\delta)\, b_{t,k+1}\big] = -\xi b\,\big[1+\beta\phi\, H_{\bar{t}-t}\big].$$

The second equality holds because the bracket equals $-(1-\beta(1-\delta))\, b$ before $\bar{t}$, $-b$ at $\bar{t}$, and zero afterwards, and $H_{m+1} - \beta(1-\delta)H_m = 1+\beta\phi H_m$. After the window, $\mathrm{TE}^{\pi}_t = 0$ (Assumption X.1). The direct effect is therefore a scale, $-\xi b$, that is the same in every period, times a factor that depends only on the number of treated periods left. The factor equals one in the last treated period ($H_0 = 0$), and in every period when $\beta = 0$.

Proof of Proposition 4. Throughout, $\mathcal{W} = \{\underline{t}, \ldots, \bar{t}\}$ is the treatment window and the assumptions of Proposition 4 hold, so Lemma B.6 applies and r is identified by Proposition 3. The structured moments to which Lemma B.5 is applied are the boundary gap $G_{\bar{t}}$ and the interior gaps $G_{\bar{t}-1}$ and $G_{\bar{t}-2}$, each a function of observed treatment effects and the identified r.

Step 1: the gaps in terms of the forcing. Because the roots of P are r and $(1+\rho)/r$, $a_1 = -a_2\,[r + (1+\rho)/r]$; with $a_0 = (1+\rho)a_2$, the left-hand side of (26) equals $a_2\,\big[G_{t+1} - \frac{1+\rho}{r}G_t\big]$, so

$$a_0\,[G_t - \beta r\, G_{t+1}] = -r\,\mathcal{B}_t.$$

Since $G_{\bar{t}+1} = 0$ by (14), evaluating at $\bar{t}$, $\bar{t}-1$, and $\bar{t}-2$ with (29)–(31) gives

$$a_0\, G_{\bar{t}} = -r(1-\delta)\, b, \tag{32}$$

$$a_0\, [G_{\bar{t}-1} - \beta r\, G_{\bar{t}}] = r\, \left[\delta + (1-\delta)^2\beta\right] b, \tag{33}$$

$$a_0\, [G_{\bar{t}-2} - \beta r\, G_{\bar{t}-1}] = r\, \delta\, [1 - \beta(1-\delta)]\, b. \tag{34}$$

Since $b \neq 0$, $\delta < 1$ and $r \neq 0$, $G_{\bar{t}} \neq 0$, so ratios to (32) are well defined and free of the unknown scale $a_0 b$. This verifies the condition of Lemma B.5: each ratio is a known function of (δ, β) given r.

Step 2: part (i), β known. Dividing (33) by (32) gives

$$\frac{G_{\bar{t}-1}}{G_{\bar{t}}} = -\frac{\delta}{1-\delta} + \beta\left(r - (1-\delta)\right) = -\frac{\delta}{1-\delta} + \beta\phi. \tag{35}$$

This is the ratio of the interior gap to the boundary gap in Section 3.4.3; at $\beta = 0$, it equals $-\delta/(1-\delta)$. Rearrange (35) to

$$\frac{1}{1-\delta} = 1 + \beta r - \beta(1-\delta) - \frac{G_{\bar{t}-1}}{G_{\bar{t}}}. \tag{36}$$

The left-hand side is strictly decreasing in $1-\delta$ with derivative $-1/(1-\delta)^2$; the right-hand side has derivative $-\beta \leq 0$ (it is constant when $\beta = 0$). Since $1/(1-\delta)^2 > 1 > \beta$, the two sides cross at most once on $(0, 1)$. The map from δ to $G_{\bar{t}-1}/G_{\bar{t}}$ is therefore one-to-one. By Lemma B.5 δ, and hence $\phi = r - (1-\delta)$, are identified.

Step 3: part (ii), β unknown. Divide (34) by (32):

$$\frac{G_{\bar{t}-2}}{G_{\bar{t}}} - \beta r\, \frac{G_{\bar{t}-1}}{G_{\bar{t}}} = -\frac{\delta}{1-\delta} + \beta\delta. \tag{37}$$

Subtracting (35) from (37) eliminates $-\delta/(1-\delta)$ and leaves

$$\frac{G_{\bar{t}-2}}{G_{\bar{t}}} - \frac{G_{\bar{t}-1}}{G_{\bar{t}}} - \beta r\, \frac{G_{\bar{t}-1}}{G_{\bar{t}}} = \beta\big(\delta - r + (1-\delta)\big) = \beta(1-r),$$

which is linear in β and solves to

$$\beta = \frac{G_{\bar{t}-2} - G_{\bar{t}-1}}{(1-r)\, G_{\bar{t}} + r\, G_{\bar{t}-1}}. \tag{38}$$

The right-hand side is a function of observed treatment effects and the identified r, so β is identified.

The denominator of (38) is nonzero if and only if $\phi \neq 0$. Substituting (35) and using $r =$

$(1-\delta)+\phi$ gives

$$(1-r)G_{\bar{t}}+rG_{\bar{t}-1}=G_{\bar{t}}\left[1-r+r\left(1+\beta\phi-\tfrac{1}{1-\delta}\right)\right]=G_{\bar{t}}\left[1+r\beta\phi-\tfrac{r}{1-\delta}\right]=\frac{G_{\bar{t}}}{1-\delta}\,\phi\,((1-\delta)\beta r-1).$$

Since $|(1-\delta)\beta r|<1$, the factor $(1-\delta)\beta r-1$ is nonzero, so the denominator vanishes if and only if $\phi=0$. Economically, at $\phi=0$, there is no reinforcement through which anticipation of the end of the window can operate, and the treatment-effect path carries no information about β.

Given β, Step 2 identifies δ, and $\phi=r-(1-\delta)$ follows. □

B.3.3 Testing Full Depreciation

The gap identity provides a test of full depreciation.

Corollary B.1 (Testing full depreciation)**.** *Suppose Assumptions X.1 and X.2 hold, the parameters are common across individuals, and the intervention takes the form* (9)*. If $\delta=1$, then $\mathrm{TE}_{\bar{t}+1}=r\mathrm{TE}_{\bar{t}}$. If $\delta<1$, this restriction holds if and only if $\mathrm{TE}^{\pi}_{\bar{t}}=0$. If, in addition, Assumption X.3 holds and $\mathbb{E}[b_i\mid D_i=1]\neq 0$, then $\mathrm{TE}_{\bar{t}+1}=r\mathrm{TE}_{\bar{t}}$ if and only if $\delta=1$.*

Proof. Set $t=\bar{t}+1$ in Equation (15). Because the intervention ends at $\bar{t}$, Equation (9), Assumption X.1 and random assignment imply $\mathrm{TE}^{\pi}_{\bar{t}+1}=0$, so the boundary gap of Subsection 3.4.3 is $G_{\bar{t}}=-(1-\delta)\,\mathrm{TE}^{\pi}_{\bar{t}}$. Under the null $\delta=1$ the right-hand side is zero, so $\mathrm{TE}_{\bar{t}+1}=r\mathrm{TE}_{\bar{t}}$.

Conversely, if $\delta<1$ the restriction fails whenever $\mathrm{TE}^{\pi}_{\bar{t}}\neq 0$. Under Assumption X.3, $\mathrm{TE}^{\pi}_{\bar{t}}=-\xi b$ (the direct effect above), which is nonzero when $b\neq 0$, so the test has power against $\delta<1$. □

C Empirical Applications Details

This appendix provides the details and additional analyses referred to in Sections 4–6.

C.1 Labor Supply: The Seattle–Denver Income Maintenance Experiment

In this Appendix Section, we present details of the Seattle–Denver Income Maintenance Experiment and the additional results discussed in Section 4. We first document the data, our sample, and the comparison to the Final Report (Appendix C.1.1). We then give the details on the identifying assumptions verified in Section 4.1.2 (Appendix C.1.2). Next, we present robustness to misreporting of hours (Appendix C.1.3). Finally, we give details on estimation (Appendix C.1.4), the validation on the five-year arm (Appendix C.1.5), and the decomposition of observed persistence (Appendix C.1.6).

C.1.1 Data, Sample, and Comparison with the Final Report

Data. We use the SIME/DIME public-use files of the Institute for Research on Poverty's Negative Income Tax archive: the monthly composite files of principal persons and of families in Seattle and Denver.[50] We build a panel of the heads of the families enrolled in the three- and five-year arms: the husband and wife of a couple, or a single head (a head not married at enrollment, whatever their later marital status). Plans were assigned at random within strata. Strata cross the household's normal-income level (the experimenters' estimate of its income absent the experiment in seven bands), with its race (white, black, or Chicano, in Denver only), its family type (a couple or a single head with dependents), and the site..[51] Our outcome is hours per quarter, the sum of the monthly hours reported in the quarter. Quarters 1–12 are the program and quarters 13–16 the follow-up. Quarter 0 is the average quarterly hours over January to June of the first year of the data (1970 in Seattle, 1971 in Denver), months that precede every household's enrollment (as in Riddell and Riddell, 2024). A head is treated if the household was assigned to one of the eleven NIT plans, and is a control otherwise.

Sample. Our estimation sample contains 2,728 heads in 1,835 households from the three-year arm of SIME/DIME, the arm we study (Section 4.1). We build it from this panel in five steps. We start from the 7,180 heads in 4,551 households whose normal-income level is recorded. First, we keep the 4,633 heads in 2,995 households enrolled in the three-year arm with pre-enrollment data. Second, we keep the strata that contain both treated and control heads among those still enrolled in every quarter up to quarter 12, so that the probability of treatment lies strictly between zero and one in every stratum. This step keeps 55 of the 65 strata, with 4,566 heads in 2,964 households. Third, we keep heads with labor-supply information through the end of the follow-up,

[50] https://www.irp.wisc.edu/research/nit/irp-negative-income-tax-archive/.

[51] The bands are: under \$1,000, \$1,000–2,999, \$3,000–4,999, \$5,000–6,999, \$7,000–8,999, \$9,000–10,999, and \$11,000–12,999 a year.

so that every head is observed in every quarter the estimator uses: 3,456 heads in 2,336 households. Fourth, we keep single heads and both heads of couples that stay intact through quarter 16, with neither spouse reporting a marital dissolution, because a dissolution changes the incentives each spouse faces (the household component in Equation (40) of Appendix C.1.2). This leaves 959 single heads, 893 husbands, and 893 wives, in 1,850 households. Fifth, trimming to the common support of the estimated propensity score (Appendix C.1.4) removes 17 treated heads: 13 single heads and two couples. The estimation sample has 946 single heads, 891 husbands, and 891 wives, of whom 1,718 are treated and 1,010 are controls.[52]

Comparison with the Final Report. The effects in Figure 2 are close to those reported in the Final Report once we account for the differences in samples, time-spans, and regression specification (Table C.1). Pooling both treatment arms over the second and third program years, the Final Report finds that the NIT reduced annual hours by 9–10% for husbands, 17–20% for wives, and more than 20% for single female heads (U.S. Department of Health and Human Services, Office of Income Security Policy, 1983, pp. 13–16) (column 4). Over the same years, our estimates (same estimator of Figure 2) gives reductions of 7.0%, 14.8%, and 9.9%, respectively (column 1).

The small discrepancies arise from three sources. First, the Final Report pools the three-year arm with the five-year arm, which reduced hours more (Robins and West, 1980); pooling the arms raises our estimates to 8.5%, 17.6%, and 11.8% (column 2). Second, the original analyses do not remove dissolved households, keep heads who left the data after quarter 12, and estimate the effect by regression with pre-enrollment controls. Following these choices raises our estimates to 9.9%, 19.6%, and 13.3% (column 3), which match the Final Report for husbands and wives. Third, our samples differ from the original ones, which cannot be reconstructed from the public-use files (Riddell and Riddell, 2024), so we cannot trace the remaining gap for single female heads.

C.1.2 Verifying Identifying Assumptions

In this Appendix subsection, we give the details for each of the three identifying assumptions that Section 4.1.2 verifies.

Random assignment. Treated and control heads in the estimation sample are balanced on baseline characteristics that our regression adjustment (AIPW) does not use (Table C.2). This balance is consistent with random assignment (Assumption X.2) and selection on observables. The differences are jointly insignificant, both for the baseline characteristics in Panel A ($p = 0.69$) and for the whole table, which adds the two design variables in Panel B ($p = 0.56$).

Known, constant intervention path. Section 4 treats each NIT plan as a constant shift in the cost of working while the program lasts. We derive this shift, discuss what it leaves out, and

[52]Hours are missing in 21 of the 46,376 person-quarters of the estimation sample, and we set them to zero.

Table C.1: Effects of the NIT on hours in program years 2–3: our estimates and the Final Report

	(1)	(2)	(3)	(4)
	Our estimates			
	Three-year arm	Both arms	Final Report specification	Final Report
Husbands	7.0	8.5	9.9	9–10
	(2.8)	(2.4)	(1.8)	
Wives	14.8	17.6	19.6	17–20
	(7.0)	(6.3)	(4.1)	
Single female heads	9.9	11.8	13.3	>20
	(4.9)	(4.2)	(3.7)	
All three groups	9.3	11.2	12.9	13.4–14.6
	(2.3)	(2.0)	(1.6)	
Heads	2,726	3,492	4,849	
Arms	Three-year	Both	Both	Both
Estimator	AIPW	AIPW	OLS	

Notes: This table reports effects of the NIT on hours worked in quarters 5–12, the second and third program years, for the three groups of the Final Report: husbands, wives, and single female heads; effects are percent reductions relative to the group's control mean over the same quarters. In columns (1)–(3), the entry for all three groups averages the group effects and control means with the group shares in the sample. Column (1) estimates each group's effect with the estimation sample and estimator of Figure 2, averaged over these quarters; the estimation sample's two single male heads belong to none of the three groups, so the count is 2,726. Column (2) repeats the quarterly estimation in column (1) adding the five-year arm. Column (3) keeps the heads of both arms that pass the first two sample steps and are observed in every quarter from 5 to 12, with groups defined at enrollment and no trimming; the effect is the coefficient on treatment in a regression of each head's average hours over these quarters on pre-enrollment hours, normal-income level, site, race, family size, children under six, education, age and its square, the counseling and training-subsidy assignments, and the arm (Robins and West, 1980). Column (4) reports U.S. Department of Health and Human Services, Office of Income Security Policy (1983, pp. 13–16); its entry for all three groups weights the group estimates with the shares and control means of column (3). Standard errors, clustered by family, in parentheses.

bound the one feature that could make it change over the program: the tax on the wage return to work history.

The NIT as a change in the cost of working. To see how the NIT changed incentives, consider individual i in household $h(i)$ under a (constant-rate) plan j. Let y_{it} denote individual i's hours of work in quarter t and S_{it} their accumulated work history. Write household earnings as $E_{h(i)t} = E_{-i,h(i)t} + w_{it}y_{it}$, where $E_{-i,h(i)t}$ denotes the earnings of other household members. Individual i's household $h(i)$ receives payment

$$T_{h(i)t}(j) = \max\left\{0,\, G_j - \tau_j\left(E_{-i,h(i)t} + w_{it}y_{it}\right)\right\}, \tag{39}$$

where G_j is the income guarantee and τ_j the tax rate under plan j.

The plan changes the cost of working in two ways. First, an additional hour of work by i

Table C.2: Baseline characteristics of treated and control heads

	Control group		Adjusted difference	
	Mean	(SD)	Estimate	(SE)
Panel A: Baseline, monthly averages				
Share of months on welfare	0.18	(0.37)	0.02	(0.02)
Welfare transfers, AFDC and food stamps ($)	64.8	(128.3)	5.7	(7.2)
AFDC transfers ($)	35.8	(79.6)	4.5	(4.6)
Any AFDC	0.22	(0.41)	0.03	(0.02)
Family hours worked	161.3	(101.3)	−2.8	(5.4)
Family earnings ($)	495.5	(332.5)	−11.4	(17.6)
Own overtime hours	1.9	(8.9)	0.2	(0.3)
Own bonuses, tips, and commissions ($)	3.7	(22.6)	−1.3	(0.7)
Own odd-job hours	0.7	(5.7)	−0.2	(0.2)
Own odd-job income ($)	2.1	(11.8)	−0.6	(0.5)
Any child aged 6–16 at enrollment	0.60	(0.49)	0.02	(0.03)
Panel B: Design				
Assigned to counseling or training	0.65	(0.48)	0.01	(0.03)
Month of enrollment	16.3	(2.8)	−0.3	(0.2)
Joint test, Panel A: $\chi^2(11)$		8.23	$[p = 0.69]$	
Joint test, all rows: $\chi^2(13)$		11.59	$[p = 0.56]$	
Individuals	2,728 (1,718 treated, 1,010 control)			

Notes: This table compares treated and control heads in the estimation sample on baseline characteristics. Panel A reports monthly averages over the first nine months of the data, before any household enrolled. Two rows are exceptions: any AFDC indicates receipt of Aid to Families with Dependent Children (AFDC) in any of these months, and children aged 6–16 are measured at enrollment. Panel B reports two features of the design: assignment to the separate counseling and training treatment, and the month of enrollment, counted from January 1970 in Seattle and January 1971 in Denver. The first two columns report the control-group mean and standard deviation; the last two report the difference between treated and control heads, weighted by the inverse of the propensity score of the main analysis (Appendix C.1.4), and its standard error, clustered by family. Joint tests are Wald tests that all differences in Panel A, or in the whole table, are zero.

raises earnings by w_{it} but lowers the payment by $\tau_j w_{it}$, changing the marginal tax rate relative to the control schedule by $\Delta\tau_{ij} = \tau_j - \tau_i^0$ (where τ_i^0 is the rate i would face absent the experiment). Second, the plan changes i's cost of working through the household. A spouse's lower earnings, for instance, move the household relative to the break-even level and hence change the rate i faces. Let Δq_{ij} denote this household component of the change in i's cost of working. Plan j raises the cost of working during the twelve treatment quarters:

$$\tilde{p}_{it}^{j} = p_{it} + \left(\Delta\tau_{ij}\, w_i + \Delta q_{ij}\right) \mathbf{1}[t \in \mathcal{W}^{\text{NIT}}], \tag{40}$$

where $\mathcal{W}^{\text{NIT}}$ denotes the set of periods in which the NIT is active. We take the shift to be constant in real terms over the program, $b_{ij} = -(\Delta\tau_{ij} w_i + \Delta q_{ij})$, with w_i the individual's real wage. This shift is the intervention in Equation (9), with $b_i = b_{ij}$.

Features the mapping sets aside. Equation (40) treats each plan as a linear tax, with a constant change in the marginal tax rate while the program lasts. It sets aside three features of the experiment. First, a household whose income exceeds the break-even level receives no payment and faces no NIT tax at the margin, so the mapping overstates its change in the cost of working. Most treated households were below it: at baseline, 75% of them had earnings below their plan's break-even level. Second, in some plans the tax rate declined with income, so $\Delta\tau_{ij}$ moves with the household's income as it responds. The mapping holds it fixed, which is close because the rate declined gradually. Third, the household component Δq_{ij} changes over the program as the spouse's own response to the NIT builds up. A marital dissolution changes it abruptly, because the other spouse's earnings no longer enter the household's payment, so the estimation sample keeps only couples that stay intact (Appendix C.1.1). Gradual changes would make the effects during the program violate the restrictions the model places on their path, which the data do not reject ($p = 0.57$; Table 1).

Taxing the wage return to work history. If work history raises wages, the shift is not exactly constant, because the tax applies to the whole wage. Write the wage as a part that does not depend on work history and a part that does, $w_{it} = w_i^0 + w^S(S_{it})$, so that

$$\Delta\tau_{ij}\, w_{it} = \Delta\tau_{ij}\, w_i^0 + \Delta\tau_{ij}\, w^S(S_{it}).$$

The first term is a constant shift. The second taxes the wage return to work history, which is part of the return $R(y, S)$; in the model, it lowers reinforcement among treated workers while the program lasts. Let $S_{it}(0)$ be the work history worker i would have without the program and $\Delta w_{it} \equiv w^S(S_{it}) - w^S(S_{it}(0))$ the wage difference created by the work history the program removed. The second term then splits in two. The tax on the wage return to the work history the worker would have had anyway, $\Delta\tau_{ij}\, w^S(S_{it}(0))$, is part of the shift b_i; the tax on the wage difference, $\Delta\tau_{ij}\,\Delta w_{it}$, is not. Relative to the tax on the wage, the shift therefore departs from its constant value by

$$\frac{|\Delta\tau_{ij}\,\Delta w_{it}|}{\Delta\tau_{ij}\, w_{it}(0)} = \frac{|\Delta w_{it}|}{w_{it}(0)}, \qquad w_{it}(0) \equiv w_i^0 + w^S(S_{it}(0)),$$

the percentage wage difference created by the lost work history.

This departure is small. To illustrate, let the wage return be linear, $w^S(S) = w_S\, S$. Suppose a year of work raises wages by a share g, so that $w_S = g\, w/H$, with H the hours in a year of work. Because work history depreciates ($\delta > 0$), the lost work history is on average at most the cumulative reduction in hours, $12 \times 20.9 = 251$ hours by the end of the program. Taking H to be the hours control workers work in a year, $4 \times 292.6 = 1{,}170$, the departure is at most

$$\frac{|\Delta w|}{w} \;\le\; g \times \frac{251}{1{,}170} = 0.215\, g.$$

Estimates of g are between about 2 and 13 percent (Eckstein and Lifshitz, 2011; Gladden and Taber, 2000; Heckman et al., 2006). Even at $g = 0.20$, the shift departs from its constant value by at most 4.3 percent, and by less earlier in the program; at $g = 0.13$, by 2.8 percent.

