Identification Under Universal Policy Exposure

Vedant Vohra
Department of Economics, UC San Diego

Job Market Paper
Click here for the latest version.

Abstract: Many important policy reforms have two key features. First, all units are exposed to the reform at the same time. Second, the reform changes units' incentives to take up an already available treatment. How can we leverage the reform to identify treatment effects in the absence of an unexposed comparison group? While there is no cross-sectional variation in policy exposure, units differ in how their treatment take-up responds to the reform. To leverage this variation for identification, I extend the local average treatment effect (LATE) framework to a panel setting where a unit's latent compliance type, defined by comparing potential treatment choices with and without the policy in each period, can change over time. I establish that there are two routes to point identify the LATE under parallel trends. The first approach restricts the mechanisms governing selection into treatment, ruling out specific transitions in compliance types over time. The second restricts heterogeneity in treatment effects across latent compliance types. To accommodate richer selection mechanisms and unobserved effect heterogeneity, I develop a marginal treatment effects approach that delivers sharp bounds on causal parameters of interest under economically motivated shape restrictions. Applying the framework to Medicare Part D, I estimate that prescription-drug coverage improved medication access and reduced labor supply among individuals induced to obtain it by the reform. *JEL:* C14, C21, C23.

Keywords: Panel Data; Local Average Treatment Effects; Principal Stratification; Partial Identification.

Date: September 26, 2026. E-mail: vevohra@ucsd.edu. Website: vedant-vohra.github.io.
I am grateful to my dissertation co-chairs, Graham Elliott and Kaspar Wuthrich, and committee members Itzik Fadlon, Craig McIntosh, and Xinwei Ma for their support and guidance. I also thank Julie Cullen, Riccardo Di Cato, Stefan Faridani, Jacob Goldin, Harrison Mitchell, and Paul Niehaus for helpful comments and conversations.

1 Introduction

Many economically important policies are implemented at the same time for an entire population. These include administrative reforms that reduce barriers to obtaining public benefits (e.g., Miller, 2012; Ko and Moffitt, 2024; Cholli and Wu, 2026); immigration restrictions, minimum-wage increases, and input subsidies that change firms' relative input costs (e.g., Morris et al., 2007; Clemens, Lewis and Postel, 2018; Haanwinckel, 2023); and regulatory reforms changing the penalties firms face for violating environmental or labor laws (e.g., Stafford, 2002; Wuellrich, 2010; Shapiro and Walker, 2018). Because these policies are implemented simultaneously, there is no cross-sectional variation in policy exposure that can be leveraged for causal identification, as in standard difference-in-differences designs. Treatment take-up, however, varies across units and over time, and the policy changes the incentives of taking up treatment when it is switched on. Consider, for example, the introduction of a health-insurance subsidy made available to an entire population. Individuals may obtain health insurance before the subsidy is introduced, but the subsidy lowers the cost of coverage and induces some individuals who would otherwise have remained uninsured to obtain coverage.[1] In this paper, I develop a framework that leverages this panel variation in treatment take-up to identify causal effects of universally implemented policies.

I begin by adapting the standard local average treatment effect (LATE) framework—which classifies units into latent compliance types based on their potential treatment choices—to a panel setting in which units make those choices in multiple periods and their compliance type may therefore change over time. I primarily focus on identification of a LATE: the average causal effect of the treatment on an outcome of interest in the post-policy window, for units induced to take up treatment by the policy.[2] In the example of a universal health insurance subsidy, this is the effect of insurance on health among individuals who obtain coverage because of the subsidy, but otherwise would have remained uninsured.

Universal policy exposure creates two key challenges to identification. First, the econometrician observes treatment decisions at baseline only with the policy switched off and at endline only with the policy switched on. An observed change in treatment status may therefore reflect either policy-induced take-up or a change that would have occurred even without the policy, but there is no unexposed control group against which to compare these changes. Suppose, for example, that an individual is uninsured before the subsidy is introduced but

[1] Similarly, a nationwide input subsidy program for farmers universally reduces the cost of purchasing fertilizers (Morris et al., 2007), while an increase in regulatory fines for environmental noncompliance may raise the cost of noncompliance for all firms (Stafford, 2002).

[2] I also study identification of the policy's first stage effect on treatment take-up (e.g., the effect of a universal subsidy on insurance coverage) and its reduced-form effect on the outcome (e.g., the effect of a universal subsidy on health).

obtains coverage afterward. Absent additional assumptions, it is unclear whether the subsidy causally induced this take-up; the individual might have purchased insurance in the post-period even without the subsidy. Second, identifying the LATE requires recovering the compliers' post-policy outcomes had they not selected into treatment. A common approach is to model this unobserved potential outcome using a parallel trends assumption. I show that even a strong form of parallel trends is insufficient to identify the LATE.[3] The reason is that some compliers may already be treated at baseline, so their untreated outcomes are never observed. Parallel trends restricts how these outcomes evolve over time, but not their baseline levels, leaving the post-policy untreated counterfactual unrestricted. Therefore, the LATE cannot be identified without additional assumptions.

I show that identification of the LATE entails a trade-off: researchers must impose restrictions on either the mechanisms governing selection into treatment or the extent of treatment-effect heterogeneity. I begin with restrictions on selection behavior. Consider a two-period setting in which the universal policy activates in the second period. In each period, a unit has two potential treatment choices: whether it would take up treatment if the policy were switched on or switched off. Comparing these two potential treatment choices within each period defines four latent compliance types: compliers take up treatment only with the policy, never-takers do not take up treatment under either policy state, always-takers take up treatment under both policy states, and defiers take it up only without the policy (Imbens and Angrist, 1994). Absent restrictions connecting potential treatment decisions across periods, the four types in each period yield 16 possible compliance trajectories. For example, a unit might act as a complier in both periods or transition from an always-taker at baseline to a never-taker at endline. Because some transitions may be economically implausible in a given setting, researchers can impose restrictions to rule them out.

To restrict the set of possible compliance trajectories, I consider three assumptions on selection behavior. Because universal policies typically shift take-up costs in the same direction for all units, I impose monotonicity: any unit who would take up treatment in a given period without the policy would also take it up with the policy. However, within-period monotonicity alone places no restrictions on how choice behavior evolves over time. Therefore, I consider two stationarity conditions that restrict how a unit's compliance type may change over time. First, I require that the aggregate share of always-takers remain stable across periods, which may be plausible when there are no aggregate trends driving treatment take-up. Second, I require that units treated at baseline would remain treated at endline had the policy counterfactually not been implemented, which may be plausible

[3]Following De Chaisemartin and d'Haultfoeuille (2020) and Marx, Tamer and Tang (2024), I consider a parallel trends assumption that requires the mean untreated potential outcomes to evolve identically across all realized treatment paths.

when taking up treatment entails significant fixed costs.[4] Under these assumptions, there is a one-to-one mapping between observed treatment paths and latent compliance types at endline. Moreover, they ensure that all endline compliers are untreated at baseline, so their untreated outcomes at baseline are observed. Parallel trends can then recover their mean untreated outcome at endline. The LATE is point-identified by comparing the change in outcomes among units untreated at baseline but treated at endline (the compliers) with the change in outcomes among continuously untreated units (the never-takers).

Although these assumptions on selection behavior may be plausible in many empirical contexts, they rule out settings in which time-varying unobservables drive treatment take-up. I show, however, that relaxing these selection restrictions while retaining point identification of the LATE requires restrictions on treatment-effect heterogeneity. I therefore consider an alternative approach that relaxes these restrictions on selection behavior but requires homogeneous average treatment effects across the relevant complier and always-taker subgroups. Under this restriction and parallel trends, the same difference-in-differences contrast as before—comparing the change in outcomes among units untreated at baseline but treated at endline with the change among continuously untreated units—identifies the LATE. If exact homogeneity is implausible in a given setting, I show how researchers can bound the degree of permissible heterogeneity using user-specified sensitivity parameters. Varying these parameters allows researchers to trace out a breakdown frontier showing how much heterogeneity is required to overturn a given causal conclusion.

The preceding point identification strategies require researchers either to rule out treatment take-up driven by time-varying unobservables or to impose treatment-effect homogeneity, thereby ruling out selection based on anticipated gains. For settings characterized by heterogeneity and selection patterns that are too rich to yield a single point estimate, I introduce a flexible marginal treatment effects framework that partially identifies the LATE. The framework has three distinctive features. First, it extends MTE analysis to a panel setting with period-specific unobservables governing treatment take-up. Second, I show that additional pre-policy data, combined with mild restrictions on how selection behavior evolves over time, identify all relevant selection primitives. Third, rather than bounding heterogeneity directly with user-specified sensitivity parameters, the framework allows treatment effects to vary with latent take-up propensities and uses economically motivated restrictions on this relationship to obtain sharp, easily computable bounds.

I conclude by using the framework to analyze exposure designs, commonly used to evaluate universal policies. These designs leverage variation in baseline treatment status, comparing

[4]With additional pre-period data, researchers can assess the plausibility of these stationarity conditions by examining stability in the share of treated units and persistence of treatment take-up among previously treated units, respectively. In the paper, I formalize the conditions under which these identifying assumptions are testable.

units untreated at baseline (conjectured to be more exposed) with those already treated (conjectured to be less exposed). I show that, under parallel trends, the resulting estimand equals the treatment effect among all units who switch treatment status (not necessarily induced to switch by the policy) minus a bias term arising from time-varying treatment effects among continuously treated units.[5] Time-invariant treatment effects are therefore necessary for the design to recover a valid causal effect.[6] This restriction is often implausible because it rules out calendar-time variation in average gains from treatment, such as variation driven by changing economic conditions. Moreover, even under time invariance, the design identifies a difficult-to-interpret weighted average of treatment effects across all compliance types; recovering the LATE requires restrictions on selection or treatment-effect heterogeneity.

I apply the framework in two settings. First, I benchmark the proposed strategy against the Oregon Health Insurance Experiment, which randomized the opportunity to apply for Medicaid (Taubman et al., 2014). Among lottery winners uninsured at baseline, I compare changes in emergency-department use between those who obtained coverage at endline and those who remained uninsured. The estimated LATE for the probability of an ED visit (7.3 percentage points) is close to the experimental benchmark (7.0 percentage points), supporting the plausibility of the proposed identification approach.

Next, I apply the framework to Medicare Part D, which in 2006 subsidized prescription-drug coverage for adults aged 65 and older. A prominent literature evaluates Part D by comparing Medicare-eligible individuals with those just below age 65 (e.g., Engelhardt and Gruber, 2011; Wettstein, 2020). Yet Part D increased prescription-drug utilization among the nonelderly through pharmaceutical advertising (Alpert, Lakdawalla and Sood, 2023), casting doubt on this comparison. Without a credible under-65 control group, Part D's universal rollout leaves no cross-sectional variation in policy exposure among eligible adults. I instead use panel variation in take-up of prescription-drug coverage among adults already aged 65 or older before the reform, sidestepping concerns about a potentially affected younger control group. Consistent with existing work, I find that Part D substantially increased coverage, which improved medication access and financial protection while reducing labor supply. Sensitivity analyses show robustness to substantial heterogeneity, while MTE bounds permitting selection on gains yield similar conclusions. Finally, I implement an exposure design comparing respondents uninsured at baseline with those already insured: the estimated effect of coverage on self-reported health is negative but insignificant, and each reported point

[5] Tazhitdinova and Vázquez-Bare (2025) similarly show that comparisons between groups with unequal baseline treatment statuses can confound the effect of switching treatment with the evolution of treatment effects over time. Sun and Shapiro (2022) show that treatment-effect heterogeneity can prevent estimators exploiting differential exposure to a common policy change from recovering even a weighted average of causal effects.
[6] Using baseline treatment status as an instrument for policy exposure requires the same time-invariance restriction (Goldsmith-Pinkham, Sorkin and Swift, 2020).

estimate lies outside the corresponding MTE bounds. This suggests that exposure designs may yield biased estimates of the LATE.

Despite empirical interest in evaluating universally implemented policies, econometric guidance on identification in these settings remains comparatively limited. Conventional difference-in-differences compares outcome changes between treated and untreated or not-yet-treated units (e.g., Ashenfelter and Card, 1985; Abadie, 2005; Roth et al., 2023). Extensions allow treatment status to change over time (De Chaisemartin and d'Haultfoeuille, 2026), or examine comparisons across treatment doses (Callaway, Goodman-Bacon and Sant'Anna, 2025). Fuzzy difference-in-differences likewise uses cross-sectional comparisons between two groups experiencing different changes in treatment rates (De Chaisemartin and d'Haultfoeuille, 2018). Xu, Zhao and Ding (2026) study events affecting all units but use an observed baseline factor to construct comparison groups. Here, all units are universally exposed to a policy change. My primary contribution is to characterize the assumptions under which panel variation in treatment take-up identifies the LATE under universal exposure.

The paper also relates work on principal stratification (Imbens and Angrist, 1994; Frangakis and Rubin, 2002) to difference-in-differences by establishing the identifying content of restrictions on latent compliance types together with parallel trends. Angrist et al. (2024) study identification with period-specific compliance using a randomized instrument, while Ghanem, Sant'Anna and Wüthrich (2026*a*) and Marx, Tamer and Tang (2024) study the compatibility of different selection models with parallel trends. Here, all units face the same policy change at the same time and endogenous treatment choices allow units to move between compliance types over time. I characterize the sharp set of restrictions on compliance transitions that, together with parallel trends, point identify the LATE.[7]

The paper also contributes to partial identification in panel-data settings (e.g., Manski and Pepper, 2018; Rambachan and Roth, 2023). Building on the MTE literature (e.g., Heckman and Vytlacil, 1999, 2005; Mogstad, Santos and Torgovitsky, 2018), I allow treatment effects to vary with period-specific latent propensities to take up treatment. The discrete variation created by the universal policy is generally insufficient to point identify selection mechanisms governed by multidimensional, period-specific unobservables (Kamat, Norris and Pecenco, 2026). I show that additional pre-policy data, combined with restrictions on the dependence of these unobservables, identify the relevant selection primitives.

The remainder of the paper proceeds as follows. Section 2 formalizes the setting and provides motivating examples. Section 3 presents the point identification and sensitivity results, while Section 4 develops the MTE framework for partial identification. Section 5 analyzes exposure designs. Section 6 provides an application. Section 7 concludes.

[7]These selection restrictions also connect the paper to correlated-random-effects models (Chamberlain, 1982) and to nonseparable panel models identified through exchangeability or time-homogeneity restrictions (Altonji and Matzkin, 2005; Chernozhukov et al., 2013).

2 Setting and Empirical Examples

2.1 Setup and Notation I begin with a two-period panel dataset, where each period is denoted by $t \in \{0, 1\}$ and each unit is denoted by i. Let $Z_t \in \{0, 1\}$ be an indicator that equals one if the policy is switched on in period t and zero otherwise. Crucially, Z_t is not indexed by i; the policy switches on for all units in the same period. In this simple two-period setup, the policy is introduced at $t = 1$, such that $Z_0 = 0$ and $Z_1 = 1$.

Let $D_{it}(z) \in \{0, 1\}$ denote the potential treatment status for unit i in period t if the policy state Z_t were set to z. The observed treatment status is given by:

$$D_{it} = D_{it}(0) \cdot (1 - Z_t) + D_{it}(1) \cdot Z_t$$

This notation assumes that treatment status depends only on the contemporaneous policy state. Next, let $Y_{it}(d)$ denote the potential outcome of interest for unit i in period t if their treatment status were set to $d \in \{0, 1\}$. The observed outcome is given by:

$$Y_{it} = Y_{it}(0) \cdot (1 - D_{it}) + Y_{it}(1) \cdot D_{it}$$

This notation for the potential outcomes imposes important restrictions. First, it rules out anticipatory effects, as the outcome in period t does not depend on future treatment status. Second, it rules out carryover effects, as the period t outcome depends only on contemporaneous treatment status, rather than the entire treatment path. Third, it imposes an exclusion restriction: the policy Z_t affects the outcome only through the treatment status D_{it}. Consequently, the potential outcome satisfies $Y_{it}(d, z) = Y_{it}(d)$ for $z \in \{0, 1\}$. Finally, the notation also imposes no interference across units, so that a unit's potential outcomes do not depend on the treatment statuses of other units. Throughout, I assume that $\mathbb{E}[|Y_{it}(d)|] < \infty$ for all $t, d \in \{0, 1\}$ and that each treatment path $(D_{i0}, D_{i1}) \in \{(0, 0), (0, 1), (1, 1)\}$ has positive probability.

For each period $t \in \{0, 1\}$, units can be categorized into four latent compliance types based on their potential treatment decisions $D_{it}(z)$ (Imbens and Angrist, 1994). Specifically, period-t always-takers are units for whom $D_{it}(0) = D_{it}(1) = 1$; period-$t$ compliers are units for whom $D_{it}(1) > D_{it}(0)$; period-t never-takers are units for whom $D_{it}(0) = D_{it}(1) = 0$; and period-$t$ defiers are units for whom $D_{it}(1) < D_{it}(0)$. A critical departure in our setting is that a unit's latent compliance type is period-specific. Because units make potential treatment decisions in both periods—represented by the random variables $\{D_{i0}(0), D_{i0}(1), D_{i1}(0), D_{i1}(1)\}$—there are 16 possible compliance types across time. However, because the aggregate policy is switched off at baseline ($Z_0 = 0$) and switched on at endline ($Z_1 = 1$), the econometrician only observes the realized treatment path $(D_{i0}, D_{i1}) = (D_{i0}(0), D_{i1}(1))$. Importantly, absent additional assumptions, each realized

treatment path is consistent with four latent compliance trajectories and therefore does not uniquely determine a unit's compliance type in either period.

I primarily focus on identification of the causal effect of the treatment on the subpopulation induced to take it up by the universal policy change. Specifically, I define the target parameter as the Local Average Treatment Effect (LATE) for period-1 compliers:

$$\lambda = \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) | D_{i1}(1) > D_{i1}(0)]$$

I assume that $\mathbb{P}(D_{i1}(1) > D_{i1}(0)) > 0$, so that λ is well defined. This parameter is the average treatment effect for units who would take up treatment at $t = 1$ under the policy but would not do so in its absence. In standard instrumental variables (IV) designs, the LATE is frequently criticized because the compliant subgroup is tied to the specific instrument (e.g., a randomized informational nudge) used for identification (e.g., Heckman and Urzua, 2010); the critique is that the identified LATE need not correspond to the treatment effect for the subgroup relevant to the policy question of interest. However, in our setting, the universal policy itself induces variation in treatment take-up, such that the compliant subgroup consists exactly of the units whom the policy induces to take up treatment at $t = 1$. Consequently, λ is an important policy-relevant parameter in our setting.[8]

2.2 Empirical Examples To help fix ideas and motivate this setting, I consider three distinct policy domains where universally implemented policies alter compliance behavior.

Example 2.1. (Universal Changes to the Cost of Accessing Social Programs) A broad class of policies universally lowers the costs of enrolling in social safety net programs. As a leading example, consider Cholli and Wu (2026), who study a reform that simultaneously reduced administrative barriers to welfare enrollment for an entire population.[9] In this context, Z_t is an indicator for the introduction of a unified, streamlined online portal to apply for welfare benefits. The binary treatment D_{it} represents individual i's decision to enroll in welfare benefits in period t. Because the portal is absent at baseline ($t = 0$) and active at endline ($t = 1$), the econometrician only observes the realized enrollment sequence $(D_{i0}, D_{i1}) = (D_{i0}(0), D_{i1}(1))$. The counterfactual enrollment decisions—whether the individual would have enrolled at $t = 0$ had the portal been launched, or at $t = 1$ had the portal not been launched—remain latent. Finally, the target parameter λ captures the causal effect of welfare enrollment on outcomes of interest, for the subpopulation induced to enroll at $t = 1$ precisely because the online portal lowered the cost of enrolling.

[8] While λ serves as the primary target for the subsequent identification results, I also provide results for the identification of the first stage effect, $\mathbb{E}[D_{i1}(1) - D_{i1}(0)]$, the reduced-form effect of the policy, $\mathbb{E}[Y_{i1}(D_{i1}(1)) - Y_{i1}(D_{i1}(0))]$, and other policy-relevant parameters of interest.

[9] Other examples include policies that change the cost of enrolling in health insurance for all individuals in a population of interest (Miller, 2012; Argys et al., 2020) and policies that make early childhood education universally more accessible (DeMalach and Schlosser, 2026).

Example 2.2. (Universal Changes to the Cost of Technology Adoption) A broad class of policies alters the relative costs of adopting production technologies, shifting the incentives to adopt inputs that firms could already procure prior to the reform. As a leading example, consider the introduction of a universal subsidy program (for example, see Morris et al., 2007) to purchase fertilizers made available to all farmers.[10] In this context, Z_t is an indicator for the availability of the fertilizer subsidy. The binary treatment D_{it} represents farmer i's decision to purchase fertilizer in period t. Because the subsidy is absent at baseline ($t = 0$) and active at endline ($t = 1$), the econometrician only observes the realized adoption sequence $(D_{i0}, D_{i1}) = (D_{i0}(0), D_{i1}(1))$. The counterfactual adoption decisions—whether the farmer would have purchased fertilizer at $t = 0$ had the subsidy been available, or at $t = 1$ had the subsidy not been available—remain latent. Finally, the target parameter λ captures the causal effect of fertilizer adoption on an outcome of interest for the subpopulation of farmers induced to purchase fertilizers at $t = 1$ precisely because the subsidy lowered their cost.

Example 2.3. (Universal Regulatory Changes) A broad class of federal policies alters firm behavior by universally changing the pecuniary penalties for non-compliance, shifting the incentives to adhere to existing regulatory standards. As an illustrative example, consider the enforcement of environmental regulations where the government universally increases the fines for not complying with specific environmental standards (e.g., Stafford, 2002).[11] In this context, Z_t is an indicator for the period in which fines are increased for non-adherence. The binary treatment D_{it} represents firm i's decision to adhere to the environmental standard in period t. Because the higher penalty is absent at baseline ($t = 0$) and active at endline ($t = 1$), the econometrician only observes the realized adherence sequence $(D_{i0}, D_{i1}) = (D_{i0}(0), D_{i1}(1))$. The counterfactual adherence decisions—whether the firm would have adhered at $t = 0$ had the fine been higher, or at $t = 1$ had the fine remained low—remain latent. Finally, the target parameter λ captures the causal effect of regulatory compliance on an outcome of interest specifically for the subpopulation induced to comply at $t = 1$ because the increased fine altered the relative cost of non-compliance.

[10] Other examples include national minimum wage hikes that can alter the relative cost of labor versus automation (Aaronson and Phelan, 2017; Haanwinckel, 2023) and federal immigration changes restricting foreign farm-workers that can change farmers' incentives to mechanize (Clemens, Lewis and Postel, 2018).

[11] Other examples include federal labor law changes (Quach, 2024) and nationwide changes in financial penalties for violating employment mandates (Wuellrich, 2010). Indeed, Shapiro and Walker (2018) emphasize that a key challenge in evaluating many environmental regulations using existing econometric tools has been that regulations are often universally implemented, which leaves researchers without a natural comparison group. Similarly, Quach (2024) is explicitly unable to evaluate the consequences of important federal regulatory changes on firm behavior specifically because they are universally implemented, and instead focuses on smaller state-level regulatory changes.

3 Point Identification of λ

In this section, I study point identification of the local average treatment effect, λ. A natural starting point is the standard parallel trends assumption, which serves as the key identifying assumption in the difference-in-differences literature. This assumption posits that the mean evolution of untreated potential outcomes is constant across all realized treatment paths (e.g., De Chaisemartin and d'Haultfoeuille, 2020; Marx, Tamer and Tang, 2024).

Assumption PT (Parallel Trends)**.** For some constant $\Delta^{PT} \in \mathbb{R}$ and all realized treatment paths $d_0, d_1 \in \{0, 1\}$:

$$\mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = d_0, D_{i1} = d_1\Big] = \Delta^{PT}$$

Assumption PT imposes a strong exogeneity restriction on untreated outcome changes, requiring their mean to be invariant to the realized treatment path. In Section 3.1, I demonstrate that while Assumption PT identifies counterfactual outcome trends, it is insufficient to identify λ. Motivated by this negative result, I provide two distinct pathways to point-identify λ. Section 3.2 explores assumptions that restrict selection into treatment over time, while Section 3.3 considers an alternative set of assumptions restricting treatment effect heterogeneity. Notably, I show that under either set of restrictions, the parallel trends condition is required only among units untreated at baseline: $\mathbb{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 0] = \mathbb{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1]$.[12] Section 3.4 provides a representation of λ nesting both point-identification results as special cases, and Section 3.5 uses this representation to conduct sensitivity analyses under bounded heterogeneity assumptions.

3.1 Lack of Identification under Parallel Trends

In this section, I describe a central challenge in identifying the target parameter, λ, and demonstrate why additional assumptions beyond Assumption PT are necessary to identify it. Specifically, I show that λ remains entirely unconstrained under Assumption PT.

To understand this failure of identification, consider the subpopulation that selects into treatment in both periods. For these continuously treated units ($D_{i0} = 1, D_{i1} = 1$), the econometrician only ever observes the treated potential outcomes, such that $Y_{i0} = Y_{i0}(1)$ and $Y_{i1} = Y_{i1}(1)$.

Assumption PT posits that the average untreated potential outcome at $t = 1$ for this continuously treated group can be written as:

$$\mathbb{E}[Y_{i1}(0) \mid D_{i0} = 1, D_{i1} = 1] = \mathbb{E}[Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 1] + \Delta^{PT}.$$

The counterfactual trend Δ^{PT} is identified by the never-treated units ($D_{i0} = 0, D_{i1} = 0$). However, the moment $\mathbb{E}[Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 1]$ is unobserved and absent additional

[12] In Section 4, I provide a framework to bound λ under further relaxations of Assumption PT considered in previous work (Manski and Pepper, 2018; Rambachan and Roth, 2023).

assumptions, this moment can be arbitrarily large or small. Crucially, period-1 compliers may include units already treated at baseline, for whom neither untreated potential outcome is observed. Shifting both of these potential outcomes by the same amount preserves the observed data and Assumption PT, while changing their treatment effects. Consequently, λ can be arbitrarily large or small absent additional assumptions. I summarize this negative result in the following proposition.

Proposition 1. *Let Λ^* denote the identified set for λ. Under Assumption PT, $\Lambda^* = \mathbb{R}$ when no restrictions are imposed on the support of potential outcomes.*

Proof. See Appendix Section A.

Remark 1. Even if a researcher knew precisely which units are period-1 compliers, λ would remain unconstrained provided $\mathbb{P}[D_{i0} = 1, D_{i1}(1) > D_{i1}(0)] > 0$. For these baseline-treated compliers, neither untreated potential outcome is observed. Shifting both by the same constant preserves the observed data and Assumption PT, while changing λ. ■

Remark 2. It is instructive to contrast the negative result in Proposition 1 with settings where treatment is entirely unavailable prior to the policy change, such that $D_{i0} = 0$ for all i. For example, when evaluating the rollout of a job-training program (e.g., LaLonde, 1986), individuals cannot enroll before the program exists. In such cases, any unit treated at endline is necessarily untreated at baseline. As a result, the baseline untreated potential outcome for the treated group is identified, i.e., $\mathbb{E}[Y_{i0}(0) \mid D_{i1} = 1] = \mathbb{E}[Y_{i0} \mid D_{i0} = 0, D_{i1} = 1]$. Combined with the trend Δ^{PT} identified from the never-treated units, this allows for point identification of the treatment effect on the treated under Assumption PT. In contrast, under the universal policy changes described in Section 2.2, units select into the underlying treatment (D_{it}) even at $t = 0$ and the policy (Z_t) switches on at $t = 1$ as a universal cost-shifter. Because units may select into treatment both pre- and post-reform, the baseline untreated potential outcome for the continuously treated subpopulation remains unidentified under Assumption PT. ■

Proposition 1 highlights the key challenge to identifying λ: the maintained assumptions permit period-1 compliers to have been treated at baseline, while leaving their post-policy untreated potential outcome unrestricted. To overcome this challenge, a researcher has two possible avenues for point identification. The first avenue restricts admissible selection mechanisms governing treatment take-up, such that units already treated prior to the reform are inframarginal to the policy change and therefore are not in the compliant subgroup at endline. I explore this avenue in Section 3.2. The second avenue permits baseline-treated units to be in the compliant subgroup at endline, but imposes restrictions on treatment effect heterogeneity. Specifically, this approach assumes that period-1 compliers have the same average treatment effect regardless of their baseline treatment status, and that period-1

compliers and always-takers who were untreated at baseline have the same average treatment effect. I explore this avenue in Sections 3.3.

3.2 Restricting Admissible Selection Mechanisms In this section, I explore the first avenue to point identification: imposing restrictions on the mechanism governing selection into treatment over time. Recall from Section 2.1 that because units make potential treatment decisions in both periods—choosing the sequence $\{D_{i0}(0), D_{i0}(1), D_{i1}(0), D_{i1}(1)\}$—there are 16 possible joint compliance trajectories. Absent additional structure, a unit's compliance type may evolve arbitrarily over time. For instance, an individual might act as an always-taker at baseline ($D_{i0}(0) = D_{i0}(1) = 1$) but transition into a complier at endline ($D_{i1}(1) > D_{i1}(0)$). Absent additional assumptions, each observed treatment path is consistent with four latent compliance trajectories and therefore does not uniquely determine a unit's compliance type in either period. The following assumption restricts this mapping and rules out specific dynamic compliance sequences.

Assumption SEL (Selection Restrictions)**.**

(a) Monotonicity:

$$\mathbb{P}[D_{it}(1) \geq D_{it}(0)] = 1 \text{ for } t \in \{0, 1\}$$

(b) Stability in Share of Always-Takers:

$$\mathbb{P}[D_{i0}(0) = D_{i0}(1) = 1] = \mathbb{P}[D_{i1}(0) = D_{i1}(1) = 1]$$

(c) Stability of Choosers over Time:

$$\mathbb{P}[D_{i1}(0) = 0 \mid D_{i0}(0) = 1] = 0$$

Assumption SEL(a) imposes monotonicity, stating that the policy weakly encourages selection into treatment, thereby ruling out defiers (i.e., units for whom $D_{it}(1) < D_{it}(0)$) in both periods. In the context of Example 1 in Section 2.2, this requires that universally reducing the administrative cost of accessing welfare—such as by introducing a streamlined online portal—does not prevent any individual who would otherwise enroll from doing so. Crucially, however, monotonicity alone places no intertemporal constraints on choice behavior, as it only restricts potential treatment statuses within a given time period. To address the dynamic nature of compliance, Assumptions SEL(b) and SEL(c) introduce novel restrictions on the time-series properties of the selection mechanism.

Assumption SEL(b) imposes stability in the aggregate share of always-takers over time. This assumption may be plausible when there are no secular trends driving treatment take-up in the absence of the policy change. In the context of Example 1 in Section 2.2, this requires that the proportion of the population who would enroll in welfare benefits without the policy—that is, without the aid of a streamlined online portal to reduce administrative

frictions—remains constant across periods. Additional pre-period data on treatment take-up can provide evidence supporting the plausibility of this assumption. Under Assumption SEL(a), any unit that selects into treatment absent the policy is necessarily an always-taker. Consequently, stable aggregate enrollment rates across multiple pre-reform periods would indicate that the share of always-takers was stable before the policy was introduced.

Assumption SEL(c) states that units who select into treatment in the period before the policy is implemented would continue to do so in the post-period had the policy counterfactually not been implemented. This assumption may be plausible when take-up entails substantial fixed costs, while continued participation is relatively inexpensive. In the context of Example 1 in Section 2.2, this assumption requires that individuals who successfully enrolled in welfare benefits at baseline—navigating the high administrative frictions without the aid of the online portal—would continue to enroll in the subsequent period had the portal counterfactually not been introduced. With unit-level panel data spanning multiple pre-policy periods, researchers can assess whether treated units persistently remain in treatment to support the plausibility of this assumption. This assumption can equivalently be stated as a monotonicity condition on treatment take-up over time, under the counterfactual that the policy was never switched on: $\mathbb{P}[D_{i1}(0) \geq D_{i0}(0)] = 1$.

Next, I formalize the identifying content of Assumption SEL. Recall that absent any restrictions on selection into treatment, a unit's compliance type may evolve arbitrarily across the two periods, generating 16 possible transition paths as illustrated in Figure 1a. Assumption SEL restricts the admissible compliance types over time as follows. First, Assumption SEL(a) rules out defiers in both periods. Consequently, any unit observed selecting into treatment at baseline (i.e., a unit for whom $D_{i0} = D_{i0}(0) = 1$) must be a period-0 always-taker. Next, under Assumption SEL(a), Assumptions SEL(b) and (c) jointly imply that potential treatment decisions in the absence of the policy are time-invariant, such that $D_{i0}(0) = D_{i1}(0)$ almost surely. This time-invariance yields two consequences. First, an always-taker at $t = 0$ must remain an always-taker at $t = 1$. Second, units who are never-takers or compliers at baseline (i.e., units for whom $D_{i0} = D_{i0}(0) = 0$) must also remain never-takers or compliers at endline, since their potential treatment status in the absence of the policy must satisfy $D_{i1}(0) = 0$. I summarize this result in Lemma 1 below.

Lemma 1. *Under Assumption SEL, the admissible unit types are:*

(i) never-takers in both periods, i.e., $D_{i0}(0) = D_{i0}(1) = 0$ and $D_{i1}(0) = D_{i1}(1) = 0$,

(ii) never-takers at $t = 0$ and compliers at $t = 1$, i.e., $D_{i0}(0) = D_{i0}(1) = 0$ and $D_{i1}(1) > D_{i1}(0)$,

(iii) compliers at $t = 0$ and never-takers at $t = 1$, i.e., $D_{i0}(1) > D_{i0}(0)$ and $D_{i1}(0) = D_{i1}(1) = 0$,

(iv) compliers in both periods, i.e., $D_{i0}(1) > D_{i0}(0)$ and $D_{i1}(1) > D_{i1}(0)$,

(v) always-takers in both periods, i.e., $D_{i0}(0) = D_{i0}(1) = 1$ *and* $D_{i1}(0) = D_{i1}(1) = 1$.

Proof. See Appendix Section B.

All admissible compliance types over time under Assumption SEL are illustrated in Figure 1b.

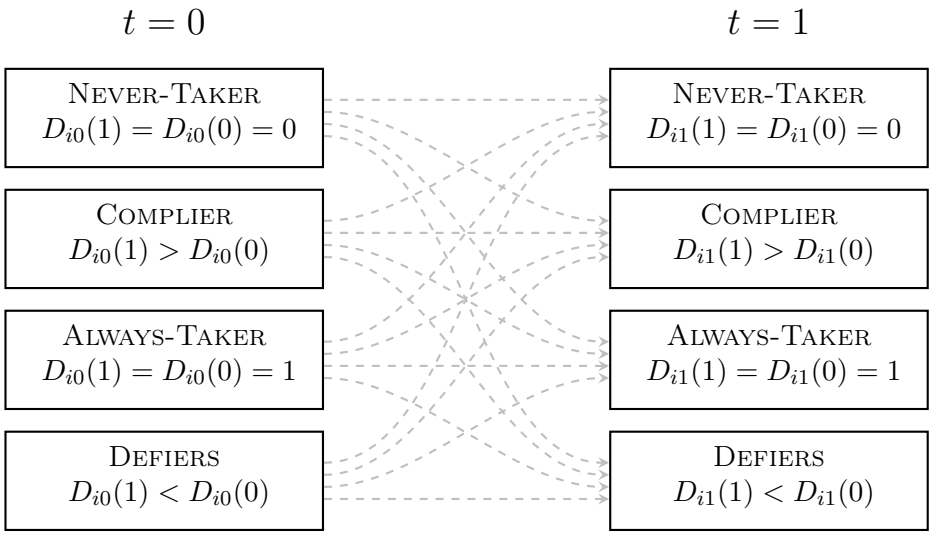

(a) Compliance types over time without imposing Assumption SEL

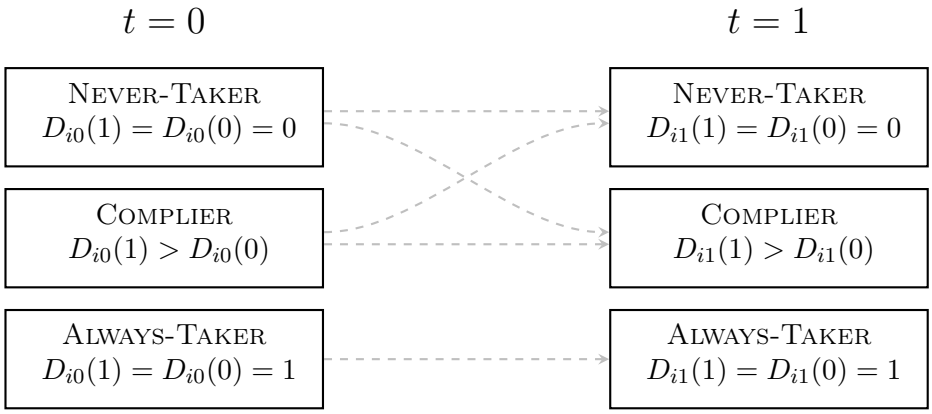

(b) Compliance types over time after imposing Assumption SEL

FIGURE 1. Identifying Content of Assumption SEL

Lemma 1 characterizes the identifying content of Assumption SEL through the lens of non-parametric compliance types. It is also instructive to examine the implications of these restrictions within a separable threshold-crossing model:

$$D_{it}(z) = \mathbb{1}(\alpha_i + \gamma_{it} \cdot z \geq U_{it}).$$

First, a sufficient condition for monotonicity is $\gamma_{it} \geq 0$ for all units and both periods, which can be interpreted as the policy weakly reducing the cost of selecting into treatment. Second, under monotonicity, the share of always-takers in period t is given by $\mathbb{P}[\alpha_i \geq U_{it}]$. Assumption SEL(b) therefore requires that $\mathbb{P}[\alpha_i \geq U_{i0}] = \mathbb{P}[\alpha_i \geq U_{i1}]$. Third, Assumption SEL(c) requires that $\mathbb{P}[\alpha_i \geq U_{i0}, \alpha_i < U_{i1}] = 0$. Under Assumption SEL(a), Assumptions SEL(b) and (c) jointly require that potential treatment decisions in the absence of the policy are time-invariant almost surely, restricting the joint distribution of $(\alpha_i, U_{i0}, U_{i1})$ such that $\mathbb{P}[\min(U_{i0}, U_{i1}) \leq \alpha_i < \max(U_{i0}, U_{i1})] = 0$. Thus, idiosyncratic shocks may vary across periods but almost surely do not induce a unit to alter their treatment decision absent the policy.

Remark 3. Suppose the idiosyncratic shocks can be decomposed as $U_{it} = \nu_i + \delta_t$. Defining $\tilde{\alpha}_i \equiv \alpha_i - \nu_i$, Assumption SEL requires that $\mathbb{P}[\min\{\delta_0, \delta_1\} \leq \tilde{\alpha}_i < \max\{\delta_0, \delta_1\}] = 0$. If $\delta_1 \neq \delta_0$, this restriction requires that the distribution of $\tilde{\alpha}_i$ assign zero probability to the nonempty interval between the two thresholds. If its distribution function is strictly increasing on an interval containing both thresholds, the restriction therefore implies $\delta_1 = \delta_0$. Consequently, $D_{it}(0) = \mathbb{1}(\tilde{\alpha}_i \geq \delta_0)$ in both periods. In this special case, treatment decisions absent the policy depend solely on time-invariant unobservables. Under such selection on fixed effects, Ghanem, Sant'Anna and Wüthrich (2026*b*) show that parallel pre- and post-trends hold in their two-way fixed effects outcome model when the conditional mean of idiosyncratic outcome shocks given the fixed effects is constant over time. ■

An important consequence of Lemma 1 is that it establishes a bijection between observable treatment paths and latent compliance types at endline ($t = 1$). Specifically, all units that remain continuously untreated ($D_{i0} = 0, D_{i1} = 0$) are never-takers ($D_{i1}(1) = D_{i1}(0) = 0$). All units that remain continuously treated ($D_{i0} = 1, D_{i1} = 1$) are always-takers ($D_{i1}(1) = D_{i1}(0) = 1$). Finally, all units that are untreated at baseline but treated at endline ($D_{i0} = 0, D_{i1} = 1$) are compliers ($D_{i1}(1) > D_{i1}(0)$).

Remark 4. Under Assumption SEL(a), units with the observed treatment path $(D_{i0}, D_{i1}) = (1, 0)$ are always-takers at baseline and never-takers at endline, violating Assumption SEL(c). These units, however, do not enter the parameter of interest λ or the estimand considered below. Therefore, they may be excluded from the data without changing the identification argument for λ, provided that Assumption SEL holds for the remaining units.

Having established the identifying content of Assumption SEL, I now turn to the identification of the causal parameters, beginning with the first stage. The following proposition demonstrates that Assumptions SEL(a) and (b) are sufficient to identify the average effect of the policy on treatment take-up. Furthermore, the addition of Assumption SEL(c) permits the point identification of each unit's latent compliance type.

Proposition 2. *Under Assumptions SEL(a)-(b), the first stage effect of the policy on treatment take-up is identified as:*

$$\mathbb{E}\Big[D_{i1}(1) - D_{i1}(0)\Big] = \mathbb{P}\Big[D_{i1} = 1\Big] - \mathbb{P}\Big[D_{i0} = 1\Big].$$

Under Assumptions SEL(a)-(c), each unit's latent type is also identified since $D_{i1}(0) = D_{i0}(0)$ almost surely.

Proof. See Appendix Section C.

The intuition behind this result is as follows. Under Assumption SEL(a), any unit observed taking up treatment at $t = 0$ is an always-taker. Therefore, the share of always-takers at $t = 0$ is identified. Because Assumption SEL(b) holds this share constant over time, the share of always-takers at $t = 1$ is also identified. Since all units who select into treatment at $t = 1$ are either always-takers or compliers, subtracting the treatment take-up rate at $t = 0$ from the treatment take-up rate at $t = 1$ isolates the share of compliers, yielding the first stage effect. Finally, by further imposing Assumption SEL(c), Lemma 1 shows that $D_{i1}(0) = D_{i0}(0)$ almost surely, allowing each unit's latent compliance type to be identified from their treatment path.

Building on the identification of the first stage and latent compliance types, I now turn to the point identification of the target parameter, λ. Recall from Section 3.1 that the fundamental barrier to identifying λ under Assumption PT was the presence of baseline-treated compliers, for whom the pre-policy untreated potential outcome, $Y_{i0}(0)$, is unobserved. However, as established by Proposition 2, Assumption SEL ensures that all endline compliers ($D_{i1}(1) > D_{i1}(0)$) are necessarily untreated at baseline ($D_{i0} = 0$). Consequently, the baseline potential outcome is observed ($Y_{i0} = Y_{i0}(0)$) for the entire compliant subpopulation, permitting the point identification of λ under parallel trends. I formalize this result in the following proposition.

Proposition 3. *Under Assumptions PT and SEL, λ is identified as follows:*

$$\tau := \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big] = \lambda$$

Proof. See Appendix Section D.

It is instructive to examine the empirical variation leveraged by the estimand τ. The first expectation, $\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1]$, captures the observed intertemporal change in outcomes for the compliant subpopulation. This first difference reflects both the causal effect of taking up treatment at endline and any confounding aggregate time shocks. The second expectation, $\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0]$, evaluates the corresponding over-time evolution for the continuously untreated never-takers. Under Assumption PT, differencing out the change in outcomes for never-takers eliminates the confounding time shocks, yielding a standard difference-in-differences contrast that isolates the LATE, λ.

A direct corollary of Propositions 2 and 3 is that the reduced-form effect of the policy is also identified.

Corollary 1. *Under Assumptions PT and SEL, the reduced-form effect of the policy on the outcome is identified as follows:*

$$\mathbb{E}\left[Y_{i1}\Big(D_{i1}(1)\Big) - Y_{i1}\Big(D_{i1}(0)\Big)\right] = \mathbb{E}[D_{i1}(1) - D_{i1}(0)] \times \lambda$$
$$= \left(\mathbb{P}\Big[D_{i1} = 1\Big] - \mathbb{P}\Big[D_{i0} = 1\Big]\right) \times \tau.$$

Having established the point identification of λ under Assumption SEL, a natural theoretical question is whether these selection conditions could be weakened without imposing restrictions on treatment effect heterogeneity. Specifically, one might ask whether it is possible to retain the identification results provided above upon expanding the set of admissible compliance transitions illustrated in Figure 1b. The answer is negative: point identification of λ fails if units are additionally allowed to transition from baseline never-takers to endline always-takers, from baseline compliers to endline always-takers, or from baseline always-takers to endline compliers.[13] I formalize this sharpness of Assumption SEL in detail in Appendix Section E. This highlights a familiar tradeoff between restricting selection into treatment and restricting treatment effect heterogeneity. In the next section, I proceed by making exactly this trade-off: I relax Assumption SEL and demonstrate how imposing restrictions on treatment effect heterogeneity can alternatively restore point identification.

Remark 5. Lemma 1 shows that any unit observed taking up treatment at baseline ($D_{i0} = 1$) is necessarily an always-taker. These units are excluded from the contrast in Proposition 3 and are not leveraged for the identification of the LATE. Because the estimand τ only leverages units that were untreated at baseline, identification relies on a weaker version of Assumption PT. Specifically, the parallel trends condition is only required to hold for units with $D_{i0} = 0$: $\mathbb{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1] = \mathbb{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 0]$. ■

3.3 Restricting Treatment Effect Heterogeneity In Section 3.2, I demonstrated that restricting the joint distribution of $\{D_{i0}(0), D_{i0}(1), D_{i1}(0), D_{i1}(1)\}$ permits the point identification of λ, and established that these selection conditions cannot be relaxed without simultaneously restricting treatment effect heterogeneity. In this section, I explore this alternative approach: I allow for unrestricted selection transitions—as illustrated in Figure 1a—and instead impose restrictions on treatment effect heterogeneity. I begin by formalizing the specific homogeneity condition required to restore point identification.

[13] Under monotonicity, units transitioning from baseline always-takers to endline never-takers are identified by the observed path $(D_{i0}, D_{i1}) = (1, 0)$. If this is the only additional admissible compliance trajectory, these units may be excluded without loss, leaving the five admissible trajectories in Figure 1b.

Assumption HET (Treatment Effect Homogeneity)**.** For $d_0, d_1 \in \{0, 1\}$:

$$\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = d_0, D_{i1}(0) = d_1, D_{i1}(1) = 1\Big] = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) = 1\Big]$$

Assumption HET imposes two key homogeneity restrictions across the latent subpopulations that select into treatment at endline. First, it requires that the average treatment effect is equal for both endline compliers and endline always-takers who were untreated at baseline. Second, it requires that the average treatment effect is equal for endline compliers who were treated at baseline and endline compliers who were untreated at baseline. I formalize exactly how this assumption helps identify λ in the following proposition.

Proposition 4. *Under Assumptions PT and HET, λ is identified as follows:*

$$\tau := \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big] = \lambda$$

Proof. See Appendix Section F.

Proposition 4 demonstrates that τ—the same estimand leveraged under the selection restrictions in Proposition 3—point-identifies λ under Assumption HET. To understand the intuition behind this result, note that under Assumption PT, τ identifies a weighted average of the treatment effects for two distinct latent groups: endline compliers and endline always-takers who were both untreated at baseline. The first restriction of Assumption HET dictates that the treatment effect is equal across both of these subgroups, permitting identification of the average treatment effect for baseline-untreated compliers. The second restriction then imposes that this treatment effect is homogeneous across compliers, regardless of whether they were treated or untreated at baseline. Consequently, τ recovers the treatment effect for the entire compliant subpopulation, λ.

While Assumption HET point-identifies the LATE, it is a strong restriction that rules out, for instance, economic settings characterized by Roy-style selection on gains. Before jointly relaxing both the selection conditions of Section 3.2 and the homogeneity conditions of Section 3.3, it is helpful to examine how these assumptions interact. To this end, Section 3.4 introduces a unified framework that nests both approaches.

3.4 A Unified Identification Framework

In this section, I provide a generalized representation of the LATE that unifies the identification approaches discussed in Sections 3.2–3.3. This generic framework nests Propositions 3 and 4 as special cases: imposing the selection conditions of Assumption SEL recovers the point identification result of Proposition 3, while imposing the treatment effect homogeneity of Assumption HET recovers the result of Proposition 4. Crucially, this unified characterization demonstrates why λ remains unbounded whenever Assumption SEL is relaxed and heterogeneity is left unrestricted, underscoring the sharpness result established earlier in Appendix E. Finally, this generic representation outlines a pathway for the sensitivity analysis pursued in Section 3.5.

I begin by defining four parameters that parameterize potential violations of Assumptions SEL and HET. First, the change in the share of always-takers from baseline to endline is defined as

$$\Delta^{AT} := \mathbb{P}[D_{i1}(0) = D_{i1}(1) = 1] - \mathbb{P}[D_{i0}(0) = D_{i0}(1) = 1]. \tag{3.1}$$

Under Assumption SEL(b), Δ^{AT} is set to 0. Second, the probability that a baseline-treated unit counterfactually selects out of treatment at endline is defined as

$$\Delta^{churn} := \mathbb{P}[D_{i1}(0) = 0 \mid D_{i0}(0) = 1] \tag{3.2}$$

Under Assumption SEL(c), Δ^{churn} is set to 0. Third, the difference in the average treatment effect between compliers who were treated at baseline and those who were untreated is denoted as:

$$\kappa_C := \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) > D_{i1}(0), D_{i0} = 1\Big] - \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) > D_{i1}(0), D_{i0} = 0\Big]. \tag{3.3}$$

Finally, the difference in the average treatment effect between baseline-untreated always-takers and baseline-untreated compliers is denoted as:

$$\kappa_{AT} := \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) = D_{i1}(0) = 1, D_{i0} = 0\Big] - \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) > D_{i1}(0), D_{i0} = 0\Big]. \tag{3.4}$$

Under Assumption HET, κ_C and κ_{AT} are set to 0. With these parameters defined, the following proposition expresses the LATE, λ, as a function of these parameters and the estimand, τ.

