

جامعة نيويورك أبوظبي

 NYU | ABU DHABI



Stage-Based Identification of Policy Effects

**Christian Alemán, Christopher Busch,
Alexander Ludwig and
Raül Santaeulàlia-Llopis**

Working Paper # 0094

October 2023

Division of Social Science Working Paper Series

New York University Abu Dhabi, Saadiyat Island P.O Box 129188, Abu Dhabi, UAE

<http://nyuad.nyu.edu/en/academics/academic-divisions/social-science.html>

Stage-Based Identification of Policy Effects*

Christian Alemán
NYU Abu Dhabi

Christopher Busch
LMU Munich
CESifo

Alexander Ludwig
Goethe University Frankfurt
ICIR, UAB, BSE and CEPR

Raül Santaaulàlia-Llopis
NYU Abu Dhabi
UAB, BSE and CEPR

October 16, 2023

Abstract

We develop a method that identifies the effects of nationwide policy, i.e., policy implemented across all regions at the same time. The core idea is to track outcome paths in terms of stages rather than time, where a stage of a regional outcome at time t is its location on the support of a reference path. The method proceeds in two steps. First, a normalization maps the time paths of regional outcomes onto the reference path—using only pre-policy data. This uncovers cross-regional heterogeneity of the stage at which policy is implemented. Second, this stage variation identifies policy effects inside a window of stages where a stage-leading region provides the no-policy counterfactual path for non-leading regions that are subject to policy inside that window. We assess our method's performance with Monte-Carlo experiments, illustrate it with empirical applications, and show that it captures heterogeneous policy effects across stages.

Keywords: Stages, Identification, Policy Effects, Nationwide Policy, Macroeconomics
JEL Classification: C01, E00

*We thank Alberto Abadie, Jaap Abbring, Francesco Agostinelli, Manuel Arellano, Dmitry Arkhangelsky, Stephane Bonhomme, Brant Callaway, Jeffrey Campbell, Fabio Canova, Davide Cantoni, Xu Cheng, Hanming Fang, Hans-Peter Kohler, Jakob Larsen, Fabrizia Mealli, Espen Moen, Jordan Norris, José-Víctor Ríos-Rull, Jonathan Roth, Frank Schorfheide and participants at seminars and conferences for helpful comments and suggestions. Raül Santaaulàlia-Llopis acknowledges financial support from the AGAUR 2020PANDE00036 "Pandemies" Grant 2020-2022, from the Spanish Ministry of Economy and Competitiveness, through the Proyectos I+D+i 2019 Retos Investigación PID2019-110684RB-I00 Grant, Europa Excelencia EUR2021-122011 and the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S).

Alemán: christian.c.aleman@gmail.com; Busch: chris.busch.econ@gmail.com; Ludwig: mail@alexander-ludwig.com; Santaaulàlia-Llopis: loraulet@gmail.com.

[GitHub site](#) for our SBI toolbox.

1 Introduction

Motivation. The empirical assessment of a policy requires a credible counterfactual. Standard empirical strategies critically rely on cross-regional heterogeneity of the time of policy implementation as source of identification—e.g., the existence of one untreated region or a staggered rollout.¹ The counterfactual is then based on the untreated region, typically assuming that the regional paths of the outcome of interest are similar with differences not exceeding a constant gap over time—the so-called parallel trends assumption.² However, many relevant policy contexts violate these conditions. First, many policies are implemented nationwide, carried out in all regions at the same time, which eliminates the source of identification for standard empirical strategies. Second, the pre-policy outcome paths can differ across regions—e.g., in their starting date, speed, or magnitude—in a way that violates the parallel trends assumption. In panel (a) of Figure 1, we illustrate one such scenario. Our goal in this paper is to provide an identification strategy for such policy contexts.

Idea. We develop the idea of tracking a regional outcome over stages rather than time. Specifically, we define a stage of a regional outcome at time t as its location on the support of a reference outcome path. Panel (b) of Figure 1 illustrates this concept by showing a reference outcome path evolving through stages. At any given time, the outcome of one region may be at a different stage compared to another region. In the figure, at time t region \mathcal{C} has progressed to a more advanced stage than region \mathcal{T} . This implies that cross-regional heterogeneity in stages may be present at the time of policy implementation, t_p . The core of our method is tracing out and exploiting this heterogeneity of stages for the identification of policy effects; hence the label *Stage-Based Identification* (SBI).

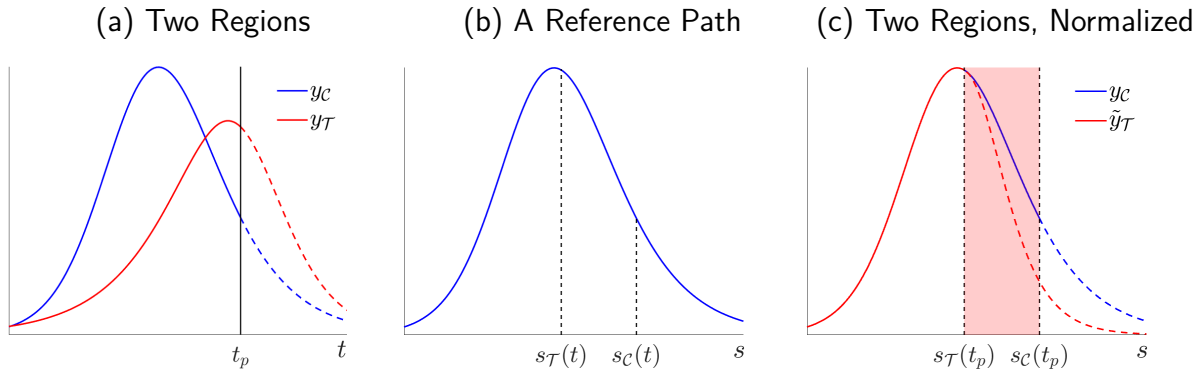
The Method. Our method comprises two steps: normalization and identification. First, we normalize the pre-policy outcome time paths of regions to a reference path. This normalization step affects, simultaneously, the transformation of time into normalized time (i.e., stages) and the transformation of the outcome level into a normalized level.³ We conduct these transformations using low degree polynomials (one for the stage-to-time transformation and one for the level-to-normalized-level conversion) and refer to the set of associated coefficients as the normalization coefficients. In particular, in our theoretical derivations and empirical applications, we focus on monotonic (linear) transformations. For a given region, the normalization coefficients are

¹E.g., [Blundell and Macurdy \(1999\)](#), [Angrist and Krueger \(1999\)](#), [Athey and Imbens \(2017\)](#) and [Card \(2022\)](#).

²The assumption of parallel trends is relaxed in [Abadie \(2005\)](#). Indeed, there is a growing and exciting body of research work regarding more flexible forms of parallel trends (e.g., [Callaway and Sant’Anna, 2021](#); [Rambachan and Roth, 2023](#); [Roth and Sant’Anna, 2023](#)).

³This transformation of time and level coordinates relates to earlier work in [Iorio and Santaeuilàlia-Llopis \(2016\)](#).

Figure 1: A Stage-Based Identification of the Effects of Nationwide Policy: An Illustration



Notes: Panel (a) shows the time paths of an outcome variable for two regions, \mathcal{C} and \mathcal{T} ; policy is implemented in both regions at the same time t_p ; dashed sections indicate post-policy paths. Panel (b) shows a reference path, where at some t , region \mathcal{C} is at a more advanced stage than region \mathcal{T} . Panel (c) shows the result of mapping $y_{\mathcal{T}}$ onto $y_{\mathcal{C}}$, resulting in the normalized path $\tilde{y}_{\mathcal{T}}$; the pink shaded area indicates the identification window.

determined as the ones that minimize the distance between the normalized pre-policy outcome path and the reference path. As a result, up to a minimization error, the regional pre-policy paths are identical in the stage domain before the stage at which policy is implemented first across regions.⁴ Panel (c) of Figure 1 shows this for an example where we use the outcome time path of region \mathcal{C} as the reference path and normalize the path of region \mathcal{T} to this reference path. When selecting one regional path as reference path, we refer to the corresponding region as *reference region*, and to the other region(s) as *non-reference region(s)*. This normalization procedure uncovers cross-regional heterogeneity of the stages at which policy is implemented.

Second, we base identification on this cross-regional heterogeneity. For example, in the illustration, policy is implemented at an earlier stage in region \mathcal{T} than in region \mathcal{C} . Then, applying the normalization coefficients—resulting from pre-policy data—on post-policy data opens a window of stages, in which a stage-leading region (in the example, region \mathcal{C}) is not subject to policy whereas the other region (in the example, region \mathcal{T}) is; see pink shaded area in panel (c) of Figure 1. Our identification assumption is that the normalization coefficients that minimize the distance between the pre-policy outcome paths across regions in the stage domain are unaffected by policy. That is, we assume that the normalization would make the post-policy paths line up as well in the absence of policy. Thus, under our identification assumption, the stage-leading region serves as *control region*, and provides the no-policy counterfactual for the other region (the *treatment region*).⁵ The difference between the outcomes of the control region and the treatment region inside the identification window captures the effects of policy.

⁴We show analytical examples where the normalization coefficients can be uniquely identified in Section 2.3. These closed-form solutions provide an interpretation for the normalization coefficients as those that reshape the structural parameters of the non-reference region into those of the reference region before the policy takes place.

⁵Note that since the stage at which each region receives the policy is a result of our normalization, SBI does not require an ex-ante assignment of control or treatment across regions, see Section 2.2.

Method Performance. We use a range of illustrative economic models for different nationwide policy scenarios, in which we simulate the effect of some policy on an outcome variable of interest. Within the model framework, we know the true counterfactual path that would occur if policy were absent throughout. Using the (simulated) data that would be available to a policy evaluator, we find that SBI can successfully identify the true (model-generated) effect of the policy.

We further assess whether and when our normalization procedure comes to its limits and plausible identification is not feasible using SBI. Precisely, we perform a Monte Carlo study that numerically characterizes the bounds within which SBI is able to recover the true (model-generated) effects of policy. We consider one benchmark region which we pair up with a large set of regions (drawing from a large set of structural parameters) one-by-one. Applying SBI to these pairs, we find that the error in the SBI policy effect relative to the true policy effect systematically increases when moving farther away from the benchmark region in the space of the normalization coefficients. Under the interpretation that the normalization reshapes the structural parameters that determine the outcome paths, SBI requires—for successful identification of policy effects—the structural parameters to not be too dissimilar across regions in the absence of policy. In addition, we assess how SBI fares in contexts where there are potential confounding factors such as time-varying latent heterogeneity and confounding policy. Using model-generated data, we find that SBI is able to identify the true policy effects in these contexts as long as the confounding factors keep the regional outcome paths sufficiently close in the space of the normalization coefficients. We further establish in a set of Placebo diagnoses that SBI successfully identifies a zero policy effect when there is none. We also show how to conduct inference on our identified policy effects with data that incorporates a stochastic component.

Three Applications. We apply SBI to study the effects of nationwide policy in three empirical applications. First, we assess the effectiveness of the stay-home policy implemented nationwide in response to the first wave of the Covid-19 pandemic in Spain. SBI assigns Madrid as the stage-leading (hence, control) region at the time of policy implementation. We find that the stay-home policy significantly reduces the amount of deaths by 24.71% in the rest of Spain inside an identification window of seven days. In other words, had the stay-home policy not been implemented, there would have been 1,734 more deaths over the course of one week. Second, we assess the effects of the Food and Drugs Administration (FDA) approval of oral contraceptives (the pill) in the United States (U.S.) in 1960. SBI assigns West Virginia as stage-leading at the time of the policy. We find that the pill reduced the crude fertility rate (number of births per 10,000) by 9.38% in rest of the U.S. in an identification window of about one decade. We also consider the effect of the pill on female education. For this outcome of interest, the proportion of college women, SBI assigns Washington DC as the stage-leading region. We find that the pill

increased the proportion of college women by 21.73% in the rest of the U.S. in an identification window of a bit over one decade that followed the FDA approval. Third, we study the effects of the German reunification in 1990 on GDP per capita in West Germany where SBI assigns Hessen as the stage-leading region rest of West Germany. Using the path of GDP per capita of Hessen as no-policy counterfactual, we find that the German reunification significantly reduced GDP per capita in the rest of West Germany by 3.51% in a window of approximately seven years.

Heterogeneous Effects. We also map the outcome paths of the non-leading regions one-by-one onto the path of the leading region. This uncovers regional heterogeneity of the stage at policy implementation and allows us to assess how the policy effects potentially differ by stage. Focusing on the Covid-19 application, we find that the policy effects systematically vary across stages. The earlier (in stages) the policy is implemented, the larger are the effects: the amount of prevented deaths is 65% in Murcia where the policy is implemented two weeks earlier in the stage domain than in the leading region, Madrid, and 12% in the Basque Country where policy arrives two days earlier in the stage domain than Madrid. Both the size of the identification window (policy implemented at earlier stages implies a larger window) and the interim policy effects (i.e., effects for the same horizon inside the identification window) contribute to generate the heterogeneous effects. We reach similar insights using a set of artificial regions from the power set of regions.⁶

Non-Nationwide Policy. Our identification method can also be applied to assess the effects of non-nationwide policy, as it is typically done in other standard empirical strategies. Importantly, in the case of non-nationwide policy our identification is still based on the stage (not the time) at which the policy is implemented. When there are regions that are treated and regions that are not, SBI delivers a right-open identification window bounded from below with the stage at which the policy is implemented in the treated region and unbounded from above. We exemplify this approach using the empirical application of the German reunification where we use alternative potential controls for West Germany built from other OECD countries as in [Abadie et al. \(2014\)](#), assuming, as in their work, that the German reunification had no effect on GDP per capita in these countries.

Related Literature. Our method is directly related to the empirical strategies designed for settings that resemble natural experiments. These strategies rely on a difference-in-differences approach in order to generate the counterfactual path (or potential outcome as in [Imbens and](#)

⁶In the same vein, we also use regional heterogeneity in the stage of policy implementation to conduct further analysis: for some regions there are several possible control regions—namely, all regions that receive the policy at a later stage.

Rubin, 2015) that serves as control for a treated region—i.e., the region subject to policy. We emphasize two main differences of SBI. First, a critical common factor in previous strategies is that the source of identification is the heterogeneity in the time of policy implementation across regions with the existence of either one untreated region (e.g., Card, 1990; Card and Krueger, 2000) or a staggered policy adoption (e.g., Athey and Imbens, 2021; Borusyak et al., 2021; Goodman-Bacon, 2021). This is not the case in SBI. Precisely, our main point of departure with respect to previous work is that SBI is able to deliver identification of policy effects for contexts where the cross-regional heterogeneity of the time of policy implementation is absent.

Second, one concern when identifying policy effects is that cross-regional differences in the pre-policy determinants of the outcome of interest may influence post-policy outcomes. One typical approach is to argue for the plausibility of the parallel trends assumption based on the presence of “parallel” pre-policy outcome time paths across regions (e.g., Bertrand et al., 2004). In contrast, our approach does not resort to parallel pre-policy outcome time paths to support the identification assumption. Instead, through a normalization of time and level, we minimize the distance of pre-policy outcome paths across regions in the stage domain. By doing so, our method aims to control for cross-regional differences in the pre-policy determinants of outcomes, regardless of their observability. Our identification assumption is that, in the absence of policy, the normalization coefficients obtained from pre-policy data would also map to the post-policy paths of non-reference regions onto the reference path. By applying these coefficients to post-policy data, we create a counterfactual path to identify policy effects. This also makes our approach complementary to a growing literature that invokes a less strict notion of parallel trends and explores identification when parallel trends do not exactly hold. Abadie (2005) conditions the parallel trends to a set of observables using propensity scores (Heckman et al., 1998). This idea is extended to staggered policy in Callaway and Sant’Anna (2021). Rambachan and Roth (2023) provide confidence sets for the identified effects building on how much the pre-policy trends potentially differ from the trends after policy. Roth and Sant’Anna (2023) characterize how the parallel trends assumption survives strictly monotonic transformations under a stronger condition on the cumulative distribution of the no-policy counterfactual.

Our work also relates to other policy evaluation approaches like the original synthetic control methods (SCM) (Abadie and Gardeazabal, 2003; Abadie et al., 2010).⁷ The SCM approach essentially constructs a counterfactual time path based on a weighted average across untreated (control) regions. Two main differences stand out. First, analogous to other empirical strategies, SCM requires the existence of a set of untreated regions to construct the synthetic control group for identification. In contrast, SBI relies on cross-regional heterogeneity in the stage—not

⁷Arkhangelsky et al. (2021) propose a SCM-like weighting scheme for more flexibility regarding parallel trends.

time—at which the policy is implemented. For this reason, we can apply SBI to a nationwide policy occurring at the same time across all regions, unlike with SCM—or other methods for that matter. We illustrate this aspect in the context of the German reunification application in Section 5.2. Second, our method does not require the use of observable determinants of cross-regional outcome differences in order to generate the counterfactual. Instead, the counterfactual in our method is constructed using solely the time paths of the outcome of interest. Also, similar to SBI, the changes-in-changes method in [Athey and Imbens \(2006\)](#) features a mapping of outcomes across regions. Their focus lies on capturing heterogeneity of the policy effect over the cross-sectional distribution of an individual level outcome. To this end, they map pre-policy cumulative cross-sectional distributions across regions and use this to construct the counterfactual distribution in the treated region.^{8,9} Instead of cross-sectional distributions, we map pre-policy regional outcome time paths. Further, the main difference described above remains, namely that the identification in changes-in-changes is also based on regional heterogeneity of the time of policy implementation, whereas SBI does not require that heterogeneity.

Turning to our notion of *stages*, in the analyses of structural transformation (e.g., [Galor and Weil, 2000](#); [Herrendorf et al., 2014](#); [Cervellati and Sunde, 2015](#)) or of the demographic transition (e.g., [Greenwood et al., 2005](#)), the level of income per capita typically summarizes the “stage” of development for an outcome of interest (e.g., agricultural share of output, urbanization rates, population growth rates, etc.) in cross-country comparisons. In contrast, rather than replacing time for an observable such as income per capita, SBI provides a normalization of the time path of the outcome of interest (possibly income per capita itself: see our evaluation of the effect of the German reunification on income per capita in West Germany). This implies that the level of the outcome of interest (e.g., income per capita) is not a sufficient statistic to define the stage of a region in our approach. This same argument is discussed in the earlier work of [lorio and Santaeulàlia-Llopis \(2010, 2016\)](#) that also conducts a normalization mapping country-specific time paths of HIV prevalence onto a reference path in order to define stages of the HIV epidemic. We depart from that work in that we use our normalization to a reference path as base for identifying the effects of policies that aim to alter the path of the outcome of interest. For this reason, our normalization coefficients are obtained by strictly using pre-policy outcome paths.

Finally, note that we consider scenarios which have been studied in the literature typically using macroeconomic models, that is, where a policy is implemented nationwide, i.e. across all regions at the same time. In terms of our first application, recent econ-epi frameworks highlight the link between aggregate economic activity and a pandemic in order to study related policy

⁸This analysis is related to [Altonji and Blank \(1999\)](#), who consider a decomposition of relative wage changes across groups into changes of the distribution of skills and the payoff for those skills.