The departure also has a known direction. As treated workers lose work history, their wages fall relative to $w_{it}(0)$, so the NIT's increase in their cost of working shrinks over the program. This dampens the build-up of the effects during the program, so if anything the effects we measure understate the build-up that a constant shift would produce. After the program, the tax no longer applies, so the post-program decay that identifies propagation is unaffected.[53]

Exclusion. A possible threat to exclusion (Assumption X.1) is income effects: by changing lifetime resources, a plan could change the marginal utility of lifetime wealth (MUW), and with it the incentive to work after the payments end. We relax the model by letting each plan change the MUW and identify this income effect from the persistence of the treatment effects after the program and from their differences across plans. Allowing for it leaves the estimated propagation rate close to its baseline value, and the income-effect terms are jointly insignificant.

Model. The model in Section 2 rules out income effects because utility is linear in the numeraire. We change two of its features. First, consumption enters utility through a strictly concave function, $v(c_{it})$ with $v' > 0 > v''$, rather than linearly. Second, individuals borrow and save in perfect credit markets and discount consumption at the market interest rate, so they smooth it. The MUW, $\lambda_{it} \equiv v'(c_{it})$, is then a martingale, $\lambda_{it} = \mathcal{E}_{it}[\lambda_{i,t+1}]$: news about lifetime resources changes it, and the change persists in expectation. This persistence is what could bias our estimates: without consumption smoothing, the change in the MUW would end with the payments, like the price change. The discount rate of consumption is separate from patience β, which governs how individuals weigh the future effects of current work on work history; our estimates of β concern only the latter. Given the path of the MUW, the choice problem is (1) with $\lambda_{it} p_{it}$ in place of the cost of working p_{it}, which is measured in units of the numeraire. We abstract from the effect of the MUW on the return R.

A NIT plan changes lifetime resources at enrollment. It therefore shifts the MUW at enrollment by a factor $1 + \kappa_i^j \equiv \lambda_i^j / \lambda_i^0$, where λ_i^j and λ_i^0 are i's MUW under plan j and under control, and the shift persists. The plan therefore changes the cost of working by

$$(1 + \kappa_i^j)\tilde{p}_{it}^j - p_{it} = \underbrace{\left(\tilde{p}_{it}^j - p_{it}\right)}_{\text{price change}} + \underbrace{\kappa_i^j \, \tilde{p}_{it}^j}_{\text{income change}} .$$

The price change ends with the program, and the income change does not.

Identification. Consider first the case $\beta = 0$, in which the cost outlook π_{it} responds only to the

[53] The NIT also had little effect on the hourly wages of those who worked, as West (1980) found. In our sample, its average effect over the program years is -0.3 log points (SE 1.3), which rules out wage reductions larger than about 3% (results available upon request).

current cost of working, with slope $\xi = 1/U_{yy}$ (Equation (21)). We let the change in the MUW be common within a plan, $\kappa_i^j = \kappa^j$, and specify it below. The plan's average direct effect in quarter t is then

$$\mathrm{TE}_t^{\pi,j} = \xi\big[\Delta p_t^j + \kappa^j\,\tilde{p}_t^j\big],$$

where Δp_t^j is the plan's mean price change, zero after the program, and $\tilde{p}_t^j$ is its mean cost of working. By Equation (15), applied plan by plan, $\mathrm{TE}_t^j - r\,\mathrm{TE}_{t-1}^j = \mathrm{TE}_t^{\pi,j} - (1-\delta)\,\mathrm{TE}_{t-1}^{\pi,j}$. Without income effects, the direct effect is zero after the program, and the effects decay geometrically at rate r. With them, the income change remains. If the cost of working is constant after the program, at $\tilde{p}^j$, then $\mathrm{TE}_t^j = r\,\mathrm{TE}_{t-1}^j + \delta\,\xi\kappa^j\tilde{p}^j$, and the effects converge to $\delta\,\xi\kappa^j\tilde{p}^j/(\delta-\phi)$ rather than to zero. An estimator that ignores the income change attributes this persistence to slower decay. The persistence after the program, and its variation across plans that changed lifetime resources by different amounts, identify the income change. With $\beta > 0$, the cost outlook also depends on expected future costs (Equation (21)). Patience then changes the size of this limit and the path of the effects during the program, which anticipate the program's end, but not the argument for identification.

Specification. Our main specification makes the change in the MUW linear in the plan's parameters,

$$\kappa^j = -\theta_G\,G_j + \theta_\tau\,\tau_j,$$

where G_j is the plan's guarantee at its households' baseline family sizes, in thousands of dollars per year, and τ_j is its tax rate, evaluated at the plan's mean baseline income for plans whose rate declines with income. An income effect implies $\theta_G > 0$, since a larger guarantee lowers the MUW, and $\theta_\tau > 0$, since a higher tax rate raises it. Plans combine guarantees and tax rates differently, so differences across plans separate θ_G from θ_τ.

Estimation. We estimate r, δ, β, the slope ξ, θ_G, and θ_τ by extending the GMM approach developed in Appendix D with a block of moments across plans. The block contains 22 plan-contrast moments, two for each of the eleven plans. They compare the deviation of each plan's treatment effects from the pooled ones, $\mathrm{TE}_t^j - \mathrm{TE}_t$, with the deviation the model implies, D_t^j. Applying Equation (15) to the deviations gives

$$D_t^j = r\,D_{t-1}^j + d_t^j - (1-\delta)\,d_{t-1}^j, \qquad D_0^j = d_0^j = 0,$$

where $d_t^j \equiv \mathrm{TE}_t^{\pi,j} - \mathrm{TE}_t^{\pi}$ is the deviation of the plan's direct effect from the pooled one, which depends on ξ, θ_G, and θ_τ through $\mathrm{TE}_t^{\pi,j} = \xi\big[\Delta p_t^j + \kappa^j\tilde{p}_t^j\big]$, the direct effect of the case $\beta = 0$ above. The two moments for plan j average $\big(\mathrm{TE}_t^j - \mathrm{TE}_t\big) - D_t^j$ over the program quarters 1–12 and over the post-program quarters 13–16. Each plan's treatment effects compare its treated heads with controls reweighted to match them on baseline employment and hours, and the pooled path averages the plans' effects with their shares of treated heads. The pooled block uses iterated

efficient weights and the contrast block diagonal weights; standard errors are clustered by family. The estimates without income terms, column (1) of Table C.3, are therefore the comparison point for those with them: their control weights and moment weights differ from those of Table 1.

Results. Table C.3 reports the estimates. Without income terms (column 1), the estimator gives $\hat{r} = 0.870$ (SE 0.119), close to the 0.862 of Table 1, but splits it differently: $\hat{\phi} = 0.301$ and $1 - \hat{\delta} = 0.569$, against 0.235 and 0.627 there. As in Section 4, the estimated patience is $\hat{\beta} = 0$, with and without income terms. With income terms in our main specification (column 2), propagation barely moves: $\hat{r} = 0.902$ (SE 0.130). The income terms are small and jointly insignificant: $\hat{\theta}_G = 0.013$ (SE 0.103) and $\hat{\theta}_\tau = 0.278$ (SE 0.820), with a joint p-value of 0.75. A guarantee \$1,000 per year higher changes contemporaneous hours by -0.4 hours per quarter (95% confidence interval from -6.7 to 5.9), against the reduction of 29 hours in the last program quarter. The split of propagation moves more, but within its standard errors: $\hat{\phi}$ rises from 0.301 to 0.371 (SE 0.254), and the long-run multiplier from 3.31 to 4.80 (SE 5.92). The estimates are similar under two alternative specifications, which make κ^j linear in dollar measures of the plan's change in lifetime resources, and when the recurrence moments start in quarter 3, 5, or 7: $\hat{r}$ lies between 0.83 and 0.90, and the joint p-values between 0.46 and 0.97 (results available upon request).

Table C.3: Allowing for income effects

	(1) No income terms	(2) Guarantee and tax rate
Propagation rate, r	0.870	0.902
	(0.119)	(0.130)
Reinforcement, ϕ	0.301	0.371
	(0.217)	(0.254)
Survival, $1 - \delta$	0.569	0.531
	(0.250)	(0.274)
Long-run multiplier	3.31	4.80
	(2.67)	(5.92)
Income-effect terms		
Guarantee, θ_G		0.013
		(0.103)
Tax rate, θ_τ		0.278
		(0.820)
Joint test of the income terms, p-value		0.747

Notes: This table presents GMM estimates of the model parameters with and without income-effect terms. The GMM estimator pools the treatment-effect path across plans and adds 22 plan-contrast moments. Column (1) has no income terms. Column (2) makes the change in the MUW linear in the plan's guarantee and tax rate. Patience is estimated at $\hat{\beta} = 0$ in both columns. Standard errors, clustered by family, are reported in parentheses; the long-run multiplier's is by the delta method.

C.1.3 Misreporting of Hours

Our estimates in Section 4 rest on the path of treatment effects on hours, from which we recover the parameters of the model (Section 3.4). Treated households had an incentive to underreport their work: Greenberg and Halsey (1983) compared interview reports with administrative earnings records, and find that they reported less of their work than controls did. The measured path of effects then is likely to overstate the reduction in hours during the program, and so biases the parameters we recover from it. We re-estimate the model on a path that attributes part of the measured effects to misreporting. We show the parameters barely change, and neither do our counterfactuals; a permanent NIT still exceeds the effect in the last program quarter by 27%.

Implementation. To account for misreporting, we consider the hypothetical case in which a share $1-k$ of the measured treatment effect in a set of quarters reflects misreporting, and re-estimate the model under it. To attribute a share $1-k$ of the measured effect in a set of quarters to misreporting, we replace each treated head's reported hours in those quarters, Y_{it}, with $Y_{it} - (1-k)\,\widehat{\text{TE}}_t$, where $\widehat{\text{TE}}_t$ is the estimate in Figure 2. This replacement makes the estimated treatment effect in those quarters exactly $k\widehat{\text{TE}}_t$ and leaves its standard error unchanged. We then re-estimate the model with the same sample, moments and weighting as in Table 1.

Results. We consider four scenarios. In the first three, we assume misreporting accounts for 5, 10 and 20% of the measured effects in the program quarters ($k = 0.95$, 0.90 and 0.80). Columns (2)–(4) of Table C.4 present the estimates under these scenarios and compare them with the main estimates of Table 1, in column (1). Panels A and B report the model's parameters, and Panel C the long-run multiplier and the response to a permanent NIT. Across these columns, the point estimates change little relative to the main estimate, both for the model's parameters and for the long-run multiplier. In the 5 and 10% scenarios, every estimate moves by less than half a standard error. Even in the most extreme scenario, where 20% of the treatment effect is bias, reinforcement moves by about one standard error and survival by less, and propagation and the long-run multiplier move little.[54]

[54] In a fourth scenario not presented for brevity, we go further and assume that misreporting continues after the program, in all sixteen quarters ($k = 0.90$). Because the effects after the program identify propagation, this scenario could bias estimates even further. It does not by construction: the model is identified from the shape of the treatment-effect path, the ratios between its effects, not from its level, so a misreporting share that is the same in every quarter cannot move the estimates.

Table C.4: Robustness to misreporting of hours

	(1)	(2)	(3)	(4)
		Misreported during the program		
	Main estimates	5%	10%	20%
Panel A: Propagation				
Propagation rate, r	0.862	0.857	0.852	0.840
	(0.114)	(0.118)	(0.122)	(0.132)
Panel B: Reinforcement, survival, and patience				
Reinforcement, ϕ	0.235	0.262	0.297	0.412
	(0.155)	(0.187)	(0.232)	(0.401)
Survival, $1-\delta$	0.627	0.595	0.555	0.429
	(0.225)	(0.257)	(0.301)	(0.464)
Patience, β	0.000	0.000	0.000	0.000
	(0.444)	(0.466)	(0.490)	(0.544)
Panel C: Permanent NIT				
Long-run multiplier, $\delta/(1-r)$	2.70	2.84	3.01	3.58
	(1.43)	(1.58)	(1.80)	(2.66)
Long-run response (hours/quarter)	37.0	35.1	33.3	29.6
Share of long run by Q12 (%)	88	88	88	89
Individuals	2,728	2,728	2,728	2,728

Notes: This table presents estimates of the model after attributing part of the measured effect of the NIT on hours to misreporting by treated households. Panel A reports the propagation rate r, and Panel B reinforcement ϕ, survival $1-\delta$ and patience β. Panel C reports the long-run multiplier, the long-run response to a permanent NIT and the share of that response reached by quarter 12, computed as in Section 4.3 from the adjusted treatment-effect path. Column (1) reproduces the main estimates of Table 1. Columns (2)–(4) attribute 5, 10 and 20% of the effects in program quarters 1–12 to misreporting ($k = 0.95$, 0.90 and 0.80). Standard errors are reported in parentheses; the long-run multiplier's is by the delta method.

C.1.4 Estimation Details

In this Appendix subsection, we give the estimation details that Section 4 leaves out: the inputs of the AIPW estimator, which enter both the treatment-effect estimates and the GMM estimator of Appendix D, and how we translate the long-run multiplier, a ratio of responses in the model, into hours.

AIPW inputs. *Propensity score.* We estimate each head's probability of treatment, $\hat{p}_i(X_i)$, by a probit of treatment on indicators for the randomization strata; on quarter-0 hours, employment and earnings; and on age, years of schooling, family size and the number of children under six at enrollment. We drop the heads outside the common support of $\hat{p}(X_i)$, from the larger of the two arms' minimum scores to the smaller of their maximum scores, $[0.179, 0.943]$; this removes 17 treated heads (step five of Appendix C.1.1).

Outcome model. The outcome model, $\hat{m}_{dt}$, predicts quarterly hours in each arm from the head's characteristics. It is a linear regression of Y_{it} on indicators for the strata and on quarter-0 employment, estimated separately in each arm and quarter.

From the long-run multiplier to hours. Figure 3 reports the effect of a permanent NIT in quarters 5 and 12 and in the long run. The model gives these effects up to a scale, the direct effect of the NIT in a single quarter, $\Delta\pi$; we first write them in units of $\Delta\pi$ and then estimate $\Delta\pi$ from the experiment.

The permanent path. A permanent NIT shifts the cost of working in every quarter from quarter 1 on, so its direct effect is the same in every quarter, $A\,\Delta\pi$, where $A = (1-\beta(1-\delta))/(1-\beta r)$ is the factor of Proposition 1(iii), equal to one when $\beta = 0$. Through work history, its effect in quarter t is

$$\mathrm{TE}^{\mathrm{P}}_t = A\,\Delta\pi\Big[1 + \phi\,\frac{1-r^{\,t-1}}{1-r}\Big],$$

which converges to $A\,\Delta\pi\,\delta/(\delta-\phi)$, the long-run multiplier times $\Delta\pi$ (Appendix Proposition B.1). The bars of Figure 3 evaluate this expression at $t = 5$, at $t = 12$, and in the limit.

The scale. We estimate $\Delta\pi$ from the observed effects of the 3-year program. Per unit of $\Delta\pi$, the model implies a direct effect in program quarter t of

$$W_t = 1 + \sum_{k=1}^{12-t} \frac{\phi}{1+\rho}\Big(\frac{r}{1+\rho}\Big)^{k-1},$$

which adds the response to the program quarters that remain (Equation (21)) and equals one when $\beta = 0$, and no direct effect after quarter 12. The program's effects in these units accumulate through work history: $x_t = \phi\,x^S_t + W_t$ and $x^S_{t+1} = (1-\delta)\,x^S_t + x_t$, with $x^S_1 = 0$. We choose $\Delta\pi$ by least squares over quarters 5–8, net of the quarter-0 estimate,

$$\widehat{\Delta\pi} = \frac{\sum_{t=5}^{8}\big(\widehat{\mathrm{TE}}_t - \widehat{\mathrm{TE}}_0\big)\,x_t}{\sum_{t=5}^{8} x_t^2}.$$

Quarters 5–8 come before the end of the program, which patient workers would anticipate. At our estimates ($\hat\beta = 0$), $\widehat{\Delta\pi} = -13.7$ hours, so the permanent NIT reduces hours by 24 in quarter 5, 32 in quarter 12, and 37.0 in the long run. With $\beta = 0.95$, we re-estimate reinforcement and survival with β fixed (0.205 and 0.673) and repeat the fit: $\widehat{\Delta\pi} = -7.9$ hours and $A = 2.17$, and the permanent NIT reduces hours by 29, 39, and 46.

C.1.5 Five-Year Validation

In this Appendix subsection, we give the details of the out-of-sample validation in Section 4.2: the construction of the five-year sample, the validation exercise, and a second exercise.

Five-year sample. We build the five-year sample with the steps of Appendix C.1.1, applied to the five-year arm. We keep the 2,362 heads enrolled in the five-year arm with pre-enrollment data, and then the strata that contain both treated and control heads among those still enrolled

in every quarter up to quarter 20, the end of the five-year program (1,169 heads). We keep single heads and both heads of couples that stay intact through quarter 20 (852 heads), and heads with labor-supply information through quarter 18 (570 heads). Trimming to the common support of a propensity score estimated as in Appendix C.1.4 removes 12 heads, leaving 558: 295 treated and 263 controls.

Validation. The validation asks whether the dynamics estimated from the three-year arm predict the path of the five-year arm's effects. We proceed in three steps. First, we estimate the five-year arm's effects with the same AIPW estimator, but with weights that account for two differences: the five-year arm assigned the plans in different proportions, and it enrolled a different population. We therefore align the weights in the five-year sample with the three-year arm's composition by entropy balancing: treated heads are weighted so that their distribution over the nine guarantee–tax-rate cells matches the three-year arm's, and all heads are weighted so that their site, head type, and quarter-0 hours match the three-year-arm sample. Second, we predict the path of these effects with the dynamic parameters fixed at the three-year estimates, for a twenty-quarter program that starts in quarter 1. The model gives the path up to its scale, the first-quarter effect, which we estimate from quarters 1–6 as in Appendix C.1.4, but by GLS with the inverse of the bootstrap covariance of the effects as the weight matrix. Third, we compare the predicted and observed effects in quarters 7–18, which the fit does not use, with the χ^2 statistic of the errors, weighted by their bootstrap covariance. The predictions match the continued build-up of the response: the root mean squared error is 9.6 hours per quarter and the mean absolute error 6.4, 8% of the average effect over those quarters, and the test does not reject the model ($\chi^2(11) = 4.57$, $p = 0.95$; Figure A.2).

A second validation. The second exercise uses no five-year outcome at all. It fixes the scale at its three-year estimate, a first-quarter effect of 13.7 hours (Appendix C.1.4), and predicts all eighteen quarters. The predicted path has the shape of the five-year response but not its level: the first-quarter effect estimated from the five-year arm, 31.9 hours, is more than twice the three-year one, consistent with the larger response of the five-year arm, which raises the pooled estimates in column (2) of Table C.1 above those in column (1). Because the five-year effects are imprecise, the joint test of the prediction errors does not reject the model ($\chi^2(18) = 17.81$, $p = 0.47$), although the errors are large (a root mean squared error of 40.6 hours per quarter). The five-year arm therefore supports the dynamics estimated from the three-year arm; its first-quarter effect is more than twice as large, although the joint test cannot reject a common one.

C.1.6 The Persistence Decomposition

In this Appendix subsection, we describe how we compute the decomposition of persistence in Section 4.2: first its four components, and then the share attributed to last quarter's hours when we impose $\delta = 1$.

Computing the covariances. Each component of observed persistence is a covariance between variables in the population, and each involves variables that we do not observe in the data: work history, S_{it}, and the other determinants of hours, $\varepsilon_{it} \equiv K\alpha_i + \pi_{it}$. We therefore first predict these variables from observed hours and the model's estimates, and then estimate the four covariances. We use only the control group, which did not receive the NIT and whose hours give our measure of observed persistence, $\hat{r}^{\mathrm{OLS}}$.

Predicted variables. Given $\hat{\delta}$, the law of motion predicts work history from observed hours, but the recursion needs a starting value: work history at the start of the panel, which depends on work before it and need not be zero, since the panel does not start at the beginning of working life. We assume that individuals were in a steady state before the panel, where $S^{SS} = (1-\delta)S^{SS} + Y^{SS}$, so that $\hat{S}_{i1} = Y_{i1}/\hat{\delta}$, and build work history forward from the observed hours $\{Y_{it}\}_{t\geq 1}$, where $t = 1$ is the first program quarter. Then, given the predicted sequence $\{\hat{S}_{it}\}$, we compute the residual $\hat{\varepsilon}_{it} = Y_{it} - \hat{\phi}\hat{S}_{it}$. Its stable part, $\bar{\hat{\varepsilon}}_{it}$, is the average of $\hat{\varepsilon}_{is}$ over the remaining quarters, $s = t, \ldots, 16$, which we require to span at least three quarters, so it is defined through quarter 14. Because the starting value is an assumption, we use only quarters 8 and later, by which its weight in work history has fallen below 5%, since $(1-\hat{\delta})^7 = 0.04$.