Proposition 5. *Under Assumption PT and* $\mathbb{P}[D_{it}(1) \geq D_{it}(0)] = 1$,

$$\begin{aligned}\lambda = \tau &- \kappa_{AT} \frac{\Delta^{AT} + \Delta^{churn}\mathbb{P}[D_{i0} = 1]}{\mathbb{P}[D_{i0} = 0, D_{i1} = 1]} \\ &+ \kappa_C \frac{\Delta^{churn}\mathbb{P}[D_{i0} = 1] - \mathbb{1}\{\Delta^{churn} > 0\}\mathbb{P}[D_{i0} = 1, D_{i1} = 0]}{\mathbb{P}[D_{i1} = 1] - (\mathbb{P}[D_{i0} = 1] + \Delta^{AT})}.\end{aligned}$$

Proof. See Appendix Section G.

Proposition 5 nests Propositions 3 and 4 as special cases. Specifically, observe that imposing the selection conditions of Assumption SEL forces both $\Delta^{AT} = \Delta^{churn} = 0$, yielding $\lambda = \tau$ as established in Proposition 3. Importantly, this point identification holds regardless of the treatment effect heterogeneity parameters, κ_{AT} and κ_C. However, if the selection restrictions are relaxed such that either $\Delta^{AT} > 0$ or $\Delta^{churn} > 0$, λ becomes unbounded, because κ_{AT} and κ_C can be arbitrarily large or small absent additional assumptions. The alternative solution, as derived in Proposition 4, is to impose the homogeneity of Assumption HET, setting $\kappa_{AT} = \kappa_C = 0$, which once again restores point identification regardless of the selection parameters Δ^{AT} and Δ^{churn}.

Setting $\kappa_{AT} = \kappa_C = 0$, as required by Assumption HET, may prove implausible in empirical settings characterized by Roy-style selection on gains. In Section 3.5, I relax this strict homogeneity by considering user-specified bounds on κ_{AT} and κ_C to limit the permissible degree of unobserved heterogeneity. By combining these bounds with additional pre-period data to estimate the selection parameters (Δ^{AT} and Δ^{churn}), I demonstrate how researchers can assess sensitivity of their causal conclusions under different degrees of unobserved heterogeneity.

3.5 Sensitivity Analysis While the homogeneity restrictions imposed by Assumption HET might be plausible in certain contexts (see, e.g., Kowalski, 2023), they may prove unpalatable in other empirical settings. Notably, if units sort into treatment according to their anticipated gains—for instance, if always-takers adopt the treatment at baseline precisely because they anticipate a systematically larger treatment effect than the compliers—then this homogeneity is violated. In this section, I relax Assumption HET and instead restrict how much treatment effects are permitted to diverge across latent compliance types. Under this bounded heterogeneity restriction, the causal parameter of interest, λ, is partially identified. I demonstrate how this partial identification framework can be leveraged to conduct sensitivity analyses, allowing researchers to assess what causal conclusions can be robustly drawn under parameterized violations of Assumption HET.[14]

I begin by formalizing the specific bounded heterogeneity assumption that facilitates this partial identification strategy.

Assumption BH (Bounded Heterogeneity)**.** For known constants M_C, $M_{AT} \geq 0$:

$$\left|\kappa_C\right| \leq M_C \quad \text{and} \quad \left|\kappa_{AT}\right| \leq M_{AT}$$

Assumption BH constrains the maximum plausible difference in average treatment effects across latent subgroups using two user-specified sensitivity parameters. Specifically, the researcher sets M_C to restrict the difference in average treatment effects between baseline-treated and baseline-untreated compliers, while setting M_{AT} to restrict the difference between baseline-untreated always-takers and baseline-untreated compliers. I formalize the resulting bounds for λ under Assumption BH in the following proposition.

Proposition 6. *Under Assumptions PT and BH,*

$$\lambda \in \left[\tau - M_{AT} \cdot w_{AT} - M_C \cdot w_C \ , \ \tau + M_{AT} \cdot w_{AT} + M_C \cdot w_C\right]$$

where $w_{AT} = \frac{\mathbb{P}[D_{i0}=0, D_{i1}(1)=D_{i1}(0)=1]}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$ *and* $w_C = \frac{\mathbb{P}[D_{i0}=1, D_{i1}(1)>D_{i1}(0)]}{\mathbb{P}[D_{i1}(1)>D_{i1}(0)]}$.

Proof. See Appendix Section H.

[14] For recent examples of similar approaches to sensitivity analysis and partial identification, see Manski and Pepper (2018), Rambachan and Roth (2023), and Kwon and Sun (2025).

Proposition 6 establishes bounds for λ under the assumption of bounded heterogeneity. If the researcher imposes strict homogeneity by setting $M_{AT} = M_C = 0$, the bounds collapse exactly to the point estimate τ, recovering the result from Section 3.3. Conversely, as $M_{AT}, M_C \to \infty$, the interval expands to the real line, recovering result of Proposition 1 where both selection and heterogeneity are left fully unrestricted. Crucially, the width of these bounds is governed by the user-specified sensitivity parameters (M_{AT} and M_C) and their respective subpopulation weights (w_{AT} and w_C). Because these weights depend on latent compliance transitions, they are not yet identified absent additional assumptions. Before demonstrating how to point-identify these weights using additional pre-period data, I first summarize how applied researchers can leverage these bounds to characterize a "breakdown frontier": the set of values for the sensitivity parameters beyond which a specific causal conclusion breaks down. For instance, an empirical researcher might ask: how severe must the unobserved heterogeneity be before the bounds on λ include zero? I formalize this in the following corollary.

Corollary 2. *Without loss of generality, assume that $\tau \geq 0$. The breakdown frontier, $\mathcal{B}^0$, is characterized by the values of M_{AT} and M_C such that the lower bound in Proposition 6 is equal to zero. This set is given by all pairs of (M_{AT}, M_C) such that $M_{AT}w_{AT} + M_C w_C = \tau$:*

$$\mathcal{B}^0 = \left\{ (M_{AT}, M_C) \;\middle|\; M_{AT}w_{AT} + M_C w_C = \tau \right\}$$

Similarly, the breakdown region, $\mathcal{B}^-$, is characterized by the values of M_{AT} and M_C such that the lower bound in Proposition 6 is strictly negative. This set is given by all pairs of (M_{AT}, M_C) such that $M_{AT}w_{AT} + M_C w_C > \tau$:

$$\mathcal{B}^- = \left\{ (M_{AT}, M_C) \;\middle|\; M_{AT}w_{AT} + M_C w_C > \tau \right\}$$

Corollary 2 formalizes how the bounds in Proposition 6 can be leveraged to evaluate the sensitivity of causal conclusions to varying degrees of unobserved heterogeneity. The breakdown frontier, $\mathcal{B}^0$, represents the level of unobserved heterogeneity where the researcher can no longer conclude that λ is strictly positive. Applied researchers may robustly conclude that λ is positive if every pair (M_{AT}, M_C) on the breakdown frontier $\mathcal{B}^0$ represents an implausibly large degree of treatment effect heterogeneity.[15]

To operationalize the approach described in this Section, it remains to identify the weights w_C and w_{AT}. In Appendix Section I, I show how additional pre-period data can be leveraged to identify these weights under monotonicity and a weak stationarity condition.

[15] In practice, researchers may benchmark the magnitude of these sensitivity parameters against the estimand τ, against the standard deviation of the outcome distribution among units who are continuously untreated, or against existing causal evidence.

4 Partial Identification of λ

In Section 3.2, I established that λ can be point-identified by imposing restrictions on the class of admissible selection mechanisms. However, by setting Δ^{AT} and Δ^{churn}—as defined in Equations 3.1 and 3.2, respectively—to zero, this approach precludes selection into treatment based on time-varying unobservables. Alternatively, Section 3.3 establishes point identification by assuming treatment effect homogeneity across latent compliance types. However, setting κ_C and κ_{AT}—as defined in Equations 3.3 and 3.4, respectively—to zero rules out economic models of selection on gains. Finally, Section 3.5 relaxes these strict conditions via bounded heterogeneity. While this framework provides a valuable tool for sensitivity analysis, the resulting bounds on λ rely on ad-hoc, user-specified sensitivity parameters.

This section introduces a flexible Marginal Treatment Effects (MTE) framework for partial identification (Heckman and Vytlacil, 1999, 2005; Mogstad, Santos and Torgovitsky, 2018). This approach simultaneously relaxes the selection restrictions of Assumption SEL and the homogeneity restrictions of Assumption HET, yielding bounds on λ that are not governed by user-specified sensitivity parameters as required by Assumption BH. This bounding approach is operationalized via a tractable linear programming framework that provides three important advantages over the preceding strategies. First, it allows researchers to flexibly adjust the maintained assumptions without requiring additional identification analysis. For instance, one can readily incorporate proposed relaxations of Assumption PT, as in Rambachan and Roth (2023). Second, it enables researchers to systematically tighten the identified set for λ by incorporating economically motivated shape restrictions as constraints—for example, imposing that units are weakly more likely to select into treatment if they anticipate a larger treatment effect. Finally, this approach permits the partial identification of alternative target parameters beyond λ, such as the Average Treatment Effect on the Treated (ATT), within the same framework.

The remainder of this section proceeds as follows. In Section 4.1, I describe a model where treatment is governed by a weakly separable selection equation. Crucially, I allow for multidimensional unobserved heterogeneity, permitting period-specific unobservables to govern treatment takeup in each period. A key theoretical contribution is demonstrating how the joint distribution of these selection unobservables can be identified under a mild stationarity condition and a restriction on their dependence structure. Section 4.2 introduces the Marginal Treatment Response functions, and formally defines the treatment effect surface and target parameter λ. Section 4.3 then operationalizes this approach by parameterizing the Marginal Treatment Response functions and imposing economically motivated shape restrictions. I conclude by demonstrating how partial identification of λ translates into a standard linear programming problem, yielding sharp, easily computable bounds.

4.1 Identification of Selection Primitives I define a threshold-crossing model of treatment selection characterized by multidimensional, period-specific unobserved heterogeneity. Let (V_{i0}, V_{i1}) denote a vector of latent scalar unobservables, where each V_{it} represents individual i's underlying resistance to taking up the treatment in period $t \in \{0, 1\}$. Let $\nu_t(z)$ denote the structural threshold for treatment selection in period t when the active policy state is $Z_t = z$. Crucially, this formulation imposes that latent resistance V_{it} is invariant to the policy state. It is required that the potential treatment decision for individual i in period t under policy state z is given by the weakly separable selection equation:

$$D_{it}(z) = \mathbb{1}(V_{it} \leq \nu_t(z)) \quad \text{for} \quad t, z \in \{0, 1\}$$

Let F_{V_0} and F_{V_1} denote the marginal cumulative distribution functions of the latent unobservables V_{i0} and V_{i1}, respectively. To map this structural selection model into identified propensity scores, I apply the probability integral transform. Define the normalized latent unobservables as $U_{it} \equiv F_{V_t}(V_{it})$ and the transformed selection thresholds as $p_t(z) \equiv F_{V_t}(\nu_t(z))$.[16] I impose the following regularity conditions on the distribution of these latent variables.

Assumption SRC (Selection Regularity Conditions)**.**

(a) The marginal distribution functions F_{V_0} and F_{V_1} are absolutely continuous and strictly increasing.

(b) The joint distribution of (U_{i0}, U_{i1}) belongs to a specified one-parameter family of absolutely continuous copulas $\{C_\theta : \theta \in \Theta\}$, and for every fixed $(u_0, u_1) \in (0, 1)^2$, the mapping $\theta \mapsto C_\theta(u_0, u_1)$ is strictly monotonic.

Under Assumption SRC(a), the structural selection equation can be equivalently expressed in terms of the normalized variables:

$$D_{it}(z) = \mathbb{1}(U_{it} \leq p_t(z)).$$

Before proceeding, it is instructive to explicitly delineate which components of the selection model are identified thus far. By the probability integral transform, the marginal distributions of the normalized selection unobservables are standard uniform, $U_{it} \sim \text{Unif}[0, 1]$. Consequently, the transformed threshold $p_t(z)$ maps to the unconditional probability of treatment in period t under policy state z, yielding $p_t(z) = \mathbb{P}[D_{it}(z) = 1]$. Because the policy is switched off at baseline ($Z_0 = 0$) and switched on at endline ($Z_1 = 1$), the realized thresholds are point-identified by the observed treatment shares: $p_0(0) = \mathbb{P}[D_{i0} = 1]$ and $p_1(1) = \mathbb{P}[D_{i1} = 1]$. This leaves two selection primitives unidentified: the counterfactual endline threshold, $p_1(0) = \mathbb{P}[D_{i1}(0) = 1]$, and the joint distribution of the unobservables,

[16] I impose Assumption SEL(a) throughout. Within this structural model, this condition requires that $p_t(1) \geq p_t(0)$ for all t.

(U_{i0}, U_{i1}).[17] In the remainder of this subsection, I demonstrate how both of these remaining selection primitives can be recovered using the selection parameters Δ^{AT} and Δ^{churn} defined in Equations 3.1 and 3.2, respectively.[18]

I begin with the identification of $p_1(0)$. Recall that, under within-period monotonicity, the change in the share of always-takers across periods is equal to the difference between the counterfactual endline threshold and the observed baseline threshold: $\Delta^{AT} = \mathbb{P}[D_{i1}(0) = 1] - \mathbb{P}[D_{i0}(0) = 1] = p_1(0) - p_0(0)$. Because the baseline threshold is directly observed in the data ($p_0(0) = \mathbb{P}[D_{i0} = 1]$), the counterfactual endline threshold is point-identified as: $p_1(0) = p_0(0) + \Delta^{AT}$.

Next, I demonstrate how the joint distribution of (U_{i0}, U_{i1}) is point-identified under the copula restriction of Assumption SRC(b). I leverage the counterfactual drop-out rate, Δ^{churn}, which represents the probability of counterfactually selecting out of treatment at $t = 1$, conditional on having selected into treatment at $t = 0$:

$$\Delta^{churn} = 1 - \frac{\mathbb{P}[U_{i0} \leq p_0(0), U_{i1} \leq p_1(0)]}{p_0(0)} \tag{4.1}$$

Because $p_0(0)$ is observed, and $p_1(0)$ and Δ^{churn} are identified, this identifies the joint cumulative distribution function of the unobservables at exactly one interior coordinate, $(p_0(0), p_1(0))$: $\mathbb{P}[U_{i0} \leq p_0(0), U_{i1} \leq p_1(0)]$. Absent further structural restrictions, the dependence structure over the remainder of the support $[0,1] \times [0,1]$ would not generally be point-identified. However, Assumption SRC(b) restricts this joint distribution to a single-parameter copula. Let the joint CDF of the normalized unobservables be denoted by $C_\theta(u_0, u_1)$, where $\theta \in \Theta$ is the unknown scalar dependence parameter. Rearranging Equation 4.1 and plugging in $\mathbb{P}[U_{i0} \leq p_0(0), U_{i1} \leq p_1(0)] = C_\theta(p_0(0), p_1(0))$:

$$C_\theta(p_0(0), p_1(0)) = p_0(0)(1 - \Delta^{churn}).$$

Because $p_0(0)$, $p_1(0)$, and Δ^{churn} are all identified, this constitutes a single equation with a single unknown parameter, θ, within the specified copula family. Crucially, Assumption SRC(b) imposes that for any fixed interior coordinates $(u_0, u_1) \in (0,1)^2$, the copula $C_\theta(u_0, u_1)$ is strictly monotonic in the dependence parameter θ. Because this mapping is strictly monotonic, identifying the joint probability at a single interior point $(p_0(0), p_1(0))$ is sufficient to uniquely identify the dependence parameter θ. Once θ is point-identified, the joint distribution $C_\theta(u_0, u_1)$ is fully characterized on $[0,1] \times [0,1]$.

Remark 6. In settings characterized by discrete variation (here, the variation induced by Z_t) and multidimensional unobserved heterogeneity (here, the period-specific resistances U_{i0}

[17] Under maintained within-period monotonicity, imposing $p_1(0) = p_0(0)$ and $U_{i0} = U_{i1}$ almost surely is sufficient for Assumption SEL, as illustrated in Remark 3.

[18] As demonstrated in Appendix I, Δ^{AT} and Δ^{churn} can be point-identified using additional pre-period data under mild stationarity conditions.

and U_{i1}), point identification of the joint distribution of selection unobservables is generally infeasible (see, e.g., Kamat, Norris and Pecenco 2026). However, in this framework, identifying the selection parameters (Δ^{AT} and Δ^{churn}) together with a specified one-parameter copula family satisfying Assumption SRC(b) is sufficient to fully recover these selection primitives. Standard copula families, including the Gaussian, Clayton, Gumbel, and Frank families (Nelsen, 1999), provide possible specifications. Appendix Section I describes how Δ^{AT} and Δ^{churn} can be identified using additional pre-period data.

4.2 Marginal Treatment Responses and Target Parameters Having fully characterized the joint distribution of the unobservables governing selection, I now map these latent variables to the potential outcomes. I define the two-dimensional Marginal Treatment Response (MTR) surface as the expected potential outcome conditional on the latent selection unobservables, U_{i0} and U_{i1}:

$$m_d^t(u_0, u_1) = \mathbb{E}[Y_{it}(d) \mid U_{i0} = u_0, U_{i1} = u_1].$$

Consequently, the Marginal Treatment Effect (MTE) surface at endline ($t = 1$) is defined by the difference in these MTR functions:

$$m_1^1(u_0, u_1) - m_0^1(u_0, u_1) = \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid U_{i0} = u_0, U_{i1} = u_1].$$

Specific causal parameters of interest can be expressed as distinct weighted averages integrated over this treatment effect surface. For instance, the target parameter λ—the Local Average Treatment Effect for period-1 compliers—is formulated by integrating the difference in these state-specific MTRs over the period-1 complier region, weighted by the copula density:

$$\lambda = \frac{1}{p_1(1) - p_1(0)} \int_{p_1(0)}^{p_1(1)} \int_0^1 \Big(m_1^1(u_0, u_1) - m_0^1(u_0, u_1) \Big) c_\theta(u_0, u_1)\, du_0\, du_1 \tag{4.2}$$

Expressing the target parameter in terms of MTR functions does not itself provide additional identifying power. Even when the selection primitives are identified, λ can remain unbounded without further restrictions on the MTR functions. In the next section, I begin by parameterizing the MTR functions, $m_d^t(u_0, u_1)$, as flexible polynomials in the selection unobservables, thereby imposing smoothness on the MTE surface.

4.3 Characterizing the Identified Set I begin by parameterizing the MTR functions as linear in the period-specific unobservables, noting that this framework naturally extends to more flexible polynomial basis functions in practice. [19]

[19] As noted by Mogstad, Santos and Torgovitsky (2018), Bernstein polynomials provide a flexible parameterization of the MTR surfaces that remains linear in the unknown coefficients. Let $b_{k,K}(u) = \binom{K}{k} u^k (1-u)^{K-k}$. I parameterize each bivariate MTR surface as $m_d^t(u_0, u_1) = \sum_{k=0}^K \sum_{\ell=0}^K \theta_{dt,k\ell} b_{k,K}(u_0) b_{\ell,K}(u_1)$. Because the basis functions are nonnegative and sum to one, constraining the coefficients to $[0, 1]$ ensures that binary-

Let the MTR for treatment status $d \in \{0,1\}$ in period $t \in \{0,1\}$ be defined by the parameter vector $\boldsymbol{\gamma}_d^t \equiv (\gamma_{d,0}^t, \gamma_{d,1}^t, \gamma_{d,2}^t)'$:

$$m_d^t(u_0, u_1) = \gamma_{d,0}^t + \gamma_{d,1}^t u_0 + \gamma_{d,2}^t u_1.$$

Crucially, under this linear specification, the target parameter λ reduces to a simple linear combination of the MTR parameters. By substituting the linear MTR functions into Equation 4.2, λ may be re-written as:

$$\lambda = (\gamma_{1,0}^1 - \gamma_{0,0}^1) + (\gamma_{1,1}^1 - \gamma_{0,1}^1) \cdot \mathbb{E}\Big[U_{i0} \,\Big|\, p_1(0) \leq U_{i1} \leq p_1(1)\Big] + (\gamma_{1,2}^1 - \gamma_{0,2}^1) \cdot \mathbb{E}\Big[U_{i1} \,\Big|\, p_1(0) \leq U_{i1} \leq p_1(1)\Big] \tag{4.3}$$

As demonstrated in Section 4.1, the selection primitives are point-identified, such that the conditional expectations $\mathbb{E}\big[U_{it} \mid p_1(0) \leq U_{i1} \leq p_1(1)\big]$ are known for $t \in \{0,1\}$. Therefore, the target parameter λ is linear in the MTR parameters. To bound λ, it remains to constrain $\boldsymbol{\gamma}_1^1$ and $\boldsymbol{\gamma}_0^1$. To do so, I define a linear program with Equation 4.3 as the objective function. In what follows, I describe the observed moments and economically motivated shape restrictions that constrain the MTR parameters. I then maximize and minimize Equation 4.3 subject to this system of constraints, yielding sharp bounds for λ.

Observed Moment Conditions: The first set of constraints is implied by the observed moment conditions that restrict the expected MTRs. Specifically, units in the population realize one of four possible treatment paths: $(d_0, d_1) \in \{(0,0),(0,1),(1,0),(1,1)\}$. Taking the expectation of the linear MTR function conditional on a specific realized path yields the following equality for each period $t \in \{0,1\}$:

$$\gamma_{d_t,0}^t + \gamma_{d_t,1}^t \mathbb{E}[U_{i0} \mid D_{i0} = d_0, D_{i1} = d_1] + \gamma_{d_t,2}^t \mathbb{E}[U_{i1} \mid D_{i0} = d_0, D_{i1} = d_1] = \mathbb{E}[Y_{it} \mid D_{i0} = d_0, D_{i1} = d_1]$$

Since the joint distribution of unobservables is point-identified via the copula density $c_\theta(u_0, u_1)$, the conditional expectations of U_{i0} and U_{i1} evaluate to known scalar constants. Consequently, matching the structural MTRs to the observed data generates a system of linear equality constraints on the unknown parameters, $\boldsymbol{\gamma}_d^t$.

Parallel Trends: I next turn to the constraints on the MTR parameters imposed by Assumption PT, which requires that

$$\mathbb{E}[m_0^1(U_{i0}, U_{i1}) - m_0^0(U_{i0}, U_{i1}) \mid D_{i0} = d_0, D_{i1} = d_1] = \Delta^{PT}$$

for all treatment paths $(d_0, d_1) \in \{(0,0),(0,1),(1,0),(1,1)\}$. By substituting the parameterized MTR above, the restriction can be equivalently stated as:

outcome MTRs remain in $[0,1]$. Economically motivated shape restrictions can also be expressed as linear constraints on the coefficients, allowing the bounds to be computed by linear programming. In the empirical application in Section 6, I set $K = 4$ in each latent variable.

$$(\gamma^1_{0,0} - \gamma^0_{0,0}) + (\gamma^1_{0,1} - \gamma^0_{0,1})\mathbb{E}[U_{i0} \mid D_{i0} = d_0, D_{i1} = d_1] + (\gamma^1_{0,2} - \gamma^0_{0,2})\mathbb{E}[U_{i1} \mid D_{i0} = d_0, D_{i1} = d_1] = \Delta^{PT}.$$

Shape Restrictions: Next, I consider two frequently considered shape restrictions in the partial identification literature. First, monotone treatment response restrictions impose that the treatment effect is non-negative for all units (Manski, 1997). Motivated by this, I require $m^t_1(u_0, u_1) - m^t_0(u_0, u_1) \geq 0$ everywhere on the unit square. Given the linear parameterization, this restricts the period-specific MTR parameter vectors $\boldsymbol{\gamma}^t_1$ and $\boldsymbol{\gamma}^t_0$ such that:

$$(\gamma^t_{1,0} - \gamma^t_{0,0}) + (\gamma^t_{1,1} - \gamma^t_{0,1})u_0 + (\gamma^t_{1,2} - \gamma^t_{0,2})u_1 \geq 0 \quad \text{for all } (u_0, u_1) \in [0,1]^2.$$

Because the treatment effect surface is affine in u_0 and u_1, this inequality holds globally over the domain if and only if it is satisfied at the four vertices of the unit square. Therefore, this shape restriction reduces to a tractable system of four linear inequalities for each period $t \in \{0, 1\}$:

$$\gamma^t_{1,0} - \gamma^t_{0,0} \geq 0, \quad (\gamma^t_{1,0} - \gamma^t_{0,0}) + (\gamma^t_{1,1} - \gamma^t_{0,1}) \geq 0$$
$$(\gamma^t_{1,0} - \gamma^t_{0,0}) + (\gamma^t_{1,2} - \gamma^t_{0,2}) \geq 0, \quad (\gamma^t_{1,0} - \gamma^t_{0,0}) + (\gamma^t_{1,1} - \gamma^t_{0,1}) + (\gamma^t_{1,2} - \gamma^t_{0,2}) \geq 0$$

Second, I consider a restriction on selection motivated by Manski and Pepper (2000): conditional mean potential outcomes are weakly decreasing in each latent unobservable. Specifically, units with a higher propensity to select into treatment (i.e., those with lower unobserved resistances u_0 and u_1) possess weakly higher potential outcomes, so that for all periods $t \in \{0, 1\}$ and treatment states $d \in \{0, 1\}$:

$$\frac{\partial m^t_d(u_0, u_1)}{\partial u_0} \leq 0 \iff \gamma^t_{d,1} \leq 0$$
$$\frac{\partial m^t_d(u_0, u_1)}{\partial u_1} \leq 0 \iff \gamma^t_{d,2} \leq 0$$

With the objective function and the full system of constraints now defined as linear in the MTR parameters, bounding λ reduces to a standard linear programming problem. Specifically, I calculate the sharp lower and upper bounds, $[\lambda_{\min}, \lambda_{\max}]$, relative to the specified MTR class and the maintained moment and shape restrictions, by minimizing and maximizing Equation 4.3 over the resulting feasible set.

5 What Do Exposure Designs Identify under Universal Policy Changes?

The preceding sections establish how restrictions on selection and treatment-effect heterogeneity identify or bound the LATE under universal policy exposure. In this section, I use this framework to examine a common empirical alternative. When a policy is implemented simultaneously for all units, researchers often leverage cross-sectional heterogeneity

in how exposed units are to the policy change in order to form treatment and comparison groups. Exposure—or policy "bite"—is generally based on the policy's expected effect on treatment take-up across groups. Revisiting Example 2.1, a researcher might conjecture that a reform reducing the administrative costs of welfare enrollment disproportionately affects individuals who were not enrolled at baseline compared to those who were already enrolled. Consequently, less exposed units—those already enrolled in the welfare program—are leveraged as a control group, while more affected units—those not enrolled at baseline—serve as the treatment group.[20] An exposure design then compares changes in outcomes between the two groups. This section characterizes the resulting estimand and establishes the conditions under which it admits a causal interpretation.

This approach relies on the following difference-in-differences estimand:

$$\beta = \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 1\Big]. \tag{5.1}$$

It is instructive to contrast the variation leveraged by the exposure design estimand, β, in Equation 5.1 with the estimand, τ, proposed in Section 3:

$$\tau = \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big].$$

β relies entirely on cross-sectional variation in baseline treatment status, comparing units who were untreated at baseline to those who were treated at baseline. It does not leverage any intertemporal variation in treatment take-up. In contrast, τ compares baseline-untreated units who subsequently select into treatment to those who remain continuously untreated. By excluding all baseline-treated units from the estimand, τ eliminates cross-sectional variation in baseline exposure, relying solely on the policy-induced variation in treatment take-up.

Next, I examine what the estimand β identifies under Assumption PT. I proceed term by term, beginning with the first term in Equation 5.1. The following lemma establishes that the expected change in outcomes among the baseline-untreated units captures the sum of the aggregate counterfactual trend, Δ^{PT}, and the average treatment effect for units with the realized treatment path $\{D_{i0} = 0, D_{i1} = 1\}$, i.e., switchers who select into treatment after the policy is implemented but remained untreated at baseline.

Lemma 2. *Under Assumption PT:*

$$\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0] = \Delta^{PT} + \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \cdot \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0].$$

Proof. See Appendix Section J.

Next, I examine what the second term in Equation 5.1, $\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1]$, isolates under Assumption PT. The motivation for subtracting this term from the expected change

[20] For an illustrative, non-exhaustive selection of recent applications of this exposure design, see Argys et al. (2020); Deza et al. (2024); Cholli and Wu (2026).

in outcomes among the baseline-untreated units is to difference out the confounding aggregate time trend, Δ^{PT}, ensuring that the variation identified in Lemma 2 captures a causal effect of the policy. However, if this moment captures unobserved dynamics beyond the aggregate time trend, it may induce bias in the identified parameter. In the following lemma, I demonstrate that the expected change in outcomes for the baseline-treated units captures the sum of the counterfactual trend and the intertemporal change in treatment effects for the continuously treated subpopulation. For ease of exposition, I additionally assume that $\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1] = 0$, though Appendix Section K provides a general decomposition without this restriction.

Lemma 3. *Under Assumption PT and* $\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1] = 0$*:*

$$\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] = \Delta^{PT} + \mathbb{E}\left[\Big(Y_{i1}(1) - Y_{i1}(0)\Big) - \Big(Y_{i0}(1) - Y_{i0}(0)\Big) \mid D_{i0} = 1, D_{i1} = 1\right].$$

Proof. See Appendix Section K.

Combining the results from Lemma 2 and Lemma 3, the exposure design estimand, β, can be decomposed as follows:

$$\beta = \underbrace{\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \cdot \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0]}_{\text{Causal Effect among Switchers}}$$
$$- \underbrace{\mathbb{E}\left[\Big(Y_{i1}(1) - Y_{i1}(0)\Big) - \Big(Y_{i0}(1) - Y_{i0}(0)\Big) \mid D_{i0} = 1, D_{i1} = 1\right]}_{\text{Bias from Time-Varying Treatment Effects}}.$$

This decomposition highlights that β identifies the causal effect of the policy for the switchers, minus a bias term driven by time-varying treatment effects of the continuously treated subpopulation. Importantly, because this intertemporal change in treatment effects may take the opposite sign of the true causal effect of the treatment, β can be theoretically mis-signed relative to the target parameter absent additional restrictions.[21] Taken together, Lemmas 2 and 3 imply that time-invariance of treatment effects is a necessary condition for the exposure design estimand to identify a (weakly) causal effect of the policy.

[21] Appendix Section M.1 characterizes the magnitude of bias required for the exposure-design estimand, β, to have the opposite sign from the period-1 LATE, λ. Under Assumption PT and homogeneous post-period treatment effects, sign reversal occurs when the average change in treatment effects among baseline-treated units, normalized by λ, exceeds the differential change in treatment take-up between the two baseline groups. For example, suppose that 10 percent of baseline-untreated units enter treatment, 5 percent of baseline-treated units exit, and $\lambda = 1$. The differential change in take-up is then 0.15. If the average treatment effect among baseline-treated units is 0.20 larger in period 1 than in period 0, then $\beta = 1(0.10+0.05)-0.20 = -0.05$. The exposure-design estimand is therefore negative even though the post-period treatment effect is positive for every unit.

Having established that the time-invariance of treatment effects is a necessary condition for identification, a natural subsequent question arises: exactly which causal parameter does β recover when this restriction holds? I formalize this result in the following proposition.

Proposition 7. *Suppose that the effect of the treatment is strictly time-invariant among baseline-treated units:* $\mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1\Big] = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1\Big]$. *Under Assumption PT:*

$$\begin{aligned}\beta = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \cdot \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] \\ + \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1, D_{i1} = 0\Big] \cdot \mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1].\end{aligned}$$

Proof. See Appendix Section L.

The parameter identified in Proposition 7 is challenging to interpret because it is a non-convex weighted average of treatment effects across distinct latent compliance types. The first term identifies the average treatment effect for units with the realized path $\{D_{i0} = 0, D_{i1} = 1\}$, who enter treatment after the policy is implemented. Absent additional restrictions, this subgroup is a mixture of period-1 compliers and always-takers. The second term identifies the average treatment effect for units with the realized path $\{D_{i0} = 1, D_{i1} = 0\}$, who exit treatment at endline despite being treated at baseline. Under monotonicity, this subgroup consists of units who are always-takers at $t = 0$ but become never-takers at $t = 1$. Therefore, without further restrictions on the evolution of compliance types, β generally does not correspond to the average treatment effect for a well-defined latent subpopulation.

Remark 7. Applied researchers frequently leverage baseline exposure as an instrument for the effective "bite" of the policy under a linear outcome model, appealing to the identification results in Goldsmith-Pinkham, Sorkin and Swift (2020) (for a recent application, see Finkelstein, Notowidigdo and Shi, 2026). In Appendix Section M.2, I relate the identification arguments developed in this paper to the exposure designs interpretation of Bartik instruments in Goldsmith-Pinkham, Sorkin and Swift (2020) and demonstrate that baseline exposure serves as a valid instrument if and only if treatment effects are time-invariant. ■

Remark 8. If one additionally assumes that the causal effect is homogeneous across switchers, such that $\mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1] = \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1, D_{i1} = 0]$, then β identifies the average treatment effect among switchers scaled by the differential change in treatment take-up. Specifically:

$$\begin{aligned}\beta &= \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} \neq D_{i1}] \cdot \Big(\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] + \mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1]\Big) \\ &= \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} \neq D_{i1}] \cdot \Big(\mathbb{E}[\Delta D_i \mid D_{i0} = 0] - \mathbb{E}[\Delta D_i \mid D_{i0} = 1]\Big),\end{aligned}$$

where $\Delta D_i := D_{i1} - D_{i0}$. The term $\mathbb{E}[\Delta D_i \mid D_{i0} = 0] - \mathbb{E}[\Delta D_i \mid D_{i0} = 1]$ captures how the "bite" of the policy is differential across baseline take-up. Consequently, dividing β

by the differential change in take-up yields the familiar Wald-DiD estimand considered by De Chaisemartin and d'Haultfoeuille (2018). ■

Remark 9. A practical advantage of exposure designs is that they can be implemented using pooled panel data and a continuous measure of baseline exposure, whereas estimating τ requires unit-level panel data. However, Sun and Shapiro (2022) show that, even within a linear panel model, exposure designs need not recover even a weighted average of heterogeneous unit-level effects. On the other hand, panel microdata on unit-level treatment paths permit researchers to accommodate treatment-effect heterogeneity, relaxing the strong homogeneity assumptions required to interpret conventional exposure designs. ■

6 Applications

This section illustrates the usefulness of the proposed framework in two real-world settings. Section 6.1 conducts a falsification exercise using the Oregon Health Insurance Experiment, where random assignment provides an independent experimental benchmark against which to compare the LATE estimated using the strategy proposed in this paper. This exercise is informative about the plausibility of the identifying assumptions proposed in this paper: the experimental benchmark allows me to assess whether those assumptions jointly recover the experimental LATE. The two estimates need not align because treatment take-up remains endogenous in the nonexperimental analysis. If individuals who take up treatment differ from those who do not in their counterfactual outcome trends or treatment gains, the panel estimand will generally differ from the experimental LATE. I then apply the method to Medicare Part D in Section 6.2, a universal reform to which the entire eligible population was exposed and for which no natural, credibly unaffected control group exists.

6.1 Oregon Health Insurance Experiment In 2008, Oregon randomly gave a group of low-income uninsured adults the opportunity to apply for Medicaid. Taubman et al. (2014) leverage this random variation to estimate the effect of Medicaid coverage on emergency department (ED) use. Medicaid enrollment nevertheless remained an individual choice. Some lottery winners did not enroll, while 15% of the control group obtained coverage through other channels. Thus, among lottery winners, those insured at endline include both individuals induced to enroll by the lottery and individuals who would have obtained coverage without it.

Within the treatment arm, the lottery operated as a universal shock that reduced the cost of obtaining Medicaid coverage for everyone selected. I therefore restrict the analysis to lottery winners, reproducing a universal rollout in which the policy reduces the cost of treatment take-up but take-up itself remains endogenous. Because all respondents were uninsured at baseline, $D_{i0} = 0$ holds mechanically. I compare the change in an indicator for any ED use between baseline (Y_{i0}) and endline (Y_{i1}) among individuals who obtained

coverage ($D_{i1} = 1$) with the corresponding change among those who remained uninsured ($D_{i1} = 0$). Specifically, conditioning on prior ED utilization, I estimate the following:

$$\tau = \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big].$$

As shown in Section 3.3, τ identifies the LATE under Assumptions PT and HET. In this setting, Assumption PT requires that individuals who obtained coverage and those who remained uninsured would have experienced the same average change in ED use without Medicaid. A sufficient condition is that this average change does not vary with the factors driving enrollment (Ghanem, Sant'Anna and Wüthrich, 2026*a*). Obtaining coverage through the lottery required completing an application and documenting eligibility (Taubman et al., 2014), suggesting that differences in take-up may partly reflect enrollment costs rather than changes in healthcare needs. Assumption HET requires that, conditional on prior ED utilization, Medicaid have the same average effect for individuals induced to enroll by the lottery and those who would have obtained coverage through other channels. Although Medicaid's effect on ED visits may vary across latent compliance types, Kowalski (2023) shows that prior ED utilization explains nearly all of the estimated heterogeneity. After conditioning on prior ED use, the estimated marginal treatment effect curve is approximately flat, providing support for Assumption HET in this setting.

Figure 2 presents the results. Taubman et al. (2014) estimate that Medicaid increased the probability of an ED visit by 7.0 percentage points among individuals induced to enroll by the lottery. The estimate obtained using τ is 7.3 percentage points, closely matching the experimental benchmark. For the number of ED visits, τ yields an increase of 0.39 visits, compared with 0.40 using lottery assignment.[22] This agreement is not mechanical: τ uses only lottery winners and could differ from the experimental LATE if those who obtain coverage would have experienced different untreated outcome trends from those who remain uninsured, or if compliers and always-takers experience different treatment effects. The results provide evidence that the identifying restrictions are plausible in this setting, and illustrate how panel variation in endogenous treatment take-up can reproduce experimental benchmarks without a comparison group unexposed to the policy.

6.2 Introduction of Medicare Part D Introduced in 2006, the Medicare Part D program provided subsidized access to prescription drug coverage for all US adults aged 65 and older. The policy allowed eligible individuals to choose among several subsidized plans, universally reducing the out-of-pocket cost of obtaining prescription drug insurance. Prior to the reform, approximately 80% of elderly US adults already held some form of prescription drug coverage. Following the introduction of Part D, coverage increased but approximately 2.6

[22]Random assignment implies that compliers among lottery winners have the same average treatment effect as compliers in the full experimental population. Under Assumptions PT and HET, the two strategies identify the same LATE.

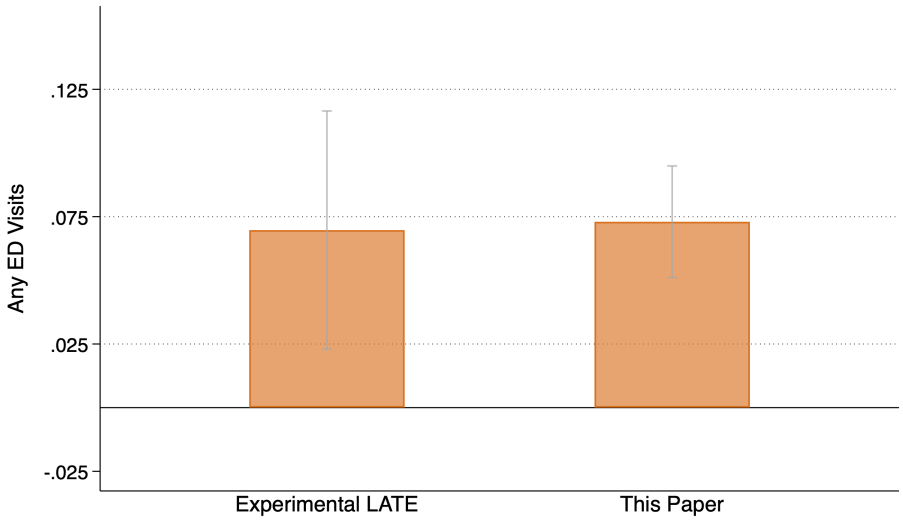

(a) Any ED Visit

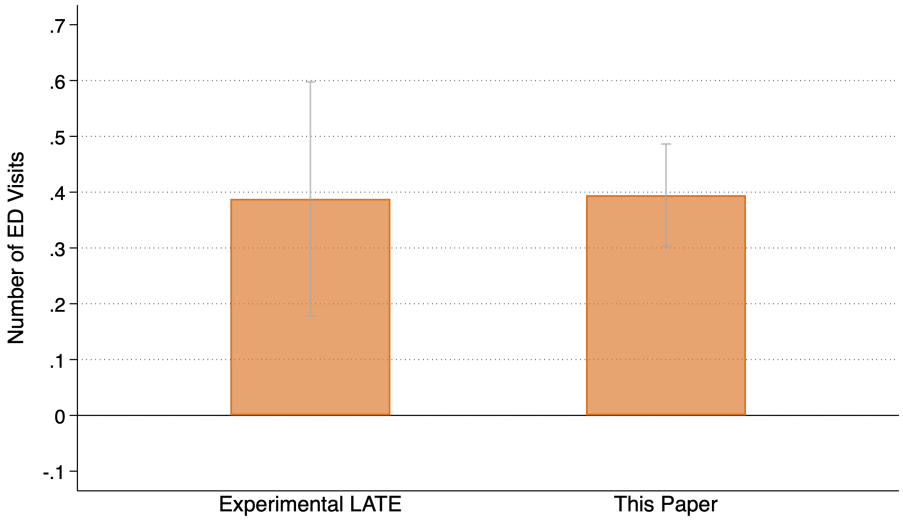

(b) Number of ED Visits

FIGURE 2. Effect of Medicaid on Emergency-Department Use

Notes. Panels (a) and (b) report estimated effects of Medicaid coverage on the probability of any ED visit and the number of ED visits, respectively. Experimental estimates instrument Medicaid coverage with lottery assignment. The proposed estimator, τ, compares changes in ED use between lottery winners who enrolled in Medicaid during the study period and those who did not. Both specifications include household lottery-list size indicators; the experimental specification also controls for the baseline outcome. The 95% confidence intervals are plotted; standard errors are clustered by household.

million individuals (or 7.5% of the eligible population) remained uninsured, in part because the residual premium costs remained too high (Heiss, McFadden and Winter, 2006).

A prominent literature leverages the Medicare eligibility threshold to estimate the effects of Part D on financial protection, health, and labor supply by comparing individuals aged 65 and older with those just below the age threshold (Engelhardt and Gruber, 2011; Huh and Reif, 2017; Wettstein, 2020). Using this design, Wettstein (2020) finds evidence that Part D reduced labor supply by alleviating "retirement lock." Such designs, however, require Part D to have had no effect on the younger comparison group. This assumption is difficult to defend. Alpert, Lakdawalla and Sood (2023) show that Part D increased prescription-drug utilization among the nonelderly through increased pharmaceutical advertising, while Card, Dobkin and Maestas (2009) and Diamond et al. (2018) provide evidence that individuals time health-care use around anticipated insurance coverage, implying that individuals approaching age 65 need not constitute an unaffected control group. The introduction of Part D is therefore more naturally viewed as a universal policy that reduced the cost of obtaining prescription-drug coverage for the entire eligible population and thereby induced coverage take-up among the previously uninsured. Using only individuals who were already Medicare-eligible before the reform, I identify the effect of prescription-drug coverage among those induced to obtain it by Part D, without relying on younger individuals as an unaffected control group.

Using the biennial waves of the Health and Retirement Study from 2002 through 2010 (Health and Retirement Study, 2025), I estimate the effects of prescription-drug coverage induced by Medicare Part D on medication use and out-of-pocket costs, self-reported health, and labor supply.[23] I represent the introduction of Part D by $Z_t = \mathbb{1}\{t \geq 2006\}$ and define D_{it} as an indicator for prescription-drug coverage.

Proposition 2 identifies the first stage effect of Part D under two restrictions. First, monotonicity requires that the reform weakly increased each individual's likelihood of obtaining prescription-drug coverage. This assumption is mild because Part D reduced the cost of obtaining coverage, making it unlikely that anyone eligible for the subsidy would lose coverage because of the reform. Second, the share of always-takers must remain constant over time. Under monotonicity, the observed pre-reform coverage rate identifies this share. Figure 3 supports this restriction: coverage remained nearly unchanged between 2002 and 2004. Under these two assumptions, Proposition 2 shows that the first stage can be estimated by subtracting the 2004 coverage rate from the observed 2006 coverage rate. Therefore, Part D causally increased prescription-drug coverage by 16.6 percentage points, from 75.6% to 92.2%.[24]

[23]To avoid changes in sample composition from individuals aging into Medicare eligibility, I restrict the sample to respondents who were at least 65 years old in 2002.

[24]The 2006 coverage rate implies an uninsured rate of 7.8%, closely matching the 7.5% reported by Heiss, McFadden and Winter (2006).

FIGURE 3. Effect of Medicare Part D on Prescription-Drug Coverage

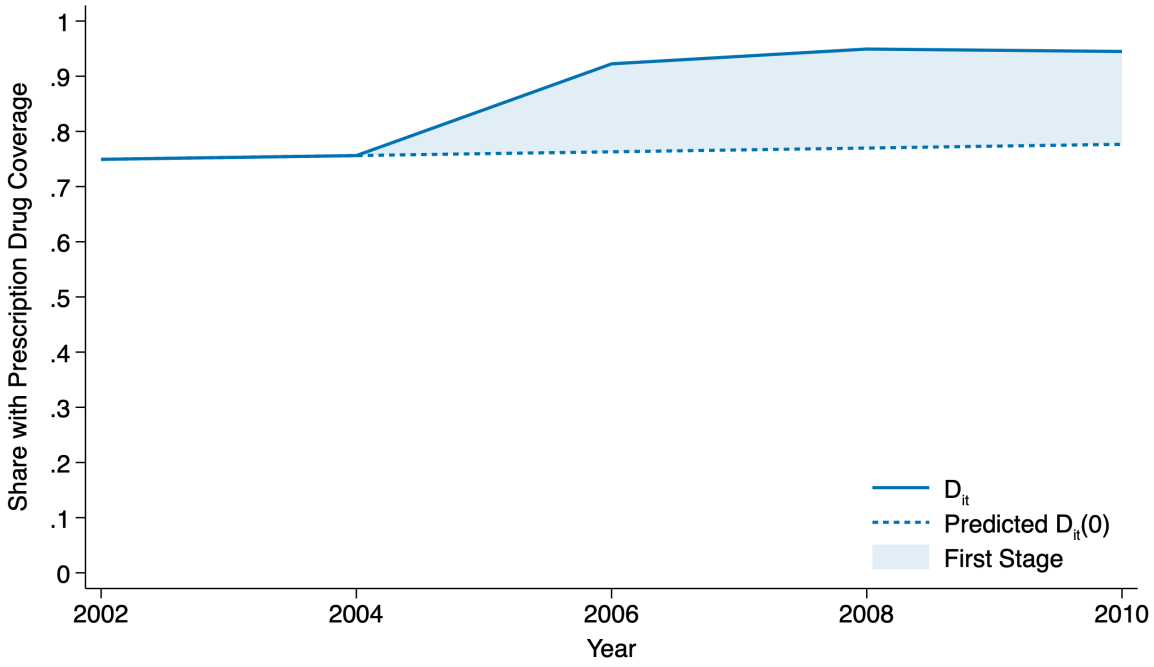

Notes. This figure plots prescription-drug insurance coverage before and after the introduction of Medicare Part D. The solid line reports the observed share of respondents with prescription-drug coverage, D_{it}, in each Health and Retirement Study wave. A respondent is classified as covered if they report that prescription drugs are covered by any health-insurance plan. Specifically, D_{it} measures any prescription-drug insurance rather than Medicare Part D enrollment. The dashed line reports predicted coverage in the absence of Medicare Part D, $D_{it}(0)$, obtained by fitting a linear trend to coverage in the 2002 and 2004 pre-policy waves and extrapolating it through the post-policy period. The shaded area between observed and predicted coverage represents the estimated first stage effect of Medicare Part D on prescription-drug insurance coverage, following the result in Proposition 2. The sample consists of respondents who were at least 65 years old in 2002.

Having established that Part D increased prescription-drug coverage, I next examine the effects of that coverage on outcomes. For each post-policy wave t, I estimate τ_t by comparing changes since 2004 between initially uninsured respondents who obtained coverage in wave t and those who remained uninsured in wave t:

$$\tau_t = \mathbb{E}\Big[Y_{i,t} - Y_{i,2004} \mid D_{i,s\leq 2004},\ D_{i,t} = 1\Big] - \mathbb{E}\Big[Y_{i,t} - Y_{i,2004} \mid D_{i,s\leq 2004},\ D_{i,t} = 0\Big],$$

where s indexes the pre-reform survey waves. Under Assumption PT, together with either Assumption SEL or Assumption HET, τ_t identifies the period-t effect of coverage among respondents induced to obtain it by Part D, λ_t.[25] I first assess the plausibility of parallel trends. Figure 4 presents event-study estimates for cost-related prescription-drug outcomes.[26] The 2002 estimates are close to zero and statistically insignificant, providing no evidence of

[25] Appendix Section N formalizes this result for $t \geq 2006$ and formalizes how to test the plausibility of parallel trends using pre-policy outcomes and subsequent coverage status.

[26] The 2002 placebo estimates compare pre-policy outcome changes between respondents who subsequently obtained coverage in 2006 and those who remained uninsured.

differential pre-policy trends between the two groups. The post-policy estimates are negative for both outcomes. By 2008, prescription-drug coverage reduced the probability of avoiding medication because of its cost by approximately 5 percentage points and the probability of incurring any out-of-pocket prescription cost by approximately 7 percentage points.

FIGURE 4. Effects of Prescription-Drug Insurance on Medication Access and Cost

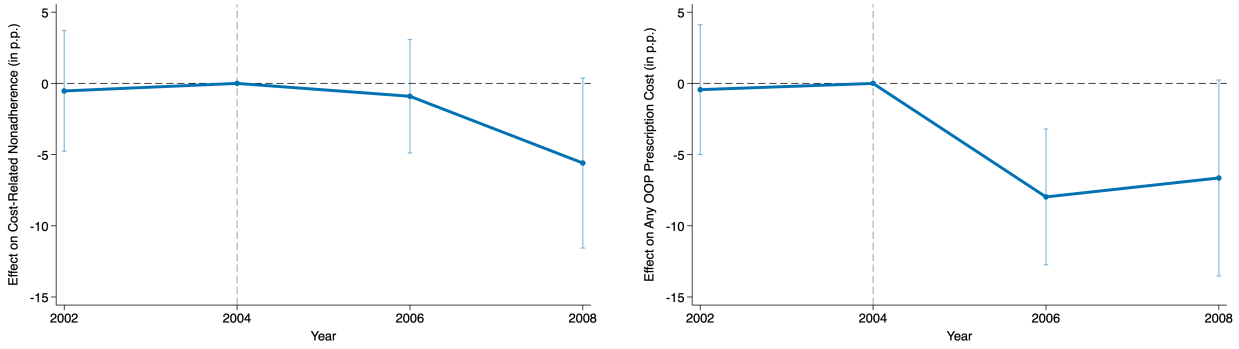

(a) Avoided Medication Due to Cost (b) Any OOP Prescription Cost

Notes. This figure reports wave-specific estimates of the effect of prescription-drug insurance coverage on cost-related outcomes. The sample is restricted to respondents who were at least 65 years old in 2002 and reported no prescription-drug coverage in either 2002 or 2004. The 2004 coefficient is normalized to zero. The 2002 coefficient compares the change in the outcome from 2002 to 2004 across respondents who subsequently obtained coverage in 2006 and those who remained uninsured. Each post-policy coefficient corresponds to τ_t, defined in Section 6.2. Panel (a) uses an indicator equal to one if the respondent reports not taking prescribed medication because of its cost. Panel (b) uses an indicator equal to one if the respondent reports positive monthly out-of-pocket spending on prescription drugs. Estimates are reported in percentage points. Vertical bars show 95 percent confidence intervals constructed using standard errors clustered at the respondent level.

For τ_t to point-identify the period-t LATE, I additionally require that either Assumption SEL or Assumption HET holds. Assumption SEL(c) requires respondents covered before the reform to have remained covered in its absence. This restriction implies $\Pr(D_{i,2004} = 0 \mid D_{i,2002} = 1) = 0$. In the data, however, 9.0% of respondents covered in 2002 lacked coverage in 2004, providing evidence against Assumption SEL. Point identification can instead be obtained by imposing Assumption HET, which requires treatment effects to be equal across post-policy compliers and always-takers, regardless of their baseline coverage status. Table I reports the average of τ_t across the post-policy waves. Under Assumption HET, each reported coefficient identifies the effect of prescription-drug coverage across post-policy waves among respondents induced to obtain coverage by Part D. Coverage reduced the probability of not taking prescribed medication because of its cost by 3.1 percentage points and the probability of any out-of-pocket prescription spending by 7.8 percentage points,declines of 26% and 10% relative to their 2004 means. Coverage also reduced the probability of working by 2.9 percentage points, or 15% relative to its 2004 mean. Together, under Assumption

HET, these estimates suggest that prescription-drug coverage induced by Part D improved access to medication, provided meaningful financial protection, and eased retirement lock.

TABLE I. Effects of Prescription-Drug Insurance on Outcomes

	(1)	(2)	(3)	(4)
	Avoided Medication Due to Cost	Any OOP Prescription Cost	Self-Reported Excellent Health	Working Part-Time or Full-Time
Prescription-Drug Coverage	-0.031	-0.078	0.022	-0.029
	(0.019)	(0.022)	(0.021)	(0.017)
Outcome Mean	0.121	0.793	0.112	0.193
Observations	1,217	1,216	1,216	1,218

Notes: This table reports estimates of the effect of prescription-drug insurance coverage on the indicated outcomes following the introduction of Medicare Part D. The analysis uses the 2002, 2004, 2006, 2008, and 2010 waves of the Health and Retirement Study. The sample is restricted to respondents who were at least 65 years old in 2002 and reported no prescription-drug coverage in either 2002 or 2004. Each reported coefficient is the average of the wave-specific estimands τ_t, defined in Section 6.2, across the 2006, 2008, and 2010 post-policy survey waves. *Avoided Medication Due to Cost*, reported in Column 1, is an indicator equal to one if the respondent reports not taking prescribed medication because of its cost. *Any OOP Prescription Cost*, reported in Column 2, is an indicator equal to one if the respondent reports positive monthly out-of-pocket spending on prescription drugs. *Self-Reported Excellent Health*, reported in Column 3, is an indicator equal to one if the respondent rates their health as excellent. *Working Part-Time or Full-Time*, reported in Column 4, is an indicator equal to one if the respondent reports working either part-time or full-time. Standard errors are clustered at the respondent level. The outcome mean is the 2004 mean of the corresponding outcome among respondents in the analysis sample for whom that outcome is observed. Observations corresponds to the number of distinct respondents in the sample.

Assumption HET may be implausible if individuals select into coverage partly according to their expected gains: always-takers may have systematically different prescription-drug needs than those induced to enroll by Part D. I therefore apply the sensitivity analysis developed in Section 3.5, replacing exact homogeneity with bounded heterogeneity. The parameter M_{AT} bounds the difference in average effects between baseline-uninsured always-takers and baseline-uninsured compliers, while M_C bounds the difference between baseline-insured and baseline-uninsured compliers. I set $M_{AT} = M_C = M^*$, thereby allowing both effect differences to be as large as the same common magnitude M^*. Table II reports the largest M^* for which the bounds on the LATE exclude zero. Across outcomes, M^* ranges from 3.2 to 11.2 percentage points. These values are 43–46% larger than the corresponding point estimates and equal 14–37% of the 2004 outcome means. Thus, the signs of the estimated LATEs are robust to departures from homogeneity: treatment effects would need to differ substantially across the relevant latent groups for the bounds to include zero.

TABLE II. Sensitivity Analysis

	(1)	(2)	(3)	(4)
	$\widehat{\tau}$	$\overline{Y}_{2004}$	M^*	$M^*/\overline{Y}_{2004}$ (%)
Avoided Medication Due to Cost	−0.031	0.121	0.0448	37.0
Any OOP Prescription Cost	−0.078	0.793	0.1118	14.1
Self-Reported Excellent Health	0.022	0.112	0.0320	28.6
Working Part-Time or Full-Time	−0.029	0.193	0.0419	21.7

Notes. This table evaluates the sensitivity of the pooled estimates reported in Table I to treatment-effect heterogeneity. For each outcome, M^* is the smallest value $M^* = M_{AT} = M_C$ for which the bounds in Proposition 6 include zero. The parameter M_{AT} bounds the difference between the average treatment effects for baseline-untreated always-takers and baseline-untreated compliers, while M_C bounds the difference between the average treatment effects for baseline-treated and baseline-untreated compliers. Because all outcomes are binary, M^* is measured in percentage points; for example, $M^* = 0.0448$ represents a maximum permitted difference of 4.48 percentage points across the relevant latent groups. Column 4 expresses M^* as a percentage of the corresponding 2004 outcome mean.

The sensitivity analysis permits departures from homogeneity, but constructing the resulting bounds requires the researcher to specify the maximum permissible degree of heterogeneity. To instead discipline treatment-effect heterogeneity using economically motivated shape restrictions, I implement the MTE approach developed in Section 4 and impose two restrictions directly on the treatment-effect surface. First, coverage weakly reduces both cost-related outcomes and the probability of working, while weakly increasing excellent health. Second, selection on gains requires individuals more likely to obtain coverage to experience weakly larger gains from it.[27] Figure 5 reports the resulting bounds. Under these restrictions, coverage reduced the probability of any out-of-pocket prescription spending by between 4.8 and 14.4 percentage points, increased the probability of reporting excellent health by between 0.5 and 3.1 percentage points, and reduced the probability of working by between 0.3 and 7.7 percentage points. Thus, even after allowing treatment effects to vary systematically with the demand for insurance, the results indicate that coverage induced by Part D provided meaningful financial protection and eased retirement lock.

[27]For each outcome, I use respondents observed in both 2004 and 2006 and parameterize the four period-treatment MTR surfaces using fourth-degree tensor-product Bernstein polynomials in the latent resistance variables (U_{i0}, U_{i1}). I model the dependence between the two latent resistance variables using a Gaussian copula. I constrain the MTR surfaces to respect binary outcome support, match the observed outcome moments, and satisfy parallel trends and the stated shape restrictions. I then use linear programming to minimize and maximize the 2006 LATE over all admissible MTR surfaces. The selection parameters are estimated from the 2002 and 2004 waves, following the approach described in Appendix Section I.

FIGURE 5. Identified Set for the Effects of Prescription-Drug Insurance

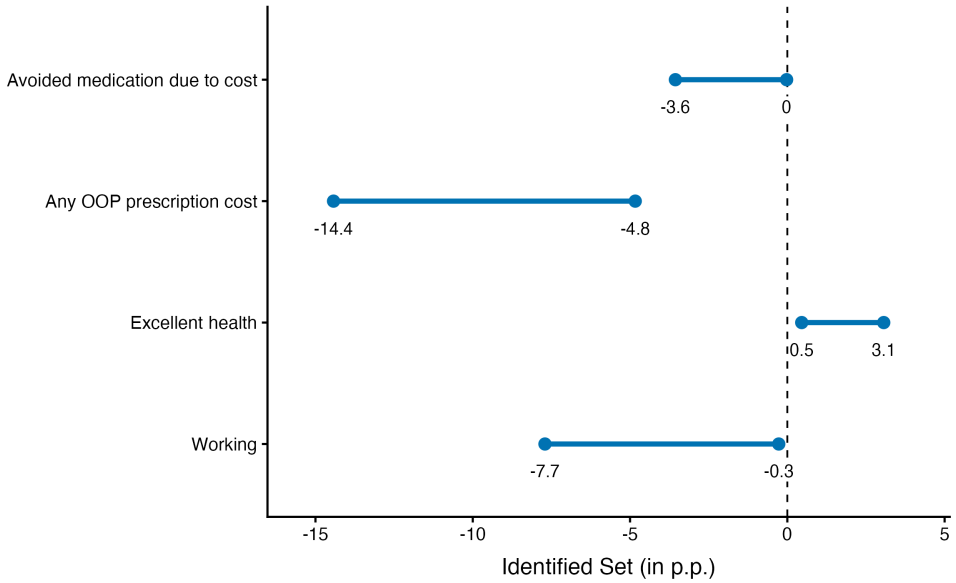

Notes. This figure reports identified sets for the effect of prescription-drug insurance among respondents induced to obtain coverage by the introduction of Medicare Part D. Each horizontal segment reports the identified set from the MTE analysis. These specifications impose parallel trends, binary-outcome support, the monotone treatment response, and selection on gains. The analysis uses a balanced panel from the 2004 and 2006 Health and Retirement Study waves. The selection parameters Δ^{AT} and Δ^{churn} are estimated using the 2002 and 2004 waves.

Finally, I implement an exposure design (discussed in Section 5) that compares respondents uninsured in both 2002 and 2004 to those insured in both years, scaling this contrast by their differential change in coverage. This comparison requires the average effect of prescription-drug coverage among the baseline-insured to remain constant over time. This restriction may be implausible in an aging population: changing medication needs can alter how much coverage improves medication access and health, even among respondents who remain continuously insured. The exposure design subtracts these changing gains along with the common untreated trend, potentially biasing the resulting estimate. The exposure design estimates effects of -4.1 percentage points on avoiding medication because of cost (95% CI: $[-6.4, -1.8]$), -1.4 percentage points on excellent health (95% CI: $[-3.5, 0.7]$), and -0.2 percentage points on working (95% CI: $[-2.3, 2.0]$). To assess their plausibility, I benchmark these estimates against the MTE bounds estimated in Figure 5; these bounds contain the values of the LATE consistent with the observed moments and maintained shape restrictions. Each point estimate lies outside the corresponding MTE bounds, suggesting that exposure designs may yield biased estimates of the LATE in empirical settings of interest.

7 Conclusion

This paper develops a unified econometric framework for identifying causal effects under universal policy implementation. Its central contribution is to establish that universal exposure to a policy change need not preclude identification and to characterize the additional assumptions required to identify causal effects of interest. In particular, I characterize when unit-level variation in treatment take-up can identify causal effects even though policy exposure provides no cross-sectional identifying variation. The results also provide practical guidance for applied researchers. Researchers can identify the LATE by restricting how selection behavior evolves over time or, alternatively, by restricting treatment-effect heterogeneity across latent groups. When neither approach is defensible, the marginal treatment effects framework provides a disciplined partial-identification alternative. Researchers using exposure designs must justify time-invariant treatment effects; without this restriction, the resulting estimand need not identify a valid causal effect. A LATE interpretation further requires restrictions such as homogeneous average treatment effects across the relevant subgroups.

A limitation of the point-identification results is their reliance on parallel trends. Under the considered selection restrictions, observed treatment paths identify compliers, but recovering their mean untreated post-policy outcome requires additional assumptions. Parallel trends provides one route to recovering this counterfactual; exploring alternative restrictions on untreated outcomes is a direction for future research. Because the framework relies on observing unit-level treatment sequences, another direction for future research is to characterize what can be identified by combining treatment histories observed for a subsample of units with pooled panel data covering a broader population.

References

Aaronson, Daniel, and Brian J Phelan. 2017. "Wage shocks and the technological substitution of low-wage jobs." *The Economic Journal*, 127(599): 1–34. 10

Abadie, Alberto. 2005. "Semiparametric Difference-in-Differences Estimators." *The Review of Economic Studies*, 72(1): 1–19. 1

Alpert, Abby, Darius Lakdawalla, and Neeraj Sood. 2023. "Prescription Drug Advertising and Drug Utilization: The Role of Medicare Part D." *Journal of Public Economics*, 221: 104860. 1, 6.2

Altonji, Joseph G., and Rosa L. Matzkin. 2005. "Cross Section and Panel Data Estimators for Nonseparable Models with Endogenous Regressors." *Econometrica*, 73(4): 1053–1102. 7

Angrist, Joshua, Bruno Ferman, Carol Gao, Peter Hull, Otavio L. Tecchio, and Robert W. Yeh. 2024. "Instrumental Variables with Time-Varying Exposure: New Estimates of Revascularization Effects on Quality of Life." National Bureau of Economic Research Working Paper 33296. Revised July 2026. 1

Argys, Laura M, Andrew I Friedson, M Melinda Pitts, and D Sebastian Tello-Trillo. 2020. "Losing public health insurance: TennCare reform and personal financial distress." *Journal of Public Economics*, 187: 104202. 9, 20

Ashenfelter, Orley, and David Card. 1985. "Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs." *The Review of Economics and Statistics*, 67(4): 648–660. 1

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2025. "Difference-in-Differences with a Continuous Treatment." 1

Card, David, Carlos Dobkin, and Nicole Maestas. 2009. "Does Medicare save lives?" *The quarterly journal of economics*, 124(2): 597–636. 6.2

Chamberlain, Gary. 1982. "Multivariate Regression Models for Panel Data." *Journal of Econometrics*, 18(1): 5–46. 7

Chernozhukov, Victor, Iván Fernández-Val, Jinyong Hahn, and Whitney Newey. 2013. "Average and quantile effects in nonseparable panel models." *Econometrica*, 81(2): 535–580. 7

Cholli, Neil, and Derek Wu. 2026. "Multiple program participation in the safety net: Incidence, impediments and implications." Working Paper. 1, 2.1, 20

Clemens, Michael A, Ethan G Lewis, and Hannah M Postel. 2018. "Immigration restrictions as active labor market policy: Evidence from the Mexican bracero exclusion." *American Economic Review*, 108(6): 1468–1487. 1, 10

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2018. "Fuzzy differences-in-differences." *The Review of Economic Studies*, 85(2): 999–1028. 1, 8, M.2

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American economic review*, 110(9): 2964–2996. 3, 3

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2026. "Difference-in-differences estimators of intertemporal treatment effects." *Review of Economics and Statistics*, 108(4): 863–880. 1

DeMalach, Elad, and Analia Schlosser. 2026. "Short-and long-term effects of universal preschool: Evidence from the Arab population in Israel." *American Economic Journal: Economic Policy*, 18(2): 404–440. 9

Deza, Monica, Thanh Lu, Johanna Catherine Maclean, and Alberto Ortega. 2024. "Losing medicaid and crime." National Bureau of Economic Research. 20

Diamond, Rebecca, Michael J Dickstein, Timothy McQuade, and Petra Persson. 2018. "Take-up, drop-out, and spending in ACA marketplaces." 6.2

Engelhardt, Gary V, and Jonathan Gruber. 2011. "Medicare Part D and the financial protection of the elderly." *American Economic Journal: Economic Policy*, 3(4): 77–102. 1, 6.2

Finkelstein, Amy, Matthew J Notowidigdo, and Steven X Shi. 2026. "Trading Goods for Lives: NAFTAâs Mortality Impacts and Implications." National Bureau of Economic Research. 7

Frangakis, Constantine E, and Donald B Rubin. 2002. "Principal stratification in causal inference." *Biometrics*, 58(1): 21–29. 1

Ghanem, Dalia, Pedro HC Sant'Anna, and Kaspar Wüthrich. 2026*a*. "Selection and parallel trends." *arXiv preprint arXiv:2203.09001*. 1, 6.1

Ghanem, Dalia, Pedro HC Sant'Anna, and Kaspar Wüthrich. 2026*b*. "When Should Pre-trends Be Parallel?" Vol. 116, 64–69, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203-2425. 3

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik instruments: What, when, why, and how." *American Economic Review*, 110(8): 2586–2624. 6, 7, M.2

Haanwinckel, Daniel. 2023. "Does regional variation in wage levels identify the effects of a national minimum wage?" *arXiv preprint arXiv:2307.01284*. 1, 10

Health and Retirement Study. 2025. "Health and Retirement Study, Public Survey

Data." *Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740)*, Available from the University of Michigan. 6.2

Heckman, James J., and Edward J. Vytlacil. 1999. "Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects." *Proceedings of the National Academy of Sciences*, 96(8): 4730–4734. 1, 4

Heckman, James J., and Edward J. Vytlacil. 2005. "Structural Equations, Treatment Effects, and Econometric Policy Evaluation." *Econometrica*, 73(3): 669–738. 1, 4

Heckman, James J, and Sergio Urzua. 2010. "Comparing IV with structural models: What simple IV can and cannot identify." *Journal of Econometrics*, 156(1): 27–37. 2.1

Heiss, Florian, Daniel McFadden, and Joachim Winter. 2006. "Who Failed To Enroll In Medicare Part D, And Why? Early Results: Survey results find gains in coverage among seniors but indicate some dissatisfaction and doubts about its adequacy in the long run." *Health Affairs*, 25(Suppl1): W344–W354. 6.2, 24

Huh, Jason, and Julian Reif. 2017. "Did Medicare part D reduce mortality?" *Journal of health economics*, 53: 17–37. 6.2

Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica*, 62(2): 467–475. 1, 2.1

Kamat, Vishal, Samuel Norris, and Matthew Pecenco. 2026. "Identification in Multiple Treatment Models under Discrete Variation." 1, 6

Kowalski, Amanda E. 2023. "Reconciling seemingly contradictory results from the Oregon health insurance experiment and the Massachusetts health reform." *Review of Economics and Statistics*, 105(3): 646–664. 3.5, 6.1

Ko, Wonsik, and Robert A Moffitt. 2024. "Take-up of social benefits." *Handbook of Labor, Human Resources and Population Economics*, 1–42. 1

Kwon, Soonwoo, and Liyang Sun. 2025. "Estimating treatment effects under bounded heterogeneity." *arXiv preprint arXiv:2510.05454*. 14

LaLonde, Robert J. 1986. "Evaluating the econometric evaluations of training programs with experimental data." *The American economic review*, 604–620. 2

Manski, Charles F. 1997. "Monotone treatment response." *Econometrica: Journal of the Econometric Society*, 1311–1334. 4.3

Manski, Charles F, and John V Pepper. 2000. "Monotone Instrumental Variables: With an Application to the Returns to Schooling." *Econometrica*, 68(4): 997–1010. 4.3

Manski, Charles F, and John V Pepper. 2018. "How do right-to-carry laws affect crime rates? Coping with ambiguity using bounded-variation assumptions." *Review of*

Economics and Statistics, 100(2): 232–244. 1, 12, 14

Marx, Philip, Elie Tamer, and Xun Tang. 2024. "Parallel trends and dynamic choices." *Journal of Political Economy Microeconomics*, 2(1): 129–171. 3, 1, 3

Miller, Sarah. 2012. "The effect of insurance on emergency room visits: an analysis of the 2006 Massachusetts health reform." *Journal of public Economics*, 96(11-12): 893–908. 1, 9

Mogstad, Magne, Andres Santos, and Alexander Torgovitsky. 2018. "Using instrumental variables for inference about policy relevant treatment parameters." *Econometrica*, 86(5): 1589–1619. 1, 4, 19

Morris, Michael, Valerie A. Kelly, Ron J. Kopicki, and Derek Byerlee. 2007. *Fertilizer Use in African Agriculture: Lessons Learned and Good Practice Guidelines.* Washington, DC:World Bank. 1, 1, 2.2

Nelsen, Roger B. 1999. "Methods of constructing Copulas." In *An Introduction to Copulas.* 45–87. Springer. 6

Quach, Simon. 2024. "The Labor Market Effects of Expanding Overtime Coverage." *Available at SSRN 3608506.* 11

Rambachan, Ashesh, and Jonathan Roth. 2023. "A more credible approach to parallel trends." *Review of Economic Studies*, 90(5): 2555–2591. 1, 12, 14, 4

Roth, Jonathan. 2022. "Pretest with caution: Event-study estimates after testing for parallel trends." *American Economic Review: Insights*, 4(3): 305–322. 30

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe. 2023. "What's trending in difference-in-differences? A synthesis of the recent econometrics literature." *Journal of econometrics*, 235(2): 2218–2244. 1

Shapiro, Joseph S, and Reed Walker. 2018. "Why is pollution from US manufacturing declining? The roles of environmental regulation, productivity, and trade." *American economic review*, 108(12): 3814–3854. 1, 11

Stafford, Sarah L. 2002. "The Effect of Punishment on Firm Compliance with Hazardous Waste Regulations." *Journal of Environmental Economics and Management*, 44(2): 290–308. 1, 1, 2.3

Sun, Liyang, and Jesse M Shapiro. 2022. "A linear panel model with heterogeneous coefficients and variation in exposure." *Journal of Economic Perspectives*, 36(4): 193–204. 5, 9

Taubman, Sarah L, Heidi L Allen, Bill J Wright, Katherine Baicker, and Amy N Finkelstein. 2014. "Medicaid increases emergency-department use: evidence from Oregon's Health Insurance Experiment." *Science*, 343(6168): 263–268. 1, 6.1

Tazhitdinova, Alisa, and Gonzalo Vázquez-Bare. 2025. "Difference-in-Differences with Unequal Baseline Treatment Status." Working paper, revised May 2025. 5

Wettstein, Gal. 2020. "Retirement lock and prescription drug insurance: Evidence from medicare part d." *American Economic Journal: Economic Policy*, 12(1): 389–417. 1, 6.2

Wuellrich, Jean-Philippe. 2010. "The effects of increasing financial incentives for firms to promote employment of disabled workers." *Economics Letters*, 107(2): 173–176. 1, 11

Xu, Yiqing, Anqi Zhao, and Peng Ding. 2026. "Factorial difference-in-differences." *Journal of the American Statistical Association*, , (just-accepted): 1–23. 1

Appendix

A Proof of Proposition 1

Fix an observed-data distribution compatible with Assumption PT and the maintained regularity conditions. To establish that $\Lambda^* = \mathbb{R}$, I construct, for each $\ell \in \mathbb{R}$, a specification of potential outcomes and treatment statuses that generates this same observed-data distribution, satisfies the maintained assumptions, and yields $\lambda = \ell$.

Let $S_i = \mathbf{1}\{D_{i0} = 1, D_{i1} = 1\}$ indicate continuously treated units. By the maintained regularity conditions, $\mathbb{P}[S_i = 1] > 0$. Starting from any compatible specification, redefine the unobserved potential treatment status under no policy at endline as

$$D_{i1}(0) = (1 - D_{i0})D_{i1},$$

leaving all other potential treatment statuses and potential outcomes unchanged. This redefinition preserves the observed data and Assumption PT. Because $D_{i1}(1) = D_{i1}$, period-1 compliers are now exactly the continuously treated units:

$$\{D_{i1}(1) > D_{i1}(0)\} = \{S_i = 1\}.$$

Let

$$b = \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid S_i = 1],$$

which is finite under the maintained moment conditions.

For any $\ell \in \mathbb{R}$, define a new specification of potential outcomes by

$$Y_{it}^{\ell}(0) = Y_{it}(0) + (b - \ell)S_i, \qquad Y_{it}^{\ell}(1) = Y_{it}(1), \qquad t \in \{0, 1\}.$$

This transformation changes only the untreated potential outcomes of continuously treated units, which are unobserved in both periods. It therefore preserves the joint distribution of observed outcomes and treatment statuses.

Moreover, the transformation leaves each unit's untreated outcome change unchanged:

$$Y_{i1}^{\ell}(0) - Y_{i0}^{\ell}(0) = Y_{i1}(0) - Y_{i0}(0).$$

Consequently, Assumption PT continues to hold. The transformed potential outcomes also have finite first moments because the adjustment is a finite constant multiplied by an indicator.

Finally, since the period-1 complier subgroup remains $\{S_i = 1\}$, the LATE under this specification is

$$\lambda^{\ell} = \mathbb{E}[Y_{i1}^{\ell}(1) - Y_{i1}^{\ell}(0) \mid S_i = 1] = b - (b - \ell) = \ell.$$

Thus, every $\ell \in \mathbb{R}$ is attainable while preserving the observed-data distribution and all maintained assumptions. Since these assumptions also ensure that λ is finite, $\Lambda^* = \mathbb{R}$. $\square$

B Proof to Lemma 1

The joint distribution of $\{D_{i0}(0), D_{i0}(1), D_{i1}(0), D_{i1}(1)\}$ may be defined as

$$\pi_{d_{00}d_{01}d_{10}d_{11}} := \mathbb{P}\Big[D_{i0}(0) = d_{00},\ D_{i0}(1) = d_{01},\ D_{i1}(0) = d_{10},\ D_{i1}(1) = d_{11}\Big]$$

for $(d_{00}, d_{01}, d_{10}, d_{11}) \in \{0,1\}^4$. Absent additional assumptions, the only restrictions on the data-generating process are:

$$\pi_{d_{00}d_{01}d_{10}d_{11}} \geq 0 \quad \text{for all} (d_{00}, d_{01}, d_{10}, d_{11}) \in \{0,1\}^4, \quad \text{and} \sum_{d_{00},d_{01},d_{10},d_{11} \in \{0,1\}} \pi_{d_{00}d_{01}d_{10}d_{11}} = 1.$$

Assumption SEL(a) states the only admissible compliance types are never-takers ($D_{it}(0) = 0, D_{it}(1) = 0$), compliers ($D_{it}(0) = 0, D_{it}(1) = 1$), and always-takers ($D_{it}(0) = 1, D_{it}(1) = 1$). Equivalently, the following probabilities must be zero:

$$\pi_{1000}, \pi_{1001}, \pi_{1010}, \pi_{1011}, \pi_{0010}, \pi_{0110}, \pi_{1110} = 0$$

Note, however, that Assumption SEL(a) does not constrain intertemporal choice behavior, only requiring within-period monotonicity. Assumption SEL(c), on the other hand, imposes restrictions on intertemporal choice behavior requiring that

$$\begin{aligned} &\mathbb{P}[D_{i1}(0) = 0 \mid D_{i0}(0) = 1] = 0 \\ \iff &\pi_{1100}, \pi_{1101} = 0 \end{aligned}$$

Finally, Assumption SEL(b) requires that

$$\begin{aligned} &\mathbb{P}[D_{i0}(0) = D_{i0}(1) = 1] = \mathbb{P}[D_{i1}(0) = D_{i1}(1) = 1] \\ \iff &\pi_{1100} + \pi_{1101} + \pi_{1111} = \pi_{0011} + \pi_{0111} + \pi_{1111} \\ \iff &\pi_{1100} + \pi_{1101} = \pi_{0011} + \pi_{0111} \end{aligned}$$

Assumption SEL(b) and SEL(c) jointly imply that

$$\pi_{0011}, \pi_{0111} = 0$$

Therefore, under the only admissible types of units under Assumption SEL are:

Type	$\Big(D_{i0}(0), D_{i0}(1), D_{i1}(0), D_{i1}(1)\Big)$	Interpretation
π_{0000}	(0,0,0,0)	never-taker
π_{0001}	(0,0,0,1)	period-1 complier
π_{0100}	(0,1,0,0)	period-0 complier
π_{0101}	(0,1,0,1)	period-0 and period-1 complier
π_{1111}	(1,1,1,1)	always-taker

□

C Proof to Proposition 2

Our identification target is the effect of the policy Z_t on treatment take-up in the post-period D_{i1}:

$$\mathbb{E}\Big[D_{i1}(1) - D_{i1}(0)\Big]$$

In order to identify the effect of the policy on treatment take-up, we make two assumptions. First, we assume that, within each time-period, units are weakly more likely to select into treatment with the policy switched on.

At $t = 0$, the policy is switched off, so the observed treatment status satisfies $D_{i0} = D_{i0}(0)$. Assumption SEL(a) implies $D_{i0}(1) \geq D_{i0}(0)$. Hence, for units with $D_{i0} = 1$, we necessarily have

$$\Big(D_{i0}(0), D_{i0}(1)\Big) = (1, 1),$$

so the individual is an always-taker at $t = 0$. Therefore

$$\mathbb{P}[D_{i0}(0) = 1,\ D_{i0}(1) = 1] = \mathbb{P}[D_{i0} = 1].$$

Under Assumption SEL(b), which states that the share of always-takers is stable over time, we can also identify the share of always-takers at $t = 1$,

$$\mathbb{P}[D_{i1}(0) = 1,\ D_{i1}(1) = 1] = \mathbb{P}[D_{i0}(0) = 1,\ D_{i0}(1) = 1] = \mathbb{P}[D_{i0} = 1].$$

At $t = 1$, the policy is on, so the observed treatment status satisfies $D_{i1} = D_{i1}(1)$. Any unit with $D_{i1}(1) = 1$ must either be a complier or an always-taker:

$$\Big(D_{i1}(0), D_{i1}(1)\Big) = (0, 1) \quad \text{or} \quad \Big(D_{i1}(0), D_{i1}(1)\Big) = (1, 1).$$

Therefore,

$$\begin{aligned}
\mathbb{P}[D_{i1} = 1] &= \mathbb{P}[D_{i1}(0) = 1,\ D_{i1}(1) = 1] + \mathbb{P}[D_{i1}(0) = 0,\ D_{i1}(1) = 1] \\
&= \mathbb{P}[D_{i0} = 1] + \mathbb{P}[D_{i1}(0) = 0,\ D_{i1}(1) = 1] \\
\iff \mathbb{P}[D_{i1} = 1] - \mathbb{P}[D_{i0} = 1] &= \mathbb{P}[D_{i1}(0) = 0,\ D_{i1}(1) = 1]
\end{aligned}$$

Under Assumption SEL(a),

$$\mathbb{E}\Big[D_{i1}(1) - D_{i1}(0)\Big] = \mathbb{P}[D_{i1}(0) = 0,\ D_{i1}(1) = 1] = \mathbb{P}[D_{i1} = 1] - \mathbb{P}[D_{i0} = 1].$$

Finally, we show how Assumption SEL(a)-(c) identifies the latent compliance type of all units. Lemma 1 shows that this assumption requires that $D_{i1}(0) = D_{i0}(0)$. $D_{i0}(0)$ is observed as D_{i0}. Similarly, $D_{i1}(1)$ is observed as D_{i1}. Therefore, $D_{i1}(0)$ and $D_{i1}(1)$ are identified for all units, such that:

$$\{D_{i1}(1), D_{i1}(0)\} = \{D_{i1}, D_{i0}\}$$

□

D Proof to Proposition 3

In Lemma 1, we show that $\{D_{i0} = 0, D_{i1} = 1\} = \{D_{i1}(1) > D_{i1}(0)\}$. Therefore, we have that:

$$\lambda = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big]$$

Under Assumption PT, we can show that:

$$\begin{aligned}
\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] &= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 0\Big] \\
&= \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big] = \tau
\end{aligned}$$

where the second equality follows from Assumption PT. Therefore, τ identifies λ under Assumptions PT and SEL:

$$\lambda = \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big] = \tau.$$

□

E Sharpness of Assumption SEL

In this Section, I establish that the restrictions imposed by Assumption SEL are sharp. That is, it is not possible to further relax these selection conditions and retain the identification results in Section 3.2 without imposing restrictions on treatment effect heterogeneity. Specifically, I demonstrate that the set of admissible latent compliance types derived in Lemma 1 cannot be expanded to include any of the remaining four compliance types that are otherwise possible under within-period monotonicity. I proceed by systematically ruling out the inclusion of each of these four compliance trajectories.

Period-0 Compliers or Never-Takers to Period-1 Always-Takers: Consider the latent compliance types corresponding to units that are never-takers or compliers at baseline ($t = 0$), but transition to always-takers at endline ($t = 1$). These correspond to the latent paths $(D_{i0}(0) = 0, D_{i0}(1) = 0, D_{i1}(0) = 1, D_{i1}(1) = 1)$ and $(D_{i0}(0) = 0, D_{i0}(1) = 1, D_{i1}(0) = 1, D_{i1}(1) = 1)$. Because the policy is switched off at baseline ($Z_0 = 0$) and switched on at endline ($Z_1 = 1$), both of these latent types generate the exact same observable treatment sequence: $D_{i0} = 0$ and $D_{i1} = 1$.

Under Assumption SEL, the observed subgroup $(D_{i0} = 0, D_{i1} = 1)$ consists exclusively of period-1 compliers. However, if either of these two alternative transition types is admissible, this observed subgroup becomes a mixture of period-1 compliers and period-1 always-takers. Applying Assumption PT, the estimand τ subsequently identifies a weighted average of the complier effect, λ, and the treatment effect for these transitioning always-takers:

$$
\begin{aligned}
\tau &= \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 0\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \\
&= \lambda \cdot \mathbb{P}[D_{i1}(1) > D_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1] \\
&\quad + \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1}(1) = D_{i1}(0) = 1\Big] \cdot \mathbb{P}[D_{i1}(1) = D_{i1}(0) = 1 \mid D_{i0} = 0, D_{i1} = 1]
\end{aligned}
$$

Absent additional assumptions restricting the treatment effect heterogeneity between the period-1 compliers (λ) and always-takers ($\mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1}(1) = D_{i1}(0) = 1]$), the parameter of interest λ cannot be point-identified from the observable estimand τ.

Period-0 Always-Takers to Period-1 Compliers: Consider the latent compliance type corresponding to units that are always-takers at baseline ($t = 0$) but transition to compliers at endline ($t = 1$). This corresponds to the latent path $(D_{i0}(0) = 1, D_{i0}(1) = 1, D_{i1}(0) = 0, D_{i1}(1) = 1)$. Because the policy is off at baseline ($Z_0 = 0$) and on at endline ($Z_1 = 1$), the observed treatment sequence for these units is $D_{i0} = 1$ and $D_{i1} = 1$.

Because these units are period-1 compliers ($D_{i1}(1) > D_{i1}(0)$), they contribute to the target LATE parameter, λ. Under Assumption PT, point-identifying this effect requires recovering their average untreated potential outcome at baseline: $\mathbb{E}[Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 1]$. However, because these units always select into treatment at baseline regardless of the policy, their realized baseline outcome is treated ($Y_{i0} = Y_{i0}(1)$). Consequently, the untreated baseline moment $\mathbb{E}[Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 1]$ is unobserved and not identified. Without imposing restrictions to bound this unobserved counterfactual, the treatment effect for baseline-treated compliers remains entirely unconstrained, thereby precluding the point identification of λ.

Period-0 Always-Takers to Period-1 Never-Takers: Consider the latent compliance type corresponding to units that are always-takers at baseline ($t = 0$) but transition to never-takers at endline ($t = 1$). This corresponds to the latent path $(D_{i0}(0), D_{i0}(1), D_{i1}(0), D_{i1}(1)) = (1, 1, 0, 0)$. Under monotonicity, these units are identified by the observable treatment sequence $(D_{i0}, D_{i1}) = (1, 0)$. If this is the only admissible additional compliance trajectory admitted, units that are period-1 compliers and units with treatment path $(D_{i0}, D_{i1}) = (0, 1)$ are the still the same set. Consequently, λ is identified by τ under the maintained parallel trends assumption. As noted in Remark 4, the $(1, 0)$ subgroup may therefore be excluded without changing λ or τ, leaving only the five admissible trajectories in Figure 1b. Admitting units with $(D_{i0}, D_{i1}) = (1, 0)$ violates Assumption SEL(b), so the first-stage is identified as:

$$\mathbb{E}[D_{i1}(1) - D_{i1}(0)] = \mathbb{P}[D_{i0} = 0, D_{i1} = 1] = \mathbb{E}[D_{i1}] - \mathbb{E}[D_{i0}] + \mathbb{P}[D_{i0} = 1, D_{i1} = 0].$$

F Proof to Proposition 4

We define λ conditional on baseline treatment status as follows:

$$\lambda_d := \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) > D_{i1}(0), D_{i0} = d\Big]$$

for $d \in \{0, 1\}$.

Assumption HET imposes two distinct forms of treatment effect homogeneity. First, it states that the treatment effect heterogeneity among compliers is mean-independent of D_{i0}, i.e., $\lambda_0 = \lambda_1$. Second, it states that treatment effect heterogeneity is mean-independent of the counterfactual treatment decision, $D_{i1}(0)$. This implies that the treatment effect is homogeneous across always-takers and compliers, such that $\lambda_d = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) = D_{i1}(0) = 1, D_{i0} = d\Big]$.

The treatment effect for units with $D_{i0} = 0$ and $D_{i1} = 1$ can be decomposed into the effect of the policy on baseline-untreated compliers and always-takers:

$$\begin{aligned}
&\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \\
&= \lambda_0 \cdot \mathbb{P}[D_{i1}(1) > D_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1] \\
&\quad + \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1}(1) = D_{i1}(0) = 1\Big] \cdot \mathbb{P}[D_{i1}(1) = D_{i1}(0) = 1 \mid D_{i0} = 0, D_{i1} = 1]
\end{aligned}$$

Under Assumption HET, $\lambda_0 = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1}(1) = D_{i1}(0) = 1\Big]$. Therefore,

$$\begin{aligned}
&\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \\
&= \lambda_0 = \lambda
\end{aligned}$$

where the final equality follows again from Assumption HET.

Under Assumption PT, we can show that:

$$\begin{aligned}
\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] &= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{i1} = 0\Big] \\
&= \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big] = \tau
\end{aligned}$$

where the second equality follows from Assumption PT. Therefore, τ identifies λ under under Assumptions PT and HET:

$$\lambda = \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0, D_{i1} = 0\Big] = \tau.$$

□

G Proof to Proposition 5

We define the joint probabilities of period-1 compliance types and baseline treatment status as follows:

Period-1 Status	$D_{i0}=0$	$D_{i0}=1$	**Total**
Always-Takers (A)	π_{A0}	π_{A1}	π_A
Compliers (C)	π_{C0}	π_{C1}	π_C
Never-Takers (N)	π_{N0}	π_{N1}	π_N
Defiers (D)	π_{D0}	π_{D1}	π_D
Total	$\mathbb{P}[D_{i0}=0]$	$\mathbb{P}[D_{i0}=1]$	1

We define λ conditional on baseline treatment status as follows:

$$\lambda_d := \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) > D_{i1}(0), D_{i0} = d\Big]$$

for $d \in \{0,1\}$.

Under Assumption PT, we can show that:

$$\begin{aligned}
\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0}=0, D_{i1}=1\Big] &= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0}=0, D_{i1}=1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0}=0, D_{i1}=1\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0}=0, D_{i1}=1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0}=0, D_{i1}=0\Big] \\
&= \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0}=0, D_{i1}=1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0}=0, D_{i1}=0\Big] = \tau
\end{aligned}$$

where the second equality follows from Assumption PT.

The units $\{D_{i0}=0, D_{i1}=1\}$ consists of baseline-untreated compliers and always-takers (π_{A0}). Therefore, τ recovers the weighted average of treatment effects for these two groups:

$$\tau = \frac{\pi_{C0}\lambda_0 + \pi_{A0}\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0}=0, D_{i1}(1)=D_{i1}(0)=1\Big]}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$$

Recall that $\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0}=0, D_{i1}(1)=D_{i1}(0)=1\Big] - \lambda_0 = \kappa_{AT}$:

$$\tau = \frac{\pi_{C0}\lambda_0 + \pi_{A0}(\lambda_0 + \kappa_{AT})}{\mathbb{P}[D_{i0}=0, D_{i1}=1]} = \lambda_0 + \kappa_{AT}\frac{\pi_{A0}}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$$

Solving for λ_0:

$$\lambda_0 = \tau - \kappa_{AT}\frac{\pi_{A0}}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$$

Recall that $\lambda_1 - \lambda_0 = \kappa_C$. Substituting this into the definition of λ:

$$\begin{aligned}
\lambda &= \frac{\pi_{C0}\lambda_0 + \pi_{C1}\lambda_1}{\pi_C} \\
&= \frac{\pi_{C0}\lambda_0 + \pi_{C1}(\lambda_0 + \kappa_C)}{\pi_C} = \lambda_0 + \kappa_C\frac{\pi_{C1}}{\pi_C}
\end{aligned}$$

Plugging in λ_0 from above:

$$\lambda = \left(\tau - \kappa_{AT}\frac{\pi_{A0}}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}\right) + \kappa_C \frac{\pi_{C1}}{\pi_C}$$

It remains to identify π_{A0}, π_{C1}, and π_C. Recall that:

$$\begin{aligned}
&\mathbb{P}[D_{i1}(0) = D_{i1}(1) = 1] - \mathbb{P}[D_{i0}(0) = D_{i0}(1) = 1] = \Delta^{AT} \\
\implies &\pi_A - \mathbb{P}[D_{i0} = 1] = \Delta^{AT} \\
\implies &\pi_A = \Delta^{AT} + \mathbb{P}[D_{i0} = 1] \\
\implies &\pi_{A0} + \pi_{A1} = \mathbb{P}[D_{i0} = 1] + \Delta^{AT}
\end{aligned}$$

where the second line follows from requiring that $D_{it}(1) \geq D_{it}(0)$ almost surely. Next,

$$\begin{aligned}
&\mathbb{P}[D_{i1}(0) = 0 \mid D_{i0}(0) = 1] = \Delta^{churn} \\
\implies &\mathbb{P}[D_{i1}(0) = 0 \mid D_{i0} = 1] = \Delta^{churn} \\
\implies &\frac{\pi_{C1} + \pi_{N1}}{\mathbb{P}[D_{i0} = 1]} = \Delta^{churn} \\
\implies &\pi_{C1} + \pi_{N1} = \Delta^{churn}\mathbb{P}[D_{i0} = 1]
\end{aligned}$$

Since $D_{it}(1) \geq D_{it}(0)$ almost surely, we can observe π_{N1} directly, since at $t = 1$, never-takers have $D_{i1}(1) = D_{i1} = 0$. Therefore, $\pi_{N1} = \mathbb{P}[D_{i0} = 1, D_{i1} = 0]$. Solving for π_{C1}:

$$\pi_{C1} = \Delta^{churn}\mathbb{P}[D_{i0} = 1] - \mathbb{P}[D_{i0} = 1, D_{i1} = 0]$$

Note that $\pi_{A1} + \pi_{C1} + \pi_{N1} = \mathbb{P}[D_{i0} = 1]$. Using the fact that $\pi_{C1} + \pi_{N1} = \Delta^{churn}\mathbb{P}[D_{i0} = 1]$, we have:

$$\begin{aligned}
\pi_{A1} + \Delta^{churn}\mathbb{P}[D_{i0} = 1] &= \mathbb{P}[D_{i0} = 1] \\
\pi_{A1} &= (1 - \Delta^{churn})\mathbb{P}[D_{i0} = 1]
\end{aligned}$$

Recall that $\pi_{A0} + \pi_{A1} = \mathbb{P}[D_{i0} = 1] + \Delta^{AT}$. Substituting π_{A1} back into this expression:

$$\begin{aligned}
\pi_{A0} + (1 - \Delta^{churn})\mathbb{P}[D_{i0} = 1] &= \mathbb{P}[D_{i0} = 1] + \Delta^{AT} \\
\pi_{A0} &= \Delta^{AT} + \Delta^{churn}\mathbb{P}[D_{i0} = 1]
\end{aligned}$$

Finally,

$$\begin{aligned}
\pi_C &= \mathbb{P}[D_{i1} = 1] - \pi_A \\
&= \mathbb{P}[D_{i1} = 1] - \Big(\Delta^{AT} + \mathbb{P}[D_{i0} = 1]\Big)
\end{aligned}$$

Substituting π_{A0}, π_{C1}, and π_C as derived above:

$$\lambda = \left(\tau - \kappa_{AT}\frac{\Delta^{AT} + \Delta^{churn}\mathbb{P}[D_{i0}=1]}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}\right) + \kappa_C \frac{\Delta^{churn}\mathbb{P}[D_{i0}=1] - \mathbb{P}[D_{i0}=1, D_{i1}=0]}{\mathbb{P}[D_{i1}=1] - (\mathbb{P}[D_{i0}=1] + \Delta^{AT})}$$

Recall that, as discussed in Section 3.2, $\Delta^{churn} = 0 \implies \mathbb{P}[D_{i0} = 1, D_{i1} = 0] = 0$. Therefore:

$$\begin{aligned}\lambda = \tau &- \kappa_{AT}\frac{\Delta^{AT} + \Delta^{churn}\mathbb{P}[D_{i0}=1]}{\mathbb{P}[D_{i0}=0, D_{i1}=1]} \\ &+ \kappa_C \frac{\Delta^{churn}\mathbb{P}[D_{i0}=1] - \mathbb{1}\{\Delta^{churn} > 0\}\mathbb{P}[D_{i0}=1, D_{i1}=0]}{\mathbb{P}[D_{i1}=1] - (\mathbb{P}[D_{i0}=1] + \Delta^{AT})}.\end{aligned}$$

Observe that, if $\Delta^{AT} = \Delta^{churn} = 0$, $\lambda = \tau$. This is the case discussed in Section 3.2. Similarly, if $\kappa_C = \kappa_{AT} = 0$, $\lambda = \tau$. This is the case discussed in Section 3.3.

□

H Proof to Proposition 6

We define the joint probabilities of period-1 compliance types and baseline treatment status as follows:

Period-1 Status	$D_{i0}=0$	$D_{i0}=1$	**Total**
Always-Takers (A)	π_{A0}	π_{A1}	π_A
Compliers (C)	π_{C0}	π_{C1}	π_C
Never-Takers (N)	π_{N0}	π_{N1}	π_N
Defiers (D)	π_{D0}	π_{D1}	π_D
Total	$\mathbb{P}[D_{i0}=0]$	$\mathbb{P}[D_{i0}=1]$	1

We define λ conditional on baseline treatment status as follows:

$$\lambda_d := \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1) > D_{i1}(0), D_{i0} = d\Big]$$

for $d \in \{0,1\}$.

We begin by identifying the treatment effect for the units $\{D_{i0}=0, D_{i1}=1\}$ under Assumption PT:

$$\begin{aligned}
\tau &:= \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0}=0, D_{i1}=1\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0}=0, D_{i1}=0\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0}=0, D_{i1}=1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0}=0, D_{i1}=0\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i0}(0) \mid D_{i0}=0, D_{i1}=1\Big] - \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0}=0, D_{i1}=1\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0}=0, D_{i1}=1\Big]
\end{aligned}$$

where the third equality follows from Assumption PT.

The units $\{D_{i0}=0, D_{i1}=1\}$ consists of baseline-untreated compliers and always-takers. Therefore, τ recovers the weighted average of treatment effects for these two groups:

$$\tau = \frac{\pi_{C0}\lambda_0 + \pi_{A0}\mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0}=0, D_{i1}(1)=D_{i1}(0)=1\Big]}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$$

Assumption BH bounds the difference between the always-taker effect and the complier effect. Specifically, $\mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i1}(1)=D_{i1}(0)=1, D_{i0}=0] \in [\lambda_0 - M_{AT}, \lambda_0 + M_{AT}]$.

First, we focus on finding the upper bound for λ. In order to do this, we begin by finding the maximum possible value for λ_0. We substitute the lower bound for the always-taker effect into the expression for τ:

$$\tau \geq \frac{\pi_{C0}\lambda_0 + \pi_{A0}(\lambda_0 - M_{AT})}{\mathbb{P}[D_{i0}=0, D_{i1}=1]} = \frac{\lambda_0(\pi_{C0} + \pi_{A0}) - \pi_{A0}M_{AT}}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$$

Because $\mathbb{P}[D_{i0}=0, D_{i1}=1] = \pi_{C0} + \pi_{A0}$, this simplifies to $\tau \geq \lambda_0 - M_{AT}w_{AT}$, where we define $w_{AT} \equiv \frac{\pi_{A0}}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$. Rearranging yields the upper bound for the baseline-untreated

complier effect:

$$\lambda_0^{max} = \tau + M_{AT} w_{AT}$$

The parameter of interest, λ, is the weighted average of the treatment for baseline-untreated compliers (λ_0) and baseline-treated compliers (λ_1):

$$\lambda = \frac{\pi_{C0}\lambda_0 + \pi_{C1}\lambda_1}{\pi_C}$$

Assumption BH also bounds the heterogeneity among compliers across baseline treatment states, requiring $\lambda_1 \in [\lambda_0 - M_C, \lambda_0 + M_C]$. Substituting the upper bound for λ_1 into the expression for λ yields:

$$\lambda \leq \frac{\pi_{C0}\lambda_0 + \pi_{C1}(\lambda_0 + M_C)}{\pi_C} = \lambda_0 + M_C w_C$$

where we define $w_C \equiv \frac{\pi_{C1}}{\pi_C}$. Because λ is increasing in λ_0, the upper bound for λ is achieved by evaluating this expression at λ_0^{max}:

$$\lambda^{max} = \lambda_0^{max} + M_C w_C$$
$$\lambda^{max} = \tau + M_{AT} w_{AT} + M_C w_C$$

Next, we focus on finding the lower bound for λ. In order to do this, we begin by finding the minimum possible value for λ_0. We substitute the upper bound for the always-taker effect into the expression for τ:

$$\tau \leq \frac{\pi_{C0}\lambda_0 + \pi_{A0}(\lambda_0 + M_{AT})}{\mathbb{P}[D_{i0} = 0, D_{i1} = 1]} = \frac{\lambda_0(\pi_{C0} + \pi_{A0}) + \pi_{A0} M_{AT}}{\mathbb{P}[D_{i0} = 0, D_{i1} = 1]}$$

Because $\mathbb{P}[D_{i0} = 0, D_{i1} = 1] = \pi_{C0} + \pi_{A0}$, this simplifies to $\tau \leq \lambda_0 + M_{AT} w_{AT}$, where $w_{AT} = \frac{\pi_{A0}}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$. Rearranging yields the lower bound for the baseline-untreated complier effect:

$$\lambda_0^{min} = \tau - M_{AT} w_{AT}$$

The parameter of interest, λ, is the weighted average of the treatment for baseline-untreated compliers (λ_0) and baseline-treated compliers (λ_1):

$$\lambda = \frac{\pi_{C0}\lambda_0 + \pi_{C1}\lambda_1}{\pi_C}$$

Assumption BH also bounds the heterogeneity among compliers across baseline treatment states, requiring $\lambda_1 \in [\lambda_0 - M_C, \lambda_0 + M_C]$. Substituting the lower bound for λ_1 into the expression for λ yields:

$$\lambda \geq \frac{\pi_{C0}\lambda_0 + \pi_{C1}(\lambda_0 - M_C)}{\pi_C} = \lambda_0 - M_C w_C$$

where $w_C = \frac{\pi_{C1}}{\pi_C}$. Because λ is increasing in λ_0, the lower bound for λ is achieved by evaluating this expression at λ_0^{min}:

$$\lambda^{min} = \lambda_0^{min} - M_C w_C$$
$$\lambda^{min} = \tau - M_{AT} w_{AT} - M_C w_C$$

Taken together,

$$\lambda \in \left[\tau - M_{AT} \cdot w_{AT} - M_C \cdot w_C \ , \ \tau + M_{AT} \cdot w_{AT} + M_C \cdot w_C\right]$$

where $w_{AT} = \frac{\pi_{A0}}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$ and $w_C = \frac{\pi_{C1}}{\pi_C}$.

□

I Identifying Transition Probabilities using Additional Pre-Period Data

In Sections 2-3.4, we required that we observe each unit's selection decision for $t \in \{0, 1\}$, with the policy switched off at $t = 0$ ($Z_0 = 0$) and the policy switched on at $t = 1$ ($Z_1 = 1$). The identification results in Sections 3.5-4, however, require us to identify the change in the share of always-takers from baseline to endline:

$$\Delta^{AT} = \mathbb{P}[D_{i1}(0) = D_{i1}(1) = 1] - \mathbb{P}[D_{i0}(0) = D_{i0}(1) = 1],$$

and the probability that a baseline-treated unit counterfactually selects out of treatment at endline:

$$\Delta^{churn} = \mathbb{P}[D_{i1}(0) = 0 \mid D_{i0}(0) = 1].$$

In this Section, we leverage additional pre-period data to show Δ^{AT} and Δ^{churn} are identified under mild stationarity assumptions. We require that we observe each unit's selection decision for $t \in \{-1, 0, 1\}$, with the policy switched off at $t \in \{-1, 0\}$ ($Z_{-1}, Z_0 = 0$) and the policy switched on at $t = 1$ ($Z_1 = 1$). Therefore, for each unit, we observe $\{D_{i,-1}, D_{i,0}, D_{i,1}\}$. We maintain Assumption SEL(a) (within-period monotonicity) for the remainder of this section: $\mathbb{P}[D_{it}(1) \geq D_{it}(0)] = 1$.

We begin by identifying Δ^{AT}. The key assumption here is that the change in share of always-takers is constant over time: $\mathbb{P}[D_{i,t}(0) = D_{i,t}(1) = 1] - \mathbb{P}[D_{i,t-1}(0) = D_{i,t-1}(1) = 1] = \mathbb{P}[D_{i,t+1}(0) = D_{i,t+1}(1) = 1] - \mathbb{P}[D_{i,t}(0) = D_{i,t}(1) = 1]$ for all t.

For $t \in \{-1, 0\}$, the policy is switched off, so the observed treatment status satisfies $D_{it} = D_{it}(0)$. Assumption SEL(a) implies $D_{it}(1) \geq D_{it}(0)$. Hence, for units with $D_{i,-1} = 1$, we necessarily have

$$\Big(D_{i,-1}(0), D_{i,-1}(1)\Big) = (1, 1),$$

so the individual is an always-taker at $t = -1$. Similarly, for units with $D_{i,0} = 1$, we necessarily have

$$\Big(D_{i,0}(0), D_{i,0}(1)\Big) = (1, 1),$$

so the individual is an always-taker at $t = 0$. Therefore,

$$\mathbb{P}[D_{i,-1}(0) = 1,\ D_{i,-1}(1) = 1] = \mathbb{P}[D_{i,-1} = 1]$$
$$\mathbb{P}[D_{i,0}(0) = 1,\ D_{i,0}(1) = 1] = \mathbb{P}[D_{i,0} = 1]$$

Then, Δ^{AT} is identified as:

$$\begin{aligned}
\Delta^{AT} &= \mathbb{P}[D_{i1}(0) = D_{i1}(1) = 1] - \mathbb{P}[D_{i0}(0) = D_{i0}(1) = 1] \\
&= \mathbb{P}[D_{i,0}(0) = D_{i,0}(1) = 1] - \mathbb{P}[D_{i,-1}(0) = D_{i,-1}(1) = 1] \\
&= \mathbb{P}[D_{i,0} = 1] - \mathbb{P}[D_{i,-1} = 1]
\end{aligned}$$

where the second equality follows from the assumption that the change in share of always-takers is constant over time and the third equality follows from within-period monotonicity.[28]

Next, we turn to identifying Δ^{churn}. The key assumption here is that the rate at which units select out of treatment absent the policy is constant over time: $\mathbb{P}[D_{i,t+1}(0) = 0 \mid D_{i,t}(0) = 1] = \mathbb{P}[D_{i,t}(0) = 0 \mid D_{i,t-1}(0) = 1]$ for all t. Since $Z_{-1} = 0$ and $Z_0 = 0$, we observe $D_{i,-1}(0) = D_{i,-1} = 0$ and $D_{i,0}(0) = D_{i0} = 0$. Therefore,

$$\begin{aligned}\Delta^{churn} &= \mathbb{P}[D_{i,1}(0) = 0 \mid D_{i,0}(0) = 1] \\ &= \mathbb{P}[D_{i,0}(0) = 0 \mid D_{i,-1}(0) = 1] \\ &= \mathbb{P}[D_{i,0} = 0 \mid D_{i,-1} = 1]\end{aligned}$$

where the second equality follows from the assumption that the rate at which units select out of treatment absent the policy is constant over time and the third assumption follows from the fact that $Z_{-1} = 0$ and $Z_0 = 0$.

Next, we show how the weights in Proposition 6, w_{AT} and w_C, can be identified as functions of Δ^{churn} and Δ^{AT}. We define the joint probabilities of period-1 compliance types and baseline treatment status as follows:

Period-1 Status	$D_{i0} = 0$	$D_{i0} = 1$	**Total**
Always-Takers (A)	π_{A0}	π_{A1}	π_A
Compliers (C)	π_{C0}	π_{C1}	π_C
Never-Takers (N)	π_{N0}	π_{N1}	π_N
Defiers (D)	π_{D0}	π_{D1}	π_D
Total	$\mathbb{P}[D_{i0} = 0]$	$\mathbb{P}[D_{i0} = 1]$	1

Recall that, $w_{AT} = \frac{\pi_{A0}}{\mathbb{P}[D_{i0}=0, D_{i1}=1]}$ and $w_C = \frac{\pi_{C1}}{\pi_C}$.

We begin by identifying π_{C1}. Recall that,

$$\begin{aligned}\Delta^{churn} &= \mathbb{P}[D_{i,1}(0) = 0 \mid D_{i,0}(0) = 1] \\ &= \frac{\mathbb{P}[D_{i,1}(0) = 0, D_{i,0}(0) = 1]}{\mathbb{P}[D_{i,0}(0) = 1]} \\ &= \frac{\mathbb{P}[D_{i,1}(0) = 0, D_{i,0} = 1]}{\mathbb{P}[D_{i,0} = 1]} \\ \implies \Delta^{churn} \cdot \mathbb{P}[D_{i,0} = 1] &= \mathbb{P}[D_{i,1}(0) = 0, D_{i,0} = 1]\end{aligned}$$

where the third equality follows from the fact that $Z_0 = 0$. The subpopulation $\{D_{i,1}(0) = 0, D_{i,0} = 1\}$ consists of units who do not take treatment when the policy is active ($D_{i1}(1) = 0$, period-1 never-takers π_{N1}) and units who do take treatment when the policy is active

[28] Because the share of always-takers at $t = 1$ must be bounded between 0 and 1, this implicitly requires that the predicted share is also bounded as follows: $0 \leq 2\mathbb{P}[D_{i,0} = 1] - \mathbb{P}[D_{i,-1} = 1] \leq 1$. This condition is verifiable using the available pre-period data.

($D_{i1}(1) = 1$, period-1 compliers π_{C1}). Thus:

$$\Delta^{churn} \cdot \mathbb{P}[D_{i,0} = 1] = \pi_{C1} + \pi_{N1}$$

Under within-period monotonicity, $\pi_{D1} = 0$, the never-taker share is identified as $\pi_{N1} = \mathbb{P}[D_{i0} = 1, D_{i1} = 0]$. Therefore,

$$\pi_{C1} = \Delta^{churn}\mathbb{P}[D_{i0} = 1] - \mathbb{P}[D_{i0} = 1, D_{i1} = 0]$$

Next, we turn to identification of π_{A0}. By definition, the change in the share of always-takers across periods is given by:

$$\begin{aligned} \Delta^{AT} &= \pi_A - \mathbb{P}[D_{i0}(0) = 1, D_{i0}(1) = 1] \\ &= \pi_A - \mathbb{P}[D_{i0} = 1] \\ \implies \Delta^{AT} + \mathbb{P}[D_{i0} = 1] &= \pi_A \\ \implies \Delta^{AT} + \mathbb{P}[D_{i0} = 1] &= \pi_{A0} + \pi_{A1} \end{aligned}$$

where the second equality follows from within-period monotonicity. In order to identify π_{A0}, it remains to identify π_{A1}. Observe that:

$$\pi_{A1} + (\pi_{C1} + \pi_{N1}) = \mathbb{P}[D_{i0} = 1]$$

Substituting $\pi_{C1} + \pi_{N1} = \Delta^{churn}\mathbb{P}[D_{i0} = 1]$ allows us to isolate π_{A1}:

$$\pi_{A1} = \mathbb{P}[D_{i0} = 1] - \Delta^{churn}\mathbb{P}[D_{i0} = 1] = (1 - \Delta^{churn})\mathbb{P}[D_{i0} = 1]$$

Substituting π_A and π_{A1} into the decomposition $\pi_{A0} = \pi_A - \pi_{A1}$ yields:

$$\begin{aligned} \pi_{A0} &= (\Delta^{AT} + \mathbb{P}[D_{i0} = 1]) - (1 - \Delta^{churn})\mathbb{P}[D_{i0} = 1] \\ \pi_{A0} &= \Delta^{AT} + \Delta^{churn}\mathbb{P}[D_{i0} = 1] \end{aligned}$$

Next, we identify π_C. At endline ($Z_1 = 1$), units with observed treatment $D_{i1} = 1$ necessarily have $D_{i1}(1) = 1$. This group consists entirely of period-1 compliers (π_C) and period-1 always-takers (π_A).

$$\mathbb{P}[D_{i1} = 1] = \pi_C + \pi_A$$

Substituting the identified share $\pi_A = \Delta^{AT} + \mathbb{P}[D_{i0} = 1]$ and rearranging yields the total share of period-1 compliers:

$$\pi_C = \mathbb{P}[D_{i1} = 1] - (\mathbb{P}[D_{i0} = 1] + \Delta^{AT})$$

Finally, we substitute π_{C1}, π_{A0}, and π_C into the expressions for the weights w_{AT} and w_C:

$$w_{AT} = \frac{\Delta^{AT} + \Delta^{churn}\mathbb{P}[D_{i0} = 1]}{\mathbb{P}[D_{i0} = 0, D_{i1} = 1]}$$

$$w_C = \frac{\Delta^{churn}\mathbb{P}[D_{i0} = 1] - \mathbb{P}[D_{i0} = 1, D_{i1} = 0]}{\mathbb{P}[D_{i1} = 1] - (\mathbb{P}[D_{i0} = 1] + \Delta^{AT})}$$

J Proof to Lemma 2

Consider the first term of the estimand β in Equation 5.1:

$$\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0]$$

This term is the change in outcomes over time among units who were untreated at baseline. We examine what this term identifies under Assumption PT. At $t = 0$, we observe $Y_{i0} = Y_{i0}(0)$. At $t = 1$, we observe $Y_{i1} = Y_{i1}(0) + D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big)$.

$$\begin{aligned}\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0] &= \mathbb{E}\Big[Y_{i1}(0) + D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) - Y_{i0}(0) \mid D_{i0} = 0\Big] \\ &= \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0\Big] + \mathbb{E}\Big[D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 0\Big].\end{aligned}$$

Under Assumption PT, the first term evaluates to Δ^{PT}. The second term may be re-written as:

$$\mathbb{E}\Big[D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 0\Big] = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \cdot \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0]$$

Therefore,

$$\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0] = \Delta^{PT} + \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \cdot \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0].$$

□

K Proof to Lemma 3

Consider the second term of the estimand β in Equation 5.1:

$$\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1]$$

This term is the change in the outcomes over time among units who were treated at baseline. We examine what this term identifies under Assumption PT. At $t = 0$, we observe $Y_{i0} = Y_{i0}(1)$. At $t = 1$, we observe $Y_{i1} = Y_{i1}(0) + D_{i1} \cdot \Big(Y_{i1}(1) - Y_{i1}(0)\Big)$.

$$\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] = \mathbb{E}\Big[Y_{i1}(0) + D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) - Y_{i0}(1) \mid D_{i0} = 1\Big].$$

Adding and subtracting $Y_{i0}(0)$ inside the expectation:

$$\begin{aligned}
\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] &= \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) + D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) - \Big(Y_{i0}(1) - Y_{i0}(0)\Big) \mid D_{i0} = 1\Big] \\
&= \mathbb{E}\Big[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 1\Big] \\
&\quad + \mathbb{E}\Big[D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 1\Big] \\
&\quad - \mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1\Big].
\end{aligned}$$

Under Assumption PT, the first term evaluates to Δ^{PT}.

$$\begin{aligned}
\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] &= \Delta^{PT} \\
&\quad + \mathbb{E}\Big[D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 1\Big] \\
&\quad - \mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1\Big].
\end{aligned}$$

Applying the Law of Total Probability to expand the remaining two terms,

$$\begin{aligned}
&\mathbb{E}\Big[D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 1\Big] - \mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1\Big] \\
&= \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1, D_{i1} = 1\Big]\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 1] \\
&\quad - \mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 1\Big]\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 1] \\
&\quad - \mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 0\Big]\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1]
\end{aligned}$$

Therefore, under Assumption PT,

$$
\begin{aligned}
&\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] \\
&= \Delta^{PT} \\
&\quad + \mathbb{E}\left[\Big(Y_{i1}(1) - Y_{i1}(0)\Big) - \Big(Y_{i0}(1) - Y_{i0}(0)\Big) \mid D_{i0} = 1, D_{i1} = 1\right]\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 1] \\
&\quad - \mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 0\Big]\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1]
\end{aligned}
$$

If we additionally assume that $\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1] = 0$, we obtain:

$$
\begin{aligned}
&\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] \\
&= \Delta^{PT} \\
&\quad + \mathbb{E}\left[\Big(Y_{i1}(1) - Y_{i1}(0)\Big) - \Big(Y_{i0}(1) - Y_{i0}(0)\Big) \mid D_{i0} = 1, D_{i1} = 1\right]
\end{aligned}
$$

□

L Proof to Proposition 7

Recall that,

$$\beta = \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 0\Big] - \mathbb{E}\Big[Y_{i1} - Y_{i0} \mid D_{i0} = 1\Big]$$

In Appendix Section J, we showed that:

$$\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0] = \Delta^{PT} + \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big] \cdot \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0].$$

In Appendix Section K, we showed that:

$$\begin{aligned}\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] = \Delta^{PT} &\\ &+ \mathbb{E}\Big[D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 1\Big] \\ &- \mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1\Big].\end{aligned}$$

Under the time-invariance of treatment effects, the third term is equal to $\mathbb{E}\Big[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1\Big] = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1\Big]$. Substituting this into the expression above yields:

$$\begin{aligned}\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] = \Delta^{PT} &\\ &- \mathbb{E}\Big[(1 - D_{i1})\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 1\Big]\end{aligned}$$

Taken together:

$$\begin{aligned}\beta &= \left(\Delta^{PT} + \mathbb{E}\Big[D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 0\Big]\right) \\ &\quad - \left(\Delta^{PT} - \mathbb{E}\Big[(1 - D_{i1})\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 1\Big]\right) \\ &= \mathbb{E}\Big[D_{i1}\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 0\Big] + \mathbb{E}\Big[(1 - D_{i1})\Big(Y_{i1}(1) - Y_{i1}(0)\Big) \mid D_{i0} = 1\Big].\end{aligned}$$

Applying the Law of Total Probability, we obtain:

$$\begin{aligned}\beta = \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1\Big]\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] &\\ + \mathbb{E}\Big[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1, D_{i1} = 0\Big]\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1].&\end{aligned}$$

□

M Additional Results on Exposure Designs

M.1 Characterizing the Bias in Exposure Designs Recall that the exposure-design estimand is

$$\beta = \mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0] - \mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1].$$

We characterize each term under Assumption PT without imposing restrictions on treatment transitions or treatment-effect heterogeneity. First, consider units who were untreated at baseline. Because $D_{i0} = 0$,

$$Y_{i0} = Y_{i0}(0),$$

whereas

$$Y_{i1} = Y_{i1}(0) + D_{i1}\big(Y_{i1}(1) - Y_{i1}(0)\big).$$

Therefore,

$$\begin{aligned}\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0] = \mathbb{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0] \\ + \mathbb{E}\big[D_{i1}\big(Y_{i1}(1) - Y_{i1}(0)\big) \mid D_{i0} = 0\big].\end{aligned}$$

Assumption PT and the law of iterated expectations imply

$$\mathbb{E}[Y_{i1}(0) - Y_{i0}(0) \mid D_{i0} = 0] = \Delta^{PT}.$$

Applying the law of total probability to the remaining term gives

$$\begin{aligned}\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 0] = \Delta^{PT} + \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1] \\ \times \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0].\end{aligned}$$

Next, consider units who were treated at baseline. Because $D_{i0} = 1$,

$$Y_{i0} = Y_{i0}(1),$$

and hence

$$Y_{i1} - Y_{i0} = Y_{i1}(0) - Y_{i0}(0) + D_{i1}\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big).$$

Taking expectations conditional on $D_{i0} = 1$ and applying Assumption PT yields

$$\begin{aligned}\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] = \Delta^{PT} + \mathbb{E}\big[D_{i1}\big(Y_{i1}(1) - Y_{i1}(0)\big) \mid D_{i0} = 1\big] \\ - \mathbb{E}[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1].\end{aligned}$$

Expanding the remaining terms by endline treatment status gives

$$\begin{aligned}\mathbb{E}[Y_{i1} - Y_{i0} \mid D_{i0} = 1] = \Delta^{PT} + \mathbb{E}\big[\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big) \mid D_{i0} = 1, D_{i1} = 1\big] \\ \times \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 1] \\ - \mathbb{E}[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 0] \\ \times \mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1].\end{aligned}$$

Combining the two expressions, the common counterfactual trend cancels, yielding

$$\begin{aligned}\beta = {} & \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1]\,\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] \\ & - \mathbb{E}\big[\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big) \mid D_{i0} = 1, D_{i1} = 1\big]\,\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 1] \\ & + \mathbb{E}[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1, D_{i1} = 0]\,\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1].\end{aligned}$$

An equivalent representation is obtained by adding and subtracting

$$\mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1].$$

This gives

$$\begin{aligned}\beta = {} & \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1]\,\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] \\ & + \mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1, D_{i1} = 0]\,\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1] \\ & - \mathbb{E}\big[\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big) \mid D_{i0} = 1\big]\,.\end{aligned}$$

Thus, under Assumption PT alone, β equals the weighted sum of period-1 treatment effects among units who switch treatment status, minus the average intertemporal change in treatment effects among baseline-treated units.

We now impose homogeneity of treatment effects across all units in the post-period:

$$Y_{i1}(1) - Y_{i1}(0) = \lambda$$

almost surely. Because the post-period treatment effect is common across units, it is equal to the period-1 LATE. This restriction eliminates cross-sectional heterogeneity in post-period treatment effects while allowing treatment effects to vary over time. Under this restriction, the exposure-design estimand has the opposite sign from the LATE when the average change in treatment effects among baseline-treated units is sufficiently large to dominate the contribution of the LATE generated by the differential change in treatment take-up. Equivalently, after scaling by the LATE, the bias from time-varying treatment effects must exceed the differential change in take-up between the two baseline groups.

Substituting the homogeneity restriction into the preceding decomposition gives

$$\begin{aligned}\beta = {} & \lambda\,\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] + \lambda\,\mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1] \\ & - \mathbb{E}\big[\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big) \mid D_{i0} = 1\big]\,.\end{aligned}$$

Therefore,

$$\begin{aligned}\beta = {} & \lambda\,(\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] + \mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1]) \\ & - \mathbb{E}\big[\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big) \mid D_{i0} = 1\big]\,.\end{aligned}$$

The sum of the two conditional transition probabilities is equal to the differential change in treatment take-up between the baseline groups:

$$\begin{aligned}&\mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] + \mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1] \\ &\quad = \mathbb{E}[D_{i1} - D_{i0} \mid D_{i0} = 0] - \mathbb{E}[D_{i1} - D_{i0} \mid D_{i0} = 1].\end{aligned}$$

Suppose that $\lambda \neq 0$. The exposure-design estimand and the LATE have opposite signs if and only if

$$\beta\lambda < 0.$$

Because $\lambda \neq 0$,

$$\beta\lambda < 0 \quad \iff \quad \frac{\beta}{\lambda} < 0.$$

Substituting the preceding decomposition yields

$$\begin{aligned}\frac{\beta}{\lambda} &= \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] + \mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1] \\ &\quad - \frac{\mathbb{E}\big[\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big) \mid D_{i0} = 1\big]}{\lambda}.\end{aligned}$$

It follows that $\beta\lambda < 0$ if and only if

$$\frac{\mathbb{E}\big[\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big) \mid D_{i0} = 1\big]}{\lambda} > \mathbb{P}[D_{i1} = 1 \mid D_{i0} = 0] + \mathbb{P}[D_{i1} = 0 \mid D_{i0} = 1].$$

M.2 Leveraging Baseline Treatment Status as an Instrument This section relates identification under a universal policy change to the exposure interpretation of Bartik instruments in Goldsmith-Pinkham, Sorkin and Swift (2020), hereafter GPSS. I first describe their construction and then show how its counterpart in this setting is equivalent to using baseline treatment status as an instrument for the change in treatment.

Consider two periods, $t \in \{0,1\}$, and two industries, $k \in \{1,2\}$. Write the outcome equation in levels as

$$y_{\ell t} = \alpha_\ell + \beta_t + \vartheta x_{\ell t} + \epsilon_{\ell t},$$

where ℓ indexes locations, $y_{\ell t}$ is log wages, and $x_{\ell t}$ is log employment. First differencing yields

$$\Delta y_\ell = (\beta_1 - \beta_0) + \vartheta \Delta x_\ell + \Delta \epsilon_\ell.$$

This gives a log-growth version of GPSS's two-industry, one-period equation. Let $s_{\ell k0}$ be location ℓ's initial employment share in industry k, with $s_{\ell 10} + s_{\ell 20} = 1$. Let g_k be the common employment-growth component for industry k between the two dates, typically measured using national industry growth. The Bartik instrument is

$$B_\ell = s_{\ell 10} g_1 + s_{\ell 20} g_2 = g_2 + (g_1 - g_2) s_{\ell 10}.$$

When $g_1 \neq g_2$, cross-sectional variation in the instrument comes from differences in initial industry shares. GPSS's exogenous-shares interpretation requires these shares to be uncorrelated with $\Delta\epsilon_\ell$, the structural error in the differenced equation.

In this paper's setting, the policy switches on simultaneously for all units, from $Z_0 = 0$ to $Z_1 = 1$, so observed treatment is $D_{i0} = D_{i0}(0)$ and $D_{i1} = D_{i1}(1)$. To clarify the identification problem, write potential outcomes as

$$Y_{it}(d) = \alpha_i + \beta_t + \lambda_{it} d + \varepsilon_{it},$$

where $d \in \{0,1\}$ denotes treatment status and $\lambda_{it} = Y_{it}(1) - Y_{it}(0)$ may vary across units and over time. Observed outcomes satisfy

$$\begin{aligned} Y_{i0} &= \alpha_i + \beta_0 + \lambda_{i0} D_{i0}(0) + \varepsilon_{i0}, \\ Y_{i1} &= \alpha_i + \beta_1 + \lambda_{i1} D_{i1}(1) + \varepsilon_{i1}. \end{aligned}$$

Define $\Delta Y_i = Y_{i1} - Y_{i0}$, $\Delta D_i = D_{i1} - D_{i0}$, and $\Delta\beta = \beta_1 - \beta_0$. First differencing and adding and subtracting $\lambda_{i1} D_{i0}$ gives

$$\begin{aligned} \Delta Y_i &= \Delta\beta + \lambda_{i1} D_{i1} - \lambda_{i0} D_{i0} + \varepsilon_{i1} - \varepsilon_{i0} \\ &= \Delta\beta + \lambda_{i1} \Delta D_i + (\lambda_{i1} - \lambda_{i0}) D_{i0} + \varepsilon_{i1} - \varepsilon_{i0} \\ &= \Delta\beta + \lambda_{i1} \Delta D_i + u_i, \end{aligned}$$

where $u_i = (\lambda_{i1} - \lambda_{i0}) D_{i0} + \varepsilon_{i1} - \varepsilon_{i0}$.

First differencing removes the unit fixed effect but does not imply $\text{Cov}(\Delta D_i, u_i) = 0$. Treatment changes may be correlated with changes in untreated outcomes or with changes in treatment effects among baseline-treated units, both of which enter u_i. This motivates seeking an instrument that predicts ΔD_i and is uncorrelated with u_i.

The universal policy change itself supplies no cross-sectional variation. An exposure design instead uses differences in treatment changes associated with baseline treatment status. Define the following Bartik-like instrument:

$$\begin{aligned} B_i &:= (1 - D_{i0})\mathbb{E}[\Delta D_i \mid D_{i0} = 0] + D_{i0}\mathbb{E}[\Delta D_i \mid D_{i0} = 1] \\ &= (1 - D_{i0})\Pr(D_{i1} = 1 \mid D_{i0} = 0) - D_{i0}\Pr(D_{i1} = 0 \mid D_{i0} = 1) \\ &= \Pr(D_{i1} = 1 \mid D_{i0} = 0) - \Big[\Pr(D_{i1} = 1 \mid D_{i0} = 0) + \Pr(D_{i1} = 0 \mid D_{i0} = 1)\Big] D_{i0}. \end{aligned}$$

Each groups' average treatment changes play the role of industry growth components in GPSS. The motivation is that a universal policy may induce differential entry among initially untreated and initially treated units. Whereas GPSS compare locations with different initial industry compositions, this exposure design compares units with different baseline treatment statuses. The observed transition rates also include changes that would have occurred without the policy, so they need not measure its causal effect on treatment.

By construction, B_i can be simplified so that $B_i = \mathbb{E}[\Delta D_i \mid D_{i0}]$ is the treatment change predicted by baseline status. Instrumenting ΔD_i with B_i therefore uses differences in average treatment changes between the baseline groups as its identifying variation. In a regression with an intercept, instrumenting ΔD_i with B_i consequently gives the same IV coefficient as instrumenting it with D_{i0}. Specifically:

$$\frac{\text{Cov}(B_i, \Delta Y_i)}{\text{Cov}(B_i, \Delta D_i)} = \frac{\text{Cov}(D_{i0}, \Delta Y_i)}{\text{Cov}(D_{i0}, \Delta D_i)}.$$

I next examine the assumptions under which this coefficient identifies a causal parameter of interest.

Using the equivalent instrument $1 - D_{i0}$, the reduced-form coefficient is the exposure-design estimand β defined in Section 5:

$$\beta = \mathbb{E}[\Delta Y_i \mid D_{i0} = 0] - \mathbb{E}[\Delta Y_i \mid D_{i0} = 1].$$

This is a difference-in-differences estimand that compares the change in mean outcomes among baseline-untreated units with the change among baseline-treated units. The corresponding first-stage coefficient is

$$\begin{aligned} &\mathbb{E}[\Delta D_i \mid D_{i0} = 0] - \mathbb{E}[\Delta D_i \mid D_{i0} = 1] \\ &\qquad = \Pr(D_{i1} = 1 \mid D_{i0} = 0) + \Pr(D_{i1} = 0 \mid D_{i0} = 1). \end{aligned}$$

Dividing the reduced form by the first stage gives

$$\frac{\mathrm{Cov}(D_{i0}, \Delta Y_i)}{\mathrm{Cov}(D_{i0}, \Delta D_i)} = \frac{\mathbb{E}[\Delta Y_i \mid D_{i0} = 0] - \mathbb{E}[\Delta Y_i \mid D_{i0} = 1]}{\mathbb{E}[\Delta D_i \mid D_{i0} = 0] - \mathbb{E}[\Delta D_i \mid D_{i0} = 1]}$$
$$= \frac{\beta}{\Pr(D_{i1} = 1 \mid D_{i0} = 0) + \Pr(D_{i1} = 0 \mid D_{i0} = 1)}.$$

This is the Wald-DiD estimand studied by De Chaisemartin and d'Haultfoeuille (2018): the DiD in outcomes divided by the DiD in treatment.

Next, I examine what causal estimand is identified under Assumption PT. The mean change in untreated potential outcomes is then the same across the two baseline treatment groups and cancels from the reduced form. Expanding observed outcomes in terms of potential outcomes gives

$$\begin{aligned}\beta = \; & \Pr(D_{i1} = 1 \mid D_{i0} = 0)\mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 0, D_{i1} = 1] \\ & + \Pr(D_{i1} = 0 \mid D_{i0} = 1)\mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1, D_{i1} = 0] \\ & - \mathbb{E}\left[\big(Y_{i1}(1) - Y_{i1}(0)\big) - \big(Y_{i0}(1) - Y_{i0}(0)\big) \mid D_{i0} = 1\right].\end{aligned}$$

The first two terms contain the endline treatment effects of observed entrants and exits, respectively. The final term is the change in the average treatment effect among all baseline-treated units. To obtain the Wald-DID estimand, divide this expression by

$$\Pr(D_{i1} = 1 \mid D_{i0} = 0) + \Pr(D_{i1} = 0 \mid D_{i0} = 1).$$

The first two terms then form a convex average of endline treatment effects, with weights equal to the two transition probabilities divided by their sum. The final term, divided by the same denominator, is the discrepancy from that average. Consequently, under Assumption PT and a nonzero first stage, the Wald-DID estimand equals this particular convex weighted average if and only if

$$\mathbb{E}[Y_{i1}(1) - Y_{i1}(0) \mid D_{i0} = 1] = \mathbb{E}[Y_{i0}(1) - Y_{i0}(0) \mid D_{i0} = 1].$$

This is time invariance of the average treatment effect among baseline-treated units. This restriction supplies a weak causal interpretation: the Wald-DID coefficient averages endline treatment effects with nonnegative weights summing to one. Under Assumption PT alone, changes in treatment effects can instead cause the coefficient to be negative even when every treatment effect is positive. For example, let untreated potential outcomes be constant, let half of the baseline-untreated units enter treatment, and let no baseline-treated units exit. Suppose treatment effects equal one for baseline-untreated units at both dates, while effects among baseline-treated units rise from one to two. Then $\beta = 1/2 - 1 = -1/2$ and the first stage is $1/2$, so the Wald-DID coefficient is -1, although all treatment effects are positive. Time invariance can fail even when economic conditions can alter the return to treatment.

For example, higher output prices increase the monetary gain from a fixed productivity improvement; a protective treatment prevents greater expected losses when the underlying risk increases; and access to subsidized credit becomes less beneficial when market borrowing costs fall.

N Identification with Multiple Pre- and Post-Policy Periods

This section extends the identification results in the paper to a setting with multiple pre- and post-policy periods. I begin by identifying the period-specific LATE under restrictions on selection and treatment-effect heterogeneity respectively. I then explain how pre-policy outcomes can be leveraged to provide evidence in support of the plausibility of the parallel-trends assumption used for identification.

Suppose panel data are available for units i and periods $t \in \{-\underline{T}, \ldots, 0, \ldots, \overline{T}\}$. The policy is inactive for $t \leq 0$ and active for $t \geq 1$: $Z_t = 0$ before and at baseline, and $Z_t = 1$ afterward. Period 0 is used as the reference period. Let $D_{it}(z) \in \{0, 1\}$ denote the potential treatment decision under policy state z, and let $Y_{it}(d)$ denote the potential outcome under treatment status d. Observed variables satisfy

$$D_{it} = D_{it}(Z_t), \qquad Y_{it} = Y_{it}(D_{it}).$$

For each post-policy period $t \in \{1, \ldots, \overline{T}\}$, the target is

$$\lambda_t = \mathbb{E}\left[Y_{it}(1) - Y_{it}(0) \mid D_{it}(1) > D_{it}(0)\right]. \tag{N.1}$$

This is the effect of treatment among units induced into treatment by the policy in period t. The compliant subpopulation may change across periods, so changes in λ_t need not describe changes in treatment effects for a fixed population. I adapt Assumption PT to a setting with multiple post-policy periods as follows:

Assumption PT (Post-Policy Parallel Trends)**.** For every $t \in \{1, \ldots, \overline{T}\}$,

$$\mathbb{E}\left[Y_{it}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{it} = 1\right] = \mathbb{E}\left[Y_{it}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{it} = 0\right].$$

Define the estimand:

$$\tau_t := \mathbb{E}\left[Y_{it} - Y_{i0} \mid D_{i0} = 0, D_{it} = 1\right] - \mathbb{E}\left[Y_{it} - Y_{i0} \mid D_{i0} = 0, D_{it} = 0\right]. \tag{N.2}$$

Under Assumption PT,

$$\tau_t = \mathbb{E}\left[Y_{it}(1) - Y_{it}(0) \mid D_{i0} = 0, D_{it} = 1\right]. \tag{N.3}$$

Next, I describe the additional assumptions under which τ_t identifies λ_t. The selection approach connects observed treatment paths to latent compliance types. The following assumption augments Assumption SEL for a setting with multiple post-policy periods.

Assumption SEL (Selection Restrictions)**.** For every $t \in \{1, \ldots, \overline{T}\}$:

(a) Monotonicity: $\mathbb{P}[D_{i0}(1) \geq D_{i0}(0)] = \mathbb{P}[D_{it}(1) \geq D_{it}(0)] = 1$

(b) Stability in the share of always-takers: $\mathbb{P}[D_{it}(0) = D_{it}(1) = 1] = \mathbb{P}[D_{i0}(0) = D_{i0}(1) = 1]$

(c) Stability of Choosers: $\mathbb{P}[D_{it}(0) = 0 \mid D_{i0}(0) = 1] = 0$

Re-applying the logic of Lemma 1, Assumption SEL implies $D_{it}(0) = D_{i0}(0)$ almost surely. Hence,

$$\{D_{it}(1) > D_{it}(0)\} = \{D_{i0} = 0, D_{it} = 1\} \quad \text{almost surely.}$$

Therefore, τ_t identifies λ_t under Assumptions PT and SEL. For completeness, Assumptions SEL(a)–(b) also identify the first stage:

$$\mathbb{E}[D_{it}(1) - D_{it}(0)] = \mathbb{P}[D_{it} = 1] - \mathbb{P}[D_{i0} = 1].$$

Assumption SEL(c) states that a unit who selects into treatment at baseline without the policy would continue to do so in period t without the policy: $\mathbb{P}[D_{it}(0) = 1 \mid D_{i0}(0) = 1] = 1$. This restricts only baseline-treated units. To see what it implies for baseline-untreated units, decompose the period-t share of always-takers by baseline status:

$$\begin{aligned}\mathbb{P}[D_{it}(0) = 1] &= \mathbb{P}[D_{it}(0) = 1 \mid D_{i0}(0) = 1]\,\mathbb{P}[D_{i0}(0) = 1] + \mathbb{P}[D_{it}(0) = 1 \mid D_{i0}(0) = 0]\,\mathbb{P}[D_{i0}(0) = 0] \\ &= \mathbb{P}[D_{i0}(0) = 1] + \mathbb{P}[D_{it}(0) = 1 \mid D_{i0}(0) = 0]\,\mathbb{P}[D_{i0}(0) = 0],\end{aligned}$$

where the second line imposes Assumption SEL(c). Assumption SEL(b) equates the left-hand side with $\mathbb{P}[D_{i0}(0) = 1]$, so that

$$\mathbb{P}[D_{it}(0) = 1 \mid D_{i0}(0) = 0] = 0 \tag{N.4}$$

whenever $\mathbb{P}[D_{i0}(0) = 0] > 0$. Because $Z_0 = 0$, the conditioning event equals $\{D_{i0} = 0\}$, so Equation (N.4) states that units untreated at baseline would remain untreated in period t absent the policy. Within the selection approach, Equation (N.4) is the restriction that guarantees any unit with $D_{i0} = 0$ and $D_{it} = 1$ is a period-t complier.

A single untreated baseline observation may provide weak grounds for imposing Equation (N.4). With access to additional pre-policy periods, one may address this concern by conditioning on a longer untreated history. For $L \in \{0, \ldots, \underline{T}\}$, define

$$\overline{D}_i^L := \max\left\{D_{i,-L}, \ldots, D_{i0}\right\},$$

so that $\overline{D}_i^L = 0$ if and only if unit i is untreated in every period from $-L$ through 0, and $\overline{D}_i^0 = D_{i0}$. Because $Z_s = 0$ for all $s \leq 0$, $\overline{D}_i^L = 0$ is equivalent to $D_{is}(0) = 0$ for all $s \in \{-L, \ldots, 0\}$. The following assumption states that units untreated throughout the pre-policy window would remain untreated had the policy not been switched on:

Assumption SEL-L (Stability of Non-Choosers within the Pre-Policy Window)**.** For every post-policy period $t \in \{1, \ldots, \overline{T}\}$:

$$\mathbb{P}\big[D_{it}(0) = 1 \mid \overline{D}_i^L = 0\big] = 0.$$

At $L = 0$, Assumption SEL-L coincides with Equation (N.4). For $L > 0$, it is imposed on a smaller subpopulation, so that Assumption SEL(c) is no longer required. Assumption

SEL-L becomes more plausible as L grows. The economic content of this argument is that demand for treatment has persistent components, so repeated decisions to remain untreated at prevailing costs is informative about future treatment status at those costs. In the Medicare Part D application in Section 6.2, demand for prescription-drug coverage depends on chronic health conditions, risk preferences, and income. An elderly individual who forgoes coverage in both 2002 and 2004 may reveal a low valuation of coverage or a persistently high cost of obtaining it, and is unlikely to obtain coverage in 2006 without a reduction in that cost. Define the estimand

$$\tau_{t,L} := \mathbb{E}\left[Y_{it} - Y_{i0} \mid \overline{D}_i^L = 0,\, D_{it} = 1\right] - \mathbb{E}\left[Y_{it} - Y_{i0} \mid \overline{D}_i^L = 0,\, D_{it} = 0\right], \tag{N.5}$$

which replaces the conditioning event $\{D_{i0} = 0\}$ in Equation (N.2) with $\{\overline{D}_i^L = 0\}$, so that $\tau_{t,0} = \tau_t$. I modify Assumption PT to condition on the pre-policy history:

Assumption PT-L (Post-Policy Parallel Trends within the Pre-Policy Window)**.** For every post-policy period $t \in \{1, \ldots, \overline{T}\}$,

$$\mathbb{E}\left[Y_{it}(0) - Y_{i0}(0) \mid \overline{D}_i^L = 0,\, D_{it} = 1\right] = \mathbb{E}\left[Y_{it}(0) - Y_{i0}(0) \mid \overline{D}_i^L = 0,\, D_{it} = 0\right].$$

At $L = 0$, Assumption PT-L coincides with Assumption PT. For $L > 0$, it is a restriction on the continuously untreated subpopulation and is neither implied by nor implies Assumption PT. Adding and subtracting $\mathbb{E}[Y_{it}(0) \mid \overline{D}_i^L = 0, D_{it} = 1]$ in Equation (N.5) and applying Assumption PT-L yields

$$\tau_{t,L} = \mathbb{E}\left[Y_{it}(1) - Y_{it}(0) \mid \overline{D}_i^L = 0,\, D_{it} = 1\right]. \tag{N.6}$$

This is the average effect of treatment among continuously untreated units who are treated in period t. Absent further restrictions, this group is a mixture of period-t compliers and period-t always-takers, since a continuously untreated unit may have $D_{it}(0) = 1$. Assumption SEL-L rules out the latter possibility. To state the resulting causal parameter, define the history-restricted LATE

$$\lambda_{t,L} := \mathbb{E}\left[Y_{it}(1) - Y_{it}(0) \mid D_{it}(1) > D_{it}(0),\, \overline{D}_i^L = 0\right], \tag{N.7}$$

the average effect of treatment among period-t compliers who were untreated throughout the pre-policy window. Under Assumption SEL-L, $D_{it}(0) = 0$ almost surely on the event $\{\overline{D}_i^L = 0\}$. Because $D_{it} = D_{it}(1)$ for $t \geq 1$, it follows that

$$\{\overline{D}_i^L = 0,\, D_{it} = 1\} = \{\overline{D}_i^L = 0,\, D_{it}(1) > D_{it}(0)\} \text{ almost surely.}$$

Substituting this equality into Equation (N.6) gives $\tau_{t,L} = \lambda_{t,L}$. Therefore, $\tau_{t,L}$ identifies $\lambda_{t,L}$ under Assumptions PT-L and SEL-L.

The parameter $\lambda_{t,L}$ differs from λ_t in Equation (N.1) by conditioning on the pre-policy history. By the law of iterated expectations,

$$\lambda_t = \lambda_{t,L}\,\mathbb{P}\big[\overline{D}_i^L = 0 \mid D_{it}(1) > D_{it}(0)\big] + \mathbb{E}\left[Y_{it}(1) - Y_{it}(0) \mid D_{it}(1) > D_{it}(0),\, \overline{D}_i^L = 1\right]\mathbb{P}\big[\overline{D}_i^L = 1 \mid D_{it}(1) > D_{it}(0)\big].$$

The two parameters therefore coincide when every period-t complier is continuously untreated over the window, so that the second term vanishes, or when the average effect among compliers does not depend on their pre-policy history. The following assumption delivers the former by extending Assumption SEL to every adjacent pair of periods in the panel, including the pre-policy periods.

Assumption SEL-M$^+$ (Selection Restrictions over the Entire Panel)**.**

(a) $\mathbb{P}[D_{is}(1) \geq D_{is}(0)] = 1$ for every $s \in \{-\underline{T}, \ldots, \overline{T}\}$
(b) $\mathbb{P}[D_{is}(0) = D_{is}(1) = 1] = \mathbb{P}[D_{i,s-1}(0) = D_{i,s-1}(1) = 1]$ for every $s \in \{-\underline{T}+1, \ldots, \overline{T}\}$.
(c) $\mathbb{P}[D_{is}(0) = 0 \mid D_{i,s-1}(0) = 1] = 0$ for every $s \in \{-\underline{T}+1, \ldots, \overline{T}\}$.

Applying the logic of Lemma 1 to each adjacent pair of periods gives $D_{is}(0) = D_{i0}(0)$ almost surely for every s. A period-t complier has $D_{it}(0) = 0$ and hence $D_{is} = D_{is}(0) = 0$ for every $s \leq 0$, so $\mathbb{P}[\overline{D}_i^L = 1 \mid D_{it}(1) > D_{it}(0)] = 0$ and $\lambda_{t,L} = \lambda_t$ for every L. By the same argument, $\{\overline{D}_i^L = 0\} = \{D_{i0} = 0\}$ almost surely, so $\tau_{t,L} = \tau_t$ and Assumption PT-L reduces to Assumption PT. Assumption SEL-M$^+$ implies Assumptions SEL and SEL-L.

An alternative approach allows unrestricted selection transitions but restricts average treatment effects across latent groups. The following assumption adapts Assumption HET to a setting with multiple post-policy periods.

Assumption HET (Treatment Effect Homogeneity)**.** For every $t \in \{1, \ldots, \overline{T}\}$ and every $d_0, d_1 \in \{0, 1\}$,

$$\mathbb{E}\left[Y_{it}(1) - Y_{it}(0) \mid D_{i0} = d_0,\, D_{it}(0) = d_1,\, D_{it}(1) = 1\right] = \mathbb{E}\left[Y_{it}(1) - Y_{it}(0) \mid D_{it}(1) = 1\right].$$

Assumption HET states that, among units treated under the policy in period t, the average treatment effect does not vary with baseline treatment status or with the counterfactual period-t treatment decision. Individual effects may vary within these groups, and average effects may vary across periods.

Under Assumption PT, Equation (N.3) shows that τ_t identifies the average effect among units with $D_{i0} = 0$ and $D_{it} = D_{it}(1) = 1$. Averaging Assumption HET over $D_{it}(0)$ within this group gives $\tau_t = \mathbb{E}[Y_{it}(1) - Y_{it}(0) \mid D_{it}(1) = 1]$. Averaging Assumption HET over D_{i0} among period-t compliers gives $\lambda_t = \mathbb{E}[Y_{it}(1) - Y_{it}(0) \mid D_{it}(1) = 1]$. Therefore, $\tau_t = \lambda_t$ under Assumptions PT and HET.

Assessing Parallel Trends Using Pre-Policy Outcomes A natural way to assess the plausibility of Assumption PT is to ask whether baseline-untreated units who were treated at endline exhibited the same outcome trends before the policy as those who were continuously untreated.[29] A feature of this setting that complicates the exercise is that the absence of the policy does not imply the absence of treatment: units may be treated in pre-policy periods, so observed pre-policy outcomes need not be untreated outcomes. For a baseline-untreated unit and $s < 0$,

$$Y_{is} - Y_{i0} = \big(Y_{is}(0) - Y_{i0}(0)\big) + D_{is}\big(Y_{is}(1) - Y_{is}(0)\big). \tag{N.8}$$

The observed pre-policy contrast between the two groups therefore equals the contrast in untreated trends plus the difference across groups in the mean of $D_{is}(Y_{is}(1) - Y_{is}(0))$. Neither Assumption SEL nor Assumption HET constrains this second term. A pre-trends test based on observed outcomes therefore conflates differences in untreated trends with differences in pre-policy treatment.

Restricting attention to continuously untreated units resolves this difficulty. Conditioning on the event $\overline{D}_i^L = 0$, every outcome observed in periods $s \in \{-L, \ldots, 0\}$ is an untreated outcome, so the second term in Equation (N.8) vanishes and the pre-policy counterpart of Assumption PT-L is directly testable

$$\mathbb{E}\left[Y_{is} - Y_{i0} \mid \overline{D}_i^L = 0,\, D_{it} = 1\right] = \mathbb{E}\left[Y_{is} - Y_{i0} \mid \overline{D}_i^L = 0,\, D_{it} = 0\right]. \tag{N.9}$$

for $s \in \{-L, \ldots, -1\}$.[30]

An alternative that avoids conditioning on the full pre-policy history is to condition only on treatment status in the two periods whose outcomes enter the comparison. For each $s < 0$, the relevant restriction imposed by parallel trends is:

$$\mathbb{E}\left[Y_{is} - Y_{i0} \mid D_{is} = D_{i0} = 0,\, D_{it} = 1\right] = \mathbb{E}\left[Y_{is} - Y_{i0} \mid D_{is} = D_{i0} = 0,\, D_{it} = 0\right]. \tag{N.10}$$

Note that for units with $D_{is} = D_{i0} = 0$, $Y_{is} = Y_{is}(0)$ and $Y_{i0} = Y_{i0}(0)$, so Equation (N.10) is a restriction on untreated potential outcomes even though units may be treated in periods between s and 0. The conditioning event varies with s, however, so each pre-policy coefficient tests parallel trends for a different subpopulation.

[29] Assumption PT restricts untreated outcome trends only between baseline and the post-policy period t, which is all that identification of λ_t requires. A pre-trends test instead examines the corresponding restriction for pre-policy outcome dates, $\mathbb{E}[Y_{is}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{it} = 1] = \mathbb{E}[Y_{is}(0) - Y_{i0}(0) \mid D_{i0} = 0, D_{it} = 0]$ for $s < 0$, holding fixed the groups defined by D_{it}.

[30] Caution is needed when interpreting these tests. They concern pre-policy restrictions and cannot establish Assumption PT, which restricts post-policy trends. Pre-trends tests also have limited power against economically meaningful violations, and conditioning the analysis on passing them can distort inference (Roth, 2022).