⁹Recently, [Gunsilius \(2023\)](#) presents a synthetic control variant of the changes-in-changes estimator.

trade-offs (e.g., in [Eichenbaum et al., 2021](#); [Farboodi et al., 2021](#); [Glover et al., 2020](#); [Kaplan et al., 2020](#)). Turning to our second application of birth-control technology, the general trade-offs faced by women are discussed in [Doepke et al. \(2023\)](#). Examples for macroeconomic frameworks in which the introduction of contraceptives has an effect on fertility can be found in [Fernández-Villaverde et al. \(2014\)](#), who study the change of societal norms, or in [Cavalcanti et al. \(2021\)](#), who focus on the context of an developing economy. The same holds for our third application on the effects of the German reunification. For example, model-based analyses of policy effects in the context of structural transformation can be found in [Rogerson \(2008\)](#), who explores the role of tax increases for hours supply in the service sector, or in [Duval-Hernández et al. \(2022\)](#), who look at female labor supply. While the mentioned macroeconomic frameworks rely on calibrated versions of their models and focus on the effects at the aggregate level, SBI provides an empirical identification of the policy effect at the regional level using cross-regional heterogeneity at the stage at which the national policy is implemented.

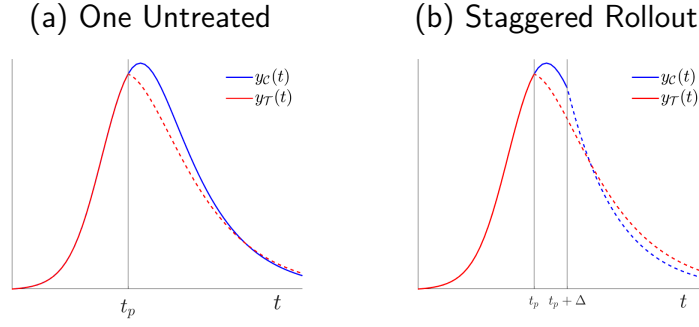
2 A Stage-Based Method to Identify Policy Effects

To contextualize our contribution, we first briefly discuss how standard empirical strategies identify policy effects. Consider a scenario in which, absent any policy intervention, the time path of an outcome $y_r(t)$ is identical across two regions $r \in \{\mathcal{C}, \mathcal{T}\}$.¹⁰ Now assume that a policy is implemented only in region \mathcal{T} at some date t_p which affects the outcome path in that region thereafter. Illustratively, we plot an outcome path of a treated region $y_{\mathcal{T}}(t)$ before policy implementation (solid red) and after policy implementation (dashed red) in panel (a) of Figure 2. We also show an outcome path for a region where the policy is not implemented, $y_{\mathcal{C}}(t)$ (solid blue). This scenario is ideal for the estimation of policy effects because the pre-policy outcome paths are identical across regions warranting the use of region \mathcal{C} as control for region \mathcal{T} . That is, the outcome path $y_{\mathcal{C}}(t)$ provides a useful no-policy counterfactual to assess the effects of policy on $y_{\mathcal{T}}(t)$ after t_p . The effects of policy are captured by the difference between $y_{\mathcal{C}}(t)$ and $y_{\mathcal{T}}(t)$ in the interval (t_p, ∞) . We can further add the implementation of the same policy to region \mathcal{C} at some later date $t_p + \Delta$ with $\Delta > 0$; see panel (b) of Figure 2. Under this staggered rollout of the policy, the effects of policy on region \mathcal{T} are identified using region \mathcal{C} as counterfactual within the interval $(t_p, t_p + \Delta]$. In that interval, region \mathcal{T} is subject to the policy whereas region \mathcal{C} is not.

The standard identification strategies just described fundamentally rely on two principles. First, there must be variation in the time of policy implementation across regions, which serves as source of identification. Second, the behavior of the outcome path before policy implementation must credibly support the parallel trends assumption. However, many policy contexts violate

¹⁰Region can be interchanged with *group* or *unit* throughout. Our empirical applications focus on *regions*.

Figure 2: Ideal Policy Scenarios with Two Regions: Standard Identification Strategies



Notes: Let $y_C(t)$ and $y_T(t)$ be the outcome paths of, respectively, region \mathcal{C} and \mathcal{T} before policy (solid lines) and after policy (dashed lines). The identified policy effects are $\int_{t_p}^h (y_C(t) - y_T(t)) dt / \int_{t_p}^h y_T(t) dt$ with $h = \infty$ for the one-untreated case and $h = t_p + \Delta$ for the staggered.

these conditions: First, a large set of policies are implemented nationwide—i.e., carried out in *all regions* at the *same time*, which eliminates the source of identification used in standard strategies. Second, the regional paths of the outcome variable before policy is implemented often differ across regions. For example, outcome paths can differ by starting date, evolve at different speed and have different magnitude. We illustrate these two challenges in panel (a) of Figure 3 where a nationwide policy is implemented in a context where the outcome path in region \mathcal{C} starts earlier, evolves at a faster speed and reaches a larger magnitude than in region \mathcal{T} .

Our strategy addresses these challenges in two steps. First, a normalization of regional outcome paths, and second, an identification based on the normalized paths.

2.1 Normalization Procedure


Again, consider two regional outcome paths $y_r(t)$ with $r \in \{\mathcal{C}, \mathcal{T}\}$. We select one region as the reference region. For this region, the stage is defined as time. For the non-reference region, the stage is the result of a normalization that maps its outcome time path onto the reference time path using only pre-policy data. We now describe our normalization of the time and level of an outcome of interest—i.e., a transformation of coordinates—and provide a formal definition of stages afterwards.

The normalization starts with postulating the existence of the composite function,

$$\tilde{y}_r(s) = (f_r \circ y_r \circ t_r)(s) = f_r(y_r(t_r(s))), \quad (1)$$

where $t_r(s) : S \rightarrow T$ is a stage-to-time transformation mapping stages $s \in S = \mathbb{R}$ into time $t \in T = \mathbb{R}$; $y_r(t) : T \rightarrow Y$ maps time into outcomes $y \in Y = \mathbb{R}$; and $f_r(y) : Y \rightarrow \tilde{Y}$ maps outcomes into normalized outcomes $\tilde{y}_r \in \tilde{Y} = \mathbb{R}$. Thus, the composite function $\tilde{y}_r(s) : S \rightarrow \tilde{Y}$ defined in (1) maps stages s —i.e., normalized time—into normalized outcomes \tilde{y} for region r .

Figure 3: Stage-Based Identification of Policy Effects: A Nationwide Policy



Notes: In panel (d), we report the policy effects γ (black marker) together with the interim cumulative effects of policy, $\gamma(s)$, as defined in Section 2.2.

Without loss of generality, we treat the outcome path of region \mathcal{T} as the reference path and that of region \mathcal{C} as the non-reference path.¹¹ For the reference region, we set s to be a fixed point of $t_{\mathcal{T}}(\cdot)$ for all s (i.e., $t = t_{\mathcal{T}}(s) = s$) and y to be a fixed point of $f_{\mathcal{T}}(\cdot)$ for all y (i.e., $\tilde{y} = f_{\mathcal{T}}(y) = y$) which implies that $\tilde{y}_{\mathcal{T}}(s) = y_{\mathcal{T}}(s) = y_{\mathcal{T}}(t = s)$ always. Instead, for the non-reference region, we approximate $t_{\mathcal{C}}(\cdot)$ and $f_{\mathcal{C}}(\cdot)$ with $t_{\mathcal{C}}(\cdot) \approx t(\cdot; \boldsymbol{\psi}) = \sum_{k=0}^K \psi_k B_k^t(\cdot)$ and $f_{\mathcal{C}}(\cdot) \approx f(\cdot; \boldsymbol{\omega}) = \sum_{m=0}^M \omega_m B_m^f(\cdot)$, respectively. $\{B^f(\cdot), B^t(\cdot)\} \in \mathcal{B}^2$ are known basis functions in the space of continuous and differentiable functions. We denote the set of $M + K + 2$ unknown normalization coefficients by $\boldsymbol{\phi} = \{\boldsymbol{\psi}, \boldsymbol{\omega}\}$. This gives the composite function $\tilde{y}_{\mathcal{C}}(s; \boldsymbol{\phi}) = (f_{\mathcal{C}}(\cdot; \boldsymbol{\omega}) \circ y_{\mathcal{C}} \circ t_{\mathcal{C}}(\cdot; \boldsymbol{\psi}))(s) = f_{\mathcal{C}}(y_{\mathcal{C}}(t_{\mathcal{C}}(s; \boldsymbol{\psi})); \boldsymbol{\omega})$ by which we approximate $\tilde{y}_{\mathcal{C}}(s)$:

$$\tilde{y}_{\mathcal{C}}(s) \approx \tilde{y}_{\mathcal{C}}(s; \boldsymbol{\phi}) = \sum_{m=0}^M \omega_m B_m^f \left(\left(y_{\mathcal{C}} \left(\sum_{k=0}^K \psi_k B_k^t(s) \right) \right) \right). \quad (2)$$

We choose monomials as benchmark for the basis functions $B^f(\cdot)$ and $B^t(\cdot)$ in (2). Further, throughout our theoretical derivations and applications we restrict our analysis to monomial basis functions of degree one, i.e., $K = M = 1$. This implies that the level and time transformation functions, respectively f_r and t_r , are linear, hence monotonic and invertible. Then, conversely to

¹¹The choice of the reference region is innocuous, see our discussion in Section 2.2.

the stage-to-time transformation, the stages in each region are defined as,

$$s = s_r(t; \boldsymbol{\psi}) = \begin{cases} t & \text{if } r = \mathcal{T} \\ t_{\mathcal{C}}^{-1}(t; \boldsymbol{\psi}) & \text{if } r = \mathcal{C}, \end{cases} \quad (3)$$

where for the reference region (here, $r = \mathcal{T}$), the stage at time t is the time itself (i.e., $s_{\mathcal{T}}(t; \boldsymbol{\psi}^*) = t = s$), whereas for the non-reference region (here, $r = \mathcal{C}$) the stage is the time at which region \mathcal{C} is at the same stage as region $r = \mathcal{T}$ is at time t (i.e., $s_{\mathcal{C}}(t; \boldsymbol{\psi}^*) = t_{\mathcal{C}}^{-1}(s_{\mathcal{T}}(t; \boldsymbol{\psi}^*); \boldsymbol{\psi}^*) = t_{\mathcal{C}}^{-1}(t; \boldsymbol{\psi}^*)$).

Hence, the approximated normalized path of the non-reference region is

$$\tilde{y}_{\mathcal{C}}(s; \boldsymbol{\phi}) = \omega_0 + \omega_1 y_{\mathcal{C}}(\psi_0 + \psi_1 s). \quad (4)$$

A nice feature of the monomial basis is that it delivers a straightforward interpretation of the normalization coefficients. The parameter ψ_0 shifts the entire outcome path of region \mathcal{C} forward (with $\psi_0 > 0$) or backwards (with $\psi_0 < 0$) in time, adjusting for different start dates. The parameter ψ_1 adjusts the speed in a constant way across periods. If $\psi_1 < 1$, then the outcome time-path of region \mathcal{C} (in time) expands, whereas with $\psi_1 > 1$ it contracts. That is, if $\psi_1 < 1$, then region \mathcal{C} is permanently faster (in time) than region \mathcal{T} —in one time-period region \mathcal{C} advances by more than one stage—and vice versa for $\psi_1 > 1$.¹²

Given observed time paths for all regions, i.e., $y_r(t)$ for $r \in \{\mathcal{C}, \mathcal{T}\}$, we determine the unknown coefficients $\boldsymbol{\phi} = \{\boldsymbol{\psi}, \boldsymbol{\omega}\}$ by minimizing the difference between the normalized path of the non-reference region, $\tilde{y}_{\mathcal{C}}(s; \boldsymbol{\phi})$, and the outcome path of the reference region, $y_{\mathcal{T}}(s)$, that is:

$$\min_{\{\boldsymbol{\phi}\}} \|\tilde{y}_{\mathcal{C}}(s; \boldsymbol{\phi}) - y_{\mathcal{T}}(s)\|_{\mathbb{C}(s)}, \quad (5)$$

where $\|\cdot\|$ is the squared Euclidean distance defined on the interval of stages,

$$\mathbb{C}(s) = [s_{\bar{r}}(t_0; \boldsymbol{\psi}), s_{\underline{r}}(t_p; \boldsymbol{\psi})] \quad (6)$$

where $s_{\bar{r}}(t_0; \boldsymbol{\psi}) = \max \{s_r(t_0; \boldsymbol{\psi})\}$ and $s_{\underline{r}}(t_p; \boldsymbol{\psi}) = \min \{s_r(t_p; \boldsymbol{\psi})\}$ for $r \in \{\mathcal{C}, \mathcal{T}\}$. That is, the interval $\mathbb{C}(s)$ ensures that the minimization (5) only uses the outcome paths up to the stage s in which the policy is implemented first across regions, i.e., $s_{\underline{r}}(t_p; \boldsymbol{\psi})$. Note that the interval $\mathbb{C}(s)$ is determined endogenously during the minimization procedure. Now, we can define the stages.

¹²More generally, allowing for the stage-to-time transformation to be quadratic (i.e., $\psi_2 \neq 0$) would capture the notion that the relative speed across the regions can change over time: for example, the outcome path of region \mathcal{C} might initially be slower than region \mathcal{T} , then catch up, and eventually move faster. Although not in the main text, we show an example with such scenario in Online Appendix [O.A](#).

Definition 1. *The stage of an outcome $y_r(t)$ of region r at time t is $s_r(t; \psi^*)$ where $\phi^* \supset \psi^*$ is the solution to the minimization of (5) subject to (4) and (6).*

In this way, stages formally emerge as the result of our normalization procedure. To gain some intuition, we exemplify our method using a nationwide policy that affects the outcome paths of two regions, $y_C(t)$ and $y_T(t)$, in Figure 3. Before policy implementation at time t_p , the outcome path of region C (solid red) differs from region T (solid blue) in that it starts earlier, grows faster and is larger; see panel (a), which also shows the outcome paths after policy implementation for the two regions (dashed lines).

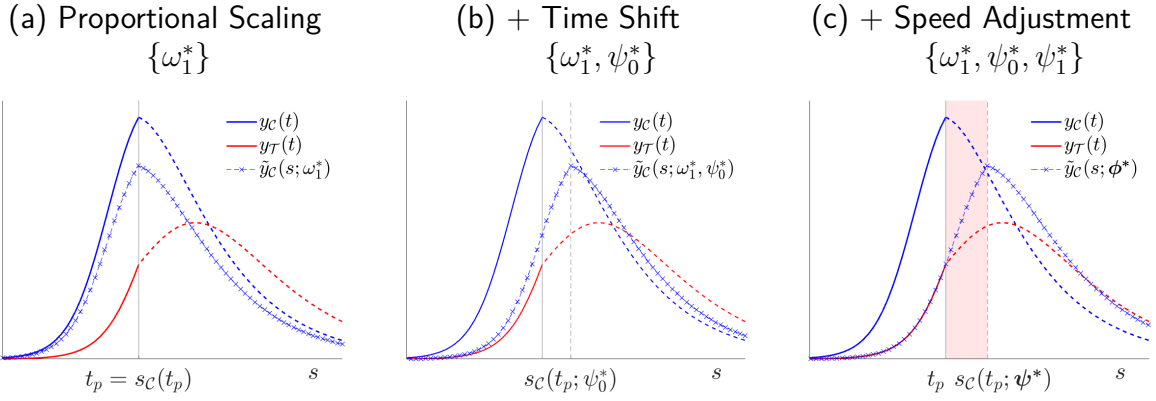
The normalization procedure—i.e., the minimization of (5) subject to (4) and (6)—achieves two goals. First, it generates a normalized outcome path for the non-reference region in the stage domain, $\tilde{y}_C(s; \phi^*)$ (cross-dashed blue), that maps—up to minimization error—onto the outcome path of the reference region before the earliest stage in which policy is implemented across regions, $s_{\underline{r}}(t_p; \psi^*)$; see panel (b) of Figure 3.¹³ Second, since $s_{\underline{r}}(t_p; \psi^*)$ is endogenous to ψ^* , the normalization uncovers heterogeneity in the stage of policy implementation across regions.¹⁴ For example, in our illustration, policy is implemented earlier—in stages—in region T than in region C , i.e., $s_{\underline{r}}(t_p; \psi) = s_T(t_p; \psi) < s_C(t_p; \psi) = s_{-\underline{r}}(t_p; \psi^*)$ with $\underline{r} = T$. Also, since we picked T to be the reference region, we obtain $s_{\underline{r}}(t_p; \psi) = t_p$.

We further decompose the effects of each of the normalization coefficients $\{\phi_i\} \in \phi^*$ on the path of the non-reference region in Figure 4. Since these coefficients are jointly determined in our minimization, we provide a non-orthogonal decomposition where we sequentially add the effects of each parameter. Note that in our illustration $y_C(t_0) = y_T(t_0) = 0$ and $\lim_{t \rightarrow \infty} y_C(t) = \lim_{t \rightarrow \infty} y_T(t) = 0$ and thus we focus on the role of the proportional level shifter ω_1 together with the stage-to-time transformation parameters ψ_0 and ψ_1 —setting the constant level shifter to zero, $\omega_0 = 0$. In panel (a), we show that the coefficient $\omega_1^* < 1$ proportionally shifts down the entire outcome path of the non-reference region C throughout its support. In panel (b), the additional time shifter, $\psi_0^* > 0$, moves the outcome path to the right delaying the outcome’s take off. In panel (c), adding the speed adjustment, $\psi_1^* < 1$, decreases the pace of the normalized outcome. See further illustrations by use of an example of a policy after the peak in Online Appendix O.A.

¹³Note that outcome variables are typically observed on discrete dates. Since the mapping can generate dates $t_C(s; \psi^*)$ that are non-integer values—i.e., non-discrete dates—we interpolate between $y_C(fl(t_C(s; \psi^*)); \omega^*)$ and $y_C(cl(t_C(s; \psi^*)); \omega^*)$, where $fl(\cdot)$ and $cl(\cdot)$ denote the integer floor or integer ceiling, respectively.

¹⁴More generally, with more than two regions, $-\underline{r}$ refers to the complement set of \underline{r} , i.e., $-\underline{r} = \underline{r}^C$

Figure 4: Decomposition by Normalization Coefficient



Notes: We sequentially add the normalization parameters $\{\phi\}$ to the non-reference path $y_C(t)$ one-by-one.

2.2 Identifying the Policy Effects

In order to identify the policy effects, we exploit the fact that our normalization uncovers heterogeneity of the stage at the time of policy implementation, i.e., $s_{\underline{r}}(t_p; \psi^*) < s_{-\underline{r}}(t_p; \psi^*)$. In particular, inside a window (interval) of stages,

$$\mathbb{W}(s; \psi^*) = [s_{\underline{r}}(t_p; \psi^*), s_{-\underline{r}}(t_p; \psi^*)], \quad (7)$$

region \underline{r} , i.e., the region where the policy is implemented first in stages, is subject to policy whereas region $-\underline{r}$ is not. In this context, we make the following assumption:

Identification Assumption 1. *The normalization parameters ϕ^* that solve the minimization of (5) subject to (4) and (6) are unaffected by policy.*

The assumption implies that, absent policy in region \underline{r} , the normalized path of the non-reference region obtained using ϕ^* and evaluated on stages $s > s_{\underline{r}}(t_p; \psi^*)$ would yield a path identical to that of the reference region for all $s \in \mathbb{W}(s; \psi^*)$.

Here, note that there is no ex-ante assignment to treatment or control for either reference or non-reference regions. Instead, the assignment of regions to treatment or control is determined endogenously (with ψ^*) by the fact that policy arrives to the regions at different stages. We refer to the region that is at a more advanced (later) stage at the policy date as the stage-leading region. This region is then endogenously assigned to be the control region. In the illustration, the stage-leading (control) region is $-\underline{r} = \mathcal{C}$, which is untreated inside the identification window $\mathbb{W}(s; \psi^*) = [t_p, s_C(t_p; \psi^*)]$ and, hence, provides the no-policy counterfactual for the stage-lagging (treated) region $\underline{r} = \mathcal{T}$ inside that window; see panel (b) in Figure 4. In the

example, the opposite roles (of reference and non-reference regions) would emerge if we picked $\underline{r} = \mathcal{C}$ as reference region.¹⁵

Policy effect. Following our illustration, where the control region is $-\underline{r} = \mathcal{C}$ and the treated region is $\underline{r} = \mathcal{T}$, we measure the policy effect for the treated region as,

$$\gamma = \frac{\int_{\mathbb{W}(s; \psi^*)} (y_{\mathcal{T}}(s) - \tilde{y}_{\mathcal{C}}(s; \phi^*)) ds}{\int_{\mathbb{W}(s; \psi^*)} \tilde{y}_{\mathcal{C}}(s; \phi^*) ds}, \quad (8)$$

which measures the cumulative effect of policy relative to the scenario without policy in the treated region inside $\mathbb{W}(s; \psi^*)$; see panel (c), Figure 3. The numerator is the area between the actual outcome path subject to policy of the treated region, i.e., $y_{\mathcal{T}}(s)$ (dashed red), and the no-policy counterfactual path for the treated region, i.e., $\tilde{y}_{\mathcal{C}}(s; \phi^*)$ (cross-dashed blue). The denominator captures the entire area below the no-policy counterfactual path for the treated region. In panel (d) of Figure 3, we zoom in on the identification window to show the policy effect γ together with the interim cumulative effects, $\gamma(s)$. Precisely, $\forall s \in \mathbb{W}(s; \psi^*)$, we define $\gamma(s) = \frac{\int_{s_{\underline{r}}(t_p; \psi^*)}^s (y_{\mathcal{T}}(s) - \tilde{y}_{\mathcal{C}}(s; \phi^*)) ds}{\int_{s_{\underline{r}}(t_p; \psi^*)}^s \tilde{y}_{\mathcal{C}}(s; \phi^*) ds}$ where $\gamma(s) = \gamma$ for $s = s_{-\underline{r}}(t_p; \psi^*)$, the stage of the leading region at the policy date (the end of the identification window).

Irrelevance of reference region. So far, we have used region \mathcal{T} as reference. Therefore, the policy effect (8) measures the impact of policy on region \mathcal{T} using as no-policy counterfactual the normalized path of region \mathcal{C} . Now consider reversing the reference region to \mathcal{C} . This implies that we redefine stages as time for region \mathcal{C} , and use the modified (to region \mathcal{T}) expressions in (2) and (3) to obtain the normalization for region \mathcal{T} . Of course, region \mathcal{C} (while now the reference region) is still the leading region, and we now obtain a policy effect by relating the normalized path of region \mathcal{T} (which is treated) to the observed path of region \mathcal{C} (which is untreated). The interpretation of the policy effect is different now, as it measures the impact that the policy would have had on region \mathcal{C} had it been treated at an earlier stage. We formally establish equivalence between the two mappings in the following

Theorem 1. *For $K = M = 1$, the policy effect (8) is invariant to the choice of the reference region.*

Proof. See Appendix A. □

¹⁵We assess our identification assumption in the context of an analytical examples with exact identification in Section 2.3 and also through a placebo test with model-generated data without exact identification in Section 3.3.

2.3 Analytical Examples with Exact Identification of Policy Effects

We now discuss a setting in which we can explicitly express the normalization of the non-reference region in terms of the structural parameters of the data generating process. Note that this serves to illustrate the method, and to provide some guidance for the interpretation of the normalization coefficients. Indeed, if the data generating process were actually known, there would be no need to apply SBI; or any other identification method for that matter.

Our method operates under the proposition that if there exists a composite function (1) such that

$$\tilde{y}_C(s) = y_{\mathcal{T}}(s), \quad (9)$$

then our normalization procedure—the minimization of (5) subject to (2) and (6)—recovers the coefficients $\phi = \{\psi, \omega\}$ up to a minimization error by approximating the functions $t_C(\cdot) \approx t_C(\cdot; \psi)$, $f_C(\cdot) \approx f_C(\cdot; \omega)$ and, hence, $\tilde{y}_C(\cdot) \approx \tilde{y}_C(\cdot; \phi)$ for all $s \in \mathbb{C}(s)$. Thus, under our identification assumption, we can identify the policy effects for all $s \in \mathbb{W}(s; \psi^*)$.

In this context, here, we are interested in cases where (9) holds and (2) holds with equality and, hence, analytical solutions for the normalization coefficients ϕ potentially exist for all $s \in \mathbb{C}(s)$. In that pursuit, consider a scenario in which the outcome path of a region $r \in \{\mathcal{C}, \mathcal{T}\}$ is,

$$y_r(t) = \left(1 + \gamma_{r,t} \mathbf{1}_{t \geq t_p}\right) g(t; \Theta_r), \quad \text{for } t \in \{t_0, \dots, t_p, \dots, T\} \quad (10)$$

where Θ_r is a set of region-specific structural parameters that determine regional outcome paths and $\gamma_{r,t}$ captures the effect of policy that emerges after its nationwide implementation at t_p .

To identify the effects of the implemented policy by SBI, we pick a region (e.g., \mathcal{T}) for the reference outcome path, $y_{\mathcal{T}}(t)$, and postulate a composite function (1) for the outcome path of the non-reference region, $\tilde{y}_C(s; \phi)$. Here, we are interested in cases where (9) holds and (2) holds with equality because then $\tilde{y}_C(s; \phi)$ and $y_{\mathcal{T}}(s)$ share exactly the same functional form before policy is implemented first in the stage domain, i.e., for all $s \in \mathbb{C}(s)$ and, hence, we can uncover—by the method of undetermined coefficients—the normalization coefficients ϕ in,

$$\Theta_{\mathcal{T}} = \tilde{\Theta}_C(\phi; \Theta_C) \quad \forall s \in \mathbb{C}(s), \quad (11)$$

which is a (potentially nonlinear) system with n equations—where n is the number of structural parameters—and with p unknowns—where $p = M + K + 2$ is the number of normalization coef-

ficients in ϕ .¹⁶ An interpretation of system (11)—and, hence, of our normalization procedure—is that the normalization coefficients ϕ reshape the structural parameters of the non-reference, Θ_C , region into those of the reference region, i.e., $\Theta_T = \tilde{\Theta}_C$ before policy implementation.

Theorem 2. *If there exists a composite function (1) such that (9) holds and (2) holds with equality for the regional outcome paths, $y_r(t)$, in (10)—i.e., if there exists a solution ϕ^* for the system (11)—then stage-based identification (SBI) exactly and uniquely identifies the true policy effects $\gamma(s)$ for all $s \in \mathbb{W}(s; \psi^*)$.*

Proof. If the system (11) holds, then the normalized outcome path of the non-reference region, $\tilde{y}_C(s; \phi^*)$, is exactly identical to the reference path, $y_T(s)$ for all $s \in \mathbb{C}(s) = [s_{\bar{r}}(t_0; \psi^*), s_{\underline{r}}(t; \psi^*)]$. Since the outcome paths follow (10)—i.e., policy affects the path $y_r(t \geq t_p)$ but not $g(t; \Theta_r)$, then ϕ^* is also a solution for the complement stage domain, i.e., for all $s \notin \mathbb{C}(s)$, in particular for $\mathbb{W}(s; \psi^*) = [s_{\underline{r}}(t_p; \psi^*), s_{-\underline{r}}(t_p; \psi^*)]$. This implies that $\tilde{y}_C(s; \phi^*)$ is exactly identical to $g(s; \Theta_T)$, that is, the true no-policy counterfactual of the reference region for all $s \in \mathbb{W}(s)$. Hence, SBI exactly and uniquely identifies the true policy effects, $\gamma(s)$, for all $s \in \mathbb{W}(s; \psi^*)$. \square

Remark 1. Note that uniqueness of the normalization coefficients ϕ^* is not necessary to recover unique policy effects.¹⁷ To see this, note that although the presence of multiple solutions of ϕ implies that there are multiple shapes for $f(\cdot; \omega)$ and $t(\cdot; \psi)$ that satisfy (9), the implied solution $\tilde{y}_C(s; \phi)$ for (9) is unique and, hence, so is the identified policy effect, $\gamma(s)$. At the same time, the overall policy effect, γ , is determined by the interim policy effects, $\gamma(s)$, and the size of the identification window, $\mathbb{W}(s; \psi^*)$, which can differ by ϕ^* ; see Online Appendix O.B.2.2.

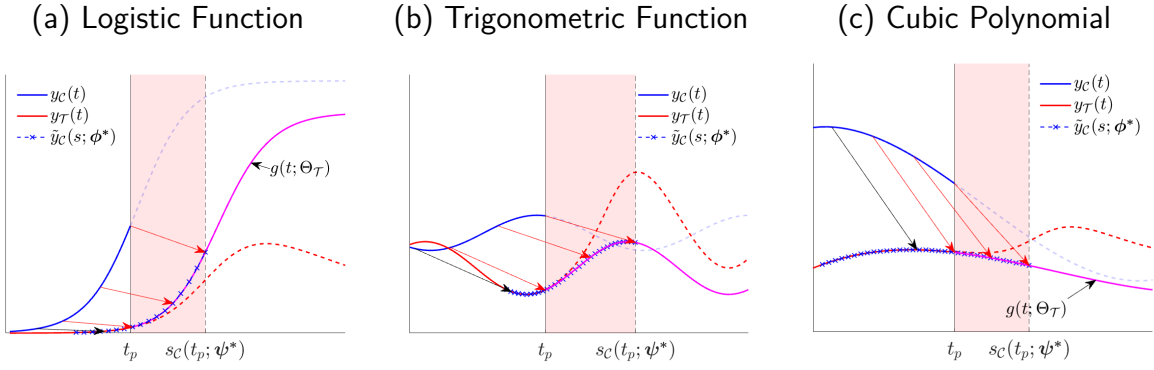
We now discuss some functional forms for the outcome paths, $y_r(t)$, for which SBI yields analytical solutions for ϕ using the approach just described. We start with logistic functions. Assume that the regional outcome paths $y_r(t)$ are determined by (10) and that, absent policy, these paths are determined by,

$$g(t; \Theta_r) = \frac{\theta_{1,r} - \theta_{0,r}}{1 + \exp(-\theta_3 t + \theta_{2,r})} + \theta_{0,r} \quad (12)$$

¹⁶That is, here, the minimization step in the normalization (in Section 2.1) is the solution to the system (11) emerging from the undetermined coefficients approach.

¹⁷Note that if (9) holds, (2) holds with equality and the inverse function $\phi = \tilde{\Theta}_C^{-1}(\Theta_T; \Theta_C)$ exists, then there exists a unique solution ϕ^* for the system (11). This sufficiency for existence and uniqueness of ϕ^* coincides with the Rouché–Frobenius Theorem in the cases where the system (11) is linear.

Figure 5: Three Examples with Exact Analytical Identification



Notes: These panels show the SBI effects in examples with closed-form solutions for ϕ^* when the data generating process is assumed to be known. In each panel, we assume that the outcome time path $g(t; \Theta_r)$ in (10) follows: (a) the logistic function in (12); (b) a trigonometric function $\theta_{0,r} + \theta_{1,r} \sin(\theta_{2,r} + \theta_{3,r}t)$; and (c) a cubic polynomial $\sum_{j=0}^3 \theta_{j,r} t^j$, respectively. The analytical derivations are in Online Appendix O.B.

with $\Theta_r = \{\theta_{0,r}, \theta_{1,r}, \theta_{2,r}, \theta_{3,r}\}$. We show an illustration of these paths for region \mathcal{C} and \mathcal{T} in panel (a) of Figure 5.¹⁸ To identify the policy effects, we apply SBI picking a region (e.g., \mathcal{T}) for the reference path and postulating a composite for the non-reference region, $\tilde{y}_C(s; \phi) = \omega_1 y_C(\psi_0 + \psi_1 s) + \omega_0$. Then, we solve for the normalization coefficients $\phi = \{\psi_0, \psi_1, \omega_0, \omega_1\}$ in (9) holding (2) with equality for all $s \in \mathbb{C}(s)$, that is,

$$\tilde{y}_C(s; \phi) = \omega_1 y_C(\psi_0 + \psi_1 s) + \omega_0 = \frac{\overbrace{\omega_1 (\theta_{1,C} - \theta_{0,C})}^{\theta_{1,\mathcal{T}} - \theta_{0,\mathcal{T}}}}{1 + \exp\left(-\underbrace{\theta_{3,C} \psi_1 s}_{\theta_{3,\mathcal{T}}} + \underbrace{(\theta_{2,C} - \theta_{3,C} \psi_0)}_{\theta_{2,\mathcal{T}}}\right)} + \underbrace{\omega_1 \theta_{0,C} + \omega_0}_{\theta_{0,\mathcal{T}}} = y_{\mathcal{T}}(s)$$

and, thus, by the method of undetermined coefficients we find $\phi = \{\psi_0, \psi_1, \omega_0, \omega_1\}$ from a linear system of four equations. It is straightforward to see that the inverse $\phi = \tilde{\Theta}_C^{-1}(\Theta_{\mathcal{T}}; \Theta_{\mathcal{C}})$ exists and, hence, there exists a unique analytical solution for ϕ^* ,

$$\begin{aligned} \omega_1^* &= \frac{\theta_{0,\mathcal{T}} - \theta_{1,\mathcal{T}}}{\theta_{0,\mathcal{C}} - \theta_{1,\mathcal{C}}}, & \omega_0^* &= \theta_{0,\mathcal{T}} - \theta_{0,\mathcal{C}} \left(\frac{\theta_{0,\mathcal{T}} - \theta_{1,\mathcal{T}}}{\theta_{0,\mathcal{C}} - \theta_{1,\mathcal{C}}} \right) \\ \psi_1^* &= \frac{\theta_{3,\mathcal{T}}}{\theta_{3,\mathcal{C}}}, & \psi_0^* &= \frac{\theta_{2,\mathcal{C}} - \theta_{2,\mathcal{T}}}{\theta_{3,\mathcal{C}}}. \end{aligned}$$

The normalization uncovers cross-regional stage heterogeneity at the time of policy implementation: in our illustration, the non-reference region is at a more advanced stage than the reference region at t_p , i.e., $s_{\mathcal{T}}(t_p; \psi^*) = t_p$. This opens a window in the stage domain in which region \mathcal{T} is subject to policy whereas region \mathcal{C} is not, i.e., $\mathbb{W}(s; \psi^*) = [t_p, s_{\mathcal{C}}(t_p, \psi^*) = \frac{\theta_{2,\mathcal{T}} - \theta_{2,\mathcal{C}}}{\theta_{3,\mathcal{T}}} + \frac{\theta_{3,\mathcal{C}}}{\theta_{3,\mathcal{T}}} t_p]$. Then, under our identification assumption, the normalized outcome path of the non-reference

¹⁸We assume that region \mathcal{C} takes off earlier, $-\frac{\theta_{2,\mathcal{C}}}{\theta_{3,\mathcal{C}}} < -\frac{\theta_{2,\mathcal{T}}}{\theta_{3,\mathcal{T}}}$, grows faster $\theta_{3,\mathcal{C}} > \theta_{3,\mathcal{T}}$, starts at level $\theta_{0,\mathcal{C}} = \theta_{0,\mathcal{T}} = 0$ (left asymptote) and shows a larger magnitude with $\theta_{1,\mathcal{C}} > \theta_{1,\mathcal{T}}$ (right asymptote) than region \mathcal{T} .

region, $\tilde{y}_C(s; \phi^*)$, serves as no-policy counterfactual for the reference region for all $s \in \mathbb{W}(s; \psi^*)$. Indeed, since the outcome paths follow (10), ϕ^* is also an analytical solution for the complement stage domain, i.e., for all $s \notin \mathbb{C}(s)$. That is, for all $s \in \mathbb{W}(s; \psi^*)$, the normalized outcome path of the non-reference region, i.e., $\tilde{y}_C(s; \phi^*)$, is identical to the reference path in the no-policy scenario, i.e., $g(t = s; \Theta_{\mathcal{T}})$ (solid magenta line); see panel (a), Figure 5. Thus, the identified policy effect is unique and identical to the true policy effect.

We repeat this analysis for outcome paths that follow trigonometric functions (panel (b), Figure 5), cubic polynomials (panel (c), Figure 5) and generalized logistic functions (Online Appendix O.B.3). In all these cases, we find a closed-form solution for ϕ^* and the identified policy effect is unique and identical to the true effect; see the derivations in Online Appendix O.B.

3 Method Performance

More generally, we are interested in assessing policy in contexts where the data generating process is unknown. Here, we first implement SBI on model-generated data (without using any knowledge about the theoretical model) and compare the identified effects with the true effects; Section 3.1. Second, we conduct a Monte Carlo analysis that provides bounds to our method performance; Section 3.2. There, we assess how our method fares in the presence of confounding factors. Third, we do a placebo diagnosis in Section 3.3. Fourth, we conduct inference in Section 3.4. As in the discussion in Section 2.1, we adopt a linear stage-to-time transformation. Further, we set the level shift parameter ω_0 to zero so that the normalization coefficients are $\phi = [\omega_1, \psi_0, \psi_1]$.


3.1 Does SBI Identify the True Policy Effects?

To address this question, we use three alternative policy contexts: a public health policy against a pandemic using a model where economic activity in the form of hours worked shapes and is shaped by a pandemic; the effects of the approval of the pill in a model of women career and fertility choices; and an economic growth policy using a model of structural transformation.

3.1.1 Public Health Policy Against a Pandemic

At the beginning of each period $t \in \{0, 1, \dots\}$, total population N_t is composed of a stock of susceptible population S_t , infected individuals I_t and recovered individuals R_t , with $N_t = S_t + I_t + R_t$ and the normalization $N_0 = 1$. An epidemic starts with an initial number of infected $I_1 > 0$ in period $t = 1$. For pre-pandemic periods $t < 1$, the population is constant with $N_0 = S_0$ and $I_0 = R_0 = 0$. The probability that a susceptible individual meets an infected

Figure 6: Stage-Based Identification of Model-Generated Policy Effects: A Nationwide Public Health Policy Against a Pandemic



Notes: We assume that $u(c_z, h_z) = \ln(c_z) - \kappa \frac{h_z^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}} + \chi$ for value of life parameter χ . Some parameters differ across regions: $\Theta_C = \{\beta = 0.509, \zeta = 0.0010, \kappa = 1.05, \xi = 0.20, I_0 = 1\}$ and $\Theta_T = \{\beta = 0.501, \zeta = 0.0008, \kappa = 1.07, \xi = 0.19, I_0 = 6\}$. The rest of the model parameters are identical across regions, $\{\delta = 0.95, \chi = 560400, z = 64, \beta = 0.501, \alpha = 0.65\}$. The parameters associated to the policy are $\bar{h} = 0.4, t_p = 38, t_f = 250$.

individual is given by $\beta \frac{I_t}{N_t}$, for $\beta \in (0, 1)$. Conditional on meeting there exists an objective probability $\lambda_{\mathcal{O}}(h_t)$ of getting infected that depends on the average hours worked h_t . Further, with probability μ infected individuals in a given period t recover or die from the disease where the conditional probability of death in turn is denoted by ζ . New infections transit to death in the same period t , i.e., h_t has an immediate effect on the survival rate between t and $t + 1$.¹⁹

In this context, we consider the problem of a social planner, who is constrained by imperfect knowledge about the infection process. In particular, the planner's beliefs of the infection probability are $\lambda_{\mathcal{P}}(h_t)$, which may differ from the objective probability. Specifically, let $\lambda_b(h_t) = \xi_b h_t^\alpha$, $\xi_b > 0$ and $\alpha \in (0, 1)$, for beliefs $b \in \{\mathcal{O}, \mathcal{P}\}$, where \mathcal{O} stands for objective and \mathcal{P} for perceived. Thus, if $\xi_{\mathcal{P}} < \xi_{\mathcal{O}}$ then the constrained planner underestimates the actual effects of average hours worked h_t on infections and vice versa if $\xi_{\mathcal{P}} > \xi_{\mathcal{O}}$.

At every period t , before making plans for all future periods $z \geq t$, the planner receives an unanticipated knowledge shock that reveals the actual state of the economy $G_{\mathcal{O},t}$ for $G = (S, I, R, D)$, which potentially differs from the perceived state $G_{\mathcal{P},t}$. We assume that the planner updates the perceived survival probability accordingly and before choosing labor supply. Precisely, letting $X_{G,b,t} = G_{b,t+1} - G_{b,t}$, the planner's perceived survival probability is revised at the beginning of every period t to $\phi_{\mathcal{P}}(h_t) = 1 - \frac{X_{D,\mathcal{P},t}}{\tilde{N}_{b,t}}$ with $\tilde{N}_{b,t} = N_{\mathcal{O},t}$ for $t = z$ and $\tilde{N}_{b,t} = N_{\mathcal{P},t}$ if $t > z$. Note that although the knowledge shock allows the planner to update the state of the economy at the beginning of every period t , she is unable to correct future forecast errors, i.e., $G_{\mathcal{O},z} - G_{\mathcal{P},z} | t$ for periods $z > t$, because the shocks are unanticipated.²⁰

¹⁹This innocuous assumption eases the exposition of the trade-off between economic activity and public health.

²⁰Note that without subjective beliefs, the population evolves essentially as in, for example, [Atkeson \(2020\)](#). Further, the forecast errors $\varepsilon_{G,z} = (G_{\mathcal{O},z} - G_{\mathcal{P},z} | t)$ can be reduced asymptotically with learning ([Adam et al., 2017](#)). For example, there could be learning about the odds of infection as in [Aleman et al. \(2022\)](#).

After updating the perceived survival probability, the planner maximizes the present-discounted stream of per period utilities for all periods $z \geq t$ with discount factor δ times the perceived unconditional probability to survive from any period t to the future, $\prod_{j=t+1}^z \phi_{\mathcal{P}}(h_{j-1})$. Importantly, since the perceived survival probability is revised at the beginning of every period t , the nature of the discounting process changes each period t and, hence, the planner needs to re-optimize—at each period t —the decision plans for all periods $z \geq t$. The per period utility, $u(c_z, h_z; \chi)$ is assumed strictly concave in consumption $c_z \geq 0$ and leisure $1 - h_z \in [0, 1]$ for a value of life parameter χ . Collecting elements, at each period t the constrained social planner solves,

$$\max_{\{c_z \geq 0, h_z \in [0, 1]\}_{z=t}^{\infty}} \sum_{z=t}^{\infty} \delta^{z-t} \prod_{j=t+1}^z \phi_{\mathcal{P}}(h_{j-1}) u(c_z, h_z; \chi), \quad (13)$$

subject to an aggregate resource constraint $N_{\mathcal{P},z} c_z = w h_z N_{\mathcal{P},z}$ where w is the implicit price (marginal product) of labor using technology $Y_z = a h_z N_{\mathcal{P},z}$. The solution is characterized by an Euler equation, which is a first-order difference equation in h_z , and thus we can easily solve for the optimal labor path during the epidemic using standard techniques.²¹

True (Model-Generated) Policy Effects. We solve the model for two regions that differ in the underlying parameter values for $\Theta = \{\delta, \chi, a, \beta, \mu, \zeta, \kappa, \nu, \{\xi_i\}_{i \in \{\mathcal{O}, \mathcal{P}\}}, \alpha, I_1\}$. In particular, we assume that the planner in region \mathcal{C} underestimates the effect that economic activity has on infections by less than the planner in region \mathcal{T} . Consequently, hours are reduced earlier and also by a larger amount in region \mathcal{C} than region \mathcal{T} in response to the epidemic. The equilibrium response of hours without policy intervention for region \mathcal{C} (solid blue) and region \mathcal{T} (solid red) are shown in panel (a) of Figure 6. The earlier and stronger response in terms of hours of region \mathcal{C} affects our outcome of interest, i.e., the epidemic path of deaths, by reducing the peak of deaths and flattening the curve in region \mathcal{C} relative to region \mathcal{T} ; see panel (b) in Figure 6. We also assume that region \mathcal{C} has higher odds of encountering infected individuals at work (i.e., higher β) which advances and increases the peak of deaths for region \mathcal{C} relative to region \mathcal{T} . Further, we assume that region \mathcal{C} has a lower disutility of work κ which implies a larger pre- and post-pandemic level of hours worked for region \mathcal{C} than region \mathcal{T} .

In this scenario, we now introduce a nationwide public health policy that imposes an upper bound on hours worked, $h < \bar{h} = 0.5$, from $t_p = 38$ to $t_f = 250$. Since, without policy, households in both regions would work more hours than \bar{h} , the policy is binding in both \mathcal{C} and \mathcal{T} —see the respective dashed lines that emerge after t_p in panel (a) of Figure 6. The lower economic activity imposed by the policy has consequences for the flow of deaths. With policy,

²¹We provide solution algorithms for all the economic models in this paper in Online Appendix O.C.

the flow of deaths peaks earlier and by a lower magnitude in both \mathcal{C} and \mathcal{T} —see the respective dashed lines that emerge after t_p in panel (b) of Figure 6. The difference between the flow of deaths with policy (dashed lines) and the flow of deaths without policy (solid lines) after t_p captures the true effects of policy generated from the model. However, the counterfactual paths of the flow of deaths without policy after policy implementation (i.e., the solid lines after t_p) are not available outside of the model.


Stage-Based Identified Policy Effects. The policy effects identified using SBI are shown in panel (c) of Figure 6. In particular, we map the path of the flow of deaths in region \mathcal{C} (solid blue line) onto the path of region \mathcal{T} (solid red line) using *only* pre-policy data; as described in Section 2. The result of the normalization step is a candidate no-policy counterfactual $\tilde{y}_{\mathcal{C}}(s; \phi^*)$ (blue line with cross markers) for region \mathcal{T} in the identification window $\mathbb{W}(s; \psi^*) = [t_p, s_{\mathcal{C}}(t_p; \psi^*)]$ (shaded pink area). In order to assess whether the identified policy effects recover the true policy effects generated by the model, we zoom in on the identification window in panel (d) of Figure 6 and compare our candidate counterfactual $\tilde{y}_{\mathcal{C}}(s; \phi^*)$ with the true counterfactual (solid red line). The main result is that the identified policy effects are not significantly different from the true effects. The identified total number of lives saved is $\int_{\mathbb{W}(s; \psi^*)} (\tilde{y}_{\mathcal{C}}(s; \phi^*) - y_{\mathcal{T}}(s)) ds = 248.545$ in a window of $s_{\mathcal{C}}(t_p; \psi^*) - t_p = 8.601$ days, whereas the true policy effects are 250.728 lives saved. Therefore, the policy prevented $\gamma = -21.496\%$ of the total deaths that would have occurred had the policy not been implemented, whereas the true effect is $\gamma_{\text{true}} = -21.644\%$, which implies a percent error of $\varepsilon(\gamma) = \left| \left(\frac{\gamma}{\gamma_{\text{true}}} - 1 \right) \times 100 \right| = 0.683\%$.

3.1.2 Oral Contraceptives and Women’s Choices

Consider a model economy where each cohort t of women derives utility from their choices on consumption $c \geq 0$, children $n \geq 0$ and sexual intercourse $x \geq 0$ and experiences disutility from pill usage o —e.g., a social norm. In addition, a woman chooses human capital investment paying q (tuition fees or job training) per unit of human capital. Earnings feature two components, a wage level w , and an endogenous human capital wage premium $z_t e(h)$ with the two components technology level z_t and a complementarity factor $e(h) \in [0, 1]$ ²² with $e_h(h) > 0, e_{hh}(h) < 0$ so that earnings per unit of time are $w(1 + z_t e(h))$. We model skill-biased technical change (SBTC) with a cohort- t specific growth factor γ_t so that $z_t = z_0 \prod_{\tau=1}^{t-1} (1 + \lambda_{\tau})$, where $z_0 > 0$ and $\lambda_t > 0$. We further assume that raising children bears a time cost of $\tau(n) \in [0, 1]$ with $\tau_n(n) > 0, \tau_{nn}(n) < 0$ so that earnings are $(1 - \tau(n))w(1 + z_t e(h))$. Sexual intercourse increases the probability of pregnancy $\phi(x) \in [0, 1]$ where $\phi_x(x) > 0, \phi_{xx}(x) < 0$ and we assume that successful

²²The mapping of the outcome variable from h to $e(h)$ is innocuous. In particular, since we model $e(h)$ as a rate we can interpret it as the fraction of educated women (e.g., college degree completion) in the population.

Figure 7: Stage-Based Identification of Model-Generated Policy Effects: Introduction of the Pill



Notes: We plot outcomes for region C (red) and T (blue) without policy (solid lines) and with policy (dashed lines). The policy is the introduction of the pill for all periods $t_p \geq 25$. The parameter values that we choose for region C are $\Theta_C = \{\zeta = 8, q = 3.2, w = 64, z_0 = 1, \lambda_{t=1}^T = 0.1\%, \theta_x = 0.5, \theta_h = 0.4, \theta_o = 0.43, \iota_{t,C}, \kappa_{t,C}\}$ and for region T they are $\Theta_T = \{\zeta = 8, q = 3.3, w = 63, z_0 = 1, \lambda_{t=1}^T = 0.1\%, \theta_x = 0.5, \theta_h = 0.4, \theta_o = 0.43, \iota_{t,T}, \kappa_{t,T}\}$.

pregnancies result in children. If women have access to the pill—which we model through policy dummy $\mathbf{1}_{t \geq t_p}$ that is equal to zero if a cohort t does not have access to the pill, and equal to one otherwise—, then the probability of pregnancy is adjusted downward by the pill effectiveness in preventing pregnancy, $g(o) \in [0, 1]$. We assume that larger use of the pill—e.g., better adherence to follow protocol—increases the effectiveness of the pill. That is, $g_o(o) > 0$ with $g_{oo}(o) < 0$.

Collecting elements, a woman solves

$$\max_{\{h, o, x\}} c + \kappa n + \zeta x - \iota_t o \quad (14)$$

subject to the budget constraint (15) and the children production technology (16):

$$c + qh = (1 - \tau(n))w(1 + z_t e(h)), \quad (15)$$

$$n = \phi(x)[1 - \mathbf{1}_{t \geq t_p} g(o)] \quad (16)$$

We derive the first order conditions for h , x and o in Online Appendix O.C.

True (Model-Generated) Policy Effects. In Figure 7, we show the equilibrium path for women’s schooling choices in panel (a) and fertility choices in panel (b). We show the model-generated paths by region in a scenario without the pill (solid lines) and in a scenario in which the government grants women legal access to the pill technology (dashed lines). Regions differ in the model parameters $\Theta = \{\kappa, \zeta, q, w, z, \{\lambda_t\}_{t=1}^T, \{\iota_t\}_{t=1}^T, \{\kappa_t\}_{t=1}^T, \theta_x, \theta_h, \theta_o\}$. In particular, we allow

for the returns to human capital to be larger and grow faster in region \mathcal{C} than in region \mathcal{T} which explains the higher human capital in region \mathcal{C} than in region \mathcal{T} . This also explains the lower fertility in region \mathcal{C} than in region \mathcal{T} . Further, we exogenously shape the SBTC parameter γ such that the endogenous human capital path is S -shaped for both regions. We choose an exogenous path for the relative utility derived from children, κ , in order for endogenous fertility to display a boom and bust. Here, we assess the effects of legalizing the pill permanently with $\mathbf{1}_{t \geq t_p} = 1$ for all cohorts of women $t \geq t_p$. The policy endogenously reduces births (n) in both regions (dashed lines panel (b), Figure 7). By reducing fertility, the pill reduces the cost of acquiring human capital, which increases the share of women entering college ($e(h)$) (dashed lines in panel (a), Figure 7).

Stage-Based Identified Policy Effects. We apply SBI using region \mathcal{T} as reference, and thus map the outcome path of region \mathcal{C} (solid blue) onto that of region \mathcal{T} (solid red) using only pre-policy data as in Section 2. Again, SBI delivers a candidate counterfactual $\tilde{y}_{\mathcal{C}}(s; \phi^*)$ (blue line with cross markers) for an identification window $\mathbb{W}(s; \psi^*) = [t_p, s_{\mathcal{C}}(t_p; \psi^*)]$ (shaded pink area); see panels (c) and (d) of Figure 7 for human capital and children, respectively. A zoomed comparison between the identified and true effects is in panel (e) and (f) of Figure 7 for human capital and children, respectively. We find that the SBI policy effects capture well the true effects. Access to the pill increases the proportion of women going to college $e(h)$ by $\gamma = 0.122\%$, whereas the true policy effects are $\gamma_{\text{true}} = 0.123\%$. The identified effect on fertility is a reduction by $\gamma = 0.830\%$, whereas the true effect is $\gamma_{\text{true}} = 0.828\%$. The error $\varepsilon(\gamma)$ of the identified policy effects relative to the true policy effects is 0.182% for human capital and of 0.232% for fertility.

3.1.3 Growth Policy and Structural Transformation

Consider a model economy with two sectors denoted by $i \in \{a, m\}$, for agriculture and manufacturing, respectively. A representative firm per sector faces competitive markets. The agricultural firm produces output y_a at relative price p_a (manufacturing is the numeraire good) employing labor n_a at wage rate w_a and land ℓ . We assume inefficient institutions in agriculture captured by a parameter τ that taxes revenue. Agricultural firms thus solve the problem

$$\max_{n_{at}} \pi_t(\ell) = (1 - \tau)p_a y_{at} - w_a n_{at} \quad \text{s.t.} \quad y_{at} = z_{at} n_{at}^{\phi} \ell^{1-\phi},$$

where ϕ is the labor share in agriculture. Since land is fixed, the agricultural technology exhibits decreasing returns to scale.²³ Manufacturing firms produce output y_{mt} with labor n_{mt} —hired at

²³The structural change—from a decreasing returns to scale technology (Malthus) to a constant returns to scale (Solow) is studied in Hansen and Prescott (2002) in the context of a one-good economy. Below, we also introduce non-homothetic preferences as an additional mechanism for structural change (e.g., Gollin et al., 2002).

wage w_{mt} —and capital k_t —rented at rate r_t —and solve the problem

$$\max_{n_{mt}, k_{t+1}} y_{mt} - w_{mt}n_{mt} - r_t k_t \quad \text{s.t.} \quad y_{mt} = z_{mt} n_{mt}^\alpha k_t^{1-\alpha},$$

where α is the labor share in manufacturing. Further, we assume that total factor productivity (TFP) differs by sector according to $z_{it} = z_{i,0}(1 + \lambda_i)^t$ for $i = \{a, m\}$ with $\gamma_a < \gamma_m$.

An infinitely-lived representative agent discounts the future at factor $\beta \in (0, 1)$ and chooses sectoral allocations of consumption $\{c_{at}, c_{mt}\}_{t=0}^\infty$, labor $\{n_{at}, n_{mt}\}_{t=0}^\infty$, and next period capital $\{k_{t+1}\}_{t=0}^\infty$. The per period utility function from agricultural goods, $u(c_{at} - \bar{c}_a)$, features a non-homotheticity through a subsistence level, \bar{c}_a . Utility from manufacturing goods, $v(c_{mt})$, is additively separable. Both $u(\cdot)$ and $v(\cdot)$ are strictly concave. The household is endowed with one unit of time in each period, i.e., $n_{at} + n_{mt} = 1 \forall t$, that is allocated to either agriculture or manufacturing and receives wage rates $\{w_{at}, w_{mt}\}$. The household receives the rents $\pi(\ell)$ from inelastically supplying (renting) land to agricultural firms. Thus, the household maximizes,

$$\max_{\{c_{at}, c_{mt}, n_{at}, n_{mt}, k_{t+1}\}_{t=0}^\infty} \sum_{t=0}^{\infty} \beta^t (u(c_{at} - \bar{c}_a) + \kappa v(c_{mt})) \quad (17)$$

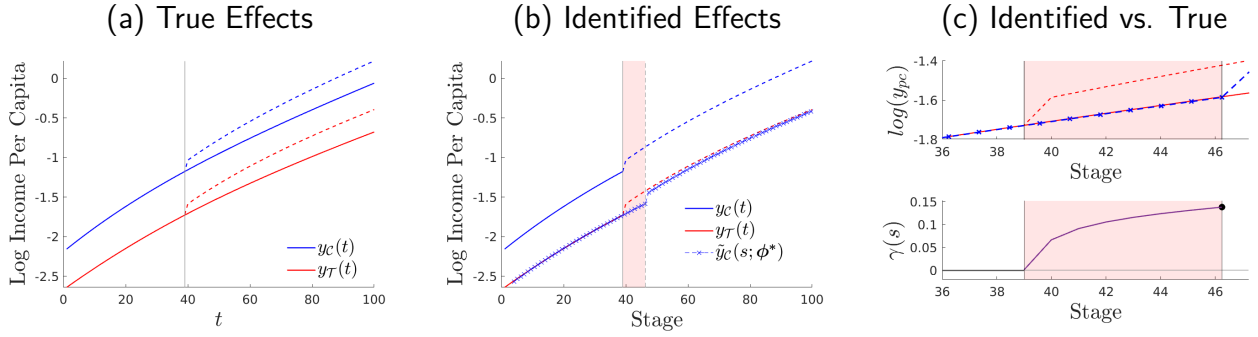
where $\kappa > 0$ is a relative utility parameter, subject to the budget constraint

$$p_{at}c_{at} + c_{mt} + k_{t+1} = \sum_{i \in \{a, m\}} w_{it}n_{it} + r_t k_t + (1 - \delta)k_t + \pi_t(\ell). \quad (18)$$

We derive the first order conditions and discuss the solution algorithm in Online Appendix [O.C](#).

True (Model-Generated) Policy Effects. We consider two regions that potentially differ in model parameters $\Theta = \{\beta, \bar{c}_a, \kappa, \delta, z_{a,0}, \lambda_a, z_{m,0}, \lambda_m, \phi, \alpha, \tau\}$. In particular, we allow for the total factor productivity in the manufacturing sector to be larger in region \mathcal{C} than in region \mathcal{T} . The larger productivity of manufacturing in region \mathcal{C} generates a larger amount of investment, lower agricultural share of labor and, ultimately, higher income per capita in region \mathcal{C} than in region \mathcal{T} at any point in time; see panel (a) in Figure 8. The model is able to generate an agricultural share that declines over time whereas, at the same time, capital and income per capita increase asymptotically reaching a balanced growth path with a trifling agricultural share. In this context, we introduce an unexpected nationwide growth policy that removes the institutional constraint τ in the agricultural sector in both regions; setting $\tau = 0$ after t_p in both regions. Removing the constraint in the agricultural sector accelerates investment (and capital) and the decline in agricultural sector. The reallocation to the non-agricultural sector increases income per capita in the economy, see (dashed lines) in panel (a) of Figure 8.

Figure 8: Stage-Based Identification of Model-Generated Policy Effects: Growth Policy



Notes: For region \mathcal{T} , we choose, $n_{a,0} = 0.75, z_{a,0} = 0.068, z_{m,0} = 0.10, \lambda_a = 0.0019, \lambda_m = 0.0055$. For region \mathcal{C} , we choose, $n_{a,0} = 0.55, z_{a,0} = 0.082, z_{m,0} = 0.13, \lambda_a = 0.0025, \lambda_m = 0.0062$. Common parameters for both regions are $\beta = 0.98, \alpha = 0.6, \phi = 0.7, \kappa = 2, \delta = 0.02$. Further, we assume that the felicity functions are logs, that is, $u(c_a - \bar{c}_a) = \ln(c_a - \bar{c}_a)$ and $v(c_m) = \ln c_m$.

Stage-Based Identified Policy Effects. We implement SBI by mapping the outcome path in region \mathcal{C} (solid blue line) onto the outcome path in region \mathcal{T} (solid red line) using only pre-policy data. We plot the resulting counterfactual candidate $\tilde{y}_C(s; \phi^*)$ (blue line with cross markers) for the identification window between t_p and $s_C(t_p; \psi^*)$ (shaded pink area); see panel (b), Figure 8. We zoom in on the identified counterfactual $\tilde{y}_C(s; \phi^*)$ and the true effects of policy in panel (c) of Figure 8. According to SBI, the growth policy increases income per capita by $\gamma = 13.781\%$ in the identification window whereas the true policy effect is $\gamma_{\text{true}} = 13.537\%$. That is, the identified policy effects catch the true policy effects with an error of $\varepsilon(\gamma) = 1.797\%$.

3.2 Bounds to Method Performance

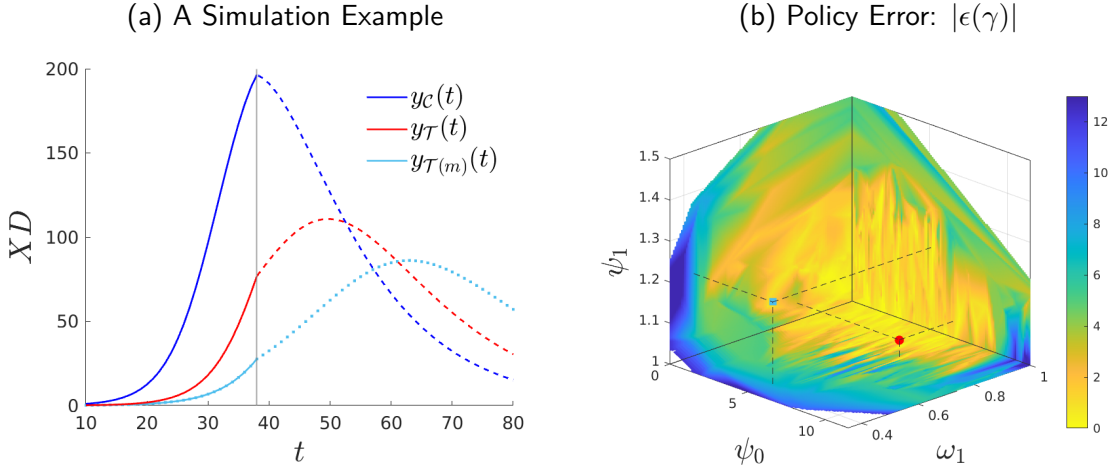
The performance analysis in Section 3.1 shows that our identification strategy can recover the true policy effects. However, it is intuitive to assume that our strategy faces some boundaries. Here, we numerically characterize the bounds within which our method is able to recover the true effects of policy with a Monte Carlo experiment in Section 3.2.1. We further discuss how we assess the performance of SBI in the presence of confounding factors in Section 3.2.2.

3.2.1 A Monte Carlo Analysis

We focus this analysis on the benchmark economic model with an endogenous pandemic described in Section 3.1.1. Specifically, we hold fixed the parameters of the non-reference region \mathcal{C} and randomize a subset— $(\beta, \zeta, \kappa, t_o)$ —of the structural parameters in that region in order to generate a large number of reference outcome paths $y_T(m)$ for regions $m \in \mathcal{M} = \{1, \dots, m, \dots, M\}$.²⁴ In panel (a) of Figure 9, we show the epidemic paths of our benchmark regions \mathcal{C} and \mathcal{T} as described in Section 3.1.1, together with one of the simulated reference regions that starts later,

²⁴We assume that the randomized parameters— β, ζ, κ and t_o —are uniformly and independently distributed. Then, we draw a total of $M = 381,000$ simulations (quadruplets).

Figure 9: Bounds to Method Performance: A Monte Carlo Analysis



Notes: To construct the set $\{y_{\mathcal{T}(m)}(t)\}_m$, we assume that $\{\beta, \zeta, \kappa, t_0\}$ are uniformly and independently distributed. The simulations are drawn from the intervals $[\beta^{lb}, \beta^{ub}] \times [\zeta^{lb}, \zeta^{ub}] \times [\kappa^{lb}, \kappa^{ub}] \times [t_0^{lb}, t_0^{ub}] = [0.5, 0.9] \times [0.001, 0.008] \times [1.05, 1.89] \times [-10, 10]$ where the superindices lb and ub denote, respectively, lower and upper bounds. We pick the bounds of the uniform distribution in a manner that our simulations generate sufficiently different outcome paths of the reference region in order to assess the performance of our method. We conduct $M = 381,000$ simulations and, in panel (b), we plot those that fall in a subset of $\Phi^o = \{\psi_0(m) | \psi_0(m) > 0.0\} \times \{\psi_1(m) | \psi_1(m) > 1.0\} \times \{\omega_1(m) | \omega_1(m) < 1.0\}$. Precisely, the hyperplane (ψ_0, ω_1) has 3,698 simulations, the hyperplane (ψ_1, ω_1) has 17,504 simulations and the hyperplane (ψ_0, ψ_1) has 3,698 simulations. Panel (b) shows values for an evenly spaced 200×200 grid on each hyperplane, approximating the values between grid points through linear interpolation.

grows slower and reaches a lower magnitude than the benchmark reference region, $y_{\mathcal{T}}(t)$, and, therefore, is further away from the non-reference region, $y_C(t)$.

In this context, in order to assess the ability of SBI to identify the true policy effect we study the policy error across all simulations. For each simulation m , we apply SBI by mapping the non-reference region, $y_C(t)$, onto the simulated reference path, $y_{\mathcal{T}(m)}(t)$. Per simulation m , we find a vector of normalization coefficients $\phi^*(m) = (\psi_0^*(m), \psi_1^*(m), \omega_1^*(m))$ that belongs to $\Phi = \Psi_0 \times \Psi_1 \times \Omega_1 \subset \mathbb{R}^3$. Then, for each simulation, we measure the policy error as the (absolute) value of the policy effect identified by SBI relative to the (model-generated) policy effect; i.e., $\varepsilon(\gamma)(m) = \left| \left(\frac{\gamma(m)}{\gamma_{\text{true}}(m)} - 1 \right) \times 100 \right|$. In panel (b) of Figure 9, we plot the policy errors of each $\phi^*(m)$ inside $\Phi^o = \{\psi_0(m) | \psi_0(m) > 0.0\} \times \{\psi_1(m) | \psi_1(m) > 1.0\} \times \{\omega_1(m) | \omega_1(m) < 1.0\}$, where $\Phi^o \subset \Phi$ is an octant with $\phi_o = (0.000, 1.000, 1.000)$ as origin. We restrict the plot to a subset of Φ^o that includes the policy error associated with the benchmark reference region $y_{\mathcal{T}}(t)$.

Now, first, note that if the vector of normalization coefficients is identical to the origin, i.e., $\phi^*(m) = \phi_o$, it implies that the outcome path of the simulated reference region, $y_{\mathcal{T}(m)}(t)$, and that of the non-reference region, $y_C(t)$, are identical.²⁵ Second, note that if the outcome path of a simulated reference region, $y_{\mathcal{T}(m)}(t)$, and the outcome path of the non-reference region, $y_C(t)$, are similar—in that SBI delivers a vector of normalization coefficients that is inside the intersection between some neighborhood of the origin, $\mathcal{N}(\phi_o)$, and Φ^o —then the policy error

²⁵If $\phi^*(m) = \phi_o$, then there is no cross-regional heterogeneity of stages and SBI cannot identify policy effects.

is small; see panel (b) of Figure 9. Indeed, policy errors with values of $\varepsilon(\gamma) \leq 5\%$ emerge (approximately) for the normalization coefficients $\phi^*(m) \in [0.000, 6.776] \times [1.000, 1.210] \times [0.764, 1.000] = \mathcal{N}(\phi_o) \cap \Phi^o$, which we depict as the yellow area. Here, note that our benchmark reference outcome path $y_{\mathcal{T}}(t)$ falls in that area with $\phi^* = (6.592, 1.041, 0.803)$ and a policy error $\varepsilon(\gamma) = 0.68\%$ (red marker). Third, moving away from the origin increases the policy error. For example, the simulated reference outcome path $y_{\mathcal{T}(m)}(t)$ in panel (a) of Figure 9 implies $\phi^*(m) = (5.083, 1.200, 0.436)$ that falls outside $\mathcal{N}(\phi_o) \cap \Phi^o$ with a policy error of 36.04%. Thus, as long as the regional outcome paths are similar enough before policy, the method can successfully identify the policy effect.²⁶

3.2.2 Confounding Factors

How does the presence of time-varying latent heterogeneity and confounding policy affect the ability of SBI to recover the true effects of policy? Before addressing this question, we would like to emphasize that the goal of our method is not to answer what would have been the effect of the policy under evaluation had confounding factors not been present. That is, we do not pursue the identification of a “pure” policy effect that nets out the presence of confounding factors.²⁷ Instead, we acknowledge that the effect of the same policy can naturally be conditioned by the context—e.g., due to the presence of different confounding factors. SBI is designed to measure these conditional policy effects.

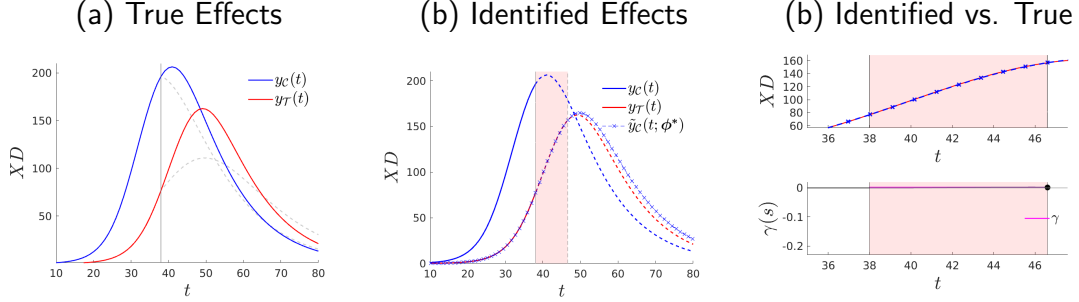
Thus, we are interested in assessing how well SBI can recover the true policy effects that emerge given different confounding factors. For this assessment, we use the benchmark econ-epi model in which there were no confounding factors (Section 3.1.1). Then, we introduce confounding factors (one by one) into the model in order to show how the true policy effects, which now explicitly depend on the specific set of confounding factors that are present, change and assess whether SBI can recover these true policy effects. We first introduce time-varying latent heterogeneity induced by an exogenous time-path of beliefs,²⁸ and then consider additional confounding policy. The results are presented in Online Appendix O.E. By and large, our results

²⁶In Online Appendix O.D, we further assess how the performance of SBI is affected by the amount of pre-policy data that is available to the policy evaluator (Roth, 2022).

²⁷Hence, from the perspective of SBI, the presence of confounding factors does not change neither the aim of the normalization step, which is to reduce the cross-regional differences in the pre-policy determinants—including latent structural parameters—of the path of the outcome of interest, nor the identification assumption.

²⁸This is one potential source of structural regional divergence, as in Rambachan and Roth (2023). In Online Appendix O.F, we separately discuss the bounds under which outcome divergence after policy due to structural parameters (and not to policy) alters the ability of SBI to capture the true policy effects.

Figure 10: Stage-Based Identification of Policy Effects: A Placebo Test



Notes: These panels show the Placebo diagnosis described in Section 3.3.

are consistent with the insights presented in Section 3.2.1: as long as the regional outcome paths maintain reasonable similarity in the presence of confounding factors, SBI can claim success.²⁹

3.3 Placebo Diagnosis

Here, we assess whether SBI identifies policy effects when these are actually non-existent. To conduct this assessment, we use as benchmark the econ-epi model from Section 3.1.1 with the relevant difference that we do not impose a policy at time t_p . Under such scenario, the paths for the flow of deaths in region \mathcal{C} (solid blue) and region \mathcal{T} (solid red) are as depicted in panel (a) of Figure 10. Then, we apply SBI as if there was a policy at some period t_p —when there is actually none. Since the normalization uses only pre-policy data, we obtain the same identification window over stages as if there was an actual policy, see panel (b) of Figure 10. Note that the normalized outcome path $\tilde{y}_{\mathcal{C}}(s, \phi^*)$ is practically identical to the outcome path $y_{\mathcal{T}}(t)$ on the identification window: the identified counterfactual matches the actual outcome path—which here is also the outcome path without policy; see panel (c) of Figure 10. That is, SBI correctly identifies that the policy effects are non-existent—or quantitatively negligible, $\gamma = 0.188\%$.

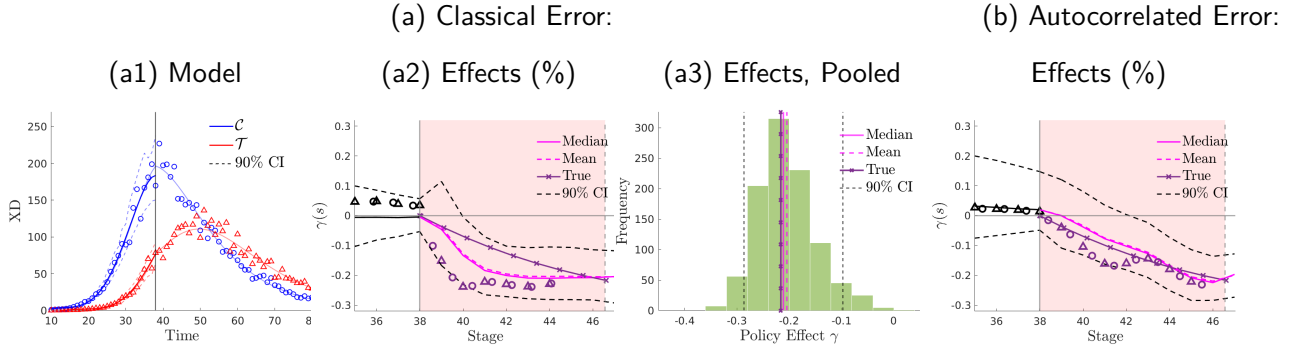
3.4 Inference

In empirical applications, the outcome paths of interest typically incorporate a stochastic component (e.g., measurement error). When facing such noisy data, we add a smoothing—or trend-extraction—step that precedes the normalization step in Section 2. The goal is to purge the observed pre-policy outcome paths of the stochastic fluctuations—of higher frequency than the object of interest—and conduct inference with them. Precisely, let the outcome paths be

$$\hat{y}_r(t) = y_r(t) + u_r(t) \quad \text{with} \quad u_r(t) \underset{\text{i.i.d.}}{\sim} N(0, \sigma_{u_r}^2), \quad (19)$$

²⁹In Online Appendix O.D, we further assess how the performance of SBI is affected by the amount of pre-policy data that is available to the policy evaluator (Roth, 2022).

Figure 11: Stage-Based Identification of Model-Generated Policy Effects: Inference



Notes: Panel (a1) shows results with classical error added to the econ-epi model (Section 3.1.1) with $\{\sigma_{\mathcal{C}}^2, \sigma_{\mathcal{T}}^2\} = \{0.008, 0.008\}$. Panel (a2) shows the policy effect on bootstrap simulations that deliver roughly the same window length (a confidence set of 10%) as the benchmark. Panel (a3) shows the unrestricted distribution of policy effects across all bootstrap draws. Panel (b) presents results for autocorrelated error with $\{\rho_{\mathcal{C}}, \rho_{\mathcal{T}}\} = \{0.13, 0.13\}$ and $\{\sigma_{\mathcal{C}}^2, \sigma_{\mathcal{T}}^2\} = \{0.008, 0.008\}$. The purple markers in panels (a2) and (b) plots the results of conducting SBI without the smoothing step.

for each region $r = \{\mathcal{C}, \mathcal{T}\}$, where $\hat{y}_r(t)$ is the outcome path observable to the policy evaluator, $y_r(t)$ is the unobservable true path and $u_r(t)$ captures classical measurement error with a normal distribution with zero mean and variance $\sigma_{u,r}^2$.³⁰ In panel (a1) of Figure 11, we show the true pre-policy outcome paths for the two regions (light solid lines) together with the observed outcome paths, for one simulation of (19), indicated by the circle markers (region \mathcal{C}) and triangle markers (region \mathcal{T}). Further, we denote by $\hat{y}_r(t < t_p)$ the estimand of $y_r(t < t_p)$ obtained by fitting Chebyshev polynomials of degree 6 to the observed data $\hat{y}_r(t < t_p)$. Thus, we recover a time-series of regional errors $u_r(t) = \hat{y}_r(t) - \hat{y}_r(t)$ for $t < t_p$. Then, for each region, we construct bootstrap draws, $\hat{y}_{r,b}(t)$ with $b \in B = 1,000$. For each bootstrap draw, we randomly draw a sequence of errors, $u_{r,b}(t < t_p)$, from the region-specific set of errors with replacement, which we add to the fitted values of the pre-policy path, $\hat{y}_r(t < t_p)$. In panel (a1) of Figure 11, we show the median (solid lines) and 90% confidence intervals (dashed lines) of the bootstrap pre-policy paths $\hat{y}_{r,b}(t < t_p)$.

Now, we apply SBI to each bootstrap sample. We first recover an estimand $\hat{y}_{r,b}(t < t_p)$, with which we perform the normalization (picking region \mathcal{T} as reference), and then measure the effect of policy for each bootstrap draw, γ_b , in the same way as for the original data sample. The heterogeneity in γ_b across bootstrap draws arises from both, differences in the policy effect by stage inside the identification window, $\gamma_b(s)$, and differences in the size of the window itself, $\mathbb{W}_b(s)$. For this reason, we focus first on bootstrap simulations of around the mean window size (with a confidence set of 10%). Two results emerge. First, the normalized outcome paths are not significantly different from each other—in the stage domain—before policy implementation;

³⁰For this illustration, we use an estimate for $\sigma_{u,r}^2$ similar to that of our Covid-19 empirical application, cf. Section 4.1.

see the non-shaded area in panel (a2), Figure 11.³¹ Second, the identified policy effects are not significantly different from the true (model-generated) policy effects defined as the benchmark without measurement error from Section 3.1.1 (purple line with crossed markers): the mean policy effect (dashed magenta line) is 20.05%—with a 90 percent confidence interval of [11.60,28.91]—which is not significantly different from the true effect, 21.60%. The median effect (solid magenta line) is of similar size, 20.06%. Not restricting the size of the window, we find similar mean policy effects (20.04%); panel (a3), Figure 11. In Online Appendix O.G we discuss an alternative approach in which we assume normality of the residuals.

Further, conducting SBI directly on the observed outcome paths, $\hat{y}_r(t)$, i.e., without the smoothing step, implies identified policy effects similar to the true ones, which suggests that smoothing, which we use in our inference procedure, does not substantially affect the identified policy effect itself; see purple markers in panel (a2), Figure 11. At the same time, it is important to acknowledge that, of course, all these results depend on the size of $\sigma_{u,r}^2$. Last, we re-conduct our analysis assuming an stochastic component that is autocorrelated: $u_r(t) = \rho u_r(t-1) + v_r(t)$ with $\rho \in [0, 1]$ and $v_r(t) \sim N(0, \sigma_{v,r}^2)$. To keep the empirical autocorrelation structure of $u_r(t)$ —including potentially temporal differences in the cross-sectional variance, we use a block bootstrap procedure that increases the sampling weight of preceding error terms in a pre-specified window (5 days) (Carlstein, 1986). Again, SBI can identify policy effects that are not significantly different from the true ones; cf. panel (b), Figure 11.

4 Applications

We use SBI to assess the effects of nationwide policies on several outcomes: the effects of stay-home policies on the flow of Covid-19 deaths in Spain (Section 4.1); the effects of the approval of oral contraceptives on fertility and women’s college education in the U.S. (Section 4.2); and the effects of the German reunification on GDP per capita (Section 4.3). Data sources are in Appendix B. As in the analytical examples of Section 2.3 and the model illustrations of Section 3, we adopt a linear stage-to-time transformation. Since the level shift parameter ω_0 turned out to be insignificant in all applications, throughout the normalization coefficients are $\phi = [\omega_1, \psi_0, \psi_1]$.

4.1 The Spanish *Confinamiento* Against Covid-19

In response to the Covid-19 pandemic, on March 14, 2020, the Spanish government announced a nationwide stay-at-home policy—enacted the following day—which locked down all non-essential

³¹This is analogous to the “pre-trend” tests typically used—in the time domain—by standard empirical strategies. Arguably, this is a necessary but not sufficient condition as post-policy trends might differ from the pre-policy trends (Rambachan and Roth, 2023). See our related discussion in Online Appendix O.F.

workers in all regions of Spain. Indicative of its strictness, the public debate referred to the policy as confinement. The strictest measures were lifted on May 2 when the first wave of the epidemic flattened out. Here, we apply SBI to assess the effects of this policy intervention on the daily flow of deaths attributed to Covid-19. We use two Spanish regions to assess the nationwide policy: Madrid and an artificially created region Rest of Spain (RoSPA) which is composed of all Spanish regions without Madrid. We label Madrid as region \mathcal{C} and RoSPA as region \mathcal{T} .³²

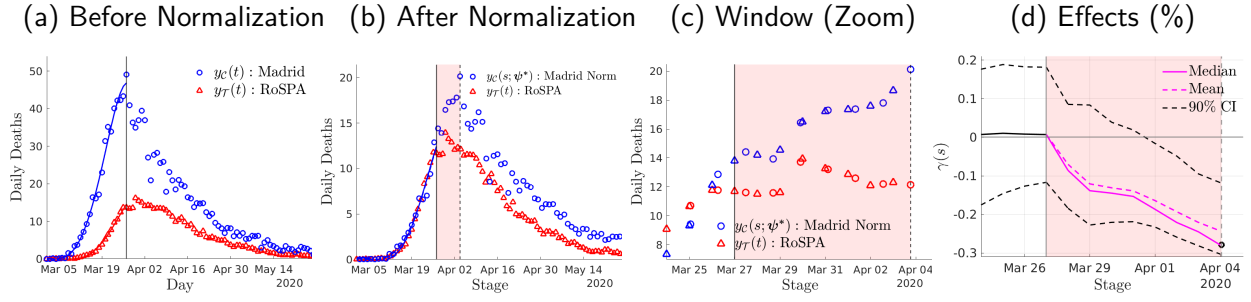
We show the daily flow of Covid-19 deaths (per million inhabitants) for Madrid (blue circles) and RoSPA (red triangles) in panel (a) of Figure 12. In order to mitigate potential measurement error on the reported deaths, we smooth the pre-policy data using as benchmark Chebyshev polynomials separately by region as described in Section 3.4. Note that we add a lag parameter to the policy date, reflecting that a policy that aims at reducing infections will show an effect on the flow of deaths with a delay. We choose a lag of 12 days, which implies that the policy is (effectively) implemented on March 27.³³ The resulting smoothed daily flow of deaths for Madrid (solid blue) and RoSPA (solid red) are also in panel (a) of Figure 12. There are clear differences in the path of the flow of deaths between Madrid and RoSPA. First, one death (per million inhabitants) is reached in March 08 for Madrid and March 14 for RoSPA. Second, by March 14 the daily flow of deaths in Madrid is 9.3 deaths (per million inhabitants) whereas this figure is 1.2 for RoSPA. Furthermore, at the (effective) time of policy implementation, the flow of deaths is reaching a peak in Madrid at 50 deaths (per million inhabitants), whereas the peak in RoSPA is smaller at 16 deaths (per million inhabitants) and occurs about a week after that in Madrid. That is, the flow of deaths starts at an earlier date, it raises more rapidly and reaches a larger peak in Madrid than in RoSPA.

Normalization. To apply SBI we pick RoSPA as reference region \mathcal{T} , and map the (smoothed) flow of deaths of the region Madrid ($y_{\mathcal{C}}(t)$, solid blue circles) onto the flow of deaths of RoSPA ($y_{\mathcal{T}}(t)$, solid red circles) using *only* pre-policy data. The normalization step delivers a normalized path for Madrid $\tilde{y}_{\mathcal{C}}(s; \phi^*)$ that is not different—up to a minimization error—from that of RoSPA, $y_{\mathcal{T}}(t)$; see panel (b) of Figure 12. We find—with bootstrapped confidence intervals reported in brackets— $\psi_0 = -0.14$ [-0.24,0.10], $\psi_1 = 1.21$ [1.18,1.26] and $\omega_1 = 0.41$ [0.34,0.45] which, respectively, delays the start, slows down the growth and lowers the peak of daily deaths in Madrid. A result of our normalization is that Madrid leads the epidemic in Spain. Precisely, the policy is implemented in Madrid at a later stage than in RoSPA, i.e., $s_{\mathcal{T}}(t_p; \psi^*) = t_p < s_{\mathcal{C}}(t_p; \psi^*)$.

³²The Covid-19 pandemic has generated lots of empirical work assessing public health policies against the pandemic; see, for example, Fang et al. (2020) for a study of the early mobility restrictions in China and Liu et al. (2021) for the provision of density forecasts with Bayesian techniques for a panel of countries and regions.

³³As reported by the Instituto de Salud Carlos III (ISCIII) it took between twelve to twenty-three days from infection to death in the first Covid-19 wave in Spain; see isci.iii.es. Here, we take the lower bound.

Figure 12: The Effects of the Spanish *Confinamiento* Against Covid-19



Notes: Panel (a) shows the daily Covid-19 deaths for Madrid, region \mathcal{C} , and for an artificial region \mathcal{T} that aggregates the rest of Spain (RoSPA). We use a Chebyshev smoother (solid lines) of degree 6. Panel (b) shows the results of our normalization using region \mathcal{T} as reference and mapping the pre-policy outcome paths of region \mathcal{C} onto region \mathcal{T} . Panel (c) zooms the identification window. Panel (d) shows the policy effect where γ is defined in equation (8). We show the mean, median and 90% confidence interval bands from bootstrapped simulations constructed as described in Section 3.4. We estimate a significant auto-correlation coefficient for the residuals ($\rho_{\mathcal{C}} = 0.48$ $\rho_{\mathcal{T}} = 0.55$, respectively) and thus perform block-bootstrap with a block window of 5 days. *Source:* Instituto de Salud Carlos III, own calculations.

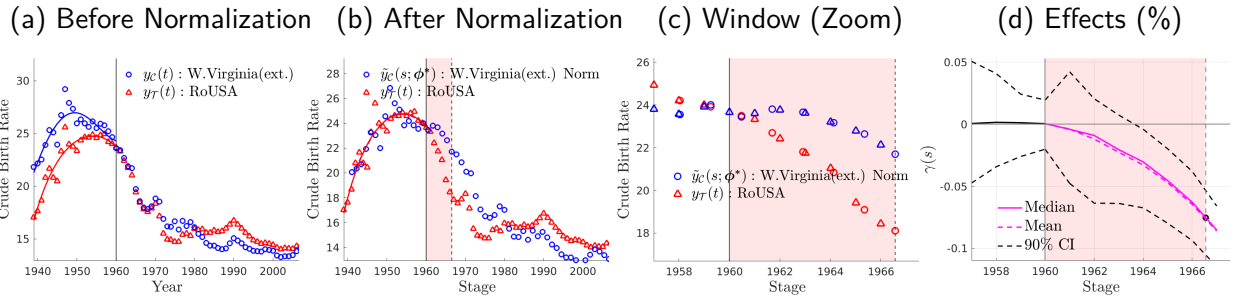
Hence, the normalization unveils a window in stages $\mathbb{W}(s; \psi^*) = [t_p, s_C(t_p; \psi^*)]$ (shaded pink area) running from March 27 to April 03 in which the stage-leading region, Madrid, is not yet subject to policy whereas RoSPA is.³⁴ Therefore, under our identification assumption, the normalized path of Madrid serves as no-policy counterfactual for RoSPA inside $\mathbb{W}(s; \psi^*)$.

Policy Effect. The implied policy effects are in panel (d), where we restrict the attention to the bootstrap simulations within the neighborhood of the median window size (plus/minus 10%). Across these bootstrap draws the (mean) identified total number of lives saved (per million inhabitants) is $\int_{\mathbb{W}(s; \psi^*)} (\tilde{y}_{\mathcal{C}}(s; \phi^*) - y_{\mathcal{T}}(s)) ds = 36.92$ within approximately one week after policy implementation, which corresponds to a total amount of lives saved by the policy in RoSPA of 1,734 during that week. That is, the stay-home policy prevented $\gamma = -24.71\%$ of the total deaths that would have occurred in RoSPA had the policy not been implemented. These effects are significant with a 90% confidence interval of $[-30.13, -11.70]$. The median effect is similar, -26.79% . Further, considering the policy effect across all bootstrap draws (i.e., without restricting the window size) we find similar significant policy effects with mean -22.43% and median -22.33% . Furthermore, we conduct a placebo diagnosis—assuming implementation of the policy before its actual implementation—to find that the identified policy effects that emerge from our method are not significantly different from zero; see Appendix C. Last, redoing our assessment without the smoothing step implies as point estimate that the policy prevented 25.61% of the deaths in RoSPA during approximately the first week; see Online Appendix O.H.³⁵

³⁴Precisely, the window $\mathbb{W}(s; \psi^*)$ runs from the effective policy date in RoSPA ($t_p =$ March 27) to the effective policy date in the stage domain for Madrid, $s_C(t_p; \psi^*) = t_p + 7.7$ days—exactly, at 6.18pm on April 03.

³⁵For all our applications, we do Placebo tests in Appendix C and robustness to trend-extraction in Online Appendix O.H.

Figure 13: The Effects of the 1960 FDA Approval of Oral Contraceptives: Crude Birth Rate



Notes: The normalization uncovers that West Virginia leads the rest of the United States (RoUSA) in that it shows the largest cross-regional stage in terms of crude birth rates at the time of policy implementation. Panel (a) shows the crude birth rate for a region \mathcal{C} which consists of West Virginia (ext.) that also incorporates the next three leading states (Idaho, Nevada and Arkansas). We use a Chebyshev smoother (solid lines) of degree 5. Panel (b) shows the results of our normalization using region \mathcal{T} as reference. Panel (c) zooms the identification window. Panel (d) shows the policy effect $\gamma(s)$ as defined in equation (8). We show the mean, median and 90% confidence interval bands from bootstrapped simulations constructed as described in Section 3.4. We find a non-significant auto-correlation coefficient for the residuals, $\rho_{\mathcal{C}} = 0.31$ $\rho_{\mathcal{T}} = 0.18$, respectively.

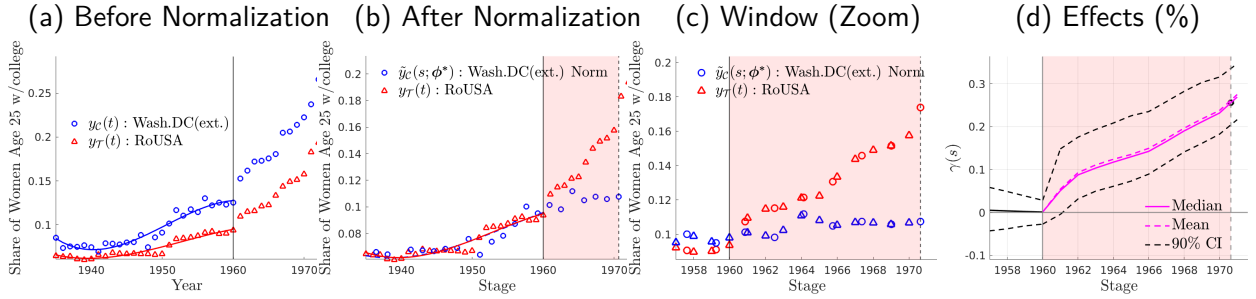
4.2 The 1960 FDA Approval of Oral Contraceptives in the U.S.

In 1960, the first hormonal birth control pill (oral contraceptive) was approved in the U.S. by the Food and Drug Administration (FDA) for use by women above the age of majority. In a seminal paper, [Goldin and Katz \(2002\)](#) use state-level variation in the age of majority in order to assess schooling and career choices of women in that threshold.³⁶ Here, since SBI does not require non-nationwide policy for identification, we assess the effects of the nationwide (federal) approval of the pill on the entire population of adult women. We focus on two outcomes: the crude birth rate and the share of women with completed college (by age 25). The crude birth rates shows an inverted-U shape pattern typically labeled as the baby boom and baby bust; cf. panel (a), Figure 13. The birth rate in the stage-leading region, West Virginia (ext.), peaks in the second half of the 1940s and in 1960 is already busting and close the 1940 levels. Instead, the birth rate in the rest of the U.S. (RoUSA) peaks in the second half of the 1950s at a lower level and has barely started to decline by year 1960. In terms of women's college education, the proportion of women of age 25 with completed college has more than tripled from 8% in 1950 to 26% in 1970 in the stage-leading region, Washington DC (ext.); panel (a), Figure 14. In the RoUSA, this figure shows a larger relative increase from 2% in 1950 to 15% in 1970.

Normalization. In terms of crude birth rates, picking RoUSA as reference region \mathcal{T} , we apply the normalization by mapping the pre-policy birth rates of West Virginia (ext.) ($y_{\mathcal{C}}(t)$, solid blue circles) onto the pre-policy crude births rates of the RoUSA ($y_{\mathcal{T}}(t)$, solid red circles). This results in a normalized path for West Virginia (ext.) $\tilde{y}_{\mathcal{C}}(s; \phi^*)$; see panel (b), Figure 13. The normalization coefficients are $\psi_0=1.80$ [1.79,1.83], $\psi_1=1.21$ [1.20,1.56] and $\omega_1=0.91$ [0.90,0.93]

³⁶Further, [Bailey \(2006\)](#) uses state-level variation in the age of majority to assess the effects of the pill on the timing of first births and women's labor force participation. [Greenwood and Guner \(2010\)](#) use an equilibrium matching model to assess the effects of oral contraceptives on premarital sex and how it is perceived in society.

Figure 14: The Effects of the 1960 FDA Approval of Oral Contraceptives: Women College



Notes: Our normalization uncovers that Washington D.C. leads the rest of the United States (RoUSA) in that it shows the largest cross-regional stage in terms of women’s college completion at the time of policy implementation. Panel (a) shows the crude birth rate for a region \mathcal{C} which consists of Washington D.C.(ext.) that also incorporates the next three leading states (Massachusetts, Colorado and Connecticut). We use a Chebyshev smoother (solid lines) of degree 5. Panel (b) shows the results of our normalization using region \mathcal{T} as reference. Panel (c) zooms the identification window. Panel (d) shows the policy effect $\gamma(s)$ as defined in equation (8). We show the mean, median and 90% confidence interval bands from bootstrapped simulations constructed as described in Section 3.4. We find a non-significant auto-correlation coefficient for the residuals, $\rho_{\mathcal{C}} = 0.21$ $\rho_{\mathcal{T}} = 0.64$, respectively.

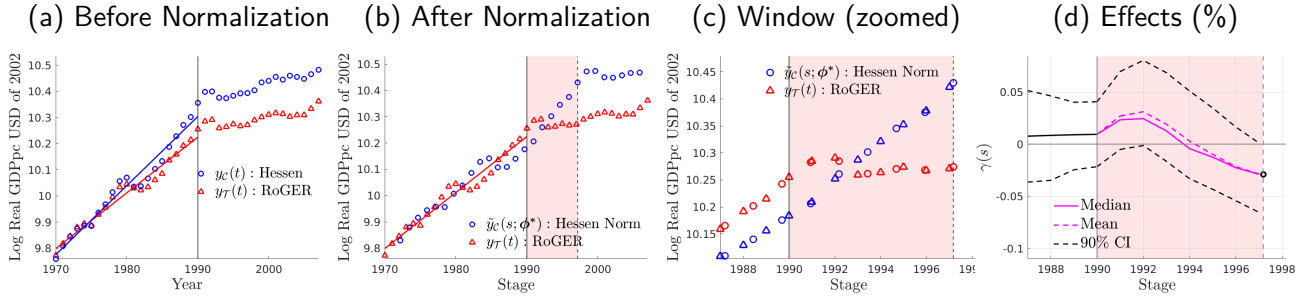
which, respectively, delays, slows down the growth, and lowers the peak of the baby boom for the leading region. The normalization unveils a window of stages $\mathbb{W}(s; \psi^*) = [t_p, s_{\mathcal{C}}(t_p; \psi^*)]$ (shaded pink area) in which West Virginia (ext.) is not subject to policy whereas RoUSA is. Analogously, for the case of women’s college education, we choose RoUSA as reference region \mathcal{T} and map the pre-policy path of Washington DC (ext.) ($y_{\mathcal{C}}(t)$, solid blue circles) onto that of RoUSA ($y_{\mathcal{T}}(t)$, solid red circles); panel (b), Figure 14. The normalization coefficients are $\psi_0 = -5.65$ $[-12.22, 0.32]$, $\psi_1 = 1.62$ $[1.28, 1.79]$ and $\omega_1 = 0.85$ $[0.75, 0.87]$ which generates a window of stages in which Washington D.C. (ext.) is not subject to policy, whereas RoUSA is.

Policy Effect. We zoom the identification window for curde birth rates and women’s college education in panel (c) of, respectively, Figure 13 and Figure 14. We find that access to the pill significantly reduced by $\gamma = -9.38\%$ the number of births (per 10,000 inhabitants) that would have otherwise occurred without the pill; panel (d), Figure 13. The median effects are similar: a -9.39% reduction. Analogously, we find that access to pill significantly increased the share of women with completed college at age 25 by $\gamma = 21.73\%$ during the decade that followed the policy compared to what would have occurred without the pill; panel (c) and (d), Figure 14. The median effects are almost identical, $\gamma = 20.69\%$.

4.3 The German Reunification

In 1990, after the fall of the Berlin wall in 1989, the German Democratic Republic was abolished and integrated fully into the Federal Republic of Germany. Given large differences between the West German states and the East German states, the political and economic reunification came at some cost—the size of which is subject to debate. Abadie et al. (2014) studies the consequences

Figure 15: The Effects of the German Reunification on GDP per capita



Notes: Panel (a) shows the GDP per capita of region \mathcal{C} , Hessen, that leads West Germany and a region \mathcal{T} that aggregates the rest of West Germany. We use a Chebyshev smoother (solid lines) of degree 3. Our normalization using region \mathcal{T} as reference is in panel (b). Panel (c) zooms the identification window. Panel (d) shows the policy effect $\gamma(s)$ as defined in equation (8), with mean, median and 90% confidence interval bands from bootstrapped simulations, see Section 3.4. We find significant auto-correlation coefficients for the residuals ($\rho_{\mathcal{C}} = 0.78$ and $\rho_{\mathcal{T}} = 0.74$, respectively) and thus perform block-bootstrap with block windows of 3 years.

for West Germany forming a counterfactual path for GDP per capita using Synthetic Control Methods (SCM). Here, we apply SBI to the same context, and construct a counterfactual for the evolution of GDP per capita in West Germany had it not been for the reunification. In contrast with Abadie et al. (2014), our counterfactual is constructed using the GDP per capita paths of West German regions only. We focus on Hessen—which our normalization uncovers as the stage-leading region—and an artificially created region for Rest of West Germany (RoGER) which is composed of all West Germany regions excluding Hessen; see panel (a), Figure 15.

Normalization. Picking RoGER as reference region, we apply the normalization by mapping the pre-policy GDP per capita of Hessen ($y_{\mathcal{C}}(t)$, solid blue circles) onto that of RoGER ($y_{\mathcal{T}}(t)$, solid red circles). This results in a normalized path for Hessen $\tilde{y}_{\mathcal{C}}(s; \phi^*)$; see panel (b) of Figure 13. The normalizing parameters are $\psi_0 = 1.95$ [1.94, 1.95], $\psi_1 = 1.24$ [1.21, 1.41] and $\omega_1 = 1.00$ [1.00, 1.01]. The normalization opens a window in stages, $\mathbb{W}(s; \psi^*) = [t_p, s_{\mathcal{C}}(t_p; \psi^*)]$ (shaded pink area), of approximately seven years in which Hessen is not subject to the German reunification but RoGER is. Therefore, under the identification assumption, the normalized path for Hessen provides a no-policy counterfactual for RoGER inside that window.

Policy Effect. We zoom in on the identification window for GDP per capita in panel (c) and the associated policy effects in panel (d) of Figure 13. We find that the German reunification significantly reduced the GDP per capita of RoGER by $\gamma = 3.51\%$ compared to what it would have been without the reunification. The median policy effects are similar, $\gamma = 3.79\%$. In the previous effects, the window size is restricted to be in the neighborhood of the median window size (plus/minus 10%, comprising a total of 541 bootstrap samples). Again, not restricting the window size, we also find significant effects of similar size $\gamma = 3.37\%$. Further, without the smoothing step, the German reunification generates a reduction of 4.82% in the GDP per capita of RoGER, which is not significantly different from our mean bootstrapped effects.

5 Further Discussion

We discuss heterogeneity of policy effects in Section 5.1 and non-nationwide policy in Section 5.2.

5.1 Heterogeneous Policy Effects


To study heterogeneous policy effects, we focus on the nationwide *confinamiento* against Covid-19 in Spain (Section 4.1). The idea is to use the regional paths of Covid-19 deaths—which might differ by stage at the time of policy implementation—in order to measure policy effects by stage. Since Madrid leads all other regions, we pick Madrid as reference. In panel (a) of Figure 16, we show the stage at which policy implementation occurs for each region emerging from separately mapping the pre-policy time path of each region onto that of Madrid. We find substantial heterogeneity in stages at the time of policy implementation. For example, policy arrives to the region of Murcia (MUR) approximately twelve days earlier than the stage at which Madrid received the policy whereas the Basque Country (PVC) precedes Madrid for two days. In addition to the “one-control-many-treatments” approach that we discuss in this Section, we also present an alternative “many-controls-one-treatment” approach in Online Appendix O.I where we exploit the fact that for non-stage leading regions multiple candidate control regions potentially exist.

Our main result is that there are heterogeneous effects by stage—hence, by region. We plot these effects (yellow markers) in panel (b) of Figure 16.³⁷ The policy effects are larger for regions that are at less advanced stages at the time of policy implementation. For example, in MUR, the policy prevented 65% of the deaths that would have otherwise occurred in MUR in a scenario without policy. In contrast, in PVC, which is closest to Madrid in terms of stages at the time of policy implementation, the policy prevented 12% of the deaths that would have otherwise occurred in PVC in a scenario without policy. We reach similar insights from hybrid regions constructed from the power set of all regions excluding Madrid, i.e., a total of $2^{16} - 1 = 65,535$ hybrid regions,³⁸ that we separately map to Madrid. We report the policy effects (tiny purple markers) associated with each of these hybrid regions in panel (b) of Figure 16. Again, the policy effects are systematically larger for regions that are at less advanced stages at policy implementation. Further, note that many hybrid regions enter policy implementation at the same stage and, hence, there is potential heterogeneity in policy effects by stage. Next, we explore this cross-regional heterogeneity by stage as means to perform additional statistical inference.

³⁷As per our Theorem 1, note that we can interpret the reported policy effects as those that would occur to the control (stage-leading) region had the policy been implemented at an earlier stage in that region or as those that actually occur to the treated region at the time of policy implementation.

³⁸A hybrid path between regions A and B is the population weighted sum of the flow of deaths per capita.

Figure 16: Heterogeneous Policy Effects by Stage, *Confinamiento* against Covid-19 in Spain



Notes: Panel (a) shows the stage at which policy implementation occurs after separately mapping the pre-policy time path of each region onto that of Madrid. We have a total of 17 region (comunidad autonoma) names: Andaluca (AND), Aragon (ARA), Asturias (AST), Baleares (BAL), Canarias (CAN), Cantabria (CNT), Castilla-La Mancha (CLM), Castilla y Leon (CLL), Catalunya (CAT), Ceuta (CEU), Valencia (VAL), Extremadura (EXT), Galicia (GAL), Madrid (MAD), Melilla (MEL), Murcia (MUR), Navarra (NAV), Pais Vasco (PVC), La Rioja (RIO), in addition to artificial region Rest of Spain (RoSPA). In panel (b), we report the policy effects γ (see Section 2.2) by region where the (yellow) marker size is the flow of deaths at the time of policy implementation. We also report the policy effects for each (hybrid) region constructed for each element in the the power set $2^{16} - 1$ of regions (tiny purple markers) in panel (b). The reported 90% CI's by stage s are constructed from the distribution of identified policy effects for all regions, including hybrids, where policy implementation occurs at the same stage s defined in rolling windows of 2 stages (or days in the stage domain of Madrid). We exclude the top 5% and bottom 5%. In panel (c) we plot the distribution of policy effects dated March 20th showing that the identified policy effect for RoSPA, -24.17%, falls within the 95% confidence interval [-34.65,-19.73] generated from all regions that enter policy at the same stage as RoSPA. In panel (d) we show the interim effect $\gamma(s)$ (see Section 2.2) by stage for day 1, day 5 and day 10 after policy implementation.

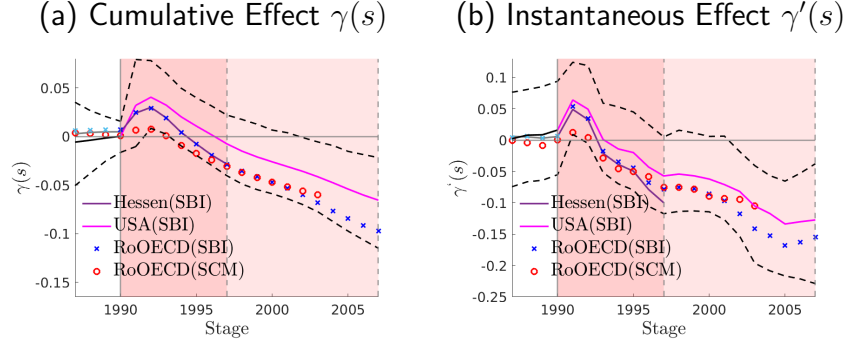
In particular, in panel (c), we plot the distribution of the identified policy effects for all regions (and hybrids) that enter policy at the stage in which Madrid is on March 20th (plus/minus one day), which purposefully includes RoSPA, i.e., our benchmark treatment in Section 4.1. We find reassuring that the identified policy effect of RoSPA is within 95% confidence interval [-34.65,-19.73] generated from all regions that enter policy at the same stage as RoSPA.

Finally, what drives the heterogeneous effects across stages? In addition to the size of the identification window, differences in policy effects can emerge within the same horizon into the policy. Indeed, the policy effects one day into the policy (i.e inside the identification window) vary in a U-shaped fashion across stages; see the magenta markers in panel (d), Figure 16. There is also cross-stage heterogeneity in effects after five (blue markers) or ten days (yellow markers) into the policy. Within panel (d) one can track the policy effect changes by region and stage. For example, consider the region of Canary Islands; the effect grows considerably from 2.5% in day 1 after policy implementation, to 6.0% in day 5, and 7.9% in day 10. The effect does not grow as much for RoSPA; from 1.9% in day 1 after policy implementation, to 2.5% in day 5.

5.2 Non-Nationwide Policy

Consider a scenario with two regions where one region, e.g., \mathcal{T} , receives the policy intervention at period t_p and the other region, e.g., \mathcal{C} , is never treated. In order to identify the policy effects, we need to modify the set on which the normalization is conducted. Picking region \mathcal{T} as reference,

Figure 17: The Effects of the German reunification: Stage-Based Identification (SBI) and Synthetic Control Methods (SCM)



Notes: The outcome variable is real GDP per capita in 2002 USD. To ease the comparison with [Abadie et al. \(2014\)](#), in panel (b), we report the instantaneous policy effects (abusing some notation), $\gamma'(s) = \frac{y_{\mathcal{T}}(s) - \tilde{y}_{\mathcal{C}}(s)}{y_{\mathcal{C}}(s)}$. The 90% CIs correspond to the U.S.

we minimize (5) subject to (3) and

$$\mathbb{C}(s) = \begin{cases} [t_0, t_p] & \text{if } r = \mathcal{T} \\ [s_{\mathcal{C}}(t_0; \psi^*), s_{\mathcal{C}}(t_f; \psi^*)] & \text{if } r = \mathcal{C} \end{cases} \quad (20)$$

for $r = \{\mathcal{C}, \mathcal{T}\}$ where t_0 denotes the first period of observed data and t_f the last. Note that since only region \mathcal{T} is treated, the identification window is,

$$\mathbb{W}(s; \psi^*) = \begin{cases} [t_p, s_{\mathcal{C}}(t_f; \psi^*)] & \text{if } s_{\mathcal{C}}(t_f; \psi^*) > t_p \\ \emptyset & \text{if } s_{\mathcal{C}}(t_f; \psi^*) < t_p \end{cases} \quad (21)$$

Hence, the policy effect can be estimated if the stage of normalized series evaluated at the last period of observed data with, $s_{\mathcal{C}}(t_f; \psi^*)$, falls beyond the stage in which the treated region enters policy. Otherwise, the treated region leads throughout—i.e., the identification window is empty. In [Online Appendix O.J](#) we illustrate such a scenario in the context of the econ-epi model framework.

Here, we exemplify how to use SBI in such a case by re-conducting our assessment of the German reunification. We pair West Germany with the United States and an aggregate consisting of the sample of OECD countries (that excludes Germany) studied in [Abadie et al. \(2014\)](#). Hence, this exercise also serves as means for comparison between SBI and SCM. We pick West Germany as reference region to conduct SBI. Our main finding is that the effects that emerge from using SBI either with the U.S. or the OECD aggregate are not significantly different from those obtained using SCM. In particular, the instantaneous policy effects imply a loss of income per capita for West Germany due to the reunification of 10.51% when compared to the United States and of 14.13% when compared to the OECD in 2003, which are not significantly different from the 10.04% reported in [Abadie et al. \(2014\)](#) for the same year.

6 Conclusion

We provide a new identification method for policy analysis, Stage-Based Identification (SBI). By uncovering heterogeneity in the stage at which policy is implemented across regions, our method allows for the analysis of nationwide policy, which expands the range of policies that can be empirically evaluated. We show the ability of SBI to accurately identify policy effects in various model simulations where the true effect is known, while also acknowledging limitations to our method's performance. Further, we show how SBI can identify heterogeneous effects of policy across stages and be applied to the assessment of non-nationwide policy.

References

- Abadie, A. (2005). Semiparametric Difference-in-Differences Estimators. *Review of Economic Studies*, 72(1):1–19.
- Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, 105(490):493–505.
- Abadie, A., Diamond, A., and Hainmueller, J. (2014). Comparative Politics and the Synthetic Control Method. *American Journal of Political Science*, 59(2):495–510.
- Abadie, A. and Gardeazabal, J. (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review*, 93(1):113–132.
- Adam, K., Marcet, A., and Beutel, J. (2017). Stock price booms and expected capital gains. *American Economic Review*, 107(8):2352–2408.
- Aleman, C., Iorio, D., and Santaaulàlia, R. (2022). A Quantitative Theory of the HIV Epidemic: Education, Risky Sex and Asymmetric Learning. Working paper, Barcelona School of Economics.
- Altonji, J. and Blank, R. (1999). Race and gender in the labor market. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3, chapter 48, pages 3143–3259. Elsevier, 1 edition.
- Angrist, J. D. and Krueger, A. B. (1999). Empirical strategies in labor economics. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3, chapter 23, pages 1277–1366. Elsevier.
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.
- Athey, S. and Imbens, G. W. (2006). Identification and inference in nonlinear difference-in-differences models. *Econometrica*, 74(2):431–497.

- Athey, S. and Imbens, G. W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, 31(2):3–32.
- Athey, S. and Imbens, G. W. (2021). Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics*.
- Atkeson, A. (2020). What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios. Working Paper 26867, National Bureau of Economic Research.
- Bailey, M. J. (2006). More Power to the Pill: The Impact of Contraceptive Freedom on Women's Life Cycle Labor Supply. *The Quarterly Journal of Economics*, 121(1):289–320.
- Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics*, 119(1):249–275.
- Blundell, R. and Macurdy, T. (1999). Labor supply: A review of alternative approaches. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3 of *Handbook of Labor Economics*, chapter 27, pages 1559–1695. Elsevier.
- Borusyak, K., Jaravel, X., and Spiess, J. (2021). Revisiting Event Study Designs: Robust and Efficient Estimation. Working papers, University College London.
- Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Card, D. (1990). The impact of the mariel boatlift on the miami labor market. *ILR Review*, 43(2):245–257.
- Card, D. (2022). Design-based research in empirical microeconomics. *American Economic Review*, 112(6):1773–81.
- Card, D. and Krueger, A. B. (2000). Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania: Reply. *The American Economic Review*, 90(5):1397–1420.
- Carlstein, E. (1986). The use of subseries values for estimating the variance of a general statistic from a stationary sequence. *The Annals of Statistics*, 14(3):1171–1179.
- Cavalcanti, T., Kocharkov, G., and Santos, C. (2021). Family Planning and Development: Aggregate Effects of Contraceptive Use. *The Economic Journal*, 131(634):624–657.
- Cervellati, M. and Sunde, U. (2015). The economic and demographic transition, mortality, and comparative development. *American Economic Journal: Macroeconomics*, 7(3):189–225.
- Doepke, M., Hannusch, A., Kindermann, F., and Tertilt, M. (2023). The Economics of Fertility: A New Era. In *Handbook of the Economics of the Family*, volume 1, chapter 4, pages 151–254. Elsevier.

- Duval-Hernández, R., Fang, L., and Ngai, L. R. (2022). Taxes, Subsidies, and Gender Gaps in Hours and Wages. Working papers.
- Eichenbaum, M. S., Rebelo, S., and Trabandt, M. (2021). The Macroeconomics of Epidemics. *The Review of Financial Studies*, 34(11):5149–5187.
- Fang, H., Wang, L., and Yang, Y. (2020). Human mobility restrictions and the spread of the novel coronavirus (2019-ncov) in china. *Journal of Public Economics*, 191:104272.
- Farboodi, M., Jarosch, G., and Shimer, R. (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory*, 196:105293.
- Fernández-Villaverde, J., Greenwood, J., and Guner, N. (2014). From shame to game in one hundred years: An economic model of the rise in premarital sex and its de-stigmatization. *Journal of the European Economic Association*, 12(1):25–61.
- Galor, O. and Weil, D. N. (2000). Population, technology, and growth: From malthusian stagnation to the demographic transition and beyond. *American Economic Review*, 90(4):806–828.
- Glover, A., Heathcote, J., Krueger, D., and Ríos-Rull, J.-V. (2020). Health versus Wealth: On the Distributional Effects of Controlling a Pandemic. Working Paper 27046, National Bureau of Economic Research.
- Goldin, C. and Katz, L. F. (2002). The Power of the Pill: Oral Contraceptives and Women’s Career and Marriage Decisions. *Journal of Political Economy*, 110(4):730–770.
- Gollin, D., Parente, S., and Rogerson, R. (2002). The role of agriculture in development. *American Economic Review*, 92(2):160–164.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.
- Greenwood, J. and Guner, N. (2010). Social change: The sexual revolution. *International Economic Review*, 51(4):893–923.
- Greenwood, J., Seshadri, A., and Vandenbroucke, G. (2005). The Baby Boom and Baby Bust. *American Economic Review*, 95(1):183–207.
- Gunsilius, F. F. (2023). Distributional Synthetic Controls. *Econometrica*, 91(3):1105–1117.
- Hansen, G. D. and Prescott, E. C. (2002). Malthus to solow. *American Economic Review*, 92(4):1205–1217.
- Heckman, J. J., Ichimura, H., and Todd, P. (1998). Matching As An Econometric Evaluation Estimator. *Review of Economic Studies*, 65(2):261–294.

- Herrendorf, B., Rogerson, R., and Valentinyi, A. (2014). Growth and Structural Transformation. In *Handbook of Economic Growth*, volume 2, chapter 06, pages 855–941. Elsevier, 1 edition.
- Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.
- Iorio, D. and Santaèulàlia-Llopis, R. (2010). Education, HIV Status, and Risky Sexual Behavior: How Much Does the Stage of the HIV Epidemic Matter? Working Paper, Wash. U in St. Louis.
- Iorio, D. and Santaèulàlia-Llopis, R. (2016). Education, HIV Status, and Risky Sexual Behavior: How Much Does the Stage of the HIV Epidemic Matter? Working papers, Barcelona School of Economics.
- Kaplan, G., Moll, B., and Violante, G. (2020). The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the U.S. Technical report.
- Liu, L., Moon, H. R., and Schorfheide, F. (2021). Panel forecasts of country-level Covid-19 infections. *Journal of Econometrics*, 220(1):2–22.
- Rambachan, A. and Roth, J. (2023). A More Credible Approach to Parallel Trends. *The Review of Economic Studies*.
- Rogerson, R. (2008). Structural transformation and the deterioration of european labor market outcomes. *Journal of Political Economy*, 116(2):235–259.
- Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–22.
- Roth, J. and Sant’Anna, P. H. C. (2023). When Is Parallel Trends Sensitive to Functional Form? *Econometrica*, 91(2):737–747.

A Proof of Theorem 1

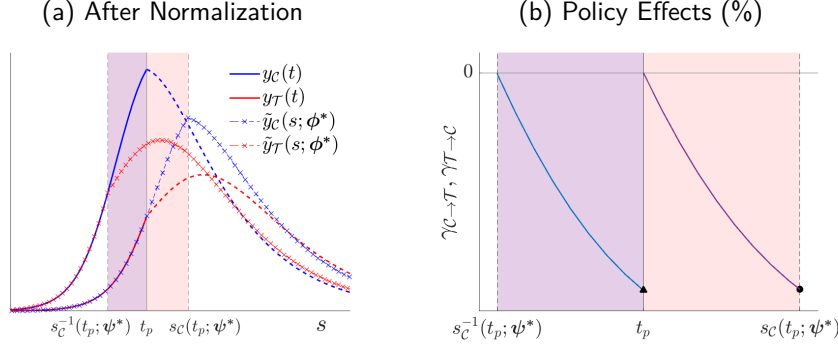
We first formalize the alternative mapping and discuss the equivalence intuitively. Stages for region \mathcal{C} are now redefined as $\mathbf{s} = \mathbf{t}_{\mathcal{C}}(\mathbf{s}) = t$. Stages for region \mathcal{T} are obtained by mapping the outcome path $y_{\mathcal{T}}(t)$ onto $y_{\mathcal{C}}(t)$ using pre-policy data, using the modified (to region \mathcal{T}) expressions in (2) and (3). The corresponding outcome path for region \mathcal{C} is $\tilde{y}_{\mathcal{C}}(\mathbf{s}) = y_{\mathcal{C}}(t = \mathbf{s})$. The normalized path of \mathcal{T} is denoted by $\tilde{y}_{\mathcal{T}}(\mathbf{s}, \bar{\phi})$.

Region \mathcal{C} (while now the reference region) is still the leading region, and we obtain the identification window $\mathbb{W}(\mathbf{s}; \bar{\psi}^*) = [\mathbf{s}_{\mathcal{T}}(t_p; \bar{\psi}^*), t_p]$, where $\bar{\psi}^*$ ($\neq \psi^*$) is the coefficient vector that is obtained when choosing \mathcal{C} as reference region. Inside this window, the policy effect is obtained by relating the normalized path $\tilde{y}_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)$ (which is treated) to the observed path of region \mathcal{C} (which is untreated).

Going back to our example, we show the relationship between the two alternative mappings in panel (a) of Figure A.1. Using region \mathcal{T} as reference yields the identification window $\mathbb{W}(s; \psi^*) = [t_p, s_{\mathcal{C}}(t_p; \psi^*)]$ (pink shaded

area), whereas using region \mathcal{C} as reference yields the identification window $\mathbb{W}(\mathbf{s}; \bar{\psi}^*) = [\mathbf{s}_{\mathcal{T}}(t_p; \bar{\psi}^*), t_p]$, which corresponds to $[s_{\mathcal{C}}^{-1}(t_p; \psi^*), t_p]$ (purple shaded area). The associated policy effects are shown in panel (b) of Figure A.1.

Figure A.1: Policy Effects with Alternative Reference Region



Notes: In panel (b), we report the policy effects γ (black markers) together with the interim cumulative effects of policy, $\gamma(s)$, as defined in Section 2.2.

In the context of the example, the policy effects obtained under the two alternative mappings are identical, assuming that the normalization holds exactly³⁹, because there are no level differences across regions at t_0 , i.e., $\omega_0 = 0$. As this insight holds generally, we summarize it in the following

Lemma 1. *If $\omega_0^* = 0$, then the choice of the reference region is irrelevant for the policy effect, i.e.,*

$$\gamma = \frac{\int_{\mathbb{W}(\mathbf{s}; \psi^*)} (y_{\mathcal{T}}(s) - \tilde{y}_{\mathcal{C}}(s; \phi^*)) ds}{\int_{\mathbb{W}(\mathbf{s}; \psi^*)} \tilde{y}_{\mathcal{C}}(s; \phi^*) ds} = \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (\tilde{y}_{\mathcal{T}}(\mathbf{s}; \bar{\phi}^*) - y_{\mathcal{C}}(\mathbf{s})) ds}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} y_{\mathcal{C}}(\mathbf{s}) ds}. \quad (\text{A.1})$$

Proof. We start with the expression on the RHS of (A.1), i.e., the policy effect expressed with region \mathcal{C} as the reference region, and go on to show that it is equal to the LHS, i.e., the policy effect expressed with region \mathcal{T} as reference region.

For the following exposition, we write it as if the normalization holds exactly. If the minimization does not give exact normalization, then the expressions hold approximately. Thus, when choosing \mathcal{T} as reference we obtain

$$y_{\mathcal{T}}(s) = \tilde{y}_{\mathcal{C}}(s, \phi^*) = \omega_1^* y_{\mathcal{C}}(t_{\mathcal{C}}(s, \psi^*)), \quad (\text{A.2})$$

where s is region \mathcal{T} -time. Likewise, when choosing \mathcal{C} as reference we obtain

$$y_{\mathcal{C}}(\mathbf{s}) = \tilde{y}_{\mathcal{T}}(\mathbf{s}, \bar{\phi}^*) = \bar{\omega}_1^* y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)), \quad (\text{A.3})$$

where \mathbf{s} is region \mathcal{C} -time.

With the above we can thus replace $\tilde{y}_{\mathcal{T}}(\mathbf{s}; \bar{\phi}^*)$ and $y_{\mathcal{C}}(\mathbf{s})$ in the RHS expression of (A.1). In particular, note that $\tilde{y}_{\mathcal{T}}(\mathbf{s}; \bar{\phi}^*) = \bar{\omega}_1^* y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*))$ using the definition of the normalization. Then, note that

$$y_{\mathcal{C}}(\mathbf{s}) = \bar{\omega}_1^* y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)) = \bar{\omega}_1^* \tilde{y}_{\mathcal{C}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*), \phi^*).$$

³⁹Otherwise, the two alternative mappings yield the same policy effect up to a minimization error.

Plugging in those two expressions into the RHS of (A.1), we thus obtain

$$\begin{aligned}
\frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (\tilde{y}_{\mathcal{T}}(\mathbf{s}; \bar{\phi}^*) - y_{\mathcal{C}}(\mathbf{s})) d\mathbf{s}}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} y_{\mathcal{C}}(\mathbf{s}) d\mathbf{s}} &= \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (\bar{\omega}_1^* y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)) - \bar{\omega}_1^* \tilde{y}_{\mathcal{C}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*), \phi^*)) d\mathbf{s}}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} \bar{\omega}_1^* \tilde{y}_{\mathcal{C}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*), \phi^*) d\mathbf{s}} \\
&= \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)) - \tilde{y}_{\mathcal{C}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*), \phi^*)) d\mathbf{s}}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} \tilde{y}_{\mathcal{C}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*), \phi^*) d\mathbf{s}} \\
&= \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (y_{\mathcal{T}}(s) - \tilde{y}_{\mathcal{C}}(s, \phi^*)) ds}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} \tilde{y}_{\mathcal{C}}(s, \phi^*) ds}, \tag{A.4}
\end{aligned}$$

which is equal to the LHS of (A.1). The last step above follows from writing the integral that is taken over \mathbf{s} in terms of $s = t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)$ directly, taking into account that the two mappings are explicitly linked because $s_{\mathcal{C}}^{-1}(t; \bar{\psi}^*) = \mathbf{s}_{\mathcal{T}}(t; \bar{\psi}^*)$ for any $t \in T$ (e.g., $t = t_p$)—that is, the mapping \mathcal{T} to \mathcal{C} undoes the mapping \mathcal{C} to \mathcal{T} , and vice versa.⁴⁰ \square

More generally, if $\omega_0^* \neq 0$, i.e., if there is a non-zero level difference between the regional outcome paths, then we can establish equivalence of the identified policy effects after taking into account the change of the reference unit. Given the interpretation as measuring the effect on the treated region, our preferred way to measure policy effects is in units of the treated region \mathcal{T} . Therefore, we adjust the numerator in the RHS of (A.1) by subtracting the level shifter ω_0 in the denominator to obtain

Lemma 2. *If $\omega_0^* \neq 0$, then the choice of the reference region is irrelevant for the policy effect after expressing the policy effect in an adjusted form when using region \mathcal{C} as the reference unit, i.e.,*

$$\gamma = \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (y_{\mathcal{T}}(s) - \tilde{y}_{\mathcal{C}}(s; \phi^*)) ds}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} \tilde{y}_{\mathcal{C}}(s; \phi^*) ds} = \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (\tilde{y}_{\mathcal{T}}(\mathbf{s}; \bar{\phi}^*) - y_{\mathcal{C}}(\mathbf{s})) d\mathbf{s}}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} y_{\mathcal{C}}(\mathbf{s}) d\mathbf{s} - \bar{\omega}_0^*}. \tag{A.5}$$

Proof. Analogous to above, note that

$$y_{\mathcal{T}}(s) = \tilde{y}_{\mathcal{C}}(s, \phi^*) = \omega_0^* + \omega_1^* y_{\mathcal{C}}(t_{\mathcal{C}}(s, \bar{\psi}^*)),$$

where s is region \mathcal{T} -time. Likewise, when choosing \mathcal{C} as reference we obtain

$$y_{\mathcal{C}}(\mathbf{s}) = \tilde{y}_{\mathcal{T}}(\mathbf{s}, \bar{\phi}^*) = \bar{\omega}_0^* + \bar{\omega}_1^* y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)),$$

where \mathbf{s} is region \mathcal{C} -time.

With the above we can thus replace $\tilde{y}_{\mathcal{T}}(\mathbf{s}; \bar{\phi}^*)$ and $y_{\mathcal{C}}(\mathbf{s})$ in the RHS expression of (A.5). In particular, note that $\tilde{y}_{\mathcal{T}}(\mathbf{s}; \bar{\phi}^*) = \bar{\omega}_0^* + \bar{\omega}_1^* y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*))$. Then, note that

$$y_{\mathcal{C}}(\mathbf{s}) = \bar{\omega}_0^* + \bar{\omega}_1^* y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)) = \bar{\omega}_0^* + \bar{\omega}_1^* \tilde{y}_{\mathcal{C}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*), \phi^*).$$

⁴⁰To see this, note that for any $t \in T$ (e.g., $t = t_p$), the stage function in the mapping \mathcal{C} to \mathcal{T} , $s_{\mathcal{C}}(t; \bar{\psi}^*) = s = t$, injects t into t whereas in the mapping \mathcal{T} to \mathcal{C} the function $\mathbf{s}_{\mathcal{T}}(t; \bar{\psi}^*) = \mathbf{s} = t$ injects t into t .

Plugging in those two expressions into the RHS of (A.5), we thus obtain

$$\begin{aligned} \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (\tilde{y}_{\mathcal{T}}(\mathbf{s}; \bar{\phi}^*) - y_{\mathcal{C}}(\mathbf{s})) d\mathbf{s}}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} y_{\mathcal{C}}(\mathbf{s}) d\mathbf{s} - \bar{\omega}_0^*} &= \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (\bar{\omega}_1^* y_{\mathcal{T}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*)) - \bar{\omega}_1^* \tilde{y}_{\mathcal{C}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*), \phi^*)) d\mathbf{s}}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} \bar{\omega}_1^* \tilde{y}_{\mathcal{C}}(t_{\mathcal{T}}(\mathbf{s}, \bar{\psi}^*), \phi^*) d\mathbf{s}} \\ &= \frac{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} (y_{\mathcal{T}}(s) - \tilde{y}_{\mathcal{C}}(s, \phi^*)) ds}{\int_{\mathbb{W}(\mathbf{s}; \bar{\psi}^*)} \tilde{y}_{\mathcal{C}}(s, \phi^*) ds}. \end{aligned}$$

□

Lemmas 1–2 thus prove Theorem 1. □□

In Online Appendix O.K, we show an alternative Lemma for the case of $\omega_0^* \neq 0$, which does not resort to the level shifter ω_0 . In this Lemma 3, we establish equivalence of the identified policy effects for a slightly modified policy effect that works with adjusted regional time paths paths, where each path is expressed relative to its value at the policy date t_p in region \mathcal{T} -time.

B Data

In this Appendix, we discuss the data construction and sources of our empirical applications conducted in Section 4.

Covid-19 Application. For the assessment of the national lockdown against the first wave of Covid-19 in Spain in Section 4.1, we use Covid-19 deaths regional series provided by the Ministerio de Sanidad. Data from the Ministerio de Sanidad can be found under the following link: www.mscbs.gob.es. All values in figure 12 are expressed per million inhabitants of each region, with Spain having 47 million inhabitants and Madrid 6.6 million.⁴¹

Oral Contraceptives Application. For our assessment of the effects of the 1960 FDA approval of oral contraceptives on fertility and women college education in Section 4.2, we use state-level data on crude birth rates and the share of women with completed college of age 25. First, To construct the crude birth rate (by state), we divide total number of births in a given year by the respective population. Birth counts by state from 1939 to 2007 are provided by IPUMS NHGIS, Vital Statistics: Natality and Mortality Data. We use population data provided by the U.S. Census Bureau annual estimates. Second, to measure the share of women of a certain age with college attainment we use decennial CENSUS data from IPUMS starting in 1940 up to 1980. In the absence of information on the year of graduation, we construct the historical series by using cohort information by CENSUS year. For example, when using CENSUS data for 1960, the share of college women of age 25 in 1959 will be the share of a woman age 26 who reported (already) having attained college by 1960.⁴² After computing the historical series per CENSUS year we compute the average across CENSUS series.

German Reunification Application. To assess the effects of the German reunification on the income per capita of West Germany in Section 4.3, we retrieve 1970-2007 GDP data for West Germany and the respective price index from the Destatis portal of the Statistisches Bundesamt. Our GDP measure of West Germany does not include

⁴¹We exclude Galicia due to the fact that we find positive (yet, non-significant) effects of the policy on the flow of deaths.

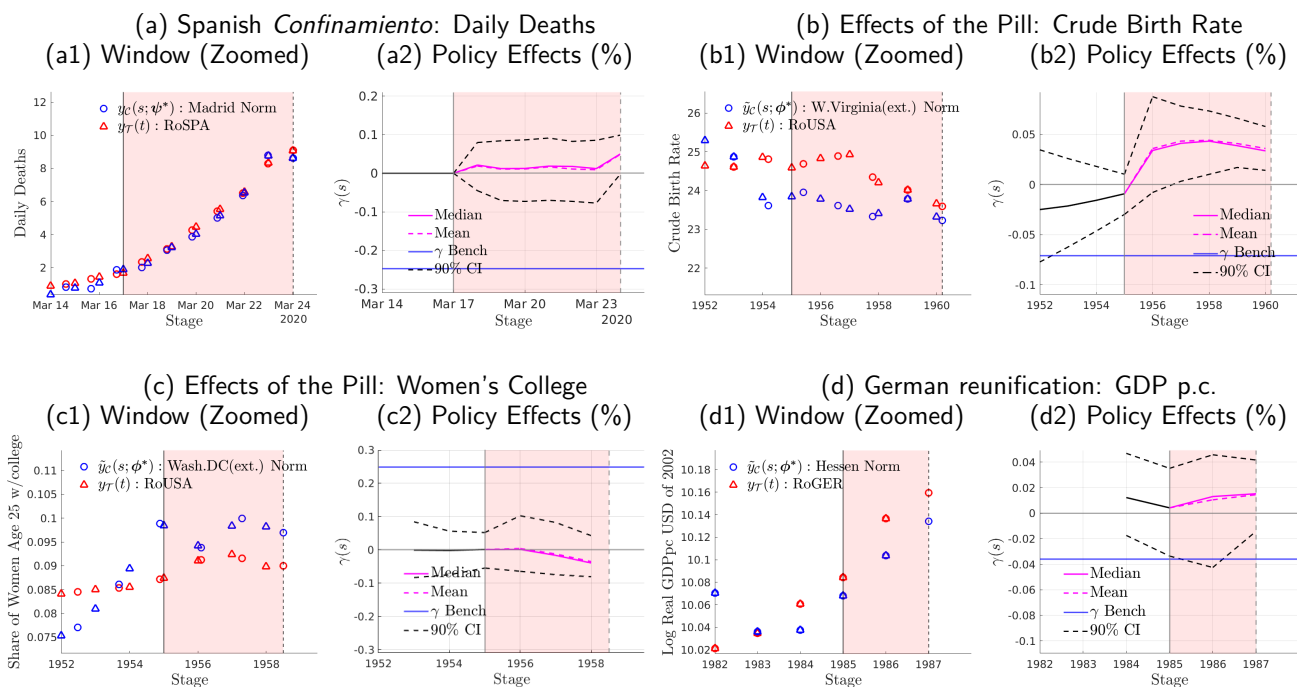
⁴²Later completion and death could hinder the precision of our measure, however after comparing the historical series from various census years, the measure does not seem to be suffering from these problems

West Berlin, and our results do not change significantly when including West Berlin. GDP data, inflation, PPP deflators and population for OECD countries for 1970-2007 is taken from the OECD data portal OECD.STAT. The list of countries that comprises our hybrid Rest of OECD (RoOECD) is the same as the one chosen in [Abadie et al. \(2014\)](#), namely Australia, Austria, Belgium, Denmark, France, Greece, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Spain, Switzerland, United Kingdom and the United States.

C Placebo Diagnosis for our Applications

Here, we conduct further inference using a placebo diagnosis in each of the applications of Section 4. That is, we show the policy effects that emerge from our method if we conduct the normalization at a some period before the actual policy took place. Precisely, for the placebo, we assume that the stay-home policy was imposed on March 10 2020, i.e., about two weeks earlier than its actual implementation; the pill was introduced 5 years before its actual market release; and the German reunification occurred 5 years before the actual reunification date. SBI will survive this test if it identifies policy effects that are not significantly different from zero. We find that this is the case for each application; see Figure C.2.

Figure C.2: Stage-Based Identification of Policy Effects: Placebo Diagnosis



Notes: In panel (a), we assume that the stay-home policy was imposed two weeks earlier than its actual effective date, March 10 2020; in panels (b) and (c), we assume that the the pill was introduced 5 years before its actual market release; in panel (d), we assume that the German reunification occurred 5 years before the actual reunification date.