Estimated covariances. We estimate each covariance in the decomposition of Section 4.2 by its sample analogue, with $\hat{S}_{it}$, $\hat{\varepsilon}_{it}$, and $\bar{\hat{\varepsilon}}_{it}$ in place of the unobserved variables. We compute the sample covariances in the control group of the estimation sample, weighting each head by the inverse of the estimated probability of being a control, after removing quarter means, and pool them over quarters 8–14, the quarters after the burn-in for which the stable part is defined. Dividing each term by the pooled covariance of current and lagged hours gives the four shares, which sum to one: 26.8% for last quarter's hours, 35.7% for earlier work history, 27.5% for stable differences, and 10.0% for changes in costs.

Decomposition imposing full depreciation. With $\delta = 1$, work history is last quarter's hours, $S_{it} = Y_{i,t-1}$, so no starting value is needed and reinforcement equals propagation, $\phi = r$. We re-estimate the model imposing $\delta = 1$, with the same sample, moments, and weighting as in Table 1, which gives $\hat{\phi} = \hat{r} = 0.783$ (SE 0.102); the $\delta = 1$ score test in that table is computed at this restricted fit. The share attributed to last quarter's hours is then the sample analogue of

$$\frac{\mathrm{Cov}(\phi\, Y_{i,t-1}, Y_{i,t-1})}{\mathrm{Cov}(Y_{it}, Y_{i,t-1})} = \phi\, \frac{\mathrm{Var}(Y_{i,t-1})}{\mathrm{Cov}(Y_{it}, Y_{i,t-1})},$$

computed in the same sample, with the same weights and quarters as above. The variance and the covariance are features of the data, with a ratio of covariance to variance of 0.878, so the share rises from $0.235/0.878 = 26.8\%$ to $0.783/0.878 = 89.2\%$. The earlier-history term is zero by construction.

C.2 Social Media Use: The Digital Addiction Experiment

In this Appendix Section, we present details of the Digital Addiction experiment and the additional results for Section 5. We first document the experiment, the data and our sample (Appendix C.2.1). We then give the details on the verification of identifying assumptions (Appendix C.2.2). Finally, we give details on estimation (C.2.3) and on how long the effect of the bonus lasts (C.2.4).

C.2.1 Experiment and Sample

Experimental design. Allcott et al. (2022) ran a field experiment on smartphone use among U.S. adults between March and July 2020. They recruited U.S. residents aged 18 to 64 who used an Android phone as their only regular smartphone, through Facebook and Instagram ads that did not mention smartphones or social media. Of 3.3 million users shown the ads, 5,320 installed Phone Dashboard, the study's app; 2,053 reached the start of the treatments, and 1,938 completed the study. Phone Dashboard records which app is in the foreground every five seconds while the screen is on. The study focuses on six apps: Facebook, Instagram, Twitter, Snapchat, web browsers, and YouTube ("FITSBY"). After three baseline weeks, participants were randomized into two independent treatments, within eight strata defined by above- or below-median baseline FITSBY use and two survey indices of self-control, and they learned their assignment at the start of week 4. Participants in the Bonus group (25%) were offered \$50 for each hour by which they reduced their average daily FITSBY use during the bonus period, weeks 7–9, below a benchmark, up to \$150, or about \$2.50 per hour of use forgone. The benchmark was their average daily use in the baseline period, rounded up to the nearest hour. They thus learned of the bonus three weeks before it took effect, and knew it would last three weeks. The Bonus Control group (75%) was told it would not receive the bonus, and 0.2% of participants had their bonus set by a multiple price list. Independently, participants in the Limit group (60%) gained access to app time limits in Phone Dashboard for the rest of the study.

Data. We use the public replication data of Allcott et al. (2022), which cover the 1,933 participants who completed the study outside the multiple-price-list group. Our outcome is FITSBY use in minutes per day, measured weekly rather than over the 20-day periods of the original analysis, to exploit more of the variation over time. The fifteen weeks split into baseline (weeks 1–3), announcement (4–6), bonus (7–9), and follow-up (10–15). A participant is treated if assigned to the Bonus group, and is a control if assigned to the Bonus Control group.

Sample. Our estimation sample contains 769 participants, 191 of them in the Bonus group. We build it from these data in two steps. First, we keep the 783 participants not assigned to the Limit treatment, so that treated and control participants differ only in the bonus, conditional on the randomization strata. Second, we keep the participants with use in every week from 7 to 15, so that every participant is observed in every week the estimator uses; this step drops 14 participants.

The conventional estimates of propagation and the decomposition of persistence use weeks 1–15, so they use the 577 controls of the estimation sample observed in all fifteen weeks.[55]

C.2.2 Verifying Identifying Assumptions

Known, constant intervention path. A participant in the Bonus group received

$$f_i(\bar{y}_i) = 50 \cdot \min\big\{3,\ \max\{0,\ \lceil \bar{y}_i^0 \rceil - \bar{y}_i\}\big\}, \tag{41}$$

where $\bar{y}_i$ is average daily FITSBY use, in hours, over the 20-day bonus period, and $\bar{y}_i^0$ is the same average in the first period, rounded up to the nearest hour to give the benchmark. The benchmark was set before participants learned of the bonus, so they could not manipulate it, and rounding it up only adds a lump-sum payment. Between the benchmark and three hours below it, an additional hour of daily use lowers the payment by \$50, or about \$2.50 per hour of use, so the bonus raises the cost of use by a constant amount during weeks 7–9, known three weeks in advance: the intervention of Equation (9). The schedule has two kinks: participants who use more than their benchmark, or who cut use by more than three hours a day, receive nothing for a further hour of reduction; in our Bonus group, 6% and 13% did so. The kinks make the shift differ across participants, which the mapping allows, but could also make it change over the bonus weeks, which Assumption X.3 rules out. They do not: the payment depends on average use over the whole period, so each day carries the same incentive, and the estimated effects barely change across the three paid weeks (60, 64, and 59 minutes per day), with no decline as participants learn they will end up at a kink.

Exclusion. After the bonus ended, nothing in the design changed the cost of use, and participants outside the Limit treatment faced no other intervention. The payments could still change use after the bonus through income effects, as for the NIT (Appendix C.1.2), but they were small, \$84 on average, so we abstract from income effects.

C.2.3 Estimation Details

In this Appendix subsection, we give additional estimation details for Section 5. We provide brief details, focusing on the differences from the implementation in Appendices C.1.4 and C.1.6: the AIPW inputs, the decomposition of persistence, and the conversion of the long-run multiplier into minutes per day.

[55] Averaged over the 20-day periods of the original analysis, our weekly effects are close to the published estimates of Allcott et al. (2022). They estimate that the bonus reduced FITSBY use by 56 minutes per day during the bonus period, 39% of the control group's use, and by 19 and 12 minutes per day in the two following periods. Averaging our weekly estimates over the same weeks gives 61 minutes per day (43%), 17, and 11. The two sets of estimates use different samples and specifications: theirs include the participants assigned to the Limit treatment and control for baseline use, while ours exclude them and condition only on the randomization strata.

AIPW inputs. The AIPW scores (43) use the same two inputs as for the NIT (Appendix C.1.4), both simpler here. We estimate the propensity score with sample analogues: $\hat{p}$, is the share of Bonus participants within each randomization stratum. The outcome model, $\hat{m}_{dt}$, is a linear regression of weekly use on indicators for the eight strata, estimated separately in each arm and week.

Decomposition. We compute the decomposition of persistence in Section 5.2 as in Appendix C.1.6, with two differences. First, we use the 577 controls observed in all fifteen weeks (Appendix C.2.1), with equal weights, because the probability of assignment is the same in every stratum. Second, fifteen weeks do not allow a burn-in: the starting value of accumulated use carries less than 5% of its weight only after twelve weeks, since $(1-\hat{\delta})^{12} = 0.05$. We therefore pool the covariances over weeks 2–13, all the weeks in which both last week's use and the stable part are defined. Because the starting value then matters more, we also start accumulated use at zero rather than at its steady-state value: the causal share falls from 34% to 25%.

From the long-run multiplier to minutes. Figure 5 reports shares of the long-run response. By Appendix Proposition B.1(ii), the share reached in week t of the incentive is $1-(\phi/\delta)\,r^{\,t-1}$, at our estimates and in the limit $\beta \to 1$ alike, so it does not depend on the level of the response to the bonus. The long-run response in minutes per day, reported in Section 5.3, does depend on his level: it is the long-run multiplier times the direct effect of the bonus in a single week, $\Delta\pi$. We estimate $\Delta\pi$ as in Appendix C.1.4, by least squares over the three bonus weeks, net of the estimate for week 6, the last week before the bonus (-4.0 minutes per day), in place of quarter 0. With $\beta = 0$, the direct effect is the same in every bonus week ($W_t = 1$) and $A = 1$, so $\widehat{\Delta\pi} = -53$ minutes per day is the one-week response. The permanent incentive then reduces use by $1.51 \times 53 = 80$ minutes per day in the long run, 36% more than the effect observed in the last bonus week (59 minutes per day) and 56% of control-group use in the bonus weeks (143 minutes per day).

C.2.4 How Long the Effect of the Bonus Lasts

Besides the response to a permanent incentive, we also report what the bonus itself achieved. For a one-week incentive, the measure would be the cumulative lifetime effect of Proposition 1(ii), together with the share of it that falls within the weeks observed. The bonus, however, lasted three weeks. We therefore measure its cumulative effect directly: we add up the observed effects and extrapolate the last one at the estimated propagation rate until it dies out.

After week 9, the effect of the bonus contains no direct response to it, so it decays at the propagation rate (Equation (14)): with $\hat{r} = 0.853$, it halves every $\ln 0.5 / \ln \hat{r} = 4.3$ weeks. In weeks 7–15, the bonus reduced use by 1,869 minutes per participant, the sum of the weekly effects times seven days. Extrapolating the week-15 effect at rate $\hat{r}$ adds a further 447 minutes, so 81% of the cumulative effect of the bonus falls within the fifteen weeks observed.

C.3 Gym Attendance: The Charness–Gneezy Experiment

In this Appendix Section, we present details of the incentives-to-exercise experiment of Charness and Gneezy (2009) and the additional results discussed in Section 6. We first document the experiment, the data and our sample (Appendix C.3.1). We then give the details on the verification of identifying assumptions (Appendix C.3.2). Next, we give details on estimation (C.3.3). Finally, we describe how we compute the effect of a permanent incentive (C.3.4).

C.3.1 Experiment and Sample

Experimental design. In their 'Study 1', Charness and Gneezy (2009) recruited 120 students at the University of Chicago and obtained the records of their visits to the campus gym, to which all students had access. Forty students, recruited in separate sessions for an unrelated experiment, form the control group. The other 80 were paid $25 to visit the gym at least once in the following week; when they returned a week later, half of them, chosen at random, were offered an additional $100 for at least eight visits in the next four weeks. We do not use their 'Study 2', which provides a different type of incentive: all its participants were paid the same amount for attending three meetings with health measurements, and its groups differed only in a requirement to go to the gym once or eight times in the following month. The requirement does not change the cost of a visit, so it is not the type of temporary intervention our model considers (Equation (2)).

Data. The experimenters obtained each participant's visits to the campus gym from the facility's swipe-card records. We use the public replication data of Charness and Gneezy (2009), which report the number of visits in each week; our outcome is weekly gym visits, the finest frequency in the public data and the one at which the payments were defined. The data cover eight weeks before the payments, the five weeks with payments (the week of the $25 payment and the four weeks of the $100 offer), and the seven weeks after them. Following Charness and Gneezy (2009), the data exclude the week of spring break, which falls just before the payments start. A participant is treated if in the five-week group, and is a control if in the control group.

Sample. Our estimation sample contains the 40 participants of the five-week group and the 40 of the control group, all observed in every week. The estimator uses weeks 1–10, the five weeks with payments and the five that follow. The conventional estimates and the decomposition use the control group over weeks 1–12, the longest run of consecutive weeks in the data, since spring break interrupts it just before the payments. We hold out the one-week group to validate the estimated dynamics (Section 6.2.2). It also offers an alternative comparison for the five-week group, randomized at the return visit, but only for the four weeks of the $100 offer, since both groups were paid in the first week; we report it as a robustness check (column 3 of Table A.2), and because it reuses the five-week group, it cannot serve as an out-of-sample check.

C.3.2 Verifying Identifying Assumptions

Known, constant intervention path. Participants knew the terms of each payment once it was offered, and Section 6 treats the payments as a constant reduction in the cost of a visit during weeks 1–5. We derive this reduction from the payment schedule, which is nonlinear in the number of visits, and then discuss the surprise \$100 offer.

The payments as a change in the cost of a visit. The five-week group was paid \$25 for visiting the gym at least once in week 1 and, in an offer made at the start of week 2, \$100 for at least eight visits in weeks 2 to 5. The payments are not a price per visit: they lower the cost of a visit only if it is the visit that reaches the target. With y_{it} participant i's visits in week t, the cost of a visit in week t is therefore

$$\tilde{p}_{it} = p_{it} - f_t(y_i), \qquad f_t(y_i) = \begin{cases} 25 \cdot \mathbf{1}\{y_{i1} = 1\}, & t = 1, \\ 100 \cdot \mathbf{1}\{\sum_{s=2}^{5} y_{is} = 8\}, & t = 2, \ldots, 5. \end{cases} \tag{42}$$

The reduction is a jump: the full payment for the visit that reaches the target, and nothing for any other. For a participant who plans to reach a target number of visits over the four weeks, whatever it is, the incentive can then be interpreted as a change in the cost of every visit by the same amount, b_i, in each week of the offer, since the payment depends only on how many visits she makes, not on when. This gives the intervention of Equation (9), with $b_i = 0$ for participants who would miss the incentive's target.

This reading assumes that participants value each visit, in each of the five weeks, at the same amount when they make it. This value is not the dollar amount per visit, which differs across weeks (\$25 for one visit in week 1 against \$100 for eight in weeks 2 to 5) and across participants; only the time profile of the shift matters for the estimates, since its level is a parameter the estimator recovers (Appendix D).

The surprise \$100 offer. The offer came a week into the payments, so in week 1 participants knew only of the \$25 payment for a single visit, and the first week may differ from the four that follow. This matters only if participants are patient, which our estimates do not suggest, and dropping the first week leaves the conclusions unchanged (column 2 of Table A.2).

C.3.3 Estimation Details

In this Appendix subsection, we give the estimation details that Section 6 leaves out, where they differ from those of the NIT (Appendices C.1.4 and C.1.6). First, we describe how we estimate the treatment effects. Second, we describe how we compute the decomposition of persistence.

Treatment effects. The experiment had no randomization strata, so treatment effects are differences in mean weekly visits between the five-week group and the control group: the AIPW scores (43) with $\hat{p}$ the sample share of the five-week group and no outcome model.

Decomposition. We compute the decomposition of persistence in Section 6.2 as in Appendix C.1.6, with two differences. First, we use the 40 controls over weeks 1–12 (Appendix C.3.1), with equal weights. Second, twelve weeks do not allow a burn-in, and because fitness survives for several weeks, the starting value matters: with fitness starting at zero, past attendance causes 58% of persistence, and at the steady state the model implies given first-week visits, 96%; we report their midpoint, 77%. Imposing $\delta = 1$, as in Appendix C.1.6, the refit gives $\hat{\phi} = \hat{r} = 0.72$, and the share attributed to last week's attendance rises from 22% to 87%.

C.3.4 The Effect of a Permanent Incentive

Figure 7 reports the effect of a permanent incentive. We estimate $\Delta\pi$ as in Appendix C.1.4, using least squares over the first three weeks with payments, net of the average estimate for the eight weeks before the payments (0.01 visits per week), in place of quarter 0, and using $\beta = 0$. This gives a one-week response of $\widehat{\Delta\pi} = 1.00$ visits per week. The permanent incentive then raises attendance by 1.7, 2.3, 3.3 and 4.2 visits per week in weeks 5, 9, 16 and 25, and by $7.42 \times 1.00 = 7.4$ in the long run. By Appendix Proposition B.1, the share of the long run reached after t weeks is $1 - (\hat{\phi}/\hat{\delta})\,\hat{r}^{\,t-1}$: 13% in the first week, with the remaining gap closing by $1 - \hat{r} = 2.9\%$ a week, so that it halves in $\ln 0.5 / \ln \hat{r} = 24$ weeks. The purple bars keep $\hat{\phi}$ and $\widehat{\Delta\pi}$ but set $\delta = 1$, so that $r = \hat{\phi}$, and the effect levels off at $\widehat{\Delta\pi}/(1 - \hat{\phi})$, about 1.2 visits per week, within two weeks.

D Estimation

This appendix describes the estimator we apply to the experiments in Sections 4–6; Appendix C reports the choices made in each application. The moments come from the recurrence that treatment effects satisfy while the incentive is in place (Lemma B.6) and from their geometric decay after it ends (Equation (14)). The treatment effects are not estimated in a separate first step; they enter the moments at the individual level, so the estimator is a single joint GMM problem.

Data. We observe $\{\{Y_{it}\}_{t=0}^{T},\, D_i,\, X_i\}_{i=1}^{N}$, where Y_{it} is the choice of individual i in period t, $D_i \in \{0,1\}$ indicates assignment to treatment, and X_i collects pre-treatment characteristics, including the randomization strata. The incentive is in place over $\mathcal{W} = \{t_0, \dots, t_1\}$, the panel extends at least two periods beyond it, $T \geq t_1 + 2$, and Assumptions X.2 and X.3 hold with common propagation.

D.1 Treatment-Effect Scores

Let $p(X_i) \equiv \Pr(D_i = 1 \mid X_i)$ and $m_{dt}(X_i) \equiv \mathbb{E}[Y_{it} \mid D_i = d, X_i]$, with estimates $\hat{p}$ and $\hat{m}_{dt}$. For each individual and period, we form the augmented inverse-probability-weighted (AIPW) score (Robins et al., 1994),

$$\tilde{Y}_{it} = \hat{m}_{1t}(X_i) - \hat{m}_{0t}(X_i) + \frac{D_i\big(Y_{it} - \hat{m}_{1t}(X_i)\big)}{\hat{p}(X_i)} - \frac{(1 - D_i)\big(Y_{it} - \hat{m}_{0t}(X_i)\big)}{1 - \hat{p}(X_i)}. \tag{43}$$

If either p or m_{dt} is estimated consistently, the sample mean of $\tilde{Y}_{it}$ is consistent for $\textsc{te}_t$. Two familiar estimators are special cases. With $\hat{m}_{dt} = 0$ and p known by design, (43) is the Horvitz–Thompson transform. With $\hat{p}$ equal to the sample share of treated individuals and $\hat{m}_{dt}$ constant, the sample mean of $\tilde{Y}_{it}$ is the difference in means between the two arms. When p is estimated, we restrict the sample to the common support of $\hat{p}$ in the two arms. The recurrence and the post-treatment decay therefore hold in expectation for $\tilde{Y}_{it}$, and we write every moment at the individual level.

D.2 Moments

We parameterize the estimator by $\psi = (r, \beta, \delta, b)$: the propagation rate r, the discount factor $\beta = 1/(1+\rho)$, the depreciation rate δ, and a scale b for the cost shift the treatment induces. A single scale imposes that the cost shift is constant across the window, as in the proof of Proposition 4; identification needs it only over the final treated periods.

Dividing (26) by a_0, and using that its characteristic roots are r and $(1+\rho)/r$, the treatment effects satisfy, over the window,

$$\mathbb{E}\Big[\tilde{Y}_{it} + \theta_1\, \tilde{Y}_{i,t+1} + \theta_2\, \tilde{Y}_{i,t+2} - R_t(b, \delta, \beta)\Big] = 0, \qquad t_0 \leq t \leq t_1, \tag{44}$$

where the second coefficient is the discount factor, $\theta_2 = \beta$, while the characteristic polynomial fixes the first,

$$\theta_1 \;=\; -\,\frac{1+\beta\, r^2}{r}. \tag{45}$$

Evaluating (27) with a constant cost shift, and absorbing the factor $1/a_0$ into b, the forcing is

$$R_t(b,\delta,\beta) = \begin{cases} \delta\big(1-\beta(1-\delta)\big)\, b, & t_0 \le t \le t_1 - 2, \\ \big(\delta + \beta(1-\delta)^2\big)\, b, & t = t_1 - 1, \\ -(1-\delta)\, b, & t = t_1. \end{cases}$$

After the program ends the forcing vanishes and the recurrence collapses to a one-period decay,

$$\mathbb{E}\Big[\tilde{Y}_{i,t+1} - r\,\tilde{Y}_{it}\Big] = 0, \qquad t_1 + 1 \le t \le T - 1, \tag{46}$$

which carries r directly. Carrying r as a primitive, rather than as a root computed from (θ_1, θ_2), is what lets a single propagation rate appear in both blocks of moments and be read off the post-experiment tail cleanly. The post-treatment pairs identify r; the moments at $t_1 - 1$ and t_1, which contain the interior and boundary gaps of Section 3.4, separate δ; and the moments at earlier treated dates identify β (Proposition 4). The estimator uses $t_1 - t_0 + 1$ recurrence moments and $T - t_1 - 1$ post-treatment pairs. When β is fixed, for example at the myopic value $\beta = 0$, where $\theta_2 = 0$ and $\theta_1 = -1/r$, the parameter vector is (r, δ, b). We recover state dependence as $\hat{\phi} = \hat{r} - (1 - \hat{\delta})$.

D.3 Estimation and Inference

We estimate ψ by GMM with equally weighted moments.[56] Because every moment is written at the individual level in $\tilde{Y}_{it}$, there is no separate first step for the treatment effects, and the ordinary GMM sandwich variance already incorporates their sampling variability. The sandwich treats $\hat{p}$ and $\hat{m}_{dt}$ as known. The AIPW score is Neyman orthogonal to both, so small errors in either change its expectation only at second order. When both are estimated consistently, their estimation therefore does not affect the asymptotic variance; when p is known by design, this holds for any outcome model. We restrict β to $[0, 1]$ and δ to $(0, 1]$. State dependence, the long-run multiplier and the other derived objects follow by the delta method.

D.4 Tests

Overidentification. The overidentification test is Hansen's J statistic for the stacked moments, with degrees of freedom equal to the number of moments less the number of free parameters.

[56] We weight all moments equally since treatment effects are highly serially correlated, generating an unstable estimated optimal weighting matrix. Equal weighting sacrifices asymptotic efficiency for stability.

Constant propagation. The constant-propagation test is the J statistic of the post-treatment pairs (46) alone, with r as the only parameter. It tests that effects decay at a common rate after the incentive ends, the restriction in Equation (14).

Full depreciation. The test of $\delta = 1$ is a one-sided score (LM) test of $\delta = 1$ against $\delta < 1$, evaluated at the fit that imposes $\delta = 1$. Because the null lies on the boundary of the parameter space, the statistic is compared with an equal mixture of χ^2_0 and χ^2_1 (5% critical value 2.71). Under $\delta = 1$, the boundary gap is zero (Corollary B.1). The score can also be evaluated at a fit in which r is first estimated from the post-treatment pairs alone and the remaining parameters are then estimated given $\hat{r}$.

Forward-looking behavior. When $\hat{\beta}$ lies at zero, a Wald statistic is not informative about β, because the restriction is at the boundary and the standard error conditional on the other parameters is too small. We test $\beta = 0$ with the score of the moments in the direction of β, orthogonalized with respect to the other parameters. With $\bar{g}$ the average moment vector at the fit that imposes $\beta = 0$, $\hat{\Omega}$ its variance, and $\hat{G}_\beta$ and $\hat{G}_{-\beta}$ the columns of the Jacobian for β and for the other parameters, let $M \equiv \hat{G}_\beta - \hat{G}_{-\beta}\big(\hat{G}'_{-\beta}\hat{G}_{-\beta}\big)^{-1}\hat{G}'_{-\beta}\hat{G}_\beta$. The statistic $z = -\sqrt{N}\, M'\bar{g}/\sqrt{M'\hat{\Omega}M}$ is asymptotically standard normal under $\beta = 0$, and the test is one-sided because $\beta \geq 0$.

Relevance. As a measure of instrument relevance, we report the squared t statistic of the treatment effect in the last treated period, $\textsc{te}_{t_1}$, the assignment-induced gap in choices whose propagation after the incentive identifies r.

E Supplementary Results

This appendix collects results that are not proofs of main-text statements, in two parts: extensions of the model (Appendix E.1) and additional identification results (Appendix E.2).

E.1 Extensions of the Model

This appendix section studies extensions to the model of Section 2. We first discuss the mapping of a life-cycle labor-supply model (Appendix E.1.1). Then, we discuss how the target parameters change when the intervention lasts more than one period (Appendix E.1.2). Next, we study the optimal policy and target parameters under non-quadratic preferences (Appendix E.1.3). Finally, we conclude with a discussion about the solution of the finite-horizon case (Appendix E.1.4).

E.1.1 Mapping to a Life-Cycle Labor-Supply Model

This subsection maps a life-cycle version of the labor-supply example of Section 2.2 into (1), with human-capital accumulation added to the search cost as a channel through which work history affects current work.

Consider a worker who chooses consumption c_t and hours y_t to maximize $\sum_{t\geq 0}(1+\rho)^{-t}\big[c_t - v(y_t, S_t)\big]$, where v is the disutility of work, subject to the lifetime budget constraint

$$\sum_{t\geq 0}\frac{c_t}{(1+\rho)^t} \leq m_0 + \sum_{t\geq 0}\frac{\big(w_t - \gamma(S_t)\big)\, y_t}{(1+\rho)^t},$$

where m_0 is initial wealth, the market interest rate equals ρ, and $\gamma(S_t)$ is the search cost of Section 2.2. Hours accumulate into work history, $S_{t+1} = (1-\delta)S_t + y_t$, and the disutility of work may depend on work history, for example through fatigue. The hourly wage has a permanent individual component α, such as ability, a time-varying component $\tilde{w}_t$, such as market conditions, and an experience premium $w^S(S_t)$:

$$w_t = \alpha + \tilde{w}_t + w^S(S_t).$$

Experience and search costs run through the same stock: hours build experience and lower search costs together, and both effects depreciate at the rate δ. Substituting the budget constraint into the objective gives

$$\max_{\{y_t\}} \sum_{t\geq 0}\frac{1}{(1+\rho)^t}\Big[\big(\alpha + w^S(S_t) - \gamma(S_t)\big)\, y_t - v(y_t, S_t) + \tilde{w}_t\, y_t\Big],$$

which is (1) with preferences $u(y,S) = -v(y,S)$, the stock-dependent return $R(y,S) = \big(\alpha + w^S(S) - \gamma(S)\big)\, y$, the wage net of the search cost, and the cost of working $p_t = -\tilde{w}_t$. The time-

varying part of the wage enters as a cost of working with the opposite sign because, like the cost of working in Section 2.2, it shifts the payoff to every hour by the same amount and does not depend on work history. The complementarity is $U_{yS} = (w^S)'(S) - \gamma'(S) - v_{yS}(y, S)$: experience raises the return to an additional hour, recent work lowers its search cost, and fatigue may lower the return.

E.1.2 Policies of Different Lengths

This subsection extends the targets of Section 2.5 to temporary interventions that last more than one period. We define the interventions and the targets in the general model of Section 2, then specialize to the linear–quadratic model of Section 2.6, where the targets have closed forms.

Setup. Suppose the transportation subsidy is in place for $L \geq 1$ consecutive periods starting at $\underline{t}$, $\mathcal{W}(L) \equiv \{\underline{t}, \ldots, \underline{t} + L - 1\}$, so that $\bar{t} = \underline{t} + L - 1$. The interventions of Section 2.5 are the special cases $\mathcal{W}(1) = \mathcal{W}^{\text{TEMP}}$ and $\mathcal{W}(\infty) = \mathcal{W}^{\text{PERM}}$. As for every intervention in Section 2.4, by $\underline{t}$ the worker knows the subsidy's magnitude and its end date, and the subsidy changes neither the worker's beliefs about baseline costs nor those about future interventions. We index periods by event time $j \geq 0$, the number of periods since the intervention began.

Targets of an intervention of any length. We normalize by the response in the last period of the intervention. Define

$$\text{CLE}_{\underline{t}}(L) \equiv \frac{\displaystyle\sum_{k \geq L} \frac{dy^\star_{\underline{t}+k}}{db}\Big|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}(L)}}{\displaystyle\frac{dy^\star_{\underline{t}+L-1}}{db}\Big|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}(L)}}, \qquad \text{LRM}_{\underline{t}}(L) \equiv \frac{\lim_{k\to\infty} \frac{dy^\star_{\underline{t}+k}}{db}\big|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}^{\text{PERM}}}}{\frac{dy^\star_{\underline{t}+L-1}}{db}\big|_{S_{\underline{t}}, b=0, \mathcal{W}=\mathcal{W}(L)}}.$$

For $L = 1$ these are $\text{CLE}_{\underline{t}}$ and $\text{LRM}_{\underline{t}}$ of Section 2.5. The cumulative lifetime effect $\text{CLE}_{\underline{t}}(L)$ answers the following question: given the hours response in the last period of an intervention that lasts L periods, how many additional hours does the intervention generate after it ends? The long-run multiplier $\text{LRM}_{\underline{t}}(L)$ answers the following question: how much larger is the long-run effect of a permanent subsidy than the response in the last period of the same subsidy for L periods? In the general model, the targets depend on $\underline{t}$, on the inherited work history and on information, as in Section 2.5. We write

$$dy_j(L) \equiv \frac{dy^\star_{\underline{t}+j}}{db}\Big|_{S_{\underline{t}},\, b=0,\, \mathcal{W}=\mathcal{W}(L)}$$

for the response of hours in event time j, and define $dS_j(L)$ analogously; $dS_0(L) = 0$ because $S_{\underline{t}}$ is held fixed.

The linear–quadratic case. From here on we impose the linear–quadratic specification (4) and the policy of Lemma 1, as Section 2.6 does.

By Lemma 1, $y^{\star}_{\underline{t}+j} = \phi S_{\underline{t}+j} + K\alpha + \pi_{\underline{t}+j}$, and the intervention changes neither α nor $S_{\underline{t}}$. Hence, with $d\pi_j(L)$ defined like $dy_j(L)$,

$$dy_j(L) = \phi\, dS_j(L) + d\pi_j(L), \qquad dS_{j+1}(L) = r\, dS_j(L) + d\pi_j(L), \qquad dS_0(L) = 0.$$

The responses, and hence the targets, do not depend on $\underline{t}$, on the inherited work history or on information, and we drop the subscript $\underline{t}$, as Section 2.6 does. The contemporaneous response to a one-period subsidy, $dy_0(1)$, is the denominator of LRM in Proposition 1. It sets the scale of every response below.

The direct effect. Under the linear–quadratic specification, the impulse response in Proposition 1(i) does not depend on t. Write $\theta^{(k)} = \phi\, r^{k-1}$ for $k \geq 1$ and $\theta^{(0)} \equiv 1$, and, as in Appendix B.3, $H_m \equiv \sum_{k=0}^{m-1} (\beta r)^k$ for $m \geq 0$, which converges as $m \to \infty$ because $|\beta r| < \sqrt{\beta} < 1$ by Lemma B.1(iii).

Lemma E.1. *Suppose the conditions of Lemma 1 hold.*

(i) A unit reduction in the cost expected $k \geq 0$ periods ahead raises the current cost outlook, and hence current hours holding work history fixed, by $dy_0(1)\, \beta^k \theta^{(k)}$.

(ii) Under $\mathcal{W}(L)$, the direct effect of the intervention is

$$d\pi_j(L) = \begin{cases} dy_0(1)\,\big[1 + \beta\phi H_{L-1-j}\big], & 0 \leq j \leq L-1, \\ 0, & j \geq L, \end{cases}$$

and $d\pi_j(\infty) = dy_0(1)\,\big(1 - \beta(1-\delta)\big)/(1-\beta r)$ for every $j \geq 0$.

Proof. (i) By (21), with $\beta = 1/(1+\rho)$, the forecast $\mathcal{E}_s[p_{s+k}]$ enters the k-th term of the sum with weight $(\beta r)^k$ and, for $k \geq 1$, the $(k-1)$-th term with weight $-\beta(1-\delta)\,(\beta r)^{k-1}$. Hence, for every $k \geq 0$,

$$\frac{\partial \pi_s}{\partial \mathcal{E}_s[p_{s+k}]} = \xi\, \beta^k r^{k-1}\big[r - (1-\delta)\big] = \xi\, \beta^k \theta^{(k)},$$

where the case $k = 0$ reads ξ, with $\xi = 1/(A + C\phi)$ the cost pass-through of Lemma B.4. Under a one-period subsidy, only $p_{\underline{t}}$ falls and $dS_0(1) = 0$, so $dy_0(1) = d\pi_0(1) = -\xi$. Substituting gives (i).

(ii) The worker knows the end date, and the subsidy conveys no information about baseline costs. At event time $j \leq L-1$, the subsidy therefore lowers the costs expected $k = 0, \ldots, L-1-j$ periods ahead by b, and no others. Summing (i) over these k gives $dy_0(1)\big[1 + \phi \sum_{k=1}^{L-1-j} \beta^k r^{k-1}\big] = dy_0(1)\big[1 + \beta\phi H_{L-1-j}\big]$. At $j \geq L$, no expected cost changes. Under the permanent subsidy, summing (i) over all $k \geq 0$ gives $dy_0(1)\big[1 + \beta\phi/(1-\beta r)\big] = dy_0(1)\big(1 - \beta(1-\delta)\big)/(1-\beta r)$. □

Part (ii) is, per unit of subsidy, the direct effect TE_t^{π} of Appendix B.3, with $t = \underline{t} + j$ and $\bar{t} = \underline{t} + L - 1$. In the last treated period one period remains, and $H_0 = 0$, so $d\pi_{L-1}(L) = dy_0(1)$ whatever the length of the intervention and the degree of patience: in that period, the response to the change in costs is the same as under a one-period subsidy. This is what makes the last period a natural normalization.

Targets. Under the linear–quadratic specification, both targets depend on the path of responses during the intervention through a single object,

$$\kappa_L \equiv \frac{dS_{L-1}(L)}{dy_{L-1}(L)} = \frac{\sum_{j=0}^{L-2}(1-\delta)^{L-2-j}\, dy_j(L)}{dy_{L-1}(L)},$$

the additional work history inherited by the last treated period per additional hour worked in that period; the second equality uses the law of motion of the stock, and $\kappa_1 = 0$ by construction.

Proposition E.1 (Targets of an intervention of any length)**.** *Suppose the conditions of Lemma 1 hold, and let $L \geq 1$. Then the responses and the target parameters do not depend on $\underline{t}$, and we write them without the subscript:*

(i) After the intervention ends, hours respond only to the work history it leaves behind: $dy_{L+k}(L) = \phi\, r^k\, dS_L(L)$ for $k \geq 0$.

(ii) If $|r| < 1$ and $dy_{L-1}(L) \neq 0$, the cumulative lifetime effect is

$$\text{CLE}(L) = \text{CLE}\left[1 + (1-\delta)\,\kappa_L\right].$$

(iii) If $|r| < 1$ and $dy_{L-1}(L) \neq 0$, the long-run multiplier is

$$\text{LRM}(L) = \text{LRM}\left[1 - \phi\,\kappa_L\right].$$

Here CLE *and* LRM *are given by Proposition 1, and, whenever $dy_{L-1}(L) \neq 0$,*

$$\kappa_L = \frac{M_L}{1 + \phi\, M_L}, \qquad M_L \equiv \frac{dS_{L-1}(L)}{dy_0(1)} = \sum_{n=2}^{L} r^{\,n-2}\big(1 + \beta\phi H_{n-1}\big),$$

with $M_1 = 0$, so that for $L = 1$ parts (ii) and (iii) are those of Proposition 1.

Proof. The responses do not depend on $\underline{t}$ by the linear–quadratic case above, and hence neither do the targets.

(i) For $j \geq L$, $d\pi_j(L) = 0$ by Lemma E.1(ii). Hence $dS_{L+k}(L) = r^k\, dS_L(L)$ and $dy_{L+k}(L) = \phi\, r^k\, dS_L(L)$ for $k \geq 0$.

(ii) Summing (i) over $k \geq 0$ and using $|r| < 1$,

$$\sum_{j \geq L} dy_j(L) = \frac{\phi}{1-r}\, dS_L(L) = \textsc{cle} \cdot dS_L(L),$$

by Proposition 1(ii). The law of motion gives $dS_L(L) = (1-\delta)\, dS_{L-1}(L) + dy_{L-1}(L)$. Dividing by $dy_{L-1}(L)$ gives $\textsc{cle}(L)$.

(iii) At $j = L-1$, $d\pi_{L-1}(L) = dy_0(1)$ by Lemma E.1(ii). Hence $dy_{L-1}(L) = \phi\, dS_{L-1}(L) + dy_0(1)$, or

$$dy_0(1) = dy_{L-1}(L)\big(1 - \phi\, \kappa_L\big).$$

By the definition of $\textsc{lrm}$, the long-run response to the permanent subsidy is $dy_0(1) \cdot \textsc{lrm}$, which is finite by Proposition 1(iii), and

$$\textsc{lrm}(L) = \frac{dy_0(1) \cdot \textsc{lrm}}{dy_{L-1}(L)} = \textsc{lrm}\,\big[1 - \phi\, \kappa_L\big].$$

Finally, iterating $dS_{j+1}(L) = r\, dS_j(L) + d\pi_j(L)$ from $dS_0(L) = 0$ and applying Lemma E.1(ii), with $n = L - j$,

$$dS_{L-1}(L) = \sum_{j=0}^{L-2} r^{\,L-2-j}\, d\pi_j(L) = dy_0(1) \sum_{n=2}^{L} r^{\,n-2}\big(1 + \beta\phi H_{n-1}\big) = dy_0(1)\, M_L,$$

and $dy_{L-1}(L) = \phi\, dS_{L-1}(L) + dy_0(1) = dy_0(1)\,(1 + \phi M_L)$. Dividing the first by the second gives the expression for κ_L; for $L = 1$ the sum is empty. □

In the linear–quadratic case, κ_L is the ratio of two objects measured in units of the response to a one-period subsidy. The numerator, M_L, is the work history inherited by the last treated period: the treated period with n periods remaining contributes its direct effect, $1 + \beta\phi H_{n-1}$ (Lemma E.1), which reaches the last treated period after $n-2$ periods of propagation at rate r. The denominator, $1 + \phi M_L$, is the response in the last treated period: the response to a one-period subsidy plus the reinforcement of that work history. Without patience ($\beta = 0$), the direct effect is the same in every treated period, so $M_L = (1 - r^{L-1})/(1-r)$ and, since $1 - r + \phi = \delta$,

$$\kappa_L = \frac{1 - r^{L-1}}{\delta - \phi\, r^{L-1}}, \qquad \textsc{lrm}(L) = \frac{\delta}{\delta - \phi}\,\frac{\delta - \phi}{\delta - \phi\, r^{L-1}} = \frac{\delta}{\delta - \phi\, r^{L-1}}.$$

As the intervention lengthens, κ_L rises from zero to $1/\delta$ and $\textsc{lrm}(L)$ falls to one. $\textsc{lrm}(L)$ is then the reciprocal of the share of the long run reached in the last treated period under a permanent subsidy (Proposition B.1(ii) at $k = L-1$), because without patience the first L periods of the two subsidies coincide.

Part (i) shows that, once the intervention ends, its effects decay at rate r from the work history it leaves behind, $dS_L(L)$, as after a one-period subsidy. Part (ii) shows that the cumulative lifetime effect per unit of that work history does not depend on the length of the intervention. Normalizing by the last treated period adds the work history inherited by that period, which survives into the post-intervention period at rate $1-\delta$; under full depreciation, $\text{CLE}(L) = \text{CLE}$ for every L. Part (iii) holds because, in the last treated period, the response to the change in costs is the same as under a one-period subsidy. What differs is the work history accumulated earlier in the intervention, which changes hours by $\phi\, dS_{L-1}(L)$; removing this component recovers $dy_0(1)$, the denominator of LRM.

E.1.3 Non-Quadratic Preferences

This subsection gives the linear–quadratic model of Section 2 a local interpretation when the payoff is not quadratic. Around a steady state, the optimal policy of the model with any smooth, strongly concave payoff $U(y, S)$ is, to first order, the policy of Lemma 1 with the curvature parameters evaluated at that steady state, whether or not costs move; and at the steady state the targets of Proposition 1 hold exactly, because they are derivatives.

Setup. The individual solves (1) with a payoff $U(y, S)$ that is three times continuously differentiable and strongly concave in (y, S). Costs are uniformly bounded, do not depend on the individual's own stock, and follow a Markov process under the individual's beliefs; choices are interior. The only structure kept from Section 2 beyond smoothness and concavity is that the cost enters the objective linearly, as $-p_t y_t$, so that the Euler equation below is linear in the cost and the response to a cost change depends on preferences only through the curvatures of U; with income effects, the marginal utility of money and its curvature would enter the responses as well. As in Appendix B, $\rho \equiv \beta^{-1} - 1$. The perturbation argument of Appendix B.1.1 gives the Euler equation

$$U_y(y_t, S_t) - p_t + \beta\, \mathcal{E}_t\Big[U_S(y_{t+1}, S_{t+1}) - (1-\delta)\big(U_y(y_{t+1}, S_{t+1}) - p_{t+1}\big)\Big] = 0, \tag{47}$$

and the optimal policy $y_t^\star = g_t(S_t; \mathcal{F}_t)$ of Section 2 is differentiable in the stock under these conditions (Section 2.1). Fix a reference cost $\bar{p}$, for example the permanent level around which the individual's costs move, and let $(\bar{y}, \bar{S})$ be the steady state the individual would reach if the cost stayed at $\bar{p}$: $\bar{y} = \delta \bar{S}$ and $\big(1 - \beta(1-\delta)\big)\big[U_y(\bar{y}, \bar{S}) - \bar{p}\big] + \beta\, U_S(\bar{y}, \bar{S}) = 0$. Costs need not stay at $\bar{p}$; write $\hat{p}_t \equiv p_t - \bar{p}$ for their deviations, and $\hat{y}_t \equiv y_t - \bar{y}$ and $\hat{S}_t \equiv S_t - \bar{S}$ for those of the choice and the stock. Bars on derivatives of U, and on other objects, denote evaluation at the steady state.

Exact targets. Fix a deterministic cost path, and let $\phi_s \equiv \partial g_s / \partial S$ and $r_s \equiv 1 - \delta + \phi_s$ along it. Differentiating (47) at s, with $dS_{s+1} = (1-\delta)\, dS_s + dy_s$ and $dy_{s+1} = \phi_{s+1}\, dS_{s+1}$, gives the

backward recursion

$$\phi_s = -\frac{B_s + (1-\delta)\, C_s\, \phi_{s+1}}{A_s + C_s\, \phi_{s+1}}, \tag{48}$$

where $A_s \equiv U_{yy,s} + \beta\big(U_{SS,s+1} - (1-\delta)U_{yS,s+1}\big)$, $B_s \equiv U_{yS,s} + \beta(1-\delta)\big(U_{SS,s+1} - (1-\delta)U_{yS,s+1}\big)$ and $C_s \equiv \beta\big(U_{yS,s+1} - (1-\delta)U_{yy,s+1}\big)$, with the derivatives of U evaluated along the path. These are the coefficients A, B and C of Appendix B.1.1 with local curvatures, and (48) is the general counterpart of the finite-horizon recursion (51) in Appendix E.1.4. As $\rho \to \infty$, $\beta \to 0$, so $C_s \to 0$ and (48) reduces to $\phi_s = -U_{yS,s}/U_{yy,s}$: without patience, the slope of the policy is the ratio of the local curvatures at the current choice, whatever the form of U, the general counterpart of the limit noted in Section 2. The impulse response is (3), $\theta_t^{(k)} = \phi_{t+k} \prod_{j=1}^{k-1} r_{t+j}$, and $\text{CLE}_t = \sum_{k\geq 1} \theta_t^{(k)}$. For the long-run multiplier, let ζ_t^{TEMP} and ζ_s^{PERM} be the responses of the choice to a unit subsidy, holding the stock fixed, under a one-period subsidy at t, and at each date $s \geq t$ under a permanent subsidy that starts at t. Differentiating (47) in the same way gives

$$\zeta_t^{\text{TEMP}} = -\frac{1}{A_t + C_t\, \phi_{t+1}}, \qquad \zeta_s^{\text{PERM}} = -\frac{1 - \beta(1-\delta) + C_s\, \zeta_{s+1}^{\text{PERM}}}{A_s + C_s\, \phi_{s+1}}. \tag{49}$$

If the path converges to a steady state with $|\bar{r}| < 1$, the stock response to the permanent subsidy converges to $\bar{\zeta}^{\text{PERM}}/(1-\bar{r})$ and the choice response to δ times this limit, so

$$\text{LRM}_t = \frac{\delta}{1-\bar{r}}\, \frac{\bar{\zeta}^{\text{PERM}}}{\zeta_t^{\text{TEMP}}}. \tag{50}$$

Local equivalence. Let $\bar{A}$, $\bar{B}$ and $\bar{C}$ be the coefficients of Appendix B.1.1 at the steady-state curvatures $(\bar{U}_{yy}, \bar{U}_{yS}, \bar{U}_{SS})$. When $\bar{U}_{yS} - (1-\delta)\bar{U}_{yy} \neq 0$, the characteristic polynomial is a nondegenerate quadratic. Let $\bar{\phi}$ be the root of (20) at these curvatures with $|1 - \delta + \bar{\phi}| < \sqrt{1+\rho}$, and $\bar{r} \equiv 1 - \delta + \bar{\phi}$.

Proposition E.2 (Local equivalence)**.** *Suppose the conditions of the setup hold and* $\bar{U}_{yS} - (1-\delta)\bar{U}_{yy} \neq 0$.

(i) (The policy to first order.) *In deviations from the steady state, the Euler equation* (47) *is the Euler equation of the linear–quadratic model with curvature parameters* $(\bar{U}_{yy}, \bar{U}_{yS}, \bar{U}_{SS})$ *and costs* $\hat{p}_t$*, up to a remainder of second order in the current and next-period deviations* $(\hat{y}_t, \hat{S}_t, \hat{y}_{t+1}, \hat{S}_{t+1})$*. To first order, the optimal policy is therefore*

$$y_t^\star = \bar{y} + \bar{\phi}\,(S_t - \bar{S}) + \hat{\pi}_t, \qquad \hat{\pi}_t = \frac{1}{\bar{A} + \bar{C}\bar{\phi}} \sum_{k\geq 0} \Big(\frac{\bar{r}}{1+\rho}\Big)^k \Big(\mathcal{E}_t[\hat{p}_{t+k}] - \frac{1-\delta}{1+\rho}\, \mathcal{E}_t[\hat{p}_{t+k+1}]\Big),$$

the policy of Lemma 1 with $\bar{y} - \bar{\phi}\bar{S}$ *in the role of* $K\alpha$ *and the cost deviations in the role of the costs, whether or not costs move. The approximation is exact along every path if and*

only if U is quadratic.

(ii) (Targets at the steady state.) *Suppose $|\bar{r}| < 1$, that absent the intervention costs stay at $\bar{p}$, and that the individual is at the steady state when the intervention arrives. Then $\phi_s = \bar{\phi}$ at every date along the intervention, and*

$$\theta^{(k)} = \bar{\phi}\,\bar{r}^{\,k-1}, \qquad \text{CLE} = \frac{\bar{\phi}}{1-\bar{r}}, \qquad \text{LRM} = \frac{\delta}{1-\bar{r}}\,\frac{1-\beta(1-\delta)}{1-\beta\bar{r}},$$

the expressions of Proposition 1, hold exactly.

Proof. (i) Expand U_y and U_S to first order around $(\bar{y}, \bar{S})$, with exact remainders R^y and R^S, which are of second order in $(\hat{y}, \hat{S})$ because U is three times continuously differentiable. Substituting into (47) and using $\hat{S}_{t+1} = (1-\delta)\hat{S}_t + \hat{y}_t$, the constant terms cancel by the steady-state condition, and

$$\bar{A}\,\hat{y}_t + \bar{B}\,\hat{S}_t + \bar{C}\,\mathcal{E}_t\hat{y}_{t+1} = \hat{p}_t - \beta(1-\delta)\,\mathcal{E}_t\hat{p}_{t+1} - \varepsilon_t, \qquad \varepsilon_t \equiv R^y_t + \beta\,\mathcal{E}_t\big[R^S_{t+1} - (1-\delta)R^y_{t+1}\big].$$

With $\varepsilon_t = 0$, this is Equation (28), the Euler equation of the linear–quadratic model with curvature parameters $(\bar{U}_{yy}, \bar{U}_{yS}, \bar{U}_{SS})$, written in deviations and with the costs $\hat{p}_t$, which are uniformly bounded. By Lemma 1 and its proof, its solution is the affine policy with slope $\bar{\phi}$ and the cost outlook (21) evaluated at the cost deviations, which is $\hat{\pi}_t$. The remainder vanishes along every path if and only if the third derivatives of U vanish, that is, if and only if U is quadratic.

(ii) Along the steady-state path, the coefficients in (48) are constant, $(A_s, B_s, C_s) = (\bar{A}, \bar{B}, \bar{C})$, and so is the slope of the policy. A constant slope solves (48) only if it solves (20) with the steady-state curvatures, and the transversality condition selects the root with $|1-\delta+\bar{\phi}| < \sqrt{1+\rho}$, as in the proof of Lemma 1. The expressions for $\theta^{(k)}$ and CLE follow. By (49), $\bar{\zeta}^{\text{TEMP}} = -1/(\bar{A} + \bar{C}\bar{\phi})$ and $\bar{\zeta}^{\text{PERM}} = -\big(1-\beta(1-\delta)\big)/\big(\bar{A} + \bar{C}(1+\bar{\phi})\big)$. By Lemma B.1(i), $\bar{A} + \bar{C}\bar{\phi} = -\bar{C}/(\beta\bar{r})$, so $\bar{\zeta}^{\text{PERM}}/\bar{\zeta}^{\text{TEMP}} = \big(1-\beta(1-\delta)\big)/(1-\beta\bar{r}) = (\rho+\delta)/(\rho+\delta-\bar{\phi})$, and (50) gives LRM. □

E.1.4 Solution in the Finite-Horizon Case

This subsection studies the solution of Problem (1) under a finite horizon and in the linear-quadratic specification in Equation (4). We show that its policy converges to the infinite-horizon policy of Lemma 1 as the horizon recedes.

Proposition E.3 (Finite horizon)**.** *Suppose the conditions of Lemma 1 hold, that the dynamic problem* (1) *ends at time T, and that the stock has no salvage value. Then: (i) The optimal policy is $y^\star_t = \phi_t S_t + K_t\alpha + \pi_t$, where the coefficients satisfy the terminal conditions $\phi_T = -U_{yS}/U_{yy}$, $K_T = -1/U_{yy}$ and $\pi_T = p_T/U_{yy}$ and, for $t < T$, with A, B and C as in Appendix B.1.1, the*

recursions

$$\phi_t = -\frac{B + (1-\delta)\, C\, \phi_{t+1}}{A + C\, \phi_{t+1}}, \tag{51}$$

$$K_t = -\frac{1 - \frac{1-\delta}{1+\rho} + C\, K_{t+1}}{A + C\, \phi_{t+1}}, \tag{52}$$

$$\pi_t = \frac{p_t - \frac{1}{1+\rho}\big[\psi + (1-\delta)\, \mathcal{E}_t p_{t+1}\big] - C\, \mathcal{E}_t \pi_{t+1}}{A + C\, \phi_{t+1}}, \tag{53}$$

where ψ is the coefficient on the stock in the payoff (4)*. (ii) As the horizon recedes, the slope converges to that of Lemma 1: with ϕ and $\phi' \equiv (1+\rho)/r - (1-\delta)$ the two roots of* (20)*,*

$$\frac{\phi_t - \phi}{\phi_t - \phi'} = \left(\frac{r^2}{1+\rho}\right)^{T-t} \frac{\phi_T - \phi}{\phi_T - \phi'}, \qquad \frac{r^2}{1+\rho} < 1,$$

and K_t and π_t converge to K and to the forward solution (21) *of Lemma 1.*

Proof. (i) The Bellman equation for the problem is

$$V_t(S) = \max_y U(y, S) - p_t y + \frac{1}{1+\rho} \mathcal{E}_t V_{t+1}((1-\delta)S + y) \tag{54}$$

with boundary condition $V_{T+1}(S) = 0$.

We show by induction that all the value functions are concave quadratics. V_{T+1} is (trivially) a concave quadratic. Suppose V_{t+1} is a concave quadratic. Then the objective in (54) is a strictly concave quadratic in (y, S), and thus remains a concave quadratic in S after optimizing over y. This shows that V_t is a concave quadratic. By induction, all the value functions are concave quadratics, and at every date (54) is a strictly concave quadratic programming problem, whose maximizer is unique and given by its first-order condition.

At T, $V_{T+1} = 0$, so the first-order condition is $\alpha + U_{yy} y_T + U_{yS} S_T - p_T = 0$, which gives the terminal conditions. For $t < T$, the perturbation argument of Appendix B.1.1 applies unchanged: raising y_t by dy and lowering y_{t+1} by $(1-\delta)dy$ changes only S_{t+1}, a date within the horizon, so the Euler equation (16) holds at t. Suppose the policy at $t+1$ is $y^\star_{t+1} = \phi_{t+1} S_{t+1} + K_{t+1}\alpha + \pi_{t+1}$. Substituting it and $S_{t+1} = (1-\delta)S_t + y_t$ into (16) and collecting terms with the definitions of A, B and C gives

$$\begin{aligned}\big(A + C\phi_{t+1}\big)\, y_t + \big(B + (1-\delta)C\phi_{t+1}\big)\, S_t + \Big(1 - \frac{1-\delta}{1+\rho} + CK_{t+1}\Big)\, \alpha \\ = p_t - \frac{1}{1+\rho}\big[\psi + (1-\delta)\, \mathcal{E}_t p_{t+1}\big] - C\, \mathcal{E}_t \pi_{t+1}.\end{aligned}$$

The coefficient on y_t is the second derivative of the objective in (54) with respect to y, $U_{yy} + V''_{t+1}/(1+\rho)$, because the envelope and first-order conditions at $t+1$ give $V'_{t+1}(S_{t+1}) = U_S(y^\star_{t+1}, S_{t+1}) +$

$(1-\delta)\big(p_{t+1} - U_y(y^\star_{t+1}, S_{t+1})\big)$, and hence $V''_{t+1} = U_{SS} - (1-\delta)U_{yS} + \big(U_{yS} - (1-\delta)U_{yy}\big)\phi_{t+1}$. It is therefore strictly negative, and solving for y_t gives the affine policy with the coefficients (51)–(53). By induction from T, the policy is affine at every date with these coefficients, and it is the unique maximizer.

(ii) The fixed points of the recursion (51) are the roots of (20), ϕ and $\phi' \equiv (1+\rho)/r - (1-\delta)$, by Lemma B.1(i), with $1-\delta+\phi = r$ and $1-\delta+\phi' = (1+\rho)/r$. Writing the recursion as $\phi_t = f(\phi_{t+1})$, each fixed point φ satisfies $f(x) - \varphi = -C\,(1-\delta+\varphi)(x-\varphi)/(A+Cx)$, so

$$\frac{\phi_t - \phi}{\phi_t - \phi'} = \frac{r^2}{1+\rho}\,\frac{\phi_{t+1} - \phi}{\phi_{t+1} - \phi'} = \cdots = \left(\frac{r^2}{1+\rho}\right)^{T-t}\frac{\phi_T - \phi}{\phi_T - \phi'}.$$

Since $|r| < \sqrt{1+\rho}$ by Lemma B.1(iii), $r^2/(1+\rho) < 1$. Thus ϕ_t converges to ϕ exponentially as $T \to \infty$. The recursions (52) and (53) are affine in K_{t+1} and $\mathcal{E}_t\pi_{t+1}$ with coefficient $-C/(A+C\phi_{t+1})$, which converges to $-C/(A+C\phi) = r/(1+\rho)$ by Lemma B.1(i), of modulus below one; their fixed points are K and the forward solution (21), so K_t and π_t converge to them as well. □

E.2 Additional Identification Results

This section collects identification results that back Section 3. First we discuss the signs of cost-outlook selection and of the bias of lagged instruments under serially correlated costs (Appendix E.2.1). Then, we study external instruments under a selection on permanent differences assumption (Appendix E.2.2). Next, we present additional identification results from studying the post-treatment path (Appendix E.2.3). Then, we present an identification argument for the decomposition of propagation using lagged choices as instruments (Appendix E.2.4). Finally, we discuss the identification of reinforcement, survival and patience under heterogeneous propagation (Appendix E.2.5).

E.2.1 Serially Correlated Costs: Cost-Outlook Selection and Lagged Instruments

This subsection relaxes the unpredictability of cost innovations to an AR(1) process and signs two biases that Section 3 reports: the cost-outlook selection term of the OLS regression (7), in Subsection 3.2, and the bias of lagged-choice instruments, in Subsection 3.3.1. Both have the sign of $\lambda - (1-\delta)$, the autocorrelation of costs net of the survival of the stock.

Lemma E.2 (Persistent costs: selection and internal instruments)**.** *Suppose types are common, $\omega_i = \omega$, with $|r| < 1$, and let $q \equiv 1-\delta$. Suppose the cost of working is $p_{it} = \bar{p} + \eta_{it}$, with*

$$\eta_{it} = \lambda\eta_{i,t-1} + \nu_{it}, \qquad |\lambda| < 1,$$

where ν_{it} is a bounded innovation with $\mathbb{E}[\nu_{it} \mid \mathcal{F}_{i,t-1}] = 0$ and constant positive variance, so that costs are uniformly bounded as Lemma 1 requires; a permanent difference in costs across individuals

would enter like a difference in the preference shifter α_i. Suppose the individual knows this process, and that hours and costs are in their stationary distribution. Within this family, Assumption I.1 is the case $\lambda = 0$.

Under the conditions of Lemma 1, after removing individual-specific means, $\pi_{it} = \xi_\lambda \eta_{it}$ with $\xi_\lambda \equiv \xi\,(1 - \beta(1-\delta)\lambda)/(1 - \beta r\lambda) \neq 0$, where ξ is the cost pass-through of Lemma B.4. Let $Q \equiv \xi_\lambda^2 \operatorname{Var}(\eta_{it}) > 0$ and $\gamma_k \equiv \operatorname{Cov}(Y_{it}, Y_{i,t-k})$, after removing individual-specific means. Then:

(i) Cost-outlook selection in the OLS regression, the last term of (7), satisfies

$$\operatorname{Cov}(Y_{i,t-1}, e_{it}) = Q\,\frac{(1-q\lambda)(\lambda - q)}{1 - r\lambda},$$

so it is positive if $\lambda > q$, zero if $\lambda = q$, and negative if $\lambda < q$.

(ii) For every $\ell \geq 2$,

$$\operatorname{Cov}(Y_{i,t-\ell}, \Delta e_{it}) = -\,Q\,\frac{(1-q\lambda)(1-\lambda)}{1 - r\lambda}\,(\lambda - q)\,\lambda^{\ell-2}.$$

Hence $Y_{i,t-2}$ is a valid instrument if and only if $\lambda = q$, and $Y_{i,t-\ell}$ with $\ell \geq 3$ is valid if and only if $\lambda \in \{0, q\}$.

(iii) Let r_ℓ denote the estimand r^{int} in (8) with instrument $Y_{i,t-\ell}$. Whenever $\gamma_{\ell-1} < \gamma_{\ell-2}$,

$$\operatorname{sign}(r_\ell - r) = \operatorname{sign}\big[(\lambda - q)\,\lambda^{\ell-2}\big].$$

The condition always holds at $\ell = 2$, where $\gamma_1 - \gamma_0 = -\frac{1}{2}\operatorname{Var}(\Delta Y_{it}) < 0$; there, $r_2 < r$ if and only if $\lambda < q$. For $\ell \geq 3$ and $\lambda \geq 0$, whenever the condition holds, $r_\ell < r$ if and only if $0 < \lambda < q$.

Proof. Remove individual-specific means throughout. Because the individual knows the process, $\mathcal{E}_{it}[\eta_{i,t+k}] = \lambda^k \eta_{it}$.

Substituting these expectations into the forward solution (21) and collecting the terms in η_{it} gives $\pi_{it} = \xi \sum_{k\geq 0} (\beta r\lambda)^k \big[1 - \beta(1-\delta)\lambda\big]\eta_{it} = \xi_\lambda \eta_{it}$, with $\xi = 1/(A + C\phi)$ as in the proof of Lemma B.4; $\xi_\lambda \neq 0$ because $\xi \neq 0$, $|\beta(1-\delta)\lambda| < 1$ and $|\beta r\lambda| < 1$. Then $e_{it} = \xi_\lambda(\eta_{it} - q\eta_{i,t-1})$, an ARMA$(1,1)$ process with autoregressive root λ and moving-average root q. Since $\operatorname{Cov}(\eta_{it}, \eta_{i,t-k}) = \lambda^k \operatorname{Var}(\eta_{it})$, for $k \geq 1$,

$$\operatorname{Cov}(e_{it}, e_{i,t-k}) = Q\left[\lambda^k(1+q^2) - q\lambda^{k-1} - q\lambda^{k+1}\right] = Q\,\lambda^{k-1}(\lambda - q)(1 - q\lambda).$$

In the stationary distribution, $Y_{it} = \sum_{m\geq 0} r^m e_{i,t-m}$, so for $j \geq 1$,

$$\operatorname{Cov}(Y_{i,t-j}, e_{it}) = \sum_{m\geq 0} r^m \operatorname{Cov}(e_{it}, e_{i,t-j-m}) = Q\,\frac{(1-q\lambda)(\lambda - q)}{1 - r\lambda}\,\lambda^{j-1}.$$

Setting $j = 1$ gives part (i). For $\ell \geq 2$, $\mathrm{Cov}(Y_{i,t-\ell}, \Delta e_{it}) = \mathrm{Cov}(Y_{i,t-\ell}, e_{it}) - \mathrm{Cov}(Y_{i,t-\ell}, e_{i,t-1})$, which gives the expression in part (ii). Because $Q > 0$, $|q\lambda| < 1$, $|\lambda| < 1$, and $|r\lambda| < 1$, it is zero exactly when $(\lambda - q)\lambda^{\ell-2} = 0$, which gives the validity statements. For part (iii), differencing (6) gives

$$r_\ell - r = \frac{\mathrm{Cov}(Y_{i,t-\ell}, \Delta e_{it})}{\mathrm{Cov}(Y_{i,t-\ell}, \Delta Y_{i,t-1})}, \qquad \mathrm{Cov}(Y_{i,t-\ell}, \Delta Y_{i,t-1}) = \gamma_{\ell-1} - \gamma_{\ell-2}.$$

When the denominator is negative, the sign of $r_\ell - r$ is the opposite of the sign of the numerator, which is $\mathrm{sign}[(\lambda - q)\lambda^{\ell-2}]$ by part (ii). At $\ell = 2$, $\gamma_1 - \gamma_0 = -\frac{1}{2}\mathrm{Var}(\Delta Y_{it})$. For $\ell \geq 3$ and $\lambda \geq 0$, the sign is negative exactly when $0 < \lambda < q$. □

E.2.2 External Instruments under Selection on Permanent Differences

Section 3.3.2 discusses identification of propagation assuming random assignment. This subsection allows treatment status to depend on permanent differences across workers, as long as it does not depend on transitory deviations from them; the condition implies parallel trends. Take the decomposition $p_{it} = \bar{p}_i + \eta_{it}$ from Assumption I.1, now without requiring the deviations η_{it} to be unpredictable, and collect the permanent components in $\Theta_i \equiv (\alpha_i, \bar{p}_i, \omega_i)$. Group workers by their permanent components, define the group-specific treatment effect $\mathrm{TE}_t(\Theta) \equiv \mathbb{E}[Y_{it} \mid D_i = 1, \Theta_i = \Theta] - \mathbb{E}[Y_{it} \mid D_i = 0, \Theta_i = \Theta]$, and write $r(\Theta)$ for the propagation rate of a group, which depends only on its propagation parameters ω.

Assumption X.4 (Selection on permanent differences)**.** *There exists an individual-specific, $\mathcal{F}_{i0}$-measurable level $\bar{p}_i$ such that, conditional on Θ_i:*

(i) without the subsidy, the process $\{S_{it}, \eta_{it}, \mathcal{E}_{it}[\eta_{i,t+k}]_{k\geq 1}\}_{t=0}^{T}$ is stationary;

(ii) D_i is independent of the process in part (i).

The model gives each part economic content. Without the subsidy, choices evolve as $Y_{it} - Y_{i,t-1} = (r_i - 1)Y_{i,t-1} + a_i + e_{it}$. Part (i) places each worker at the stationary level of hours implied by her permanent components, so her expected choice has no trend. Part (ii) ensures that treated and control workers with the same permanent components do not differ in their transitory deviations. The two groups may still differ in levels: in preference shifters, in cost levels, and even in propagation parameters. Together, the two parts imply that, absent the subsidy, average choices in the two groups would have evolved in parallel.

The assumption rules out 'Ashenfelter's dip' (Ashenfelter, 1978): workers targeted for low recent hours would recover even without the subsidy, as their hours propagate back to their stationary level. It allows, by contrast, a subsidy offered in neighborhoods with long commutes, as long as commuting costs there differ only in level, captured by $\bar{p}_i$. Choices before the subsidy provide a direct check.

Under Assumption X.4, treatment effects are measured by difference-in-differences. Let $t_p < \underline{t}$ be a period before treated workers learn about the subsidy, and define $\mathrm{TE}_s^{\mathrm{DiD}} \equiv \mathrm{TE}_s - \mathrm{TE}_{t_p}$. Within a group, treatment status is as good as random; difference-in-differences aggregates across groups the researcher does not observe. Let $F_{\Theta|1}$ denote the distribution of Θ_i among treated workers.

Proposition E.4 (External instruments under selection on permanent differences)**.** *Suppose Assumptions X.1 and X.4 hold, every group with treated workers also contains control workers, $t-1 > \bar{t}$, and $\mathrm{TE}_{t-1}^{\mathrm{DiD}} \neq 0$. Then*

$$\frac{\mathrm{TE}_t^{\mathrm{DiD}}}{\mathrm{TE}_{t-1}^{\mathrm{DiD}}} = \frac{\int r(\Theta)\mathrm{TE}_{t-1}(\Theta)\, dF_{\Theta|1}(\Theta)}{\int \mathrm{TE}_{t-1}(\Theta)\, dF_{\Theta|1}(\Theta)}.$$

If either $\mathrm{TE}_{t-1}(\Theta) \geq 0$ for all groups with treated workers or $\mathrm{TE}_{t-1}(\Theta) \leq 0$ for all of them, this is a convex, response-weighted average of propagation rates among treated workers. If, in addition, $r_i = r$ for all i, the ratio equals r.

Proof. Fix a group Θ that contains treated workers. By part (ii) of Assumption X.4, treatment status is independent, within the group, of the process that determines choices without the subsidy, so $\mathrm{TE}_t(\Theta)$ is the average effect of the subsidy on treated workers in the group. Treated workers have not learned about the subsidy at t_p, so $\mathrm{TE}_{t_p}(\Theta) = 0$. By part (i), choices without the subsidy have no expected trend within the group, and by part (ii) this holds among control workers: $\mathbb{E}[Y_{is} - Y_{it_p} \mid D_i = 0, \Theta_i = \Theta] = 0$. Hence $\mathbb{E}[Y_{is} - Y_{it_p} \mid D_i = 1, \Theta_i = \Theta] = \mathrm{TE}_s(\Theta)$ and $\mathbb{E}[Y_{is} - Y_{it_p} \mid D_i = 0] = 0$, so that $\mathrm{TE}_s^{\mathrm{DiD}} = \int \mathrm{TE}_s(\Theta)\, dF_{\Theta|1}$. For $t - 1 > \bar{t}$, the argument in the proof of Proposition 3, with part (ii) of Assumption X.4, within the group, in place of random assignment, gives $\mathrm{TE}_t(\Theta) = r(\Theta)\mathrm{TE}_{t-1}(\Theta)$ within each group. Substituting gives the ratio. Under the common-sign condition, the weights $\mathrm{TE}_{t-1}(\Theta)/\int \mathrm{TE}_{t-1}\, dF_{\Theta|1}$ are nonnegative and integrate to one. If $r(\Theta) = r$ for every group, the numerator equals r times the denominator, so the ratio equals r. □

Propositions 3 and E.4 differ only in their weights: the whole population under random assignment, and the treated workers under Assumption X.4.

E.2.3 The Post-Treatment Path

This subsection proves the second implication of Proposition 3 discussed in Section 3.3.2: the ratios of adjacent post-treatment effects rise unless propagation is common.

Corollary E.1 (Post-treatment ratios under heterogeneous propagation)**.** *Suppose Assumptions X.1 and X.2 hold, $r_i \geq 0$ almost surely and either $\mathrm{TE}_{\bar{t}+1}(\omega) \geq 0$ for all types or $\mathrm{TE}_{\bar{t}+1}(\omega) \leq 0$ for all types. Then, whenever the ratios are defined,*

$$\frac{\mathrm{TE}_{\bar{t}+2+k}}{\mathrm{TE}_{\bar{t}+1+k}}$$

is nondecreasing in k, and strictly increasing unless the distribution of r_i, weighted by first post-treatment effects, is degenerate on $r_i > 0$.

Proof. After normalizing the common sign, $\mathrm{TE}_{\bar{t}+1}(\omega) \geq 0$ for every type. Let Λ be the nonnegative measure on $[0, \infty)$ defined by $\Lambda(A) \equiv \mathbb{E}[\mathbf{1}\{r_i \in A\} \mathrm{TE}_{\bar{t}+1}(\omega_i)]$, the distribution of propagation rates weighted by first post-treatment effects. By Equation (12),

$$\mathrm{TE}_{\bar{t}+1+k} = \int r^k \, d\Lambda(r) \equiv m_k.$$

By the Cauchy–Schwarz inequality,

$$m_{k+1}^2 = \left(\int r^{k/2} \, r^{(k+2)/2} \, d\Lambda(r) \right)^2 \leq m_k \, m_{k+2},$$

so $m_{k+1}/m_k \leq m_{k+2}/m_{k+1}$ whenever the ratios are defined. Equality holds if and only if $r^{(k+2)/2}$ is proportional to $r^{k/2}$ Λ-almost everywhere. For $k = 0$ this requires Λ to be a point mass; for $k \geq 1$ it requires the restriction of Λ to $r > 0$ to be a point mass, since $r^{k/2} = 0$ at $r = 0$. The ratios are therefore strictly increasing unless the weighted distribution of r_i is degenerate on $r_i > 0$. □

E.2.4 Decomposition with Lagged Choices

This subsection shows that, when cost innovations are unpredictable and have constant variance, the residuals left after removing propagation identify survival from lagged choices alone, the result Section 3.4 notes in passing. The following additional restriction makes the scale of cost innovations comparable across periods.

Assumption E.1 (Constant innovation variance)**.** *Under Assumption I.1, cost innovations satisfy*

$$\mathbb{E}[\eta_{it}^2] = \sigma_\eta^2 > 0 \qquad \forall t.$$

Proposition E.5 (Decomposing propagation with lagged choices)**.** *Suppose types are common, $\omega_i = \omega$, so that observed choices follow Equation* (5) *with common coefficients; that Assumptions I.1 and E.1 hold; that $\delta \in (0, 1)$; that $T \geq 4$, so that five consecutive observations $Y_{i,t-4}, \ldots, Y_{it}$ exist; and that r is identified, as in Proposition 2(ii). Then*

$$-\frac{\mathrm{Cov}(\Delta e_{it}, \Delta e_{i,t-2})}{\mathrm{Cov}(\Delta e_{it}, \Delta e_{i,t-1})} = \frac{1-\delta}{(2-\delta)^2}, \qquad \Delta e_{it} = \Delta Y_{it} - r \Delta Y_{i,t-1}.$$

The denominator is nonzero because $\xi \neq 0$ (Lemma B.4), and the right-hand side is one-to-one in $\delta \in (0, 1)$. Hence δ is identified, and therefore ϕ is identified.

Proof. We construct a structured moment of the form characterized in Lemma B.5 and show that it is one-to-one in δ, so that the lemma applies.

By (6), the differenced error $\Delta e_{it} = \Delta Y_{it} - r\Delta Y_{i,t-1}$ is observed once r is identified, and by Lemma B.4,

$$\Delta e_{it} = \xi \sum_{j=0}^{2} w_j \, \eta_{i,t-j}, \qquad (w_0, w_1, w_2) = \big(1, \, -(2-\delta), \, 1-\delta\big).$$

Because information accumulates (Section 2.1), η_{is} is known at $t-1$ for $s < t$, so the objective half of Assumption I.1, $\mathbb{E}[\eta_{it} \mid \mathcal{F}_{i,t-1}] = 0$, implies $\mathbb{E}[\eta_{it}\eta_{is}] = 0$ for $s \neq t$; under Assumption E.1, $\mathbb{E}[\eta_{it}^2] = \sigma_\eta^2$ for all t. Writing $Q \equiv \xi^2\sigma_\eta^2 > 0$ and matching overlapping innovation dates,

$$\begin{aligned}
\operatorname{Cov}\big(\Delta e_{it}, \Delta e_{i,t-1}\big) &= Q\,(w_1 w_0 + w_2 w_1) = Q\, w_1(w_0 + w_2) = -\,Q\,(2-\delta)^2, \\
\operatorname{Cov}\big(\Delta e_{it}, \Delta e_{i,t-2}\big) &= Q\, w_2 w_0 = Q\,(1-\delta),
\end{aligned}$$

where the first line uses $w_0 + w_2 = 2 - \delta$. Since $\delta \in (0,1)$, $(2-\delta)^2 \geq 1$, so the first covariance is nonzero because $\xi \neq 0$. Taking the ratio,

$$-\,\frac{\operatorname{Cov}\big(\Delta e_{it}, \Delta e_{i,t-2}\big)}{\operatorname{Cov}\big(\Delta e_{it}, \Delta e_{i,t-1}\big)} = \frac{1-\delta}{(2-\delta)^2}, \tag{55}$$

and the nuisance scale Q cancels. The right-hand side does not depend on β, so the argument of Lemma B.5 for known β applies whatever β is: this ratio identifies δ without knowledge of patience.

The right-hand side is strictly decreasing on $(0,1)$, with derivative $-\delta/(2-\delta)^3 < 0$ and range $(0, 1/4)$, so it is one-to-one in δ. Lemma B.5 then identifies δ, and $\phi = r - (1-\delta)$ with it. □

E.2.5 Decomposition with Heterogeneous Propagation

This subsection asks what a temporary experiment identifies about survival, patience and reinforcement when propagation differs across individuals, the case that Section 3.4 sets aside. Identification of the propagation rates themselves under heterogeneity is Proposition 3 and is not repeated here. When only curvature varies across individuals, every treatment effect is a moment of a response-weighted distribution of propagation rates (Lemma E.3). The post-treatment path identifies that distribution up to a factor that depends on survival and patience. This yields a test of full depreciation (Proposition E.6) and, given patience, lets the last treated period identify survival and, with it, the distribution of reinforcement (Proposition E.7). A parametric version for short post-treatment panels is available upon request.

Notation. By the remark on individual parameters in Appendix B.2.1, each worker follows the policy of Lemma 1 with worker-specific reinforcement ϕ_i, propagation rate $r_i = 1 - \delta_i + \phi_i$ and pass-through ξ_i. As in Section 3.3.2, we group workers by their type ω, with type-specific treat-

ment effects $\mathrm{TE}_t(\omega)$, direct effects $\mathrm{TE}^{\pi}_t(\omega)$ and propagation rate $r(\omega)$. At the last treated date, a type's direct effect $\mathrm{TE}^{\pi}_{\bar{t}}(\omega)$ is its response to that period's subsidy. We weight the distribution of propagation rates by this response,

$$G(A) \equiv \mathbb{E}[\mathbf{1}\{r_i \in A\}\, \mathrm{TE}^{\pi}_{\bar{t}}(\omega_i)]\,.$$

When δ and ρ are common, with $\beta = 1/(1+\rho)$, define for $t \geq \underline{t}$ the polynomials

$$\varpi_t(r) \equiv \sum_{k\geq 0} (\beta r)^k \big[\mathbf{1}\{t+k \in \mathcal{W}\} - \beta(1-\delta)\,\mathbf{1}\{t+k+1 \in \mathcal{W}\}\big],$$
$$\chi_t(r) \equiv \varpi_t(r) + \big(r - (1-\delta)\big)\, S_t(r), \qquad S_t(r) \equiv \sum_{s=\underline{t}}^{t-1} r^{\,t-1-s} \varpi_s(r), \tag{56}$$

and set $\chi_t \equiv 0$ for $t < \underline{t}$. Each sum is finite, because $\varpi_t = 0$ for $t > \bar{t}$, so the coefficients of χ_t depend only on (δ, ρ) and $\mathcal{W}$; moreover $\chi_{\bar{t}+1+k}(r) = r^k \chi_{\bar{t}+1}(r)$ for $k \geq 0$. For the window $\mathcal{W} = \{\underline{t}, \ldots, \bar{t}\}$, $\varpi_t(r) = 1 + \beta\big(r - (1-\delta)\big) H_{\bar{t}-t}(r)$ for $t \in \mathcal{W}$, the factor of the direct effect in Appendix B.3, with $H_m(r) \equiv \sum_{k=0}^{m-1} (\beta r)^k$. In words, for a type with propagation rate r and per unit of its direct effect in the last treated period, $\varpi_t(r)$ is its direct effect at t, $S_t(r)$ the stock gap it has built by t, and $\chi_t(r)$ its treatment effect; $\varpi_{\bar{t}}(r) = 1$, $S_t(r) \geq 0$ for $r > 0$, and $\chi_{\bar{t}+1}(r) = \big(r - (1-\delta)\big)\big[1 + r\, S_{\bar{t}}(r)\big]$, which vanishes exactly for a type without reinforcement.

Lemma E.3 (Treatment effects as moments)**.** *Suppose that*

(i) the curvature parameters $(U_{yy,i}, U_{yS,i}, U_{SS,i})$ vary across workers, while δ and ρ are common;

(ii) the subsidy takes the form (9)*, and workers learn about it at $\underline{t}$;*

(iii) Assumptions X.1, X.2 and X.3 hold.

Then for every t, with χ_t the polynomial in (56)*,*

$$\mathrm{TE}_t = \int \chi_t(r)\, dG(r).$$

Proof. All workers of a type share $r(\omega)$, $\phi(\omega)$ and $\xi(\omega)$. Differencing (21) across arms, as in the proof of Lemma F.1 (Appendix F.1), gives $\mathrm{TE}^{\pi}_t(\omega) = -\xi(\omega)\, b(\omega)\, \varpi_t(r(\omega))$ for $t \geq \underline{t}$, with $b(\omega) \equiv \mathbb{E}[b_i \mid D_i = 1, \omega_i = \omega]$. Since $\varpi_{\bar{t}} = 1$, $\mathrm{TE}^{\pi}_t(\omega) = \varpi_t(r(\omega))\, \mathrm{TE}^{\pi}_{\bar{t}}(\omega)$. The stock gap is zero at $\underline{t}$ and evolves as $\mathrm{TE}^{S}_{t+1}(\omega) = r(\omega)\mathrm{TE}^{S}_t(\omega) + \mathrm{TE}^{\pi}_t(\omega)$, so $\mathrm{TE}_t(\omega) = \phi(\omega)\mathrm{TE}^{S}_t(\omega) + \mathrm{TE}^{\pi}_t(\omega) = \chi_t(r(\omega))\, \mathrm{TE}^{\pi}_{\bar{t}}(\omega)$. Taking expectations over ω_i gives the result. □

With a common propagation rate, G is a point mass, and the lemma reproduces the path of Section 3.4.

Full depreciation. Let $\Lambda(A) \equiv \mathbb{E}[\mathbf{1}\{r_i \in A\}\, \mathrm{TE}_{\bar{t}+1}(\omega_i)]$, as in the proof of Corollary E.1, so that $\mathrm{TE}_{\bar{t}+1+k} = \int r^k \, d\Lambda(r)$ for every $k \geq 0$.

Proposition E.6 (Testing full depreciation under heterogeneous propagation)**.** *Suppose that*

(i) the curvature parameters $(U_{yy,i}, U_{yS,i}, U_{SS,i})$ vary across workers, while δ is common;

(ii) the subsidy takes the form (9);

(iii) Assumptions X.1 and X.2 hold;

(iv) $r(\omega) \in [\underline{r}, \bar{r}] \subset (0,1)$ for every type, and either $\mathrm{TE}^{\pi}_{\bar{t}}(\omega) \geq 0$ for all types or $\mathrm{TE}^{\pi}_{\bar{t}}(\omega) \leq 0$ for all types, with $\mathrm{TE}^{\pi}_{\bar{t}}(\omega) \neq 0$ for a set of types with positive probability;

(v) the whole post-treatment path $\{\mathrm{TE}_t\}_{t>\bar{t}}$ is observed.

Then Λ is identified, and

$$\mathrm{TE}_{\bar{t}} - \int r^{-1} \, d\Lambda(r) = (1-\delta) \int r^{-1} \, dG(r).$$

Hence $\delta = 1$ if and only if $\mathrm{TE}_{\bar{t}} = \int r^{-1} \, d\Lambda(r)$.

Proof. A finite measure on a compact interval is determined by its moments, so the post-treatment path identifies Λ. Equation (15) holds type by type, and $\mathrm{TE}^{\pi}_{\bar{t}+1}(\omega) = 0$ once the subsidy ends, so $\mathrm{TE}_{\bar{t}+1}(\omega) - r(\omega)\mathrm{TE}_{\bar{t}}(\omega) = -(1-\delta)\mathrm{TE}^{\pi}_{\bar{t}}(\omega)$. Dividing by $r(\omega) > 0$ and taking expectations over ω_i gives the identity. Under (iv), $\int r^{-1} \, dG(r) \neq 0$, so the right-hand side is zero if and only if $\delta = 1$. □

The last treated effect must lie on the post-treatment path extrapolated one period back. With a common rate, Λ is a point mass, and the restriction is that of Corollary B.1. The test uses neither ρ nor the date at which workers learn about the subsidy, and the argument is type by type: if δ also varies, the right-hand side becomes $\mathbb{E}[(1-\delta_i)\, r_i^{-1} \mathrm{TE}^{\pi}_{\bar{t}}(\omega_i)]$, which is zero only if $\delta_i = 1$ for every type that responds to the subsidy. With a short post-treatment panel, the extrapolation requires restrictions on Λ.

Survival and reinforcement without a parametric family. The post-treatment path identifies Λ, the distribution of propagation rates weighted by the first post-treatment effect, not G, the distribution weighted by the last direct effect; the two differ by the boundary factor $\chi_{\bar{t}+1}$, which depends on survival and patience. The treated periods pin that factor down.

Proposition E.7 (Decomposition without a parametric family)**.** *Suppose assumptions (i)–(iii) of Lemma E.3 and (iv)–(v) of Proposition E.6 hold, and let ϖ_t, S_t and χ_t be the polynomials in (56), so that $\mathrm{TE}_t = \int \chi_t \, dG$ for every t.*

(i) (The weighted distribution, given survival and patience.) Λ *is identified from the post-treatment path, and* $d\Lambda(r) = \chi_{\bar{t}+1}(r)\, dG(r)$. *Hence, if* G *puts no mass at* $r = 1-\delta$, *then for given* (δ, ρ) *the measure* G *is identified,* $dG = d\Lambda/\chi_{\bar{t}+1}$, *and the treated periods satisfy*

$$\mathrm{TE}_t = \int \frac{\chi_t(r)}{\chi_{\bar{t}+1}(r)}\, d\Lambda(r), \qquad t \in \mathcal{W},$$

$\bar{t} - \underline{t} + 1$ *equations whose right-hand sides are known functions of* (δ, ρ) *once* Λ *is known.*

(ii) (Survival, given patience.) *If* ρ *is known and* $r(\omega) > 1-\delta$ *for every type, the right-hand side of the equation for* $t = \bar{t}$ *is strictly monotone in* $1-\delta$ *(increasing when* $\mathrm{TE}^{\pi}_{\bar{t}}(\omega) \geq 0$ *for all types), so* δ *is identified from* $\mathrm{TE}_{\bar{t}}$ *and the post-treatment path. The distribution of reinforcement weighted by the last direct effect is then* G *shifted by survival,* $G_\phi(A) = G(A + 1 - \delta)$*; if, in addition, the direct response is unrelated to propagation,* $\mathbb{E}[\mathrm{TE}^{\pi}_{\bar{t}}(\omega_i) \mid r_i] = \bar{c} \neq 0$, *it is the population distribution of* ϕ_i *up to the factor* $\bar{c}$.

(iii) (Patience.) *Suppose* ρ *is unknown,* $\bar{t} - \underline{t} \geq 1$ *and* $\underline{r} > 1-\delta$. *Then* (δ, ρ) *is locally identified at any value at which the map* $(\delta, \rho) \mapsto \big(\int \chi_t/\chi_{\bar{t}+1}\, d\Lambda\big)_{t \in \mathcal{W}}$ *in (i) has a Jacobian of rank two.*

Proof. (i) The post-treatment path gives every moment of Λ, which is a finite signed measure on a compact interval, so Λ is identified, as in the proof of Proposition E.6. By the proof of Lemma E.3, $\mathrm{TE}_{\bar{t}+1}(\omega) = \chi_{\bar{t}+1}(r(\omega))\, \mathrm{TE}^{\pi}_{\bar{t}}(\omega)$, so $\Lambda(A) = \int_A \chi_{\bar{t}+1}\, dG$ for every set A. Since $\varpi_{\bar{t}+1} = 0$, $\varpi_{\bar{t}} = 1$ and $S_{\bar{t}+1}(r) = r\, S_{\bar{t}}(r) + 1$, the boundary factor is $\chi_{\bar{t}+1}(r) = (r - (1-\delta))[1 + rS_{\bar{t}}(r)]$, which is zero only at $r = 1-\delta$ because $S_{\bar{t}}(r) \geq 0$ for $r > 0$. Where G has no mass at that point, $dG = d\Lambda/\chi_{\bar{t}+1}$, and substituting into $\mathrm{TE}_t = \int \chi_t\, dG$ gives the display.

(ii) Write $q \equiv 1 - \delta$, $x \equiv r - q > 0$ and $S \equiv S_{\bar{t}}(r)$, which depends on q through the term $\beta(1-\delta)$ in ϖ_s, with $\partial S/\partial q = -\beta \sum_{s=\underline{t}}^{\bar{t}-1} r^{\bar{t}-1-s} \sum_{k \geq 0} (\beta r)^k \mathbf{1}\{s + k + 1 \in \mathcal{W}\} \leq 0$. The integrand at $t = \bar{t}$ is

$$\mathcal{R}(r; q) \equiv \frac{\chi_{\bar{t}}(r)}{\chi_{\bar{t}+1}(r)} = \frac{1 + xS}{x\,(1 + rS)} = \frac{1/x + S}{1 + rS}, \qquad \frac{\partial \mathcal{R}}{\partial q} = \frac{(1 + rS)/x^2 + (\partial S/\partial q)\,(1 - r/x)}{(1 + rS)^2} > 0,$$

because $0 < x < r$ makes $1 - r/x < 0$, so both terms of the numerator are nonnegative and the first is positive. Under condition (iv) of Proposition E.6, G has a constant sign, and so does Λ on $\{r > q\}$ because $\chi_{\bar{t}+1} > 0$ there; hence $\int \mathcal{R}(r; q)\, d\Lambda(r)$ is strictly monotone in q and $\mathrm{TE}_{\bar{t}}$ determines q. Since $\phi(\omega) = r(\omega) - q$, the distribution of ϕ_i weighted by $\mathrm{TE}^{\pi}_{\bar{t}}(\omega_i)$ is G shifted by q; under the additional condition in (ii), iterated expectations give $G = \bar{c}\, F$, with F the population distribution of r_i.

(iii) The map in (iii) is continuously differentiable near the true (δ, ρ): its integrands are ratios of polynomials in r and β, and the denominator $\chi_{\bar{t}+1}(r) = \big(r - (1-\delta)\big)\big[1 + rS_{\bar{t}}(r)\big]$ is at least $\underline{r} - (1-\delta) > 0$ on the support. A continuously differentiable map whose derivative is injective at

a point is one-to-one on a neighborhood of it, so (δ, ρ) is locally identified; a Jacobian of rank two needs at least two equations, that is, at least two treated periods. With a one-period window, by contrast, $\underline{t} = \bar{t}$, so $\varpi_{\bar{t}} = 1$, $S_{\bar{t}} = 0$, $\chi_{\bar{t}} = 1$ and $\chi_{\bar{t}+1+k}(r) = r^k\big(r - (1-\delta)\big)$ for $k \geq 0$; none of these depends on ρ, and neither do the treatment effects, $\text{TE}_t = \int \chi_t \, dG$ with $dG = d\Lambda/\chi_{\bar{t}+1}$, while the right-hand side of $\text{TE}_{\bar{t}} = \int \big(r - (1-\delta)\big)^{-1} d\Lambda(r)$ is strictly monotone in $1 - \delta$ because Λ has a constant sign on $\{r > 1 - \delta\}$. □

Patience enters the treatment effects only through anticipation of subsidized periods still to come. With a one-period subsidy, $\varpi_{\bar{t}} = 1$ and no treatment effect depends on ρ, which is why part (iii) needs a longer window; survival is then identified from $\text{TE}_{\bar{t}} = \int \big(r - (1-\delta)\big)^{-1} d\Lambda(r)$ whatever ρ is, as the last step of the proof shows. Proposition 4(ii) needs a third treated period because it does not assume that workers learn about the subsidy at $\underline{t}$, as assumption (ii) of Lemma E.3 does.

Part (ii) also requires reinforcement to be positive for every type; a type without reinforcement responds in the window but leaves no trace after it, so the post-treatment path cannot weigh it.

F Appendix on Experimental Designs

This appendix states and proves the results behind Section 7. We start with the general setup the three designs share and preliminary results useful for the proofs (Appendix F.1). Then we study the three designs: the lead announcement (Appendix F.2), variable duration (Appendix F.3), and the staggered start (Appendix F.4). Finally, Appendix F.5 simulates each design under a data generating process designed to mimic the data studied in Section 4.

F.1 Setup

Notation. Each design adds a second treated arm to the single arm of Section 3.4: write $D_i = d \in \{1, 2\}$ for the two treated arms and $D_i = 0$ for the control group. Arm d receives the incentive over its window $\mathcal{W}_d = \{\underline{t}_d, \dots, \bar{t}_d\}$, and its participants learn of it $\ell_d \geq 0$ periods before it starts, at the arm's *information date* $\underline{t}_d - \ell_d$: from that date on, their information is $\tilde{\mathcal{F}}_{it} = \mathcal{F}_{it} \cup \{b_i, \mathcal{W}_d\}$, as in Section 2.4, and before it, $\tilde{\mathcal{F}}_{it} = \mathcal{F}_{it}$. Arm 1 is the baseline arm: $\mathcal{W}_1 = \{\underline{t}, \dots, \bar{t}\}$ and $\ell_1 = 0$.

Each arm is compared with the control group: $\text{TE}_t^d \equiv \mathbb{E}[Y_{it} \mid D_i = d] - \mathbb{E}[Y_{it} \mid D_i = 0]$ is its treatment effect at date t, and $\text{TE}_t^{S,d}$ and $\text{TE}_t^{\pi,d}$ are the corresponding differences in mean stocks and mean cost outlooks (the arm's direct effect, as in Section 3.4.2). Gaps are indexed by the date they propagate from, as in Appendix B.3: $G_t^d \equiv \text{TE}_{t+1}^d - r\,\text{TE}_t^d$, and $G^d \equiv G_{\bar{t}_d}^d$ is the arm's boundary gap. Also as in Appendix B.3, $H_m \equiv \sum_{k=0}^{m-1} (\beta r)^k$ and $\xi \equiv 1/(A + C\phi)$.

Assumption F.1 (Experiments with two treated arms)**.**

(i) The conditions of Lemma 1 are satisfied and the parameters are common across individuals ($\omega_i = \omega$).

(ii) Each treated arm d faces the cost path (9) over $\mathcal{W}_d$, and the mean cost reduction $b \equiv \mathbb{E}[b_i \mid D_i = d]$ is the same in both arms and nonzero.

(iii) Assumptions X.1 and X.2 hold for each treated arm, and Assumption X.3 holds from the arm's information date on: $\tilde{\mathcal{E}}_{it}[\tilde{p}_{i,t+k}] = \mathcal{E}_{it}[p_{i,t+k}] - b_i \mathbf{1}[t + k \in \mathcal{W}_d]$ for every $t \geq \underline{t}_d - \ell_d$ and $k \geq 1$.

(iv) For each arm, the panel runs from no later than $\underline{t}_d - \ell_d$ to at least $\bar{t}_d + 2$, so that r is identified from either arm's post-treatment decay when $\text{TE}_{\bar{t}_d+1}^d \neq 0$ (Proposition 3).

Preliminaries. The proofs below use two facts about the parameters and three properties of each arm's treatment effects. The facts are that $r \neq 0$, because r is a root of P (Lemma B.1(i)) and $P(0) = a_0 \neq 0$ by condition (ii) of Lemma 1, and hence, by the same lemma, $\xi = -r/a_0 \neq 0$. The following lemma collects the three properties: how the effects split into a stock gap and a direct effect, the direct effect itself, and how the effects behave after the incentive ends.

Lemma F.1 (Treatment effects of each arm)**.** *Under Assumption F.1, each arm d satisfies the following.*

(i) (Decomposition.) *For every t, $\mathrm{TE}_t^d = \phi\,\mathrm{TE}_t^{S,d} + \mathrm{TE}_t^{\pi,d}$ and $\mathrm{TE}_{t+1}^{S,d} = r\,\mathrm{TE}_t^{S,d} + \mathrm{TE}_t^{\pi,d}$, and $\mathrm{TE}_t^{S,d} = 0$ for every $t \le \underline{t}_d - \ell_d$.*

(ii) (Direct effect.) *For every $t \ge \underline{t}_d - \ell_d$,*

$$\mathrm{TE}_t^{\pi,d} = -\xi b \sum_{k\ge 0} (\beta r)^k \big(\mathbf{1}[t+k \in \mathcal{W}_d] - \beta(1-\delta)\,\mathbf{1}[t+k+1 \in \mathcal{W}_d]\big). \tag{57}$$

It therefore satisfies the recursion

$$\mathrm{TE}_t^{\pi,d} = -\xi b\big(\mathbf{1}[t \in \mathcal{W}_d] - \beta(1-\delta)\,\mathbf{1}[t+1 \in \mathcal{W}_d]\big) + \beta r\,\mathrm{TE}_{t+1}^{\pi,d}. \tag{58}$$

For every $t \in \mathcal{W}_d$, $\mathrm{TE}_t^{\pi,d} = -\xi b\,[1 + \beta\phi\,H_{\bar{t}_d - t}]$ and

$$r\,\mathrm{TE}_t^{\pi,d} + (1-\delta)\,\xi b = -\xi b\,\phi\,H_{\bar{t}_d - t + 1}. \tag{59}$$

(iii) (After the incentive.) *The boundary gap is*

$$G^d = (1-\delta)\,\xi b, \tag{60}$$

and the effects fade out at the propagation rate: $\mathrm{TE}_{t+1}^d = r\,\mathrm{TE}_t^d$ for every $t \ge \bar{t}_d + 1$.

Proof. (i) By Equation (5) and the law of motion, $Y_{it} = \phi S_{it} + K\alpha_i + \pi_{it}$ and $S_{i,t+1} = rS_{it} + K\alpha_i + \pi_{it}$. Taking means in arm d and in the control group and differencing gives the two equations, since the terms in α_i cancel by random assignment. Before its information date, an arm's costs and information do not differ from the control group's (Assumption F.1(ii)–(iii)), so, by random assignment, its choices before that date, and hence its stocks up to that date, have the same mean as the control group's.

(ii) For $t \ge \underline{t}_d - \ell_d$, the arm difference in the mean date-t forecast of the cost k periods ahead (for $k = 0$, the current cost) is $-b\,\mathbf{1}[t+k \in \mathcal{W}_d]$ by Assumption F.1(ii)–(iii), and random assignment balances everything else. Differencing (21) across arms gives (57), and splitting off its $k = 0$ term gives (58). The closed form follows as for the direct effect in Appendix B.3, with $\bar{t}_d$ in place of $\bar{t}$. With $m \equiv \bar{t}_d - t$, the closed form gives $r\,\mathrm{TE}_t^{\pi,d} + (1-\delta)\,\xi b = -\xi b\,[r(1+\beta\phi H_m) - (1-\delta)] = -\xi b\,\phi\,H_{m+1}$, since $r - (1-\delta) = \phi$ and $H_{m+1} = 1 + \beta r H_m$.

(iii) By (ii), $\mathrm{TE}_{\bar{t}_d}^{\pi,d} = -\xi b$, and $\mathrm{TE}_t^{\pi,d} = 0$ for every $t > \bar{t}_d$, since no cost change remains. By Equation (15), $G_t^d = \mathrm{TE}_{t+1}^{\pi,d} - (1-\delta)\,\mathrm{TE}_t^{\pi,d}$, which equals $(1-\delta)\,\xi b$ at $t = \bar{t}_d$ and zero for every $t \ge \bar{t}_d + 1$. □

Since the two arms share δ, ξ, and b, part (iii) gives restriction (3) of Table 4 in every design: each arm's effects fade out at the propagation rate after its own end, and the two arms' boundary gaps are equal.

F.2 Lead Announcement

In this design, arm 2 receives the same incentive as arm 1 over the same window, $\mathcal{W}_2 = \mathcal{W}_1 = \{\underline{t}, \dots, \bar{t}\}$, but its participants learn of it $\ell_2 = \ell \geq 1$ periods before it starts. In the announcement periods $\underline{t} - \ell, \dots, \underline{t} - 1$, arm 2 faces the same costs as the control group and differs from it only in what its participants know, so its treatment effects in those periods are pure anticipation. Proposition F.1 shows that these anticipatory responses identify patience when reinforcement is nonzero, and that the design adds restriction (4) of Table 4.

Proposition F.1 (Lead announcement)**.** *Under Assumption F.1, let arm 2 receive the incentive over $\mathcal{W}_2 = \mathcal{W}_1$ with lead $\ell_2 = \ell \geq 1$.*

(i) (The first announcement period.) *If $\phi \neq 0$, patience is identified from arm 2's first anticipatory response and the effects of arm 1:*

$$\beta^{\ell} = \frac{\mathrm{TE}^{2}_{\underline{t}-\ell}}{r^{\ell-1}\big(r\,\mathrm{TE}^{1}_{\underline{t}} + G^{1}\big)}.$$

If $\phi = 0$, the first anticipatory response $\mathrm{TE}^{2}_{\underline{t}-\ell}$ is zero whatever the value of β.

(ii) (Later announcement periods.) *For every t with $\underline{t} - \ell \leq t \leq \underline{t} - 2$,*

$$G^{2}_{t} = \beta\big[1 - (1-\delta)\beta r\big](\beta r)^{\underline{t}-2-t}\big(r\,\mathrm{TE}^{1}_{\underline{t}} + G^{1}\big).$$

In particular, consecutive gaps satisfy $G^{2}_{t} = \beta r\, G^{2}_{t+1}$.

(iii) (Restriction.) *For every $t \geq \underline{t}$, $\mathrm{TE}^{2}_{t+1} - \mathrm{TE}^{1}_{t+1} = r\left(\mathrm{TE}^{2}_{t} - \mathrm{TE}^{1}_{t}\right)$.*

Proof. The proof applies the recursion (58) for the direct effect in three steps, and then reads off the three parts.

Step 1: periods from $\underline{t}$ on. From $\underline{t}$ on, both arms face the same current and future cost changes and know them, so by (57) their direct effects coincide: $\mathrm{TE}^{\pi,2}_{t} = \mathrm{TE}^{\pi,1}_{t}$ for every $t \geq \underline{t}$. The arms then differ only in their stocks. Differencing the two arms' equations in Lemma F.1(i) shows that, from $\underline{t}$ on, the difference between their stock gaps decays at rate r, and the difference between their effects is ϕ times it.

Step 2: the last announcement period. At $\underline{t} - 1$, arm 2's cost has not yet changed but falls by b next period. By (58) and Step 1,

$$\mathrm{TE}^{\pi,2}_{\underline{t}-1} = \beta(1-\delta)\,\xi b + \beta r\,\mathrm{TE}^{\pi,1}_{\underline{t}} = \beta\big(r\,\mathrm{TE}^{1}_{\underline{t}} + G^{1}\big),$$

where the last equality uses (60) and, since arm 1's stock gap is zero at $\underline{t}$ (Lemma F.1(i)), $\mathrm{TE}^{\pi,1}_{\underline{t}} = \mathrm{TE}^{1}_{\underline{t}}$.

Step 3: earlier announcement periods. For $\underline{t} - \ell \le t \le \underline{t} - 2$, arm 2's cost changes neither at t nor at $t+1$, so (58) reduces to

$$\mathrm{TE}^{\pi,2}_{t} = \beta r \, \mathrm{TE}^{\pi,2}_{t+1} :$$

each announcement period's direct effect is βr times the next one's.

Part (i). At its information date $\underline{t} - \ell$, arm 2's stock gap is zero (Lemma F.1(i)), so its treatment effect equals its direct effect, and Steps 2 and 3 give

$$\mathrm{TE}^{2}_{\underline{t}-\ell} = (\beta r)^{\ell-1} \, \mathrm{TE}^{\pi,2}_{\underline{t}-1} = \beta^{\ell} r^{\ell-1} \big(r \, \mathrm{TE}^{1}_{\underline{t}} + G^{1} \big).$$

By (59) for arm 1 at $t = \underline{t}$, $r \, \mathrm{TE}^{1}_{\underline{t}} + G^{1} = -\xi b \, \phi \, H_{\bar{t}-\underline{t}+1}$, which is nonzero if and only if $\phi \neq 0$, because ξ and b are nonzero and $H_m > 0$ for $m \ge 1$ (Appendices F.1 and B.3; Assumption F.1(ii)). Since also $r \neq 0$, dividing gives β^{ℓ} when $\phi \neq 0$; when $\phi = 0$, both sides of the display are zero.

Part (ii). By Equation (15), arm 2's gaps depend only on its direct effects: $G^{2}_{t} = \mathrm{TE}^{\pi,2}_{t+1} - (1-\delta) \, \mathrm{TE}^{\pi,2}_{t}$. For $\underline{t} - \ell \le t \le \underline{t} - 2$, Step 3 gives $\mathrm{TE}^{\pi,2}_{t} = \beta r \, \mathrm{TE}^{\pi,2}_{t+1}$, so $G^{2}_{t} = \big[1 - (1-\delta)\beta r \big] \, \mathrm{TE}^{\pi,2}_{t+1}$, and Steps 2 and 3 give $\mathrm{TE}^{\pi,2}_{t+1} = (\beta r)^{\underline{t}-2-t} \, \beta \big(r \, \mathrm{TE}^{1}_{\underline{t}} + G^{1} \big)$. Combining the two gives the expression for G^{2}_{t}, and comparing it at t and $t+1$ gives $G^{2}_{t} = \beta r \, G^{2}_{t+1}$. The range stops at $\underline{t} - 2$ because $G^{2}_{\underline{t}-1}$ involves $\mathrm{TE}^{\pi,2}_{\underline{t}}$, to which Step 3 does not apply. These gaps are nonzero when $\beta > 0$ and $\phi \neq 0$: the factor $1 - (1-\delta)\beta r$ is positive because $|\beta r| < 1$ (Appendix B.3) and $0 \le 1 - \delta < 1$, $r \neq 0$ (Appendix F.1), and $r \, \mathrm{TE}^{1}_{\underline{t}} + G^{1} \neq 0$ by part (i).

Part (iii). By Step 1, $\mathrm{TE}^{2}_{t+1} - \mathrm{TE}^{1}_{t+1} = \phi \, r \left(\mathrm{TE}^{S,2}_{t} - \mathrm{TE}^{S,1}_{t} \right) = r \left(\mathrm{TE}^{2}_{t} - \mathrm{TE}^{1}_{t} \right)$ for every $t \ge \underline{t}$. □

Part (i) shows that arm 2's first anticipatory response identifies patience whenever reinforcement is nonzero. Without reinforcement, a stock built in anticipation does not raise later choices, so there is nothing to gain from building it early, and the response is zero however patient participants are. Part (ii) shows that each later announcement period adds a moment in which patience enters, so a lead of ℓ periods gives ℓ such moments. With $\ell \ge 3$, consecutive gaps identify patience from r alone. A longer lead, however, shrinks the first anticipatory response by a factor $|\beta r|$ for each extra period. Part (iii) shows that, for any lead, the initial difference between the arms fades out at the propagation rate, even while both are treated: this is restriction (4) of Table 4. With a one-period lead ($\ell = 1$), part (i) reads $\beta = \mathrm{TE}^{2}_{\underline{t}-1} / \big(r \, \mathrm{TE}^{1}_{\underline{t}} + G^{1} \big)$. Arm 2's stock gap at $\underline{t}$ is then its anticipatory response at $\underline{t} - 1$ (Lemma F.1(i)), so the initial difference in part (iii) is $\mathrm{TE}^{2}_{\underline{t}} - \mathrm{TE}^{1}_{\underline{t}} = \phi \, \mathrm{TE}^{2}_{\underline{t}-1}$: reinforcement times the anticipatory response.

F.3 Variable Duration

In this design, both arms learn of the incentive when it starts, at $\underline{t}$, but arm 2 receives it for longer: $\mathcal{W}_1 = \{\underline{t}, \ldots, \bar{t}\}$ and $\mathcal{W}_2 = \{\underline{t}, \ldots, \bar{t}_2\}$, with $\underline{t} + 2 \le \bar{t} < \bar{t}_2$. While both arms are treated, they

face the same costs and differ only in what they expect after $\bar{t}$, so the difference between them reflects anticipation of arm 2's extra periods. Proposition F.2 shows that this difference identifies patience when reinforcement is nonzero.

Proposition F.2 (Identifying patience with variable duration)**.** *Under Assumption F.1, let both arms learn of the incentive at $\underline{t}$ and receive it from $\underline{t}$ on, arm 1 through $\bar{t}$ and arm 2 through $\bar{t}_2$, with $\underline{t}+2 \le \bar{t} < \bar{t}_2$. If $\phi \neq 0$ and $\beta > 0$, then for every t with $\underline{t} \le t \le \bar{t}-2$,*

$$\frac{G^2_{t+1} - G^1_{t+1}}{G^2_t - G^1_t} = \frac{1}{\beta r}.$$

Proof. The argument is that of Steps 2–3 and part (ii) in the proof of Proposition F.1, applied to the difference between the arms' direct effects. For $\underline{t} \le t \le \bar{t}-1$, both arms are paid at t and at $t+1$, so by (58) this difference is βr times next period's. By Equation (15), $G^2_t - G^1_t$ is the difference at $t+1$ minus $(1-\delta)$ times the difference at t, that is, $1-(1-\delta)\beta r$ times the difference at $t+1$; hence consecutive gap differences are in ratio βr for $\underline{t} \le t \le \bar{t}-2$. They are nonzero when $\phi \neq 0$ and $\beta > 0$: the factor $1-(1-\delta)\beta r$ is positive because $|\beta r| < 1$ (Appendix B.3), $r \neq 0$ (Appendix F.1), and at $\bar{t}$, where arm 1 is paid for the last time and arm 2 continues, (58) for each arm and (59) for arm 2 at $\bar{t}+1$ give $\mathrm{TE}^{\pi,2}_{\bar{t}} - \mathrm{TE}^{\pi,1}_{\bar{t}} = \beta\big[(1-\delta)\,\xi b + r\,\mathrm{TE}^{\pi,2}_{\bar{t}+1}\big] = -\beta\xi b\,\phi\,H_{\bar{t}_2-\bar{t}}$. The ratio stops at $\bar{t}-2$ because $G^2_{\bar{t}} - G^1_{\bar{t}}$ involves $\bar{t}+1$, when only arm 2 is treated. □

The proposition shows that, while both arms are treated, consecutive gap differences identify patience from r alone; without patience or without reinforcement, the arms do not differ until arm 1's incentive ends. Besides restriction (3), which every design satisfies (Appendix F.1), restriction (5) uses the rest of the model: by Lemma F.1, each arm's treatment effects are determined by ϕ, δ, β, its window, and the scale ξb, which the arms share and which equals $G^1/(1-\delta)$ by (60), so once arm 1 identifies ϕ, δ, and β (Proposition 4), it predicts arm 2's whole path, extrapolating the model to a longer incentive.

F.4 Staggered Start

In this design, arm 2 receives the same incentive as arm 1 for the same number of periods, but starts $s \ge 1$ periods later: $\mathcal{W}_2 = \{\underline{t}+s, \ldots, \bar{t}+s\}$, and its participants learn of it when it starts. Each arm is compared with the control group at its own dates. Proposition F.3 shows that the later arm's effects equal the baseline arm's, shifted by s periods: this is restriction (6) of Table 4, and the design adds no information about patience. The restriction can fail when responses depend on calendar time, which a single arm cannot separate from the time since the incentive started.

Proposition F.3 (Testing time invariance with a staggered start)**.** *Under Assumption F.1, let arm 2 learn of the incentive at $\underline{t}+s$, for some $s \ge 1$, and receive it from $\underline{t}+s$ through $\bar{t}+s$. Then $\mathrm{TE}^2_{t+s} = \mathrm{TE}^1_t$ for every t.*

Proof. Arm 2's window and information date are arm 1's shifted by s periods, so $\mathbf{1}[t+s \in \mathcal{W}_2] = \mathbf{1}[t \in \mathcal{W}_1]$ for every t. By (57), $\text{TE}^{\pi,2}_{t+s} = \text{TE}^{\pi,1}_t$ from the information dates on; before them, both direct effects are zero, because neither an arm's costs nor its participants' information differs from the control group's (Assumption F.1(ii)–(iii)). By Lemma F.1(i), both stock gaps are zero up to the information dates and follow the same recursion, so $\text{TE}^{S,2}_{t+s} = \text{TE}^{S,1}_t$, and hence $\text{TE}^{2}_{t+s} = \text{TE}^{1}_t$ for every t. The argument uses only that ϕ, r, ξ, and β do not depend on calendar time and that, by Lemma 1 and (57), treatment effects do not depend on the path the control group follows. □

F.5 Simulation Evidence

This subsection evaluates the experimental designs of Section 7 using Monte Carlo simulations. Each design adds the moments of a second arm to those of the single arm. Added moments cannot lower the precision of the estimates, but how much they raise it is an empirical question. We simulate experiments with features similar to those of the NIT experiments in Section 4 (timing, sample size and the persistence of hours) and set reinforcement and survival to the NIT estimates at a patience of $\beta = 0.95$. On these simulated data, we implement each experimental design of Table 4, holding fixed the number of participants and the total incentive payments. In each of $M = 1{,}000$ simulated samples, we run every design on the same participants, as a researcher who has recruited them and is choosing which experiment to run would, and compare each design's estimates with the true parameters. To see what the overidentifying restrictions detect, we also simulate data in which the model fails in a specific way and record how often each design's tests reject. The simulations show that the gains in precision are meaningful, mostly for patience, in a setting that resembles our data.

F.5.1 Simulation design

Data-generating process. Table F.1 collects the numerical choices of the simulation. Each simulated sample is drawn once and then used for every experimental design. To do so, we generate each participant's counterfactual choices under every treatment condition the designs use; each design then reveals the choices under the conditions it assigns. A treatment condition is the date at which a participant learns of the incentive together with the quarters in which they receive it; the control condition offers no incentive.

Control group. Under the control condition, participant i chooses $Y_{it}(0) = \phi S_{it} + K\alpha_i + \varepsilon_{it}$, with $S_{i,t+1} = (1-\delta)S_{it} + Y_{it}(0)$ (Lemma 1). Stable preferences $K\alpha_i$ are normal across participants, and ε_{it}, the effect on hours of an unanticipated shock to the cost of working, is normal with mean zero and independent over time. Stocks start sixty quarters before the panel at each participant's steady state, $K\alpha_i/(\delta-\phi)$, so that at the start of the panel they vary as in a stationary population. Patience is $\beta = 0.95$, and reinforcement and survival are the NIT estimates at that value (Appendix C.1.4). We set the mean and standard deviation of $K\alpha_i$ and the standard deviation of ε_{it} so that untreated

hours have the mean and quarter-to-quarter correlation of the NIT control group and an eighth of its standard deviation. With the NIT's own noise, the estimator described below puts patience at the no-patience lower bound in 24 to 46% of samples, depending on the design, as it is in the NIT (Section 4).

Treatment arms. Under a treated condition d, the participant faces the cost path of Equation (9) over the condition's quarters $\mathcal{W}_d$, with a common cost reduction b, and knows the incentive's size and timing from the condition's information date on. The participant's counterfactual choices are then $Y_{it}(d) = Y_{it}(0) + \text{TE}_t^d$, with TE_t^d the treatment effect of Appendix F.1, built from the direct effect of Lemma F.1. We set the scale ξb so that an incentive paid for a single quarter, and known from its start, changes hours in that quarter by as much as the same estimates imply.

Table F.1: Parameters of the simulation

Parameter	Value	Chosen to match
Model		
Reinforcement, ϕ	0.205	NIT estimates with $\beta = 0.95$
Survival, $1 - \delta$	0.673	
Patience, β	0.95	Patient case of Section 4
Scale of the incentive, ξb	7.94	Response to a one-quarter incentive, same estimates
Untreated hours		
Mean of $K\alpha_i$	102.4	Mean of hours, 274
Standard deviation of $K\alpha_i$	11.0	Quarter-to-quarter correlation of hours, 0.884
Standard deviation of ε_{it}	12.1	Standard deviation of hours divided by 8, 32.25
Quarters before the panel	60	Stationary stocks at the start of the panel
Simulation		
Participants, N	2,728	NIT estimation sample
Samples, M	1,000	

Notes: Patience is set to $\beta = 0.95$, and reinforcement and survival are the NIT estimates at that value (Appendix C.1.4). The targets for hours are those of the NIT control group (Section 4); the two standard deviations are set jointly to match the correlation of hours and an eighth of their standard deviation. Dividing the standard deviation by eight is equivalent to a sample 64 times as large.

Experimental designs. Table F.2 lists what varies across designs: the second arm's information date and treated quarters, the group sizes, and the panel. Each design randomly assigns $N = 2{,}728$ participants to the control condition and to the treated conditions of its arms. The researcher then observes each participant's choices under their assigned condition, $Y_{it}(D_i)$, over the design's panel. The group sizes keep the designs budget neutral: every design has the same participants and pays for the same number of person-quarters, $1{,}718 \times 12$.

Estimation. For each experimental design, we extend the estimator of Appendix D with the moments of the second arm, counted from its own information date; these moments carry the information of Propositions F.1, F.1, F.2 and F.3. The moments of both arms share the parameters (r, β, δ, b) and are weighted efficiently, so that added moments cannot lower precision. Their covariance, estimated at a first-step fit, accounts for the shared control group. We restrict β

Table F.2: Simulated experimental designs

	Second arm		Participants			
Design	Informed	Treated	Control	Arm 1	Arm 2	Panel
Single arm	—	—	1,010	1,718	—	$\underline{t}-4$ to $\underline{t}+15$
Lead announcement, $\ell = 1, 2, 4$	$\underline{t}-\ell$	$\underline{t}$ to $\underline{t}+11$	1,010	859	859	$\underline{t}-4$ to $\underline{t}+15$
Variable duration	$\underline{t}$	$\underline{t}$ to $\underline{t}+19$	1,440	644	644	$\underline{t}-4$ to $\underline{t}+23$
Staggered start	$\underline{t}+4$	$\underline{t}+4$ to $\underline{t}+15$	1,010	859	859	$\underline{t}-4$ to $\underline{t}+19$

Notes: Arm 1 is the same in every design: informed at $\underline{t}$ and treated from $\underline{t}$ to $\underline{t}+11$, as in the three-year NIT. Quarters are counted from the start of the incentive, $\underline{t}$. Every design has 2,728 participants and pays for $1{,}718 \times 12$ treated person-quarters. In each sample, participants are placed in one random order, and every design fills its control group and arms in that order.

to $[0, 1)$, minimize from a grid of starting values, and compute standard errors with the GMM sandwich formula.

Departures from the model. The overidentifying restrictions of Table 4 hold when the model is correct. To see what they detect, we also generate data from three departures, each built to violate the restrictions of one experimental design. A departure changes the counterfactual choices under some treated conditions and leaves the samples, designs and assignments unchanged. We simulate each departure at four strengths a. Under an *announcement effect*, being told in advance makes participants respond more strongly to the incentive, as if it were a fraction a larger: choices under an announced condition are $Y_{it}(0) + (1+a)\,\mathrm{TE}_t^d$. The response in the announcement quarter then overstates patience: with $a = 10\%$, the formula of Proposition F.1 gives 1.045 instead of 0.95. Under a *fading response*, the direct effect $\mathrm{TE}_t^{\pi,d}$ falls by a fraction a from the thirteenth treated quarter on before accumulating in the stock; only a condition treated for more than 12 quarters is affected. A permanent incentive then has a smaller long-run response than a 12-quarter experiment predicts. Under *calendar drift*, every condition's direct effect is multiplied by $1 + a\,(t-\underline{t})/4$ before accumulating in the stock, so the response grows by a fraction a per year; a condition that starts later has a different path in event time. An incentive introduced at another date then has a different response from the one the experiment measures. For the single arm and the three two-arm designs, with a one-quarter lead for the lead announcement, we compute the overidentification statistic of the stacked fit and, for the two-arm designs, a test of each restriction the second arm adds.

F.5.2 Results

Precision. Table F.3 reports, for each design, the median, standard deviation and root mean squared error (RMSE) of the estimates of patience across the M samples, the share of samples whose 95% confidence interval contains the true value, and the RMSE of the other parameters relative to the single arm. The medians, between 0.93 and 0.95, show that every design centers its estimates close to the true value, so the RMSE mostly reflects their spread. In the RMSE column,

every two-arm design estimates patience more precisely than the single arm, whose RMSE is 0.052. The gain grows with the lead, as Proposition F.1 suggests: the RMSE of patience falls by 9% with a one-quarter lead and by 61% with a four-quarter lead, as it does with a second arm treated for 20 quarters. An arm that starts four quarters later brings no new source of information about patience (Proposition F.3), but its added moments still lower the RMSE by 26%. The coverage column shows that the confidence intervals contain the true value in 89 to 95% of samples, except under variable duration (78%), whose standard errors understate the spread of its estimates. In the last three columns, the two-arm designs also sharpen the estimates of reinforcement, survival and propagation, whose RMSEs are 0.60 to 0.87 times those of the single arm. At the same budget, every two-arm design thus estimates every parameter more precisely than the single arm, and an announcement four quarters ahead or a second arm treated for 20 quarters cuts the error in patience by more than half.

Figure F.1 shows the estimates of patience themselves; each box holds the middle half of a design's estimates. With the single arm, the box spans 0.90 to 0.97, and 9% of the estimates sit at the upper bound of 1. With a four-quarter lead or a second arm treated for 20 quarters, it spans 0.93 to 0.96, and almost no estimate reaches the bound. The other two-arm designs have boxes of intermediate width, ordered as their RMSEs. An early announcement or a longer second arm thus tells a patience of 0.95 apart from one close to 1, which the single arm fails to do in about one sample in ten.

What the restrictions detect. Table F.4 reports how often each test rejects at the 5% level. Each panel pairs a departure with the experimental design built to detect it, and its rows are that design's tests, numbered as in Table 4. The column headed "The model" uses data from the model, and the columns under "The departure" use data from the departure at four strengths a, with critical values that make each test reject in 5% of samples when the model holds. In the first column, the tests reject in 5 to 13% of samples at the nominal 5% level. In the first panel, restriction (4) and the test of all restrictions reject an announcement effect of 20% in 98 and 99% of samples: while both arms are treated, their difference no longer fades, because it now carries a fifth of the announced arm's effect. Restriction (3) rejects the same departure in only 30% of samples, because the two boundary gaps differ by about an hour. In the second panel, the test of all restrictions rejects a fading response of 25%, a departure that leaves a 12-quarter arm's path unchanged, in over 99% of samples; the boundary-gap test of restriction (3) alone rejects it in 29%. In the third panel, restriction (6) rejects a calendar drift of 20% per year in 96% of samples, whereas the single arm's test of all restrictions rejects it in 6%, because with one arm, event time and calendar time are collinear. At the precision of these simulations, each design's second arm thus detects the departure it targets once that departure reaches a strength of 20 to 25%, including calendar drift, which the single arm rarely detects.

Table F.3: Simulated precision of the experimental designs

	Patience (true value 0.95)				RMSE relative to the single arm		
Design	Median	SD	RMSE	Coverage	Reinforcement	Survival	Propagation
Single arm	0.933	0.048	0.052	0.95	1.00	1.00	1.00
Lead announcement, 1 quarter	0.932	0.042	0.048	0.91	0.72	0.75	0.87
Lead announcement, 2 quarters	0.939	0.032	0.035	0.93	0.65	0.68	0.82
Lead announcement, 4 quarters	0.946	0.020	0.020	0.93	0.60	0.62	0.76
Variable duration, 20 quarters	0.945	0.019	0.020	0.78	0.66	0.63	0.64
Staggered start, 4 quarters later	0.939	0.036	0.038	0.89	0.68	0.70	0.82

Notes: Each row is a design of Table F.2, with the lead announcement at each of its three leads. The middle columns describe the estimates of patience across the M simulated samples: their median, standard deviation and RMSE, and the share of samples whose 95% confidence interval contains the true value. The last three columns divide the RMSE of each parameter by that of the single arm.

Figure F.1: Simulated estimates of patience by experimental design

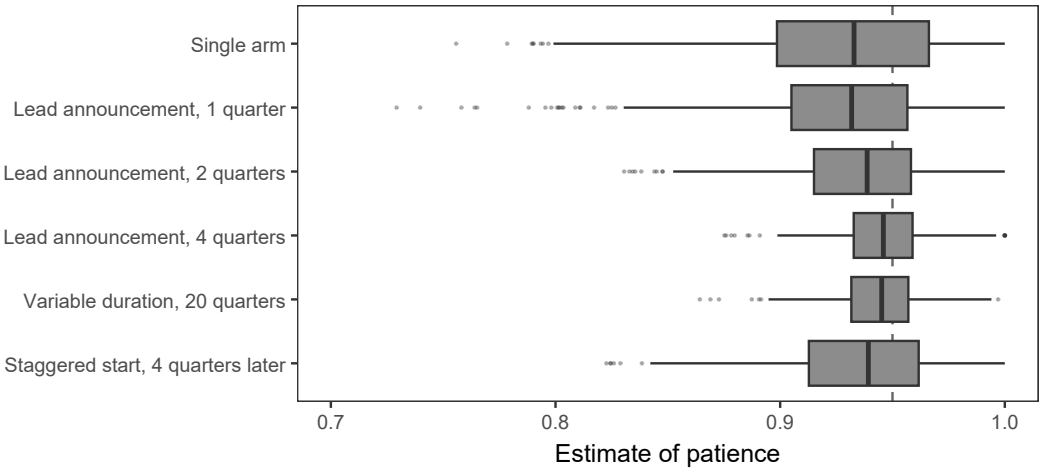

Notes: Each box shows the estimates of patience across the M simulated samples of one design (Table F.2, with the lead announcement at each of its three leads). Boxes span the interquartile range, the vertical line in each box is the median, whiskers extend to the most extreme estimates within 1.5 times the interquartile range beyond the box, and dots mark the estimates beyond the whiskers. The dashed line marks the true value, $\beta = 0.95$.

Table F.4: Simulated rejection rates of the overidentifying restrictions

	The model	The departure			
Lead announcement under an announcement effect					
Strength, a	0	10%	20%	35%	50%
(3) The arms' difference fades at the propagation rate	0.07	0.31	0.98	1.00	1.00
(4) Equal boundary gaps	0.05	0.11	0.30	0.71	0.94
All restrictions	0.08	0.37	0.99	1.00	1.00
Variable duration under a fading response					
Strength, a	0	10%	25%	40%	60%
(5) Equal boundary gaps	0.05	0.09	0.29	0.56	0.89
All restrictions	0.13	0.30	1.00	1.00	1.00
Staggered start and single arm under calendar drift					
Strength, a	0	5%	10%	20%	30%
Staggered start: (6) equal paths in event time	0.06	0.09	0.34	0.96	1.00
Staggered start: all restrictions	0.12	0.07	0.23	0.91	1.00
Single arm: all restrictions	0.08	0.05	0.04	0.06	0.11

Notes: Each entry is the share of the M simulated samples in which a test rejects at the 5% level. Each panel pairs a departure with the experimental design built to detect it (for the lead announcement, a lead of one quarter); the last panel adds the single arm. Rows are the design's tests, numbered as in Table 4: (3) a Wald test that the two arms' boundary gaps are equal, with r estimated from both arms' post-treatment pairs; (4) the overidentification statistic, with r as the only parameter, of the pairs of the difference between the two arms from $\underline{t}$ on; (6) a Wald test that the two arms' treatment effects coincide in every event quarter; and all restrictions, the overidentification statistic of the design's stacked fit. The column headed "The model" uses data from the model and the nominal χ^2 critical values; the columns under "The departure" use data from the departure at strength a and the 95th percentile of each statistic across the same samples when the model holds. The announcement effect and the fading response leave the data of the other tested designs unchanged (the two- and four-quarter leads are not tested), so those designs' tests reject in 5% of samples by construction. Calendar drift changes every design's data; the full fits of the lead announcement and variable duration reject a drift of 20% in 62 and 100% of samples.