

Counterfactual Analysis for Structural Dynamic Discrete Choice Models*

Myrto Kalouptsi, Yuichi Kitamura, Lucas Lima, and Eduardo Souza-Rodrigues[†]

June 2021

Abstract

Discrete choice data allow researchers to recover differences in utilities, but these differences may not suffice to identify policy-relevant counterfactuals of interest. This fundamental tension is important for dynamic discrete choice models because agents' behavior depends on value functions, which require utilities in levels. We propose a unified approach to investigate how much one can learn about counterfactual outcomes under mild assumptions, for a large and empirically relevant class of counterfactuals. We derive analytical properties of sharp identified sets under alternative model restrictions and develop a valid inference approach based on subsampling. To aid practitioners, we propose computationally tractable procedures that bypass model estimation and directly obtain the identified sets for the counterfactuals and the corresponding confidence sets. We illustrate in Monte Carlos, as well as an empirical exercise of firms' export decisions, the informativeness of the identified sets, and we assess the impact of (common) model restrictions on results.

KEYWORDS: Dynamic Discrete Choice, Counterfactual, Partial Identification, Structural Model.

*We would like to thank Victor Aguirregabiria, Phil Haile, Bo Honoré, Arthur Lewbel, Rob McMillan, Francesca Molinari, Ismael Mourifié, Salvador Navarro, Ariel Pakes, Aureo de Paula, John Rust, Elie Tamer, Xun Tang, Petra Todd, and Yuanyuan Wan for helpful discussions, and workshop participants at Cornell University, Ohio University, Saint Louis FED, UCLA, University of Toronto, the Gary Chamberlain Online Seminar in Econometrics, the ASSA Meetings 2019, the 3rd Conference on Structural Dynamic Models, the Counterfactuals with Economic Restrictions Conference at Western Ontario, the 5th Banff Empirical Microeconomics Workshop, the Interactions Conference: Bringing Together Econometrics and Applied Microeconomics at the University of Chicago, and the First IO Workshop at Javeriana University for additional comments. This paper circulated previously with the title "Partial Identification and Inference for Dynamic Models and Counterfactuals." Financial support from the University of Toronto Mississauga is gratefully acknowledged. All remaining errors are our own.

[†]Affiliations: Myrto Kalouptsi, Harvard University, CEPR and NBER; Yuichi Kitamura, Yale University and Cowles Foundation for Research in Economics; Lucas Lima, Harvard University; Eduardo Souza-Rodrigues, University of Toronto.

1 Introduction

Structural models have been used to answer a wide range of counterfactual questions in various fields of economics, including industrial organization, labor, public finance, and trade. In the important class of discrete choice models, it is well-known that choice data allow researchers to recover differences in individuals’ valuations: in static models, we can identify differences in flow utilities (McFadden, 1974; Train, 2009); in dynamic models, we recover differences in expected discounted streams of utilities (Rust, 1994; Magnac and Thesmar, 2002). In the latter case, the knowledge of utility differences may not suffice to identify many counterfactual outcomes of interest, as they may require knowledge of utility levels. The tension between identifying differences in valuations from choice data and the (potential) need for utility levels for counterfactuals can pose challenges to the credibility of structural estimation.

The tension is particularly relevant for structural dynamic discrete choice (DDC) models because agents’ behavior depends on future expected utilities (the value function), whose calculation requires knowledge of cardinal utilities. Applied researchers have typically addressed this issue by imposing restrictions, including assumptions often referred to as “normalizations,” that select among observationally equivalent models. However, economic theory does not always offer guidance as to the correct assumptions necessary to recover utility in levels. Further, a recent literature has shown that only a narrow class of counterfactual experiments results in counterfactual behavior that requires just utility differences recoverable from the data; see Aguirregabiria and Suzuki (2014), Norets and Tang (2014), and Kalouptsi, Scott, and Souza-Rodrigues (2021). Thus, a critical question concerns how much one can learn about counterfactual outcomes of interest for a large (and empirically relevant) class of counterfactual experiments, for which knowing utilities in levels may matter.

In this paper, we develop a tractable framework to address this question. We avoid assumptions that select (i.e., point-identify) flow utilities, but allow for milder model restrictions that can be more grounded on economic theory and on the specifics of particular empirical problems. As a result, model parameters and counterfactual outcomes are partially identified. We focus on bounds for low-dimensional counterfactual objects that are amenable to economic interpretation and relevant to researchers’ main conclusions, such as average welfare. We show how to characterize and compute sharp bounds for these objects (that is, bounds that exhaust all implications of the model and data), along with an asymptotically valid inference procedure. We also show, in the context of two applied examples, how alternative model restrictions can lead to informative bounds in practice. Our procedure is valid for a broad class of counterfactual experiments, model restrictions, and outcomes of interest, as we explain below. Our primary motivation is to offer a solution practitioners can use for a class of models used widely in empirical work.

To fix ideas, consider our running example: a firm that decides every period whether to enter (exit) a market subject to entry costs (scrap values), with the goal of maximizing its expected discounted stream

of payoffs consisting of variable profits minus fixed costs, under uncertainty. Typically, to estimate this model, researchers assume the payoff of staying out of the market (the ‘outside option’) is zero, and also impose that scrap values and/or fixed costs do not depend on state variables, and are equal to zero. While these assumptions suffice to identify flow payoffs, assuming scrap values or fixed costs are invariant over states may be strong for some industries; perhaps more importantly, setting them equal to the payoff of the outside option (zero) may be difficult to justify: economic theory does not provide guidance as to how to set these values, and cost or scrap value data are extremely rare. Mistakenly setting the scrap value to zero, for example, lowers the firm’s expected lifetime payoffs (since nothing is obtained each time the firm exits), which in turn generates a (possibly severe) downward bias in the estimated entry cost to rationalize the data. Most important, these assumptions are not always innocuous for important counterfactuals (Aguirregabiria and Suzuki, 2014; Norets and Tang, 2014; Kalouptsi, Scott, and Souza-Rodrigues, 2021). Consider, for instance, a counterfactual exploring the impact of an entry cost subsidy: mistakenly setting the scrap value to zero not just potentially leads to a quantitatively wrong prediction of the firm’s behavior, but even to a wrong sign of the subsidy’s effect (see Kalouptsi, Scott, and Souza-Rodrigues, 2021).

Given these limitations, we avoid such assumptions here. We bypass estimating the model and focus directly on the identified set of counterfactual objects (e.g., the welfare impact of a counterfactual entry subsidy) under much milder restrictions, such as that entry costs and fixed costs are positive, that variable profits are monotonic in demand shocks, or that entry is eventually profitable. In a numerical example (and in a Monte Carlo study), we illustrate that the identified sets are informative even under the mildest assumptions. We also demonstrate how the sets shrink as we add alternative model restrictions, providing intuition about the source and strength of identification. Finally, we explore which results survive under alternative model restrictions and show that we can reject common assumptions, such as zero scrap values.

For dynamic discrete choice models more generally, we first show that for a broad class of counterfactuals involving almost any change in the primitives, the sharp identified set for the counterfactual conditional choice probabilities (CCP) is a connected manifold with a dimension that can be determined from the data.¹ The set is therefore either empty (which occurs when the model is rejected by the data), or a singleton (implying point-identification), or a continuum. The dimension of the set can be calculated by checking the rank of a specific matrix, which is known by the econometrician. This dimension is typically much smaller than the dimension of the conditional probability simplex, which implies that the identified set is informative. Specific combinations of model restrictions and counterfactual experiments can reduce the dimension of the identified set further, leading to point identification in some cases (i.e., model restrictions may suffice to identify counterfactual CCPs even when they do not point-identify the model parameters). To our knowledge, while partial identification and estimation of model parameters

¹We consider a fully general class of changes in the primitives except for nonlinear transformations in payoffs (uncommon in practice).

in DDC models have been considered previously (see, e.g., Bajari, Benkard, and Levin, 2007; Norets and Tang, 2014; Berry and Compiani, 2020), these are the first analytical results characterizing the identified set of counterfactual behavior.

We then turn to the low-dimensional outcomes of interest – our main focus. Here, we show that the sharp identified set is also connected and, under additional mild conditions, compact. This is convenient as in practice it is sufficient to trace the boundary of the set. In addition, when the outcome of interest is a scalar, the identified set becomes a compact interval, in which case it suffices to calculate the lower and upper endpoints. The endpoints can be computed by solving well-behaved constrained minimization and maximization problems. The optimizations can be implemented using standard software (e.g., Knitro), and remain feasible even in high-dimensional cases involving large state spaces or a large number of model parameters. Standard solvers perform best when the researcher provides the gradient of the objective function; to help practitioners, we show how to calculate the gradient in certain nontrivial cases (e.g. long-run average effects). For cases where computing the gradient is (prohibitively) costly, we develop and propose an alternative, stochastic search procedure that takes advantage of the structure of the problem (discussed in detail in the Appendix).

The alternative computational approaches that we propose are both flexible and tractable, and they bypass the need to estimate the model parameters. Overall, an attractive feature of this procedure is that the researcher can flexibly choose (i) the set of model restrictions, (ii) the counterfactual experiment, and (iii) the target outcome of interest, all without having to derive additional analytical identification results for each possible specification.

Our approach leads naturally to an inference procedure. We develop an asymptotically valid inference approach based on subsampling, and construct confidence sets based on test inversion. As demonstrated in the Monte Carlo study, the procedure is tractable even when the state space is large, and yields tight confidence sets. We also propose a computational algorithm that substantially alleviates the burden of performing inference. Taken together, these are the first positive results on set-identification and computationally tractable valid inference procedures for counterfactual outcomes in structural dynamic multinomial choice models. These are the core contributions of our paper.

Finally, we illustrate the policy usefulness of our approach by revisiting the seminal contribution by Das, Roberts, and Tybout (2007), who perform a horserace between different types of export subsidies (export revenues, fixed cost, and entry cost subsidies). Similar to our firm entry/exit example, here firms decide whether to enter into/exit from exporting under uncertainty. Placing several restrictions on the model primitives (e.g., fixed costs and scrap values are equal to zero), they find that export revenue subsidies generate the highest net returns (in terms of the ratio of export gains over subsidy expenditures), while entry cost subsidies result on the lowest returns. We obtain the identified sets for the net returns of the three subsidy measures under mild restrictions. We show that, although the ranking of Das, Roberts, and Tybout (2007) can be confirmed under weaker restrictions than originally imposed, it does hinge on

the assumption that scrap values do not vary over states. Without this assumption, entry cost subsidies can potentially overperform the other types of subsidies because of the potential correlation (induced by the state variables) between export revenues and scrap values. This correlation can lead high-performing firms (with large export revenues) to enter into exporting frequently, and staying for long periods, while low-performing firms would not do so. In contrast, when scrap values are assumed to not vary over states, the correlation between export revenues and scrap values is eliminated, rendering entry cost subsidies particularly inefficient.

Related Literature. A large body of work studies the identification and estimation of dynamic discrete choice models. Rust (1994) showed that DDC models are not identified nonparametrically, and Magnac and Thesmar (2002) characterized the degree of underidentification. Important advances that followed include (but are not limited to) Heckman and Navarro (2007), Pesendorfer and Schmidt-Dengler (2008), Blevins (2014), Bajari, Chu, Nekipelov, and Park (2016), and Abbring and Daljord (2020). In terms of estimation, Rust (1987) introduced the nested fixed point maximum likelihood estimator in his seminal contribution, and Hotz and Miller (1993) pioneered a computationally convenient two-step estimator that was then further analyzed by a host of important studies (Hotz, Miller, Sanders, and Smith, 1994; Aguirregabiria and Mira, 2002, 2007; Bajari, Benkard, and Levin, 2007; Pakes, Ostrovsky, and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008).² We build on this literature on point-identification and estimation, and extend them to partial identification of model parameters and, more importantly, counterfactuals.

Several important papers have considered partial identification and estimation of structural parameters, namely Bajari, Benkard, and Levin (2007), Norets and Tang (2014), Dickstein and Morales (2018), Morales, Sheu, and Zahler (2019), and Berry and Compiani (2020).³ With the exception of Norets and Tang (2014) (discussed further below), these papers consider classes of models that differ from, and do not necessarily nest, ours. A common issue in this literature concerns the fact that existing inference methods for partially identified models are computationally costly – if not infeasible – when the parameter space is not small (they require repeated inversion of hypothesis testing over the parameter space); thus, prior empirical work has only estimated the most parsimonious specifications. Substantial computational costs may also limit the set of counterfactuals implemented, given that simulations for each parameter value in the identified set are required. In contrast, our approach focuses inference directly on low-dimensional final objects of interest – which are typically nonlinear functions of CCPs and model primitives – thus

²Important early contributions include Miller (1984), Wolpin (1984), and Pakes (1986).

³Bajari, Benkard, and Levin (2007) two-step estimator is the first to allow for partially-identified model parameters. Dickstein and Morales (2018) and Morales, Sheu, and Zahler (2019) pioneered the use of Euler-equation-like estimators for DDC models using moment inequalities, requiring minimal distributional assumptions on the error term. Berry and Compiani (2020) allow for serially correlated unobserved states and propose the use of lagged state variables as instrumental variables for (econometrically) endogenous states, for models with both continuous and discrete actions, and obtain partial identification of structural parameters in a discrete choice setting.

allowing for a large number of model parameters and richer empirical applications. As such, our approach complements, and can be combined with, the previous contributions.

A small but growing literature investigates the identification of counterfactuals in DDC models. The main contributions in this area are by Aguirregabiria (2010), Aguirregabiria and Suzuki (2014), Norets and Tang (2014), Arcidiacono and Miller (2020), and Kalouptsi, Scott, and Souza-Rodrigues (2017, 2021).⁴ We build on Kalouptsi, Scott, and Souza-Rodrigues (2021) (henceforth ‘KSS’), which provides the necessary and sufficient conditions for point identification of a broad class of counterfactuals. We extend KSS by showing how additional model restrictions can help point-identify counterfactual behaviour that would otherwise not be identified; by establishing the exact dimension of the counterfactual CCP identified set; and by characterizing the sharp bounds for low-dimensional objects of interest, along with practical inference procedures. Aside from KSS, the closest paper to ours is Norets and Tang (2014), who partially identify the structural parameters and the (high-dimensional) counterfactual CCPs in binary choice models with an unknown distribution of the idiosyncratic shocks. They propose a Bayesian approach to inference, based on Markov Chain Monte Carlo (MCMC), for two types of counterfactuals – pre-specified additive changes in payoffs and changes to state transitions. Compared to Norets and Tang (2014), we consider *multinomial* choice models and show that partial identification arises even when the econometrician knows the distribution of the idiosyncratic shocks. While we do not allow that distribution to be nonparametric, we allow for flexible distributions that are either known by the researcher – common in empirical work – or that belong to a parametric family with parameters not necessarily known or point-identified.⁵ We also consider a broad class of counterfactual experiments, characterize the properties of the identified sets analytically, and develop tractable procedures to compute bounds in practice (our primary motivation) – to the best of our knowledge, this is the first feasible computational approach for counterfactual bounds in dynamic multinomial choice models.⁶ In addition, our inference procedure – based on subsampling – is guaranteed to be uniformly valid asymptotically, contrary to Bayesian inference, which is known to be only pointwise valid asymptotically (see, e.g., the discussion in Canay and Shaikh, 2017). Finally, we focus on confidence sets for low-dimensional outcomes of interest, which are often what researchers are most interested in. This is not trivial as such objects involve nonlinear functions of model parameters and counterfactual CCPs.

⁴Aguirregabiria and Suzuki (2014), Norets and Tang (2014), and Arcidiacono and Miller (2020) have established the identification of two important categories of counterfactuals in different classes of DDC models: counterfactual behavior is identified when flow payoffs change additively by pre-specified amounts; counterfactual behavior is generally not identified when the state transition process changes. Kalouptsi, Scott, and Souza-Rodrigues (2017) discuss identification of counterfactual best-reply functions and equilibria in dynamic games.

⁵Recently, and independently, Christensen and Connault (2021) proposed a clever way to perform sensitivity analysis of counterfactuals in structural models allowing for the distribution of unobservables to (nonparametrically) span neighborhoods of the researcher’s assumed specification. Their approach, when applied to DDC models, complements ours; combining both is a promising venue for future work.

⁶Norets (2011) shows how to extend Norets and Tang (2014) theoretical results to multinomial choice models, but the corresponding identified sets are infeasible to compute in practice. Our (set of) parametric distributions, in contrast, preserve the smoothness – and feasibility – of our constrained optimizations.

Our inference approach builds on the formulation developed in Kitamura and Stoye (2018), where the implications of economic models are expressed in terms of the minimum value of a quadratic form. The associated quadratic form-based algorithm offers computational advantages when compared to more conventional settings. We consider a test statistic based on the minimum distance to a kinked (i.e., non-regular), *random* (estimated), and possibly *nonconvex* set. We avoid standard convexity conditions on such objects because they are typically incompatible with our model restrictions.⁷ We establish that an appropriate application of subsampling to the quadratic-form-based distance measure yields an asymptotically valid algorithm for inference.⁸

Finally, a recent and increasingly influential line of research emphasizes that (partial) identification of potential effects of policy interventions does not necessarily require identification of all the model parameters. Major contributions outside the class of structural dynamic models include Ichimura and Taber (2000, 2002) and Mogstad, Santos, and Torgovitsky (2018) for selection models; Manski (2007) for static choice models under counterfactual choice sets; Blundell, Browning, and Crawford (2008), Blundell, Kristensen, and Matzkin (2014), Kitamura and Stoye (2019), and Adams (2020) for bounds on counterfactual demand distributions and welfare analysis; Adao, Costinot, and Donaldson (2017) for international trade models; and Bejara (2020) for macroeconomic models. All these approaches, including ours, are consistent with Marschak’s (1953) advocacy of solving well-posed economic problems with minimal assumptions. See Heckman (2000, 2010) for excellent discussions of Marschak’s approach to identification in structural models.

The rest of the paper is organized as follows: Section 2 sets out the dynamic discrete choice framework; Section 3 presents the partial identification results for the model parameters, then illustrates the identified set under alternative restrictions in the context of a firm entry/exit problem; Section 4 contains our main results regarding the set-identification of counterfactuals; Section 5 discusses estimation and inference; Section 6 presents the empirical application involving export supply and subsidies; and Section 7 concludes.⁹

⁷Note that Kitamura and Stoye (2018) deal with the case where a random vector is projected on a non-smooth but fixed object with some desirable geometric features. They then show that a bootstrap procedure combined with what they call the tightening technique leads to a computationally efficient algorithm with asymptotic uniform validity.

⁸Asymptotic validity of subsampling in nonregular models with more conventional settings, such as standard moment inequality models, have been shown in the literature: see Romano and Shaikh (2008) and Romano and Shaikh (2012).

⁹The Appendix and the Online Appendix complement the main paper. The Appendix contains (a) all proofs of the propositions and theorems presented in the main text; (b) the analytical gradient of the counterfactual object of interest when it involves long-run average effects; (c) our proposed stochastic search approach to calculate the lower and upper bounds of the identified set of relevant outcomes, without analytic gradients; (d) the computational algorithm for inference based on subsampling; and (e) our Monte Carlo study. The Online Appendix presents (f) several useful examples of commonly employed restrictions in applied work (using our notation); (g) detailed information about our running example (the firm entry/exit problem); and (h) our replication of Das, Roberts, and Tybout (2007). The Online Appendix is available on the authors’ webpages.

2 Dynamic Model

Time is discrete and the time horizon is infinite. Every period t , agent i observes the state s_{it} and chooses an action a_{it} from the finite set $\mathcal{A} = \{0, \dots, A\}$ to maximize the expected discounted payoff,

$$\mathbb{E} \left(\sum_{\tau=0}^{\infty} \beta^{\tau} \bar{\pi}(a_{it+\tau}, s_{it+\tau}) \mid a_{it}, s_{it} \right),$$

where $\bar{\pi}(\cdot)$ is the per period utility or payoff function, and $\beta \in [0, 1)$ is the discount factor. The agent's state s_{it} follows a controlled Markov process. We follow the literature and assume that s_{it} is split into two components, $s_{it} = (x_{it}, \varepsilon_{it})$, where x_{it} is observed by the econometrician and ε_{it} is not. We assume $x_{it} \in \mathcal{X} = \{1, \dots, X\}$, $X < \infty$; while $\varepsilon_{it} = (\varepsilon_{0it}, \dots, \varepsilon_{Ait})$ is i.i.d. across agents and time, and has joint distribution G that is absolutely continuous with respect to the Lebesgue measure and has full support on \mathbb{R}^{A+1} .¹⁰

The transition distribution function for $(x_{it}, \varepsilon_{it})$ factors as

$$F(x_{it+1}, \varepsilon_{it+1} \mid a_{it}, x_{it}, \varepsilon_{it}) = F(x_{it+1} \mid a_{it}, x_{it}) G(\varepsilon_{it+1}),$$

and the utility or payoff function is additively separable in the unobservables,

$$\bar{\pi}(a, x_{it}, \varepsilon_{it}) = \pi_a(x_{it}) + \varepsilon_{ait},$$

where $\pi_a(x)$ is a bounded function. Let $V(x_{it}, \varepsilon_{it})$ be the expected discounted stream of payoffs under optimal behavior by the agent. By Bellman's principle of optimality,

$$V(x_{it}, \varepsilon_{it}) = \max_{a \in \mathcal{A}} \left\{ \pi_a(x_{it}) + \varepsilon_{ait} + \beta \mathbb{E} [V(x_{it+1}, \varepsilon_{it+1}) \mid a, x_{it}] \right\}.$$

Following the literature, we define the *ex ante value function*, $V(x_{it}) \equiv \int V(x_{it}, \varepsilon_{it}) dG(\varepsilon_{it})$, and the *conditional value function*:

$$v_a(x_{it}) \equiv \pi_a(x_{it}) + \beta \mathbb{E} [V(x_{it+1}) \mid a, x_{it}].$$

The ex ante value function takes the expectation of the value function with respect to ε_{it} . The conditional value function is the sum of the agent's current payoff, net of the idiosyncratic shocks ε_{it} , and the expected lifetime payoff given the choice of action a this period and optimal choices from next period onwards.

The *conditional choice probability* (CCP) function is given by

$$p_a(x_{it}) = \int \mathbf{1} \{v_a(x_{it}) + \varepsilon_{ait} \geq v_j(x_{it}) + \varepsilon_{jit}, \text{ for all } j \in \mathcal{A}\} dG(\varepsilon_{it}),$$

¹⁰Our results cover static discrete choice models, and can be extended to dynamic models with continuous states, nonstationarity, and that are finite-horizon.

where $1\{\cdot\}$ is the indicator function. We define the $(A+1) \times 1$ vector of conditional choice probabilities $p(x) = (p_0(x), \dots, p_A(x))'$, and the corresponding $(A+1)X \times 1$ vector $p = (p'(1), \dots, p'(X))'$, where $'$ denotes transpose.

It is useful to note that for any (a, x) there exists a real-valued function $\psi_a(\cdot)$ derived only from G such that

$$V(x) = v_a(x) + \psi_a(p(x)). \quad (1)$$

Equation (1) states that the ex ante value function V equals the conditional value function of any action a , v_a , plus a correction term, ψ_a , because choosing action a today is not necessarily optimal. When ε_{it} follows the type I extreme value distribution, $\psi_a(p(x)) = \kappa - \ln p_a(x)$, where κ is the Euler constant. Chiong, Galichon, and Shum (2016) propose a computationally tractable approach, based on linear-programming, that can calculate ψ_a for any given distribution G . (See also Dearing 2019.)¹¹

As we make extensive use of matrix notation below, we define the vectors $\pi_a, v_a, V, \psi_a \in \mathbb{R}^X$, which stack $\pi_a(x), v_a(x), V(x)$, and $\psi_a(p(x))$, for all $x \in \mathcal{X}$. We often use the notation $\psi_a(p)$ to emphasize the dependence of ψ_a on the choice probabilities p . We also define F_a as the transition matrix with (m, n) element equal to $\Pr(x_{it+1} = x_n | x_{it} = x_m, a)$. The payoff vector $\pi \in \mathbb{R}^{(A+1)X}$ stacks π_a for all $a \in \mathcal{A}$, and, similarly, F stacks (a vectorized version) of F_a for all $a \in \mathcal{A}$.

3 Model Restrictions and Identification

In this section, we characterize the identified set of the model parameters, allowing for additional model restrictions that the researcher may be willing to impose, and we illustrate the sets in the context of a firm entry/exit problem.

The primitives of the model are $(\mathcal{A}, \mathcal{X}, \beta, G, F, \pi)$ and generate the agent's optimal policy $\{p_a, a \in \mathcal{A}\}$. Typically, the researcher has access to panel data on agents' actions and states, $\{a_{it}, x_{it} : i = 1, \dots, N; t = 1, \dots, T\}$. Under some standard regularity conditions, the researcher can identify and estimate the agents' choice probabilities $p_a(x)$, as well as the transition distribution function F , directly from the data. We therefore take p and F as known for the identification arguments. We also follow the literature and assume for now that the econometrician knows the discount factor β and the distribution of the idiosyncratic shocks G ; we relax these assumptions in Section 4.2.¹² The main objective in this section is to (partially) identify the payoff function π .

The model is identified if there is a unique payoff that can be inferred from the observed choice probabilities and state transitions. Intuitively, π has $(A+1)X$ parameters, and there are only AX observed

¹¹Equation (1) is shown in Arcidiacono and Miller (2011, Lemma 1). It makes use of the Hotz-Miller inversion (Hotz and Miller, 1993), which, in turn, establishes that the difference of conditional value functions is a known function of the CCPs: $v_a(x) - v_j(x) = \varphi_{aj}(p(x))$, where $\varphi_{aj}(\cdot)$ is again derived only from G . When ε_{it} follows the type I extreme value distribution, $\varphi_{aj}(p(x)) = \log p_a(x) - \log p_j(x)$.

¹²Magnac and Thesmar (2002) and Abbring and Daljord (2020) discuss identification of β ; Blevins (2014), Chen (2017), and Buchholz, Shum, and Xu (2020) consider identification of G in binary choice models under different model assumptions.

CCPs; thus there are X free payoff parameters and X restrictions will need to be imposed to point-identify π (Rust, 1994; Magnac and Thesmar, 2002). We follow KSS to represent the underidentification problem as follows: for all $a \neq J$, where $J \in \mathcal{A}$ is some reference action, π_a can be represented as an affine transformation of π_J :¹³

$$\pi_a = M_a \pi_J + b_a(p), \quad (2)$$

where

$$M_a = (I - \beta F_a)(I - \beta F_J)^{-1}, \quad (3)$$

$$b_a(p) = M_a \psi_J(p) - \psi_a(p), \quad (4)$$

and I is a (conformable) identity matrix. In the logit model, $b_a(p) = \ln p_a - M_a \ln p_J$, where $\ln p_a$ is the $X \times 1$ vector with elements $\ln p_a(x)$. To simplify notation, we omit the dependence of both M_a and $b_a(p)$ on the transition probabilities F .

Example. Consider an example where the agent faces a binary choice $a \in \{0, 1\}$ and ε_{it} follows the type I extreme value distribution. Intuitively, we want to obtain the payoffs π_0, π_1 ; however, only p_1 is informative (since $p_0 = 1 - p_1$). Equation (2) becomes

$$\pi_1 = M_1 \pi_0 + \underbrace{(\log p_1 - M_1 \log p_0)}_{=b_1(p)}, \quad (5)$$

where $M_1 = (I - \beta F_1)(I - \beta F_0)^{-1}$.

It is instructive to compare this to a static model: in that case, since $\beta = 0$ (or, alternatively, $F_0 = F_1$ since choices do not affect future states), we have $M_1 = I$ and (2) becomes,

$$\pi_1 = \pi_0 + \underbrace{(\log p_1 - \log p_0)}_{=b_1(p)}. \quad (6)$$

In words, we obtain the well-known result that only the difference of payoffs $\pi_1 - \pi_0$ can be identified from choice data; under the logit assumption, it equals the log odds ratio of the choice probabilities. A typical solution here is to set $\pi_0 = 0$ and obtain $\pi_1 = \log p_1 - \log p_0$.

The difference between the dynamic model (5) and the static model (6) is the matrix M_1 , which depends on the discount factor and the state transitions. This matrix distorts both the payoff difference and the log odds ratio of the CCPs (compared to the static version) to capture the impact of agents'

¹³To see why, fix the vector $\pi_J \in \mathbb{R}^X$. Then,

$$\pi_a = v_a - \beta F_a V = V - \psi_a - \beta F_a V = (I - \beta F_a) V - \psi_a,$$

where for $a = J$, we have $V = (I - \beta F_J)^{-1}(\pi_J + \psi_J)$. After substituting for V , we obtain the result. As an aside, note that $(I - \beta F_J)$ is invertible because F_J is a stochastic matrix and hence the largest eigenvalue is smaller than or equal to one.

expectations about the future (essentially the value functions).

We rely heavily on equation (2). Given the data at hand, one can compute both the $X \times X$ matrix M_a and the $X \times 1$ vector b_a directly for each action $a \neq J$. The payoffs π_a , $a \neq J$, are not identified because the free parameter π_J is unknown. Equation (2) therefore explicitly lays out how we might estimate the payoff function if we are willing to fix the payoffs of one action at all states *a priori* (e.g. $\pi_J = 0$). However, this is not the only way to obtain identification: we simply need to add X extra restrictions. Other common possibilities involve reducing the number of payoff function parameters to be estimated using parametric assumptions and/or exclusions restrictions, as discussed below.

It will be useful to represent (2) for all actions $a \neq J$ at once using two different compact notations. First,

$$\pi_{-J} = M_{-J}\pi_J + b_{-J}, \quad (7)$$

where π_{-J} stacks π_a for all $a \neq J$, and the matrix M_{-J} and vector b_{-J} are defined similarly. Further, define $\mathbf{M} = [I, -M_{-J}]$, and arrange π in the following way: $\pi = [\pi'_{-J}, \pi'_J]'$. Then (7) becomes

$$\mathbf{M}\pi = b_{-J}. \quad (8)$$

Note that identification of π fails because \mathbf{M} is rank-deficient; indeed, \mathbf{M} is an $AX \times (A + 1)X$ matrix, and so $\text{rank}(\mathbf{M}) = AX < (A + 1)X$.

Importantly, equation (8) summarizes *all* assumptions imposed by the basic dynamic framework: any $\pi \in \mathbb{R}^{(A+1)X}$ satisfying (8) is compatible with the data.¹⁴

Model Restrictions. We consider two types of model restrictions. The first is a set of $d \leq X$ linearly independent equalities,

$$R^{eq}\pi = r^{eq}, \quad (9)$$

with $R^{eq} \in \mathbb{R}^{d \times (A+1)X}$, or in block-form, $R^{eq} = [R_{-J}^{eq}, R_J^{eq}]$, where R_{-J}^{eq} defines how π_{-J} enters into the constraints and, similarly, R_J^{eq} for π_J . This formulation is general enough to incorporate several assumptions used in practice. Examples include exclusion restrictions (setting some elements of π equal to each other), prespecifying some π_J (set $R_J^{eq} = I$, $R_{-J}^{eq} = 0$ and r^{eq} accordingly), and parametric assumptions such as $\pi_a(x) = z_a(x)\gamma$, where z_a is some known function of actions and states, and $\gamma \in \Gamma \subset \mathbb{R}^L$ is a parameter vector in the parameter space Γ , with dimension L usually much smaller than $(A + 1)X$.

¹⁴This model imposes a scale normalization. In general, the payoff function is given by $\bar{\pi}(a, x_{it}, \varepsilon_{it}) = \pi_a(x_{it}) + \sigma\varepsilon_{ait}$, where $\sigma > 0$ is a scale parameter. This means equation (8) is given by $\mathbf{M}(\pi/\sigma) = b_{-J}$. As usual in discrete choice models, when we set $\sigma = 1$ (as we do here), the scale of the payoff is measured relative to the standard deviation of one of the components of ε_{it} .

The second set of restrictions are m linear inequalities:

$$R^{iq}\pi \leq r^{iq}, \quad (10)$$

with $R^{iq} \in \mathbb{R}^{m \times (A+1)X}$, or in block-form, $R^{iq} = [R_{-j}^{iq}, R_j^{iq}]$. The inequalities (10) can incorporate shape restrictions, such as monotonicity, concavity, and supermodularity. In Online Appendix F, we explicitly lay out how several examples of assumptions used in applied work can be expressed as (9) or (10).

We assume the restrictions (9) and (10) are not redundant. Equations (8), (9), and (10) summarize therefore all model restrictions.

Model Identification. The sharp identified set for the payoff function is characterized by all payoffs satisfying all model restrictions:¹⁵

$$\Pi^I = \left\{ \pi \in \mathbb{R}^{(A+1)X} : \mathbf{M}\pi = b_{-J}, R^{eq}\pi = r^{eq}, R^{iq}\pi \leq r^{iq} \right\}. \quad (11)$$

The identified set Π^I is a convex polyhedron of dimension $X - d$, where $0 \leq d \leq X$, since it is the intersection of finitely many closed halfspaces. (Note that Π^I can be characterized in practice by linear programming or convex programming methods.) In the absence of inequalities (10), the identified set becomes a linear manifold with dimension $X - d$; and it collapses to a singleton (i.e., π is point-identified) when the matrix $[\mathbf{M}', R^{eq}]'$ is full rank (Magnac and Thesmar, 2002; Pesendorfer and Schmidt-Dengler, 2008).¹⁶

3.1 Example: Firm Entry/Exit Model

To fix ideas, we illustrate the payoff identified set in the context of a simple firm entry/exit problem. Suppose firm i faces the choice set $\mathcal{A} = \{\text{out}, \text{in}\} = \{0, 1\}$. Decompose the state space into $x_{it} = (k_{it}, w_{it})$, where $k_{it} \in \mathcal{K} = \{0, 1\}$ is the lagged decision a_{it-1} , and $w_{it} \in \mathcal{W} = \{1, \dots, W\}$ is an exogenous profit shifter (e.g. market size). Assume for convenience (unless otherwise stated) that w_{it} can take two values, low and high: $\mathcal{W} = \{w^l, w^h\}$, with $w^l < w^h$. The size of the state space is therefore $X = KW = 4$, where K and W are the number of values that k and w can take. Transition probabilities are decomposed as $F(k_{it+1}, w_{it+1} | a_{it}, k_{it}, w_{it}) = F(k_{it+1} | a_{it}, k_{it})F(w_{it+1} | w_{it})$.

Let $\pi_a(k)$ denote the $W \times 1$ payoff vector the firm obtains when it chooses action a given k and w , so

¹⁵ A sharp identified set is the smallest set of parameter values that can generate the data.

¹⁶ In the presence of unobserved heterogeneity, equations (8)–(10) hold for each unobserved type. This implies that, after type-specific choice probabilities and transition functions of finitely many unobserved types are identified (e.g., following the strategies proposed by Kasahara and Shimotsu (2009) or Hu and Shum (2012)), identified sets given by (11) hold, and can be calculated, for each type.

that $\pi_a = [\pi'_a(0), \pi'_a(1)]'$. We impose the following structure on π :

$$\pi_0 = \begin{bmatrix} oo \\ s \end{bmatrix}, \pi_1 = \begin{bmatrix} vp - fc - ec \\ vp - fc \end{bmatrix}. \quad (12)$$

The payoff the firm obtains when it was out of the market in the previous period and stays out in the current period is the vector $\pi_0(0) = oo$ (the value of the outside option); and the payoff when the firm was active and decides to exit is given by the vector of scrap values, $\pi_0(1) = s$. Note that both the outside option and the scrap values can vary with the exogenous state w . The vectors vp , fc , and ec are the variable profits, the fixed costs, and the entry costs, respectively (all of which can vary with w as well). The vector $\pi_1(0) = vp - fc - ec$ measures the profits the firm gets when it enters the market, and $\pi_1(1) = vp - fc$ are the profits when it stays.

In this example, both π_0 and π_1 are 4×1 vectors (and so π has $2X = 8$ elements). To point-identify π we need $X = 4$ restrictions. Typically, researchers identify an entry model by setting $oo = 0$ (for 2 restrictions); and further setting either $s = 0$ or assuming $vp - fc$ is known (e.g., by assuming variable profits vp can be recovered “offline,” using price and quantity data, and setting $fc = 0$). When $oo = s = 0$, then $\pi_0 = 0$, and point identification of π follows directly from (2); it is essentially a restriction on a reference action. When instead $\pi_0(0) = oo = 0$ and $vp - fc$ is known, we identify the remaining elements of π by combining (2) and (9).

Assuming the outside option equals the scrap value or the fixed costs (and all are equal to zero) may be difficult to justify in practice, as cost or scrap value data are extremely rare (Kalouptsi, 2014). When the researcher is not willing to impose such restrictions, π is not point-identified. Yet, the payoff function can be set-identified under weaker conditions. Consider, for instance, the following set of assumptions:

1. $oo = 0$, $fc \geq 0$, $ec \geq 0$, and vp is known.
2. $\pi_1(1, w^h) \geq \pi_1(1, w^l)$, and $vp - fc \leq ec \leq \frac{\mathbb{E}[vp - fc]}{1 - \beta}$, where the expectation is taken over the ergodic distribution of the state variables.
3. s does not depend on w .

Restriction 1 assumes that the outside option is zero (as usual); fixed costs and entry costs are both positive; and variable profits are known (estimated “offline”). This set of restrictions imposes $d = W = 2$ equality and $m = 4$ inequality constraints. From (11), it is clear that the identified set Π^I is a two-dimensional set ($X - d = 2$) in the eight-dimensional space.

Restriction 2 imposes $m = 5$ inequality constraints: profits are increasing in w when the firm is in the market (a monotonicity assumption); entry costs are greater than variable profits minus fixed costs (implying that entry is always costly in the first period of entry); and ec is smaller than the expected

present value of future profits when the firm stays forever in the market (meaning that, on average, it eventually pays off to enter).

Restriction 3 assumes an exclusion restriction: scrap values are state-invariant. This corresponds to $d = W - 1 = 1$ equality restriction. Note that, by combining Restrictions 1 and 3, we obtain $d = 3$ linear equalities, which makes the identified set Π^I one dimensional. In Online Appendix G, we provide explicit characterizations for this example.

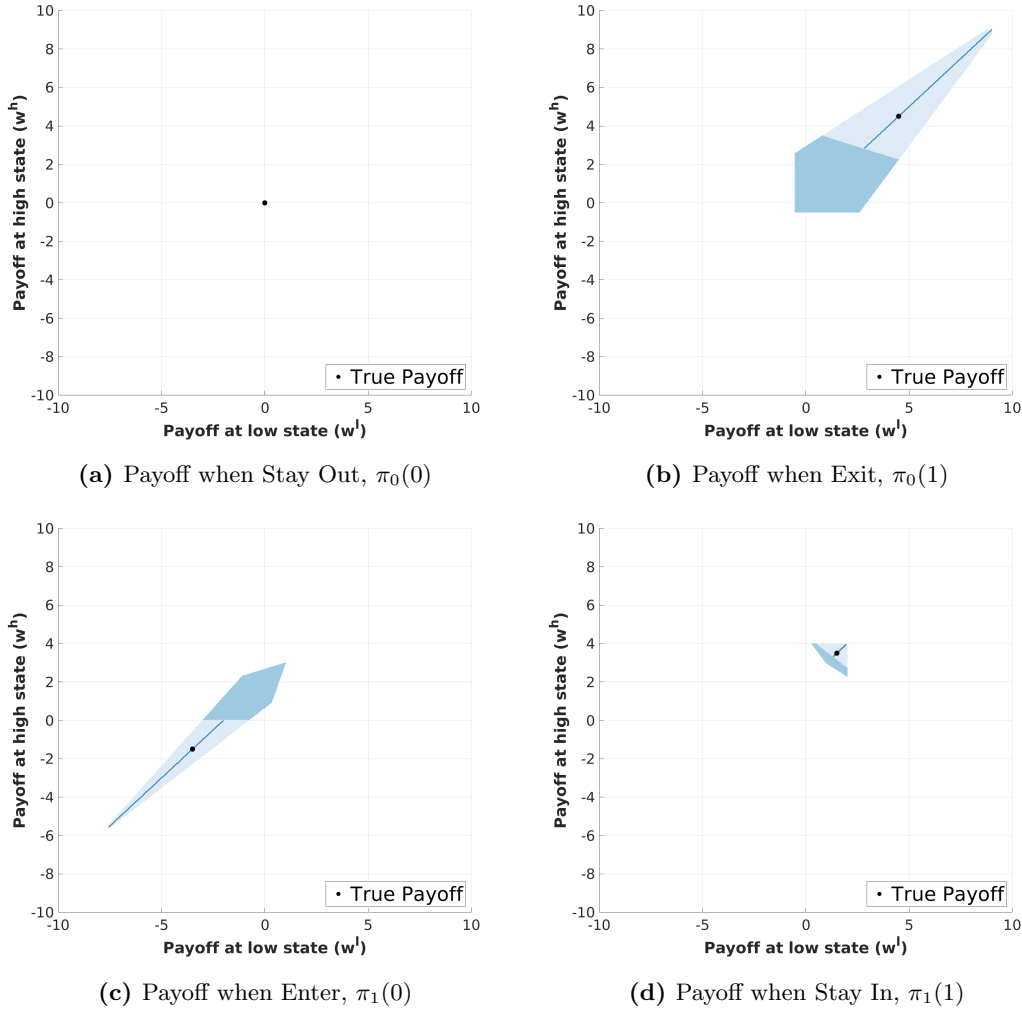


Figure 1: Firm Entry/Exit Model: Payoff Identified Set Π^I under Restrictions 1, 2, and 3. The larger polyhedron (including the dark blue areas) correspond to Π^I under Restriction 1. The light blue areas correspond to Π^I under Restrictions 1 and 2. The identified set Π^I under Restrictions 1–3 is represented by the blue lines within the light blue polyhedron. The true π is represented by the black dots.

Figure 1 presents the identified set for payoffs, Π^I , for a particular parameter configuration.¹⁷ The larger polyhedron corresponds to Π^I under Restriction 1. The identified set is informative despite the

¹⁷We assume scrap values, entry and fixed costs do not depend on w and take the following values: $s = 4.5$, $ec = 5$, and $fc = 0.5$. We also impose $vp(w^l) = 2$ and $vp(w^h) = 4$, so that $\pi_0 = (0, 0, 4.5, 4.5)'$ and $\pi_1 = (-3.5, -1.5, 1.5, 3.5)'$. The discount factor is $\beta = 0.9$, the transition process for w is $Pr(w_{t+1} = w^l | w_t = w^l) = Pr(w_{t+1} = w^h | w_t = w^h) = 0.75$, and the idiosyncratic shocks ε_{it} follow a type 1 extreme value distribution (the scale parameter is set at $\sigma = 1$). Under these assumptions, $\frac{E[vp - fc]}{1 - \beta} = 25$.

fact that the assumptions imposed are not overly restrictive. In Online Appendix G, we explicitly derive the shape of Π^I using equation (8) and the corresponding restrictions. For a brief intuition of how the linear (in)equalities interact to produce Figure 1, consider the set corresponding to scrap values (panel (b)). In this model, equation (8) alone implies that the difference between scrap values and entry costs, $s - ec$, is point-identified (see Online Appendix G). As a consequence, the inequality $ec \geq 0$ implies a lower bound on scrap values (for each state w), shifting the origin. Similarly, equation (8) implies that the sum of scrap values and the present value of fixed costs, $s + (I - \beta F^w)^{-1} fc$, is point-identified. The inequality $fc \geq 0$ entails an upper bound on scrap values, eliminating from the identified set all values for s above the intersecting lines shown in the figure. The increasing lines reflect the fact that equation (8) relates s and the *present value* of fc , so that $fc \geq 0$ leads to restrictions on scrap values *across* states.

Restrictions 1 and 2 together lead to substantial identifying power: Π^I now corresponds to the smaller polyhedron (in light blue), which is substantially smaller. Assuming that entry is costly in the first period of entry, $vp - fc \leq ec$, is the main restriction responsible for the reduction in the identified set. This assumption results in another lower bound on s (see panel (b)), but differently from $ec \geq 0$, it involves restrictions on fc and so imposes restrictions on s across states; the other assumptions in Restriction 2 are not as informative in this example; see Online Appendix G. Interestingly, the payoff function with scrap values that are equal to zero does not belong to Π^I under these two sets of restrictions. As mentioned previously, setting scrap values to zero is a common way to point-identify π , but, given that $s = 0$ is at odds with Restrictions 1 and 2, such assumption would be rejected by the data.

Finally, Restriction 3 (exclusion restriction on scrap values) also has substantial identifying power as it reduces the dimension of the identified set to one. In the figures, the identified set under Restrictions 1–3 is represented by the blue lines within the light blue polyhedron.

4 Counterfactuals and Outcomes of Interest

The applied literature has implemented several types of counterfactuals that may change one or several of the model’s primitives $(\mathcal{A}, \mathcal{X}, \beta, F, G, \pi)$. For instance, a counterfactual may change the action and state spaces (e.g. Gilleskie (1998) restricts access to medical care; Crawford and Shum (2005) do not allow patients to switch medications; Keane and Wolpin (2010) eliminate a welfare program; Keane and Merlo (2010) eliminate job options for politicians; Rosenzweig and Wolpin (1993) add an insurance option for farmers). Some counterfactuals may transform the state transitions (e.g. Collard-Wexler (2013) explores the impact of demand volatility in the ready-mix concrete industry; Hendel and Nevo (2006) study consumers’ long-run responsiveness to prices using supermarket scanner data; Kalouptsi (2014) explores the impact of time to build on industry fluctuations). Other counterfactuals change payoffs through subsidies or taxes (e.g. Keane and Wolpin (1997) consider hypothetical college tuition subsidies; Schiraldi (2011) and Wei and Li (2014), automobile scrap subsidies; Duflo, Hanna, and Ryan (2012),

bonus incentives for teachers; Das, Roberts, and Tybout (2007), export subsidies; Lin (2015) and Varela (2018), entry subsidies). Changes in payoffs may also involve a change in the agent’s “type” (e.g. Keane and Wolpin (2010) replace the primitives of minorities by those of white women; Eckstein and Lifshitz (2011) substitute the preference/costs parameters of one cohort by those of another; Ryan (2012) replaces firm entry costs post an environmental policy by those before; Dunne, Klimek, Roberts, and Xu (2013) substitute entry costs in some areas by those in others). Finally, a counterfactual may also change the discount factor (e.g., Conlon (2012) studies the evolution of the LCD TV industry when consumers become myopic).

A counterfactual is defined by the tuple $\{\tilde{\mathcal{A}}, \tilde{\mathcal{X}}, \tilde{\beta}, \tilde{G}, h^s, h\}$. The sets $\tilde{\mathcal{A}} = \{0, \dots, \tilde{A}\}$ and $\tilde{\mathcal{X}} = \{1, \dots, \tilde{X}\}$ denote the new set of actions and states respectively. The new discount factor is $\tilde{\beta}$, and the new distribution of the idiosyncratic shocks is \tilde{G} . The function $h^s : \mathbb{R}^{A \times X^2} \rightarrow \mathbb{R}^{\tilde{A} \times \tilde{X}^2}$ transforms the transition probability F into \tilde{F} . Finally, the function $h : \mathbb{R}^{AX} \rightarrow \mathbb{R}^{\tilde{A}\tilde{X}}$ transforms the payoff function π into the counterfactual payoff $\tilde{\pi}$, so that $\tilde{\pi} = h(\pi)$. Here, we restrict transformations on payoffs to affine changes

$$\tilde{\pi} = \mathcal{H}\pi + g, \tag{13}$$

where the matrix \mathcal{H} and the vector g are specified by the econometrician. I.e., the payoff $\tilde{\pi}_a(x)$ at an action-state pair (a, x) is obtained as the sum of a scalar $g_a(x)$ and a linear combination of all baseline payoffs. In practice, most applied papers consider counterfactuals that affect one primitive.

The counterfactual $\{\tilde{\mathcal{A}}, \tilde{\mathcal{X}}, \tilde{\beta}, \tilde{G}, h^s, h\}$ generates a new set of model primitives $(\tilde{\mathcal{A}}, \tilde{\mathcal{X}}, \tilde{\beta}, \tilde{G}, \tilde{F}, \tilde{\pi})$. The new set of primitives in turn leads to a new optimal behavior, denoted by \tilde{p} (the counterfactual CCP), and a new lifetime utility, denoted by \tilde{V} (the counterfactual welfare).

As the state space \mathcal{X} can be large in practice (making both \tilde{p} and \tilde{V} high-dimensional vectors), researchers are often interested in low-dimensional objects, such as the average effects of policy interventions. For instance, in the firm entry/exit application, one may be interested in predicting the average effects of an entry subsidy on: (i) how often the firm stays in the market; (ii) prices; (iii) consumer surplus; (iv) the value of the firm; and/or (v) total government expenditures, among others. Denote the low-dimensional counterfactual outcome of interest by $\theta \in \Theta \subset \mathbb{R}^n$, where Θ is a compact set (the parameter space for θ), and n is much smaller than the size of the state space X (i.e., $n \ll X$). Then, we have

$$\theta = \phi(\tilde{p}, \pi; p), \tag{14}$$

where ϕ depends on the counterfactual CCP \tilde{p} , on the baseline payoff function π , on the data (p, F) , though here and thereafter we omit its dependence on F for notational convenience, and it implicitly incorporates other quantities that may be necessary to calculate θ , such as $\tilde{\mathcal{A}}$ or \tilde{F} . For instance, take an outcome variable of interest that depends on actions and states, $Y_a(x, \varepsilon)$ (e.g., consumer surplus, or

the probability of entry), with a corresponding counterfactual given by $\tilde{Y}_a(x, \varepsilon)$. The average treatment effect of the policy intervention on Y is then $\theta = \mathbb{E}[\tilde{Y}_a(x, \varepsilon)] - \mathbb{E}[Y_a(x, \varepsilon)]$, where $\mathbb{E}[\tilde{Y}_a(x, \varepsilon)]$ integrates over the distribution of actions and states in the counterfactual scenario, while $\mathbb{E}[Y_a(x, \varepsilon)]$ integrates over the factual distribution. One may consider the long-run distribution, or may condition on an initial state and estimate short-run effects.

4.1 Identification of Counterfactual Behavior

We now investigate the identified set for the counterfactual CCP \tilde{p} . To do so, we leverage the counterfactual counterpart to (2) for any action $a \in \tilde{\mathcal{A}}$, with $a \neq J$. I.e.,

$$\tilde{\pi}_a = \tilde{M}_a \tilde{\pi}_J + \tilde{b}_a(\tilde{p}), \quad (15)$$

where

$$\begin{aligned} \tilde{M}_a &= (I - \tilde{\beta}\tilde{F}_a)(I - \tilde{\beta}\tilde{F}_J)^{-1}, \\ \tilde{b}_a(\tilde{p}) &= \tilde{M}_a \tilde{\psi}_J(\tilde{p}) - \tilde{\psi}_a(\tilde{p}), \end{aligned}$$

the functions $\tilde{\psi}_J$ and $\tilde{\psi}_a$ depend on the new distribution \tilde{G} , and, without loss of generality, the reference action J belongs to both \mathcal{A} and $\tilde{\mathcal{A}}$. As before, we omit the dependence of both \tilde{M}_a and \tilde{b}_a on the transition probabilities \tilde{F} to simplify notation.

By stacking equation (15) for all actions and rearranging it (as done previously for the baseline case), we obtain $\tilde{\mathbf{M}}\tilde{\pi} = \tilde{b}_{-J}$, where $\tilde{\mathbf{M}} = [I, -\tilde{M}_{-J}]$, I is an identity matrix, and \tilde{M}_{-J} and \tilde{b}_{-J} stack \tilde{M}_a and \tilde{b}_a , for all $a \neq J$, respectively. Next, using the fact that $\tilde{\pi} = \mathcal{H}\pi + g$, we get

$$(\tilde{\mathbf{M}}\mathcal{H})\pi = \tilde{b}_{-J}(\tilde{p}) - \tilde{\mathbf{M}}g. \quad (16)$$

Equation (16) characterizes counterfactual behavior, relating \tilde{p} and model parameters directly, with no continuation values involved.¹⁸ Importantly, the function \tilde{b}_{-J} is continuously differentiable with an everywhere invertible Jacobian (see Lemma 1 in KSS).

Our first proposition follows (all proofs are in the Appendix).

¹⁸First note that the CCP vector p generated by the model primitives is the unique vector that satisfies (8): since the Bellman is a contraction mapping, V is unique; from the definition of the conditional value function, we conclude that so are v_a and thus so is p (see the argument presented in footnote 13, which leads to equation (8)). The same reasoning applies to \tilde{p} in (16).

Proposition 1. *The sharp identified set for the counterfactual CCP \tilde{p} is*

$$\tilde{\mathbf{P}}^I = \left\{ \begin{array}{l} \tilde{p} \in \tilde{\mathbf{P}} : \exists \pi \in \mathbb{R}^{(A+1)X} \text{ such that} \\ \mathbf{M}\pi = b_{-J}(p), \\ R^{eq}\pi = r^{eq}, \quad R^{iq}\pi \leq r^{iq}, \\ (\tilde{\mathbf{M}}\mathcal{H})\pi = \tilde{b}_{-J}(\tilde{p}) - \tilde{\mathbf{M}}g \end{array} \right\}, \quad (17)$$

where $\tilde{\mathbf{P}}$ is the simplex of conditional choice probabilities. The set $\tilde{\mathbf{P}}^I$ has the following properties:

- (i) It is empty if and only if Π^I is empty.
- (ii) It is a connected manifold with boundary, and dimension in the interior given by the rank of the matrix CP_Q , where

$$C = \tilde{\mathbf{M}}\mathcal{H} \begin{bmatrix} M_{-J} \\ I \end{bmatrix}, \quad (18)$$

and P_Q is a known matrix that depends on the model restrictions (defined in the proof; see equation (A2) in the Appendix). Furthermore, $\text{rank}(CP_Q) \leq X - d$.

- (iii) It is compact when Π^I is bounded.
- (iv) In the absence of equality restrictions (9), the dimension of $\tilde{\mathbf{P}}^I$ is given by the rank of C .

In words, a vector \tilde{p} lying in the conditional probability simplex $\tilde{\mathbf{P}}$ belongs to the identified set $\tilde{\mathbf{P}}^I$ if there exists a payoff π that is compatible with the data (i.e., $\mathbf{M}\pi = b_{-J}$), satisfies the additional model restrictions (i.e., $R^{eq}\pi = r^{eq}$ and $R^{iq}\pi \leq r^{iq}$), and can generate \tilde{p} in the counterfactual scenario (i.e., $(\tilde{\mathbf{M}}\mathcal{H})\pi = \tilde{b}_{-J}(\tilde{p}) - \tilde{\mathbf{M}}g$).¹⁹

Intuitively, equation (16) implicitly defines \tilde{p} as a continuously differentiable function of π . The sharp identified set $\tilde{\mathbf{P}}^I$ is therefore the image of Π^I under this function. It is clear that $\tilde{\mathbf{P}}^I$ is empty whenever Π^I is empty (i.e., whenever the model is rejected in the data); $\tilde{\mathbf{P}}^I$ is connected because Π^I is convex; and $\tilde{\mathbf{P}}^I$ is compact when Π^I is bounded (recall that Π^I is closed). An implication of the connectedness of the identified set is that a non-empty $\tilde{\mathbf{P}}^I$ is either a singleton (in which case \tilde{p} is point-identified) or a continuum.

Proposition 1 also states that $\tilde{\mathbf{P}}^I$ is a manifold whose dimension is given by the rank of the matrix CP_Q , which is smaller than or equal to $X - d$. This is intuitive: since $\tilde{\mathbf{P}}^I$ is the image set of a function defined on a $(X - d)$ -dimensional polyhedron, \tilde{p} is specified by at most $X - d$ degrees of freedom rather than by $\tilde{X}\tilde{A}$: once $X - d$ elements are specified, the remaining are found from (16). So, whenever $X - d < \tilde{X}\tilde{A}$, the dimension of $\tilde{\mathbf{P}}^I$ is strictly smaller than the dimension of the conditional probability simplex $\tilde{\mathbf{P}}$, which implies that the (Lebesgue) measure of $\tilde{\mathbf{P}}^I$ on $\tilde{\mathbf{P}}$ is zero. In other words, the identified

¹⁹Counterfactuals involving nonlinear transformations on π change the identified set $\tilde{\mathbf{P}}^I$ defined in (17) by replacing equation (16) by $\tilde{\mathbf{M}}h(\pi) = \tilde{b}_{-J}(\tilde{p})$. We ignore such counterfactuals because they are uncommon in empirical work (such counterfactuals are considered in KSS); however, extensions to nonlinear transformations are straightforward.

set $\tilde{\mathbf{P}}^I$ is informative.

In addition, the rank of CP_Q can be strictly smaller than $X - d$. The exact value depends on (i) the counterfactual transformation (which affects the matrix C , through $\tilde{\mathbf{M}}$ and \mathcal{H} , defined by the econometrician), (ii) the model restrictions (which affect P_Q , through the linear restrictions R^{eq}), and (iii) the data (specifically, the state transitions F , which are part of C , through the matrices M_{-J} and possibly $\tilde{\mathbf{M}}$ through $\tilde{F} = h^s(F)$, and are part of P_Q). The interaction of these factors can reduce the dimension of the identified set further beyond $X - d$. Of note, once the econometrician establishes the counterfactual of interest and the model restrictions, the rank of CP_Q can be verified directly from the data.

When $\text{rank}(CP_Q) = 0$, the identified set $\tilde{\mathbf{P}}^I$ collapses to a singleton – i.e., \tilde{p} is point-identified. This means that all points $\pi \in \Pi^I$ map onto the same counterfactual CCP. Therefore, even though the model restrictions do not suffice to point identify the model parameters, they may suffice to identify counterfactual behavior.

Proposition 1 extends the literature in two directions. First, previous work (Aguirregabiria and Suzuki, 2014; Norets and Tang, 2014; Kalouptsi, Scott, and Souza-Rodrigues, 2021) focus on determining whether \tilde{p} is point-identified for specific counterfactual transformations, in the absence of any model restrictions (beyond the basic dynamic setup); in our notation, whether $\text{rank}(C) = 0$. Here, we show how additional model restrictions can help point-identify counterfactual behaviour that would otherwise not be identified. Specifically, model restrictions (9) can make $\text{rank}(CP_Q) = 0$ even when $\text{rank}(C) \neq 0$. Second, and more important, Proposition 1 establishes the exact dimension and shape of the counterfactual CCP identified set, which may be of interest in itself to practitioners, and can help compute the identified set of low-dimensional objects θ in practice, as we discuss below.

4.2 Identification of Counterfactual Outcomes of Interest

We now investigate the identified set for low-dimensional outcomes of interest $\theta \in \Theta \subset \mathbb{R}^n$.

Proposition 2. *The sharp identified set for θ is*

$$\Theta^I = \left\{ \begin{array}{l} \theta \in \Theta : \exists (\tilde{p}, \pi) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X} \text{ such that} \\ \theta = \phi(\tilde{p}, \pi; p), \mathbf{M}\pi = b_{-J}(p), \\ R^{eq}\pi = r^{eq}, R^{iq}\pi \leq r^{iq}, \\ (\tilde{\mathbf{M}}\mathcal{H})\pi = \tilde{b}_{-J}(\tilde{p}) - \tilde{\mathbf{M}}g \end{array} \right\}. \quad (19)$$

When ϕ is a continuous function of (\tilde{p}, π) , Θ^I is a connected set. In addition, when Π^I is bounded, Θ^I is compact. Finally, if θ is a scalar, then Θ^I is an interval.

Proposition 2 states that a vector θ belongs to Θ^I if and only if there exists a payoff π that is compatible with the data (i.e., $\mathbf{M}\pi = b_{-J}$), satisfies the model restrictions (i.e., $R^{eq}\pi = r^{eq}$ and $R^{iq}\pi \leq r^{iq}$), can

generate \tilde{p} in the counterfactual scenario (i.e., $(\widetilde{\mathbf{M}}\mathcal{H})\pi = \tilde{b}_{-J}(\tilde{p}) - \widetilde{\mathbf{M}}g$), and the corresponding pair (\tilde{p}, π) can generate θ (i.e., $\theta = \phi(\tilde{p}, \pi; p)$).

When ϕ is continuous, Θ^I is connected because it is the image set of a (composite) continuous function defined on the convex polyhedron Π^I . If the model restrictions make Π^I bounded, Θ^I becomes a compact and connected set, which is convenient as it suffices to trace the boundary of Θ^I to characterize this set in practice. In addition, when θ is a scalar, Θ^I reduces to a compact interval, which is even simpler to characterize: in that case we just need to compute the lower and upper endpoints of the interval Θ^I .

The upper endpoint of this interval can be calculated by solving the following constrained maximization problem:

$$\theta^U \equiv \max_{(\tilde{p}, \pi) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X}} \phi(\tilde{p}, \pi; p) \quad (20)$$

subject to

$$\begin{aligned} \mathbf{M}\pi &= b_{-J}(p), \\ R^{eq}\pi &= r^{eq}, \\ R^{iq}\pi &\leq r^{iq}, \\ (\widetilde{\mathbf{M}}\mathcal{H})\pi &= \tilde{b}_{-J}(\tilde{p}) - \widetilde{\mathbf{M}}g. \end{aligned} \quad (21)$$

The lower bound of the identified set θ^L is defined similarly (but replacing max by min). For ease of exposition, we focus on the maximization problem hereafter.

The problem (20)–(21) is a nonlinear maximization problem with linear constraints on π and smooth nonlinear constraints on \tilde{p} . When ϕ is differentiable, the optimization is well-behaved and can be solved using standard software (e.g., Knitro), even when the state space is large.

In our experience, standard algorithms are highly efficient in solving (20)–(21) in empirically-relevant high-dimensional problems when the researcher provides the gradient of ϕ . In some cases, however, the gradient of ϕ may be nontrivial to compute; for instance, this is the case when the target parameter θ involves average effects based on both factual and counterfactual ergodic distributions of the states. For such cases, we show in Appendix B how to calculate the gradient of ϕ analytically to help the numerical search.

In other cases, numerical gradients are costly to evaluate, and then standard solvers can be slow to converge. We thus develop a new a stochastic algorithm that exploits the structure of the problem (20)–(21) and combines the strengths of alternative stochastic search procedures. We discuss and describe our proposed algorithm in Appendix C.²⁰

²⁰Intuitively, one can search for θ^U by finding admissible values and updated directions for π , which is not computationally difficult as it only depends on linear constraints. But finding the corresponding \tilde{p} by solving the nonlinear equation (16) repeatedly (and calculating the gradient of ϕ numerically) can be demanding. Alternatively, searching stochastically over \tilde{p} and then finding a compatible π satisfying linear constraints is simpler. However, $\tilde{\mathbf{P}}^I$ may be a “thin” set in $\tilde{\mathbf{P}}$ with an unknown shape (since its dimension can be much smaller than the dimension of the simplex; see Proposition 1). As a result,

Extension: Unknown β and G . We now extend our results to allow for unknown discount factor β and distribution G . Formally, we assume that $\beta \in \mathcal{B} \subset [0, 1)$, and that the distribution of the idiosyncratic shocks ε_{it} belongs to a set of parametric distributions that are absolutely continuous with respect to the Lebesgue measure and have full support on \mathbb{R}^{A+1} . We denote that distribution by $G(\varepsilon_{it}; \lambda)$, where $\lambda \in \Lambda \subset \mathbb{R}^q$. The ex-ante value function, denoted here by $V(x; \lambda)$, is adjusted accordingly. The next proposition follows:

Proposition 3. *Assume that $\beta \in \mathcal{B} \subset [0, 1)$, where \mathcal{B} is a compact interval, and that $G(\varepsilon_{it}; \lambda)$, with $\lambda \in \Lambda \subset \mathbb{R}^q$, where Λ is compact and convex. Consider the counterfactual $\{\tilde{\mathcal{A}}, \tilde{\mathcal{X}}, \tilde{\beta}, \tilde{G}, h^s, h\}$, where $\tilde{\beta}$ and \tilde{G} are continuous functions of β and G , respectively, and assume that $V(x; \lambda)$ is continuously differentiable with respect to λ .²¹ Then, the sharp identified set for the counterfactual CCP \tilde{p} is a connected manifold. Moreover, if ϕ is a continuous function of $(\tilde{p}, \pi, \beta, \lambda)$, then the sharp identified set for θ is a connected set (and it is also compact when Π^I is bounded).*

Naturally, the identified sets for \tilde{p} and θ in Proposition 3 are larger than the corresponding sets when β and G are known. To find the bounds for a scalar θ in practice, we now need to solve an (augmented) constrained optimization problem, which is similar to (20)–(21) except we optimize over β, λ in addition to \tilde{p}, π . The optimization is well-behaved and can be solved using standard software, as before. That said, there is the additional computational challenge of performing the search on a larger space, as well as some computational difficulties arising from the dependence of \mathbf{M} and $\tilde{\mathbf{M}}$ on β , and the calculation of ψ_a for flexible distributions $G(\varepsilon; \lambda)$. For the formulation of the problem and a discussion on these challenges see Appendix A.

The additional computational costs reflect the difficulties in calculating counterfactuals in DDC models more generally when the discount factor and the distribution of unobservables are unknown. Norets (2011) characterizes the identified set for \tilde{p} in dynamic multinomial choice models when G is completely nonparametric, but the identified set is infeasible to compute in practice. Here, in contrast, we maintain the parametric assumption, allowing for a set of flexible distributions while preserving the smoothness (and the computational feasibility) of the constrained optimization (A7)–(A8).

4.3 Example: Firm Entry/Exit Model (Continued)

To illustrate the shape and size of the identified sets $\tilde{\mathbf{P}}^I$ and Θ^I , we now return to the firm entry/exit example. The counterfactual experiment we consider in this example is a subsidy that decreases entry

it is difficult to find points within that set randomly, and it is easy for perturbation methods to “exit” the set, increasing the cost of finding the maximum θ . We therefore develop a new algorithm in which we move in the “ \tilde{p} -world” (to avoid solving (16) repeatedly), but we keep a close eye on the “ π -world” (to keep track of the model restrictions and search in relevant directions). Searching in relevant directions without solving (16) and computing the numerical gradient of ϕ in every step improves substantially how fast θ moves on each iteration to the maximum.

²¹Formally, we need that, for any measurable and $\|\cdot\|_\infty$ -bounded function $\bar{v}_a(x)$ on $\mathcal{A} \times \mathcal{X}$, the function $\bar{V}(\bar{v}(x); \lambda) := \int_\varepsilon \max_{a \in \mathcal{A}} \{\bar{v}_a(x) + \varepsilon_a\} dG(\varepsilon; \lambda)$ is continuously differentiable with respect to λ . See Theorem 4.1 of Rust (1988).

cost by 20%. Formally, $\tilde{\pi} = \mathcal{H}\pi + g$, with $g = 0$, and \mathcal{H} block-diagonal with the diagonal blocks \mathcal{H}_{00} and \mathcal{H}_{11} given by

$$\mathcal{H}_{00} = I, \text{ and } \mathcal{H}_{11} = \begin{bmatrix} \tau I & (1 - \tau)I \\ 0 & I \end{bmatrix},$$

where $\tau = 0.8$. This implies

$$\tilde{\pi}_0 = \mathcal{H}_{00}\pi_0 = \pi_0, \text{ and } \tilde{\pi}_1 = \mathcal{H}_{11}\pi_1 = \begin{bmatrix} vp - fc - \tau \times ec \\ vp - fc \end{bmatrix}.$$

Figure 2 presents the results. First, note that the baseline and counterfactual CCPs, p and \tilde{p} , are represented by the black empty circle and the black full dot, respectively. In the baseline scenario, there is a higher probability of entry (and staying) in the high state than in the low state because higher values of w lead to greater profits and because w follows a persistent Markov process. In the counterfactual, the subsidy increases the probability of entry compared to the baseline in both low and high states w (as it should). Moreover, the subsidy decreases the probability of staying in the market, as it becomes cheaper to re-enter in the future.

We now characterize the identified set $\tilde{\mathbf{P}}^I$ under Restrictions 1–3. Similar to our representation of Π^I in Figure 1, the larger sets (including the dark blue areas) correspond to $\tilde{\mathbf{P}}^I$ under Restriction 1. The identified set is highly informative: it is a two-dimensional set in a four-dimensional space (recall from Proposition 1 that $\tilde{\mathbf{P}}^I$ is at most at the same dimension as Π^I), excluding most points in $\tilde{\mathbf{P}}$ from being possible counterfactual CCPs.²² Yet, because the baseline CCP p is at the boundary of $\tilde{\mathbf{P}}^I$, one cannot rule out the possibility that the entry subsidy has no impact on the firm’s behavior. Adding Restriction 2 reduces the size of $\tilde{\mathbf{P}}^I$ substantially (corresponding to the light blue areas in the figure). This is a direct consequence of the smaller set Π^I obtained after imposing Restriction 2 in addition to Restriction 1 (see Figure 1). The baseline CCP does not belong to $\tilde{\mathbf{P}}^I$ once we add Restriction 2; in fact, the location of p and $\tilde{\mathbf{P}}^I$ allows us to conclude that the probability of entry increases in the counterfactual and that the probability of staying decreases. In other words, the sign of the treatment effect is identified. The exclusion restriction on scrap values (Restriction 3) has substantial identification power, making $\tilde{\mathbf{P}}^I$ one-dimensional (because Π^I becomes one-dimensional as well) – see the blue lines in the figure. Note that all identified sets are connected, as expected (Proposition 1), but not necessarily convex.

We now turn to some low-dimensional outcomes θ , in particular, the long-run average impact of the entry subsidy on (i) the probability of staying in the market (labelled θ_P), (ii) consumer surplus (θ_S), and (iii) the value of the firm (θ_V). Table 1 presents the identified sets for each of these outcomes under Restrictions 1–3.²³

²²Restriction 1 determines the shape of $\tilde{\mathbf{P}}^I$, but not its dimension. Indeed, $\tilde{\mathbf{P}}^I$ is two-dimensional in the absence of any model restriction (beyond the basic setup); see Online Appendix G.

²³Assuming a (residual) linear inverse demand $P_{it} = w_{it} - \eta Q_{it}$, where P_{it} is the price and Q_{it} is the quantity demanded, and assuming a constant marginal cost mc , the variable profit is given by $vp = (w_{it} - mc)^2/4\eta$. The consumer surplus

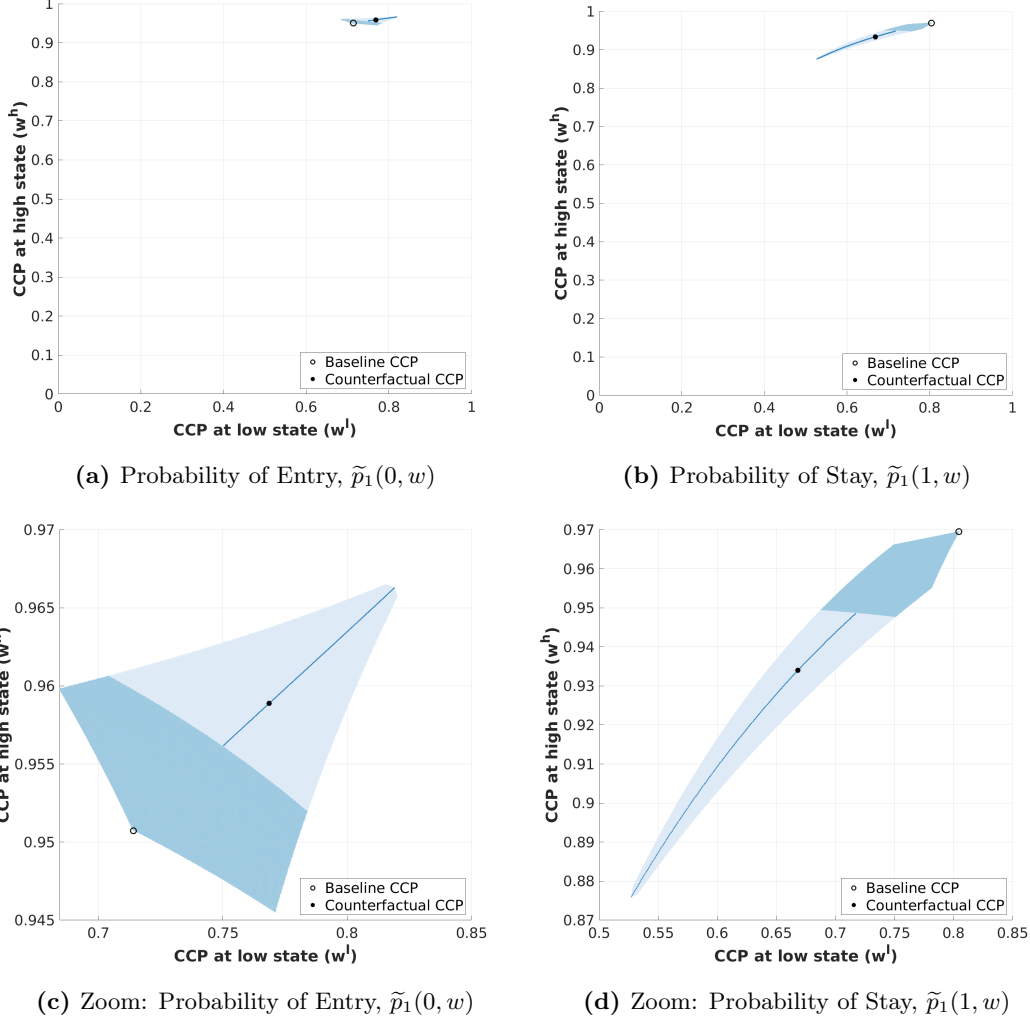


Figure 2: Identified Set for Counterfactual CCPs, $\tilde{\mathbf{P}}^I$, under Restrictions 1, 2, and 3. The larger sets (including the dark blue areas) correspond to $\tilde{\mathbf{P}}^I$ under Restriction 1. The light blue areas correspond to $\tilde{\mathbf{P}}^I$ under Restrictions 1 and 2. The identified set $\tilde{\mathbf{P}}^I$ under Restrictions 1–3 is represented by the blue lines within the light blue areas. The baseline and counterfactual CCPs, p and \tilde{p} , are represented by the black empty circle and the black full dot, respectively. The bottom panels present the “zoomed-in” versions of the top panels.

Perhaps surprisingly, the entry subsidy decreases the long-run average probability of the firm staying in the market, by approximately 6.4 percentage points. That is because, while the subsidy induces more entry, it also induces more exit. In the current case, increasing both firm’s entry and exit rates results in less time spent in the market in the long run. This in turn reduces the long-run average consumer surplus, and raises the average long-run value of the firm.

is $S = 0$ when the firm is inactive ($a = 0$), and $S = (w_{it} - mc)^2/8\eta$ when it is active ($a = 1$). Note that consumer surplus is the same in the baseline and counterfactual scenarios; the *average* S changes in the counterfactual because the firm changes its entry behavior when it receives an entry subsidy. The value of the firm in the baseline is given by the vector $V = (I - \beta F_J)^{-1} (\pi_J + \psi_J(p))$, where we take $J = 0$ (see footnote 13), and a similar expression holds for the counterfactual value: $\tilde{V} = (I - \beta \tilde{F}_J)^{-1} (\tilde{\pi}_J + \tilde{\psi}_J(\tilde{p}))$. The *average* firm value (across states) changes in the counterfactual both because the steady state distribution changes, and because the value of the firm is affected by the subsidy in *all* states. See Online Appendix G for explicit formulas for $\theta = (\theta_P, \theta_S, \theta_V)$.

Table 1: Sharp Identified Sets for the Long-run Impact of the Entry Subsidy on Outcomes of Interest, Θ^I

Outcome of Interest	Target parameter			
	True	Restriction 1	Restrictions 1-2	Restrictions 1-3
Change in Prob. of Being Active	-0.0638	[-0.1235, 0.0000]	[-0.1235, -0.0341]	[-0.1235, -0.0421]
Change in Consumer Surplus	-0.0875	[-0.1735, 0.0000]	[-0.1735, -0.0474]	[-0.1735, -0.0573]
Change in the Value of the Firm	0.9513	[0.0000, 1.8229]	[0.4489, 1.8229]	[0.6388, 1.8229]

Notes: This table shows the true and the sharp identified sets for the long-run average effect of the 20% entry subsidy on three outcomes of interest in the firm entry/exit problem: the probability of staying active, the consumer surplus, and the value of the firm. The averages are taken with respect to the state variables, using the steady-state distribution. The value of the model parameters and Restrictions 1, 2, and 3 are all specified in Section 3. See Online Appendix G for details.

As expected, the identified sets are all compact intervals (Proposition 2), and they all contain the true θ . Under Restriction 1, the upper bound of the identified set for θ_P is zero, leading to the conclusion that the long-run average probability of being active does not increase in the counterfactual. The lower bound implies that the probability of staying active can be reduced by at most 12 percentage points. Similarly, the researcher can conclude that the long-run average consumer surplus does not go up (and decreases by at most \$0.17), while the long-run average value of the firm does not go down (and increases at most by \$1.8) in response to the subsidy. These are informative identified sets despite the fact that Restriction 1 is mild.

Adding Restriction 2 makes all identified sets more informative. The upper bound on θ_P is now lower, implying that the average probability of being active is now reduced by a number between 3.4 and 12 percentage points, which clearly identifies the sign of the impact. The endpoints of the intervals for θ_S and θ_V change similarly. Adding Restriction 3 does not narrow the intervals much further, despite the fact that this restriction has substantial identifying power related to the model parameters π and counterfactual behavior \tilde{p} . That is because, while Restriction 3 reduces the dimension of Π^I and $\tilde{\mathbf{P}}^I$, it does not affect substantially the extreme points of these sets that in turn generate the endpoints of Θ^I . In Online Appendix G, we present the three-dimensional identified set Θ^I .

In Appendix E, we present results from a Monte Carlo study based on this example of firm entry and exit. Our findings there are analogous: the sets are informative even under the mildest restrictions and always contain the true parameter values. Moreover, calculating the bounds for θ is computationally fast, even in cases where the state space is large.

5 Estimation and Inference

We now present the inference procedure for the main outcomes of interest $\theta \in \Theta \subset \mathbb{R}^n$. In particular, we are interested in constructing confidence sets (CS's) for the true value of θ (rather than for the identified set Θ^I). Our approach is similar in spirit to the Hotz and Miller (1993) two-step estimator: we estimate choice probabilities p and transitions of state variables F in the first step, and then we perform inference on θ in the second step.

We assume the econometrician has access to a panel data on agents' actions and states: $\{a_{it}, x_{it} : i = 1, \dots, N; t = 1, \dots, T\}$. We consider asymptotics for the large N and fixed T case, as is typical in microeconomic applications of single-agent models, and assume i.i.d. sampling in the cross-section dimension.²⁴ Given that actions and states are finite, we consider frequency estimators for both p and F . (When states are continuous, one can use kernel estimators for p and F ; we leave extensions to continuous states for future research.) Specifically, for all $a \in \mathcal{A}$, and all $x, x' \in \mathcal{X}$,

$$\hat{p}_{aN}(x) = \frac{\sum_{it} 1 \{x_{it} = x, a_{it} = a\}}{\sum_{it} 1 \{x_{it} = x\}}, \quad (22)$$

$$\hat{F}_{aN}(x', x) = \frac{\sum_{it} 1 \{x_{it+1} = x', x_{it} = x, a_{it} = a\}}{\sum_{it} 1 \{x_{it} = x, a_{it} = a\}}, \quad (23)$$

and the vectors of sample frequencies are denoted by \hat{p}_N and \hat{F}_N .²⁵ We collect the terms \hat{p}_N and \hat{F}_N into the L -vector $\hat{\mathbf{p}}_N = [\hat{\mathbf{p}}_{1N}, \dots, \hat{\mathbf{p}}_{LN}]'$. Similarly, we collect p and F into $\mathbf{p} = [\mathbf{p}_1, \dots, \mathbf{p}_L]'$:= $E[\mathbf{e}]$, where \mathbf{e} is a vector of observed indicators. Recall that each matrix $M_a, a \in \mathcal{A}$, is a function of F , which is a subvector of \mathbf{p} , therefore we define $M_a(\mathbf{p}), a \in \mathcal{A}$, as the value of M_a evaluated at \mathbf{p} and also define $\mathbf{M}(\mathbf{p})$ accordingly. We use the same notation for $b_{-J}(\mathbf{p})$, as well as for $\widetilde{\mathbf{M}}(\mathbf{p}), \widetilde{b}_{-J}(\tilde{p}, \mathbf{p})$, and $\phi(\tilde{p}, \pi; \mathbf{p})$ when appropriate.²⁶

We construct a confidence set by inverting a test. To test the null $H_0 : \theta = \theta_0$, we reformulate the problem in the following way. For a fixed value $\theta = \theta_0$, we take the following equality constraints on π :

$$R^{eq}\pi = r^{eq}, \quad (\widetilde{\mathbf{M}}(\mathbf{p})\mathcal{H})\pi = \widetilde{b}_{-J}(\tilde{p}, \mathbf{p}) - \widetilde{\mathbf{M}}(\mathbf{p})g, \quad \text{and } \theta_0 = \phi(\tilde{p}, \pi; \mathbf{p}), \quad \text{for some } \tilde{p},$$

and collect them into

$$\mathcal{R}(\theta_0, \pi, \tilde{p}; \mathbf{p}) = 0.$$

²⁴If the data is ergodic and an appropriate mixing condition is satisfied then our procedure remains valid when $T \rightarrow \infty$ and N is fixed.

²⁵In certain cases some elements of the transition matrix F are degenerate when the corresponding states are known to evolve deterministically; see equation (G1) in the Online Appendix. We do not estimate these elements, and thus the expressions in (23) are applied only to the rest of the elements of F that need to be estimated.

²⁶Note that $\widetilde{\mathbf{M}}$ may also depend on baseline transitions F (and so may have to be estimated in the data). That is because $\widetilde{\mathbf{M}} = [I, -\widetilde{M}_{-J}]$, where \widetilde{M}_{-J} stacks \widetilde{M}_a for all $a \neq J$, with $\widetilde{M}_a = (I - \beta\widetilde{F}_a)(I - \beta\widetilde{F}_J)^{-1}$, and $\widetilde{F} = h^s(F)$. The same applies to \widetilde{b}_{-J} , which also depends on \widetilde{F} .

This leads to the criterion function

$$J(\theta_0) := \min_{\substack{(\tilde{\mathbf{p}}, \pi) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X} : R^{iq} \pi \leq r^{iq}, \\ \mathcal{R}(\theta_0, \pi, \tilde{\mathbf{p}}; \mathbf{p}) = 0}} [b_{-J}(\mathbf{p}) - \mathbf{M}(\mathbf{p})\pi]' \Omega [b_{-J}(\mathbf{p}) - \mathbf{M}(\mathbf{p})\pi], \quad (24)$$

where Ω is a (user-chosen) positive definite weighting matrix. If θ_0 belongs to Θ^I then all restrictions are satisfied and $J(\theta_0) = 0$, otherwise $J(\theta_0) > 0$. The identified set Θ^I can therefore be represented as the set of θ 's in Θ such that $J(\theta) = 0$. This implies that the null $H_0 : \theta = \theta_0$ is equivalent to $H'_0 : J(\theta_0) = 0$.

The test statistic is based on the empirical counterpart of $J(\theta_0)$, which is given by

$$\hat{J}_N(\theta_0) := \min_{\substack{(\tilde{\mathbf{p}}, \pi) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X} : R^{iq} \pi \leq r^{iq}, \\ \mathcal{R}(\theta_0, \pi, \tilde{\mathbf{p}}; \hat{\mathbf{p}}_N) = 0}} [b_{-J}(\hat{\mathbf{p}}_N) - \widehat{\mathbf{M}}_N \pi]' \widehat{\Omega}_N [b_{-J}(\hat{\mathbf{p}}_N) - \widehat{\mathbf{M}}_N \pi], \quad (25)$$

where $\widehat{\mathbf{M}}_N = \mathbf{M}(\hat{\mathbf{p}}_N)$, and $\widehat{\Omega}_N$ is a consistent estimator for Ω . For the rest of this paper we consider a general specification of Ω so that it can be a (known) continuous function of \mathbf{p} . Denoting the function by $\Omega(\cdot)$, we let $\widehat{\Omega}_N = \Omega(\hat{\mathbf{p}}_N)$ in (25).

The rejection region of the test with significance level α is $N\hat{J}_N(\theta_0) > \hat{c}_{1-\alpha}(\theta_0)$, where $\hat{c}_{1-\alpha}(\theta_0)$ is a data-dependent critical value. While a naive bootstrap for $\hat{J}_N(\theta_0)$ fails to deliver critical values that are asymptotically uniformly valid (see, e.g. Kitamura and Stoye, 2018), subsampling works under weak conditions, as we shall show shortly. Let h_N be the subsample size, with $h_N \rightarrow \infty$ as $N \rightarrow \infty$. A subsample version of $\hat{J}_N(\theta_0)$ is

$$\hat{J}_{h_N}^*(\theta_0) := \min_{\substack{(\tilde{\mathbf{p}}, \pi) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X} : R^{iq} \pi \leq r^{iq}, \\ \mathcal{R}(\theta_0, \pi, \tilde{\mathbf{p}}; \hat{\mathbf{p}}_{h_N}^*) = 0}} [\hat{b}_{-J}^* - \widehat{\mathbf{M}}_{h_N}^* \pi]' \widehat{\Omega}_{h_N}^* [\hat{b}_{-J}^* - \widehat{\mathbf{M}}_{h_N}^* \pi], \quad (26)$$

where $\hat{\mathbf{p}}_{h_N}^*$ is a subsample estimator of \mathbf{p} , $\widehat{\mathbf{M}}_{h_N}^* = \mathbf{M}(\hat{\mathbf{p}}_{h_N}^*)$, $\widehat{\Omega}_{h_N}^* = \Omega(\hat{\mathbf{p}}_{h_N}^*)$ and

$$\hat{b}_{-J}^* = b_{-J}(\hat{\mathbf{p}}_{h_N}^*) - b_{-J}(\hat{\mathbf{p}}_N) + \hat{b}_{-J}(\hat{\mathbf{p}}_N),$$

with $\hat{b}_{-J}(\hat{\mathbf{p}}_N)$ being the value of $\widehat{\mathbf{M}}_N \pi$ solving the minimization problem (25). Note that with this definition of \hat{b}_{-J}^* we implement subsampling with centering.²⁷

The testing procedure is simple: We use the empirical distribution of $h_N \hat{J}_{h_N}^*(\theta_0)$ to obtain the critical value $\hat{c}_{1-\alpha}(\theta_0)$. When the value of the test statistic is smaller than the critical value, $N\hat{J}_N(\theta_0) \leq \hat{c}_{1-\alpha}(\theta_0)$, we do not reject the null $H'_0 : J(\theta_0) = 0$, otherwise we reject it. The $1 - \alpha$ confidence set will be the collection of θ_0 's for which the tests do not reject the null.

In Appendix D, we provide several details regarding the implementation. In particular, to make the

²⁷The uncentered version takes $\hat{b}_{-J}^* = b_{-J}(\hat{\mathbf{p}}_{h_N}^*)$, and it is asymptotically valid as well. However, in our numerical experience, it has worse finite sample behavior than the centered version.

procedure operational and computationally fast, we do not actually solve problems (25) and (26) directly. We note instead that it is easier to solve relaxed versions of (20)–(21) to obtain good approximations for $\widehat{J}_N(\theta_0)$ and $\widehat{J}_{h_N}^*(\theta_0)$. We also describe how one can exploit continuity in the sequences of optimizations involved. In the Monte Carlo study, presented in Appendix E, we illustrate that confidence sets are tight and have the correct coverage probability, and that the computation remains remarkably tractable, even in cases with a large state space

Remark 1. A comment on some approaches that are alternative to ours as outlined above is in order. First, if we treat (π, \widetilde{p}) as a parameter (while θ_0 is fixed), then our system becomes one of set identified moment *equalities* (composed of equations (8), (14), and (16)), with (inequality) constraints on the parameter space for (π, \widetilde{p}) (i.e., restrictions (9) and (10)). It is then possible to test the validity of these equalities at each value of (π, \widetilde{p}) . This controls size, but would be extremely conservative; obviously the same can be done to standard moment inequality models but it is not implemented in practice for this reason. Moreover, implementing such a procedure in our context is practically impossible, as the parameter space for (π, \widetilde{p}) is too big. Second, we can fix \widetilde{p} , but not π , and rewrite the system into a moment inequality form by eliminating π (i.e. solving for other variables). As noted in Kitamura and Stoye (2018), this amounts to transforming, in the language of discrete geometry, a \mathcal{V} -representation of a polytope to an \mathcal{H} -representation, and it is generally known to be expensive to compute, and impractical even for a moderately sized system. Third, one may try to eliminate both π and \widetilde{p} from the system to get some form of moment inequalities; but this is even harder to implement, especially because of the nonlinear constraints that involve \widetilde{p} , and so it is not a practically feasible option either. For example, a recent paper by Kaido, Molinari, and Stoye (2019) is, like ours, concerned with a low dimensional object, though it is not directly applicable as their algorithm requires a moment inequality representation.

The next theorem follows:

Theorem 1. *Under Condition 1 presented in Appendix A,*

$$\liminf_{N \rightarrow \infty} \inf_{(\mathbf{p}, \theta) \in \mathcal{P}} \Pr\{N\widehat{J}_N(\theta) \leq \widehat{c}_{1-\alpha}(\theta)\} = 1 - \alpha,$$

where $\widehat{c}_{1-\alpha}(\theta)$ is the $1 - \alpha$ quantile of $h_N \widehat{J}_{h_N}^*(\theta)$, with $0 \leq \alpha \leq \frac{1}{2}$, and, for some positive constants c_1 and c_2 ,

$$\mathcal{P} := \left\{ \begin{array}{l} (\mathbf{p}, \theta) : \mathbf{p}_\ell \in (0, 1), E \left[\left| \frac{\mathbf{e}_\ell}{\sqrt{\mathbf{p}_\ell(1-\mathbf{p}_\ell)}} \right|^{2+c_1} \right] < c_2, 1 \leq \ell \leq L, \theta \in \Theta, \\ \exists (\widetilde{p}, \pi) \in \widetilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X} \text{ such that } \mathbf{M}(\mathbf{p})\pi = b_{-J}(\mathbf{p}), \\ R^{iq}\pi \leq r^{iq}, \mathcal{R}(\theta, \pi, \widetilde{p}; \mathbf{p}) = 0, \det(\Omega(\mathbf{p})) \geq c_1 \end{array} \right\}.$$

The asymptotically uniformly valid $1 - \alpha$ confidence set for θ is $CS = \{\theta \in \Theta : N\widehat{J}_N(\theta) \leq \widehat{c}_{1-\alpha}(\theta)\}$.

The set \mathcal{P} imposes regularity conditions on the data generating process for every counterfactual value

of interest $\theta \in \Theta$.²⁸ Our test statistic (25) is the squared minimum distance between the random vector $b_{-J}(\widehat{\mathbf{p}}_N)$ and a random manifold. It is therefore crucial to take sampling uncertainty in both objects into account. Importantly, the manifold does not have to be convex. We avoid such standard convexity conditions as they are typically incompatible with our model restrictions, in particular the general equality restrictions $\mathcal{R}(\theta_0, \pi, \tilde{\mathbf{p}}; \mathbf{p}) = 0$. Theorem 1 establishes the asymptotic validity of our procedure, addressing these issues.²⁹

6 Empirical Application

In this section, we illustrate our approach in the context of a dynamic model of export behavior. To that end, we consider the setup of Das, Roberts, and Tybout (2007), henceforth ‘DRT’, who perform a horserace between different kinds of export subsidies. As the authors point out, industrial exporters are highly prized in developing countries for generating gains from trade, sustaining production and employment during domestic recessions, and facilitating the absorption of foreign technologies. As a consequence, exporters often receive governmental support. Yet, seemingly similar subsidies may generate different export responses in different industries and time periods, making it difficult for policy makers to know which type of support is optimal. To shed light on these issues, DRT develop a structural dynamic model of firm export decisions and simulate returns of different subsidies. Here, we adopt their specification and explore the identifying power of alternative model restrictions to assess which of their restrictions deliver their main findings.

Data. We use DRT’s plant-level panel data from Colombian manufacturing industries and focus on the knitting mills industry. The dataset is composed of 64 knit fabric producers observed annually between 1981–1991; the sample has 704 plant-year observations. Like DRT, we focus on firms that operated continuously in the domestic market, given that they were responsible for most of the exports over this period. The share of exporting firms increased from 12% in 1981 to 18% by the end of the sample period, possibly a result of the 33% depreciation of Colombia’s real exchange rate. This industry also depicts significant turnover: the average probability of re-entry into export markets is 61%. On average, export revenues of exporting firms are approximately 1.4 times the domestic revenues.

Model. DRT assume that export markets are monopolistically competitive; this leads to a specification similar to the firm entry/exit model presented in Sections 3 and 4. In particular, every period t a firm

²⁸The first restriction in the definition of \mathcal{P} is a standard condition imposed to guarantee the Lindeberg condition. The second, the third, and the fourth collect the model restrictions (the equalities in the fourth restriction include the constraints that arise as we fix the value of the counterfactual θ). The final restriction guarantees that the minimizations (25) and (26) are asymptotically well-behaved.

²⁹Formally, the test statistic projects $b_{-J}(\widehat{\mathbf{p}}_N)$ on the manifold $S(\widehat{\mathbf{p}}_N, \theta_0)$, where $S(\mathbf{p}, \theta) := \{\mathbf{M}(\mathbf{p})\pi, \pi \in \mathbb{R}^{(A+1)X} : \mathcal{R}(\theta, \pi, \tilde{\mathbf{p}}; \mathbf{p}) = 0, \text{ and } R^{iq}\pi \leq r^{iq} \text{ hold for some } \tilde{\mathbf{p}} \in \tilde{\mathbf{P}}\}$. Condition 1 does *not* require that $S(\widehat{\mathbf{p}}_N, \theta_0)$ be convex; it requires instead that the tangent cone of $S(\widehat{\mathbf{p}}_N, \theta_0)$ be convex (see Appendix A for details).

i chooses whether to export or not, $a_{it} \in \mathcal{A} = \{0, 1\}$. The state variables are (i) the lagged decision ($k_{it} = a_{it-1}$), and (ii) exchange rates and demand/supply shocks in export markets (w_{it}). The exogenous shocks w_{it} follow (discretized) independent normal-AR(1) processes. The payoff function is given by equation (12) in Section 3. To point identify the model, DRT restrict to zero the payoffs of not exporting (i.e., both the outside value and the scrap value are set to zero). They also impose state-invariant entry and fixed costs, making their model overidentified. We relax these assumptions and instead explore the identifying power of Restrictions 1, 2, and 3 presented in the entry/exit model. In principle, scrap values may differ from zero because they may involve idleness costs (given that exiting is often temporary) or depreciation costs. Similarly, fixed costs and entry costs may depend on the aggregate states, as they involve finding trading partners, setting up distribution networks, maintaining labor and capital abroad, etc.³⁰ For ease of exposition, we leave the model details to the Online Appendix, Section H.

Counterfactuals and Outcomes of Interest. DRT focus on three counterfactual policies: (i) direct subsidies to plants' export revenues, such as a tax rebate that is proportional to foreign sales; (ii) subsidies to the cost of entering into exporting, such as grants for information or technology acquisition for export development; and (iii) payments designed to cover the annual fixed costs of operating in the export market. We follow DRT and consider a 2% export revenue subsidy, a 25% entry cost subsidy, and a 28% fixed cost subsidy.

The main outcome of interest is a benefit–cost ratio based on the long-run average annual gain in export revenues divided by the long-run average government subsidy expenditures. We denote the ratios for the revenue, fixed cost, and entry cost subsidies by θ_R , θ_F , and θ_E , respectively, and take $\theta = (\theta_R, \theta_F, \theta_E)$ – see Online Appendix H for explicit formulas for θ .

Evaluating ex-ante the impact of different model restrictions on θ is not trivial. Note first that while export revenues are observed in the data, the long-run *average* change in revenues depends on firms' decisions to export given the type of subsidy. This means that all numerators in θ depend on the counterfactual CCPs. Next, note that all denominators in θ equal the long-run average government expenditures, which depend on the fraction of firms exporting in the counterfactual steady-state; i.e., they all depend on \tilde{p} as well. In addition, θ_F and θ_E depend on the unknown parameters, fc and ec , respectively, since the government expenditures are direct functions of these costs. In the case of the entry cost subsidy, a further complication is that the (subsidized) entry cost is paid only when firms enter, implying that \tilde{p} affects the direct payments in each state (in addition to affecting the steady-state distribution). In short, θ depends on both \tilde{p} and π highly nonlinearly.

In terms of identification, the benefit-cost ratio of the revenue subsidy θ_R is point-identified. That is both because the averages in the numerator and denominator depend on observed revenues, and because

³⁰The payoff when not exporting (the outside option) may also be different from zero since it includes domestic profits. However, following DRT, we do not explore this possibility given the limitations in the data.

\tilde{p} is identified (since it involves known changes to known quantities, i.e., the identified variable profits; see KSS), implying that the counterfactual steady-state distribution is also point-identified. The other two target objects, θ_F and θ_E , are partially identified both because (i) the counterfactual behavior \tilde{p} is not point-identified (as the entry subsidy in our example in Section 4), and (ii) the denominators in the benefit–cost ratios depend *directly* on model parameters that are partially identified (i.e., on fc and ec , respectively). In sum, both θ_F and θ_E involve ratios of set-identified objects.

Results. DRT find, under their imposed restrictions ($s = 0$, as well as ec, fc invariant over states), that revenue subsidies yield the highest return, followed by the fixed cost subsidies, and then by the entry cost subsidies; i.e. $\theta_R > \theta_F > \theta_E$. We explore the robustness of this finding under milder restrictions.

We implement our procedure as explained in Sections 4 and 5, and in Appendix D.³¹ Table 2 presents the benefit–cost ratios under Restrictions 1–3. The revenue subsidy generates an estimated benefit–cost ratio, θ_R , of approximately 15 pesos of revenue per unit cost. Its impact is statistically significant and economically large, and it is fairly consistent with the estimates in DRT. Because θ_R is point identified, it does not depend on any additional model restriction (other than the basic framework (8)).

We now discuss θ_F and θ_E , which are partially identified. Restriction 1 (i.e., $fc \geq 0$ and $ec \geq 0$) is not sufficiently informative here: the fixed cost subsidies ratio, θ_F , is between 8 and 30, and the entry cost subsidies ratio, θ_E , ranges from 4 to 24. These sets are wide because there are still many model parameter values that can rationalize observed behavior. The identified sets overlap and we cannot conclude which policy generates the highest return.

Adding Restriction 2 increases the identification power substantially: the ratio for the fixed cost subsidies is now between 11 and 13. This identified set is highly informative and its upper bound is smaller than θ_R , suggesting that the revenue subsidy is more potent than the fixed cost subsidy. Intuitively, under Restriction 2, entry is profitable in the long-term, which imposes an upper limit on the values that fc can take. This upper limit, in turn, reduces the potential impact of the fixed cost subsidies.

Hence, under Restrictions 1 and 2 we can only confirm part of DRT’s findings (that $\theta_R > \theta_F$). In contrast, there is substantial uncertainty regarding the benefit–cost ratio for the entry cost subsidy θ_E : its identified set is between 7.8 and 17, containing both the estimated θ_R and the identified set of θ_F .

Incorporating exclusion restrictions on scrap values (Restriction 3) narrows the identified set for θ_E substantially: the benefit-cost ratio now ranges from 8.9 to 9.4, which is highly informative. Hence, under

³¹The transition process for exchange rates is taken from a long-time series as in DRT. Given the small sample size, we discretize the support of each exogenous state in three bins. We estimate CCPs using frequency estimators. To compute confidence intervals, we implement 1000 replications of a standard i.i.d. subsampling, resampling 20 firms over the sample time period, so that the size of each subsample is $h_N = 200 \approx 8 \times \sqrt{NT}$. To minimize the quadratic distances in (25) and (26), we take a diagonal weighting matrix Ω with diagonal elements given by the square-root of the ergodic distribution of the state variable – thus, deviations on more visited states are considered more relevant and receive greater weights. Given that θ_R is known (ex ante) to be point identified, we use the plug-in estimator proposed by Kalouptsi, Lima, and Souza-Rodrigues (2021) to estimate it, and 1000 standard i.i.d. bootstrap replications at the firm level to construct the confidence intervals for θ_R . To make our results comparable to DRT, we have also estimated the model parameters under their restrictions and obtained very similar results as theirs. See details in Section H of the Online Appendix.

Table 2: Export Revenue/Cost Ratio for Different Subsidies under Alternative Model Restrictions

	Restriction 1	Restrictions 1–2	Restrictions 1–3
Revenue Subsidies			
Estimated Identified Set	15.13	15.13	15.13
90% Confidence Interval	(11.15, 18.90)	(11.15, 18.90)	(11.15, 18.90)
Fixed Costs Subsidies			
Estimated Identified Set	[8.41, 30.82]	[11.10, 13.34]	[11.92, 12.60]
90% Confidence Interval	(7.47, 34.98)	(9.65, 14.46)	(9.92, 13.87)
Entry Costs Subsidies			
Estimated Identified Set	[4.40, 24.04]	[7.85, 17.28]	[8.88, 9.36]
90% Confidence Interval	(3.52, 34.36)	(7.03, 23.49)	(7.34, 14.33)

Notes: This table shows the estimated sharp identified sets for the average gains in total export revenues divided by the average government subsidy expenditures, both averaged over states in the long-run. The top panel shows the gains of a 2% export revenue subsidy; the middle panel, the gains of a 28% fixed cost subsidy; and the bottom panel, the gains of a 25% entry cost subsidy. The (nonsingleton) identified sets are in brackets. The data set is composed of 704 plant-year observations in the Colombian knitting mills industry. The 90% confidence intervals are in parenthesis and were calculated based on 1000 bootstrap replications for the revenue subsidies, and 1000 subsample replications for both fixed and entry costs subsidies (with subsample sizes of 200). Restrictions 1, 2, and 3 are all specified in the main text (Section 3). See Online Appendix H for details.

Restrictions 1-3, we can confirm DRT’s subsidy ranking ($\theta_R > \theta_F > \theta_E$): revenue subsidies generate the highest export revenues per unit cost, followed by fixed cost subsidies, and then by the entry cost subsidies. We thus conclude that, although this ranking can be obtained under milder restrictions than those imposed by DRT, it does seem to hinge on the assumption that scrap values do not depend on state variables. We now provide some intuition for this finding.

Why does the exclusion restriction on scrap values (Restriction 3) render entry cost subsidies less effective? Intuitively, under Restrictions 1–2, the state variables can induce a correlation between export revenues and scrap values. When that correlation is negative, entry subsidies would provide incentives for low-productivity firms (i.e., those with low export revenues) to enter and exit the export markets repeatedly, while high-productivity firms (with high export revenues) would stay longer in exporting. This makes the entry subsidy effective in terms of the benefit-cost ratio. When we impose the exclusion restriction on scrap values (Restriction 3), we eliminate that negative correlation and, as a result, the low-productivity firms stay more often in the export markets. This reduces the average gain in export revenues per peso of subsidy, pushing the upper bound on θ_E downwards, and making the entry subsidy worse than both the revenues and fixed costs subsidies. A similar result holds, but is less relevant in magnitude, for the lower bound on θ_E , reflecting possible positive correlations between export revenues and scrap values.

Of note, the uniformly valid confidence intervals indicate substantial sampling uncertainty, which is not surprising given the size of the data set.

7 Conclusion

In this paper, we develop a unified framework to investigate how much one can learn about counterfactual outcomes in dynamic discrete choice models for a large and empirically relevant class of counterfactual experiments for which the level of flow utilities (or payoffs) may matter. We derive analytical properties of the identified sets under alternative model restrictions. We also propose computationally tractable procedures for estimation and develop an asymptotically valid inference approach based on subsampling. The empirical implications of our results are illustrated based on the study of Das, Roberts, and Tybout (2007) on exporting decisions and subsidies.

Our primary motivation is to offer a solution for practitioners that is applicable to a widely used class of empirical models. We hope our results can aid practitioners in assessing which empirical findings survive under minimal restrictions, in understanding the impact of commonly imposed restrictions on counterfactuals, and in including confidence sets around their counterfactual outcomes. We leave extensions to identification of counterfactual optimal policy interventions and to dynamic games for future research.

References

- ABBRING, J. H., AND O. DALJORD (2020): “Identifying the Discount Factor in Dynamic Discrete Choice Models,” *Quantitative Economics*, 11(2), 471–501.
- ADAMS, A. (2020): “Mutually Consistent Revealed Preference Demand Predictions,” *American Economic Journal: Microeconomics*, 12(1), 42–74.
- ADAO, R., A. COSTINOT, AND D. DONALDSON (2017): “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade,” *American Economic Review*, 107(3), 633–89.
- AGUIRREGABIRIA, V. (2010): “Another Look at the Identification of Dynamic Discrete Decision Processes: An application to Retirement Behavior,” *Journal of Business & Economic Statistics*, 28(2), 201–218.
- AGUIRREGABIRIA, V., AND P. MIRA (2002): “Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models,” *Econometrica*, 70(4), 1519–1543.
- (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75(1), 1–53.
- AGUIRREGABIRIA, V., AND J. SUZUKI (2014): “Identification and Counterfactuals in Dynamic Models of Market Entry and Exit,” *Quantitative Marketing and Economics*, 12(3), 267–304.
- ARCIDIACONO, P., AND R. A. MILLER (2011): “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity,” *Econometrica*, 79(6), 1823–1867.

- (2020): “Identifying Dynamic Discrete Choice Models off Short Panels,” *Journal of Econometrics*, 215(2), 473–485.
- BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75(5), 1331–1370.
- BAJARI, P., C. S. CHU, D. NEKIPELOV, AND M. PARK (2016): “Identification and semiparametric estimation of a finite horizon dynamic discrete choice model with a terminating action,” *Quantitative Marketing and Economics*, 14(4), 271–323.
- BEJARA, M. (2020): “A Semi-structural Methodology for Policy Counterfactuals,” Discussion paper, MIT.
- BERRY, S., AND G. COMPIANI (2020): “An Instrumental Variables Approach to Dynamic Models,” Discussion paper, Yale University.
- BLEVINS, J. R. (2014): “Nonparametric identification of dynamic decision processes with discrete and continuous choices,” *Quantitative Economics*, 5(3), 531–554.
- BLUNDELL, R., M. BROWNING, AND I. CRAWFORD (2008): “Best Nonparametric Bounds on Demand Responses,” *Econometrica*, (76), 1227–1262.
- BLUNDELL, R., D. KRISTENSEN, AND R. MATZKIN (2014): “Bounding quantile demand functions using revealed preference inequalities,” *Journal of Econometrics*, 179(2), 112 – 127.
- BUCHHOLZ, N., M. SHUM, AND H. XU (2020): “Semiparametric Estimation of Dynamic Discrete Choice Models,” Discussion paper, Princeton University.
- CANAY, I. A., AND A. M. SHAIKH (2017): “Practical and Theoretical Advances in Inference for Partially Identified Models,” in *Advances in Economics and Econometrics: Eleventh World Congress*, ed. by B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson, vol. 2 of *Econometric Society Monographs*, pp. 271–306. Cambridge University Press.
- CHEN, L.-Y. (2017): “Identification of Discrete Choice Dynamic Programming Models with Nonparametric Distribution of Unobservables,” *Econometric Theory*, 33(3), 551–577.
- CHERNOFF, H. (1954): “On the Distribution of the Likelihood Ratio,” *The Annals of Mathematical Statistics*, pp. 573–578.
- CHIONG, K. X., A. GALICHON, AND M. SHUM (2016): “Duality in dynamic discrete-choice models,” *Quantitative Economics*, 7(1), 83–115.

- CHRISTENSEN, T., AND B. CONNAULT (2021): “Counterfactual Sensitivity and Robustness,” Discussion paper, New York University.
- COLLARD-WEXLER, A. (2013): “Demand Fluctuations in the Ready-Mix Concrete Industry,” *Econometrica*, 81(3), 1003–1037.
- CONLON, C. T. (2012): “A Dynamic Model of Costs and Margins in the LCD TV Industry,” Discussion paper, New York University.
- CRAWFORD, G. S., AND M. SHUM (2005): “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 73(4), 1137–1173.
- DAS, S., M. J. ROBERTS, AND J. R. TYBOUT (2007): “Market Entry Costs, Producer Heterogeneity, and Export Dynamics,” *Econometrica*, 75(3), 837–873.
- DEARING, A. (2019): “Pseudo-Value Functions and Closed-Form CCP Estimation of Dynamic Discrete Choice Models,” Discussion paper, Ohio State University.
- DICKSTEIN, M. J., AND E. MORALES (2018): “What Do Exporters Know?,” *The Quarterly Journal of Economics*, 133(4), 1753–1801.
- DUFLO, E., R. HANNA, AND S. P. RYAN (2012): “Incentives work: Getting teachers to come to school,” *The American Economic Review*, 102(4), 1241–1278.
- DUNNE, T., S. D. KLIMEK, M. J. ROBERTS, AND D. Y. XU (2013): “Entry, exit, and the determinants of market structure,” *The RAND Journal of Economics*, 44(3), 462–487.
- ECKSTEIN, Z., AND O. LIFSHITZ (2011): “Dynamic Female Labor Supply,” *Econometrica*, 79(6), 1675–1726.
- GILLESKIE, D. B. (1998): “A Dynamic Stochastic Model of Medical Care Use and Work Absence,” *Econometrica*, 66(1), 1–45.
- HECKMAN, J. J. (2000): “Causal Parameters and Policy Analysis in Economics: A Twentieth Century Retrospective,” *The Quarterly Journal of Economics*, 115(1), 45–97.
- (2010): “Building Bridges between Structural and Program Evaluation Approaches to Evaluating Policy,” *Journal of Economic Literature*, 48(2), 356–98.
- HECKMAN, J. J., AND S. NAVARRO (2007): “Dynamic discrete choice and dynamic treatment effects,” *Journal of Econometrics*, 136(2), 341–396.
- HENDEL, I., AND A. NEVO (2006): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74(6), 1637–1673.

- HOTZ, V. J., AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60(3), 497–529.
- HOTZ, V. J., R. A. MILLER, S. SANDERS, AND J. SMITH (1994): “A Simulation Estimator for Dynamic Models of Discrete Choice,” *Review of Economic Studies*, 61(2), 265–89.
- HU, Y., AND M. SHUM (2012): “Nonparametric identification of dynamic models with unobserved state variables,” *Journal of Econometrics*, 171(1), 32–44.
- ICHIMURA, H., AND C. TABER (2000): “Direct Estimation of Policy Impacts,” *NBER Working Paper* 254.
- (2002): “Semiparametric Reduced-Form Estimation of Tuition Subsidies,” *American Economic Review*, 92(2), 286–292.
- KAIDO, H., F. MOLINARI, AND J. STOYE (2019): “Confidence intervals for projections of partially identified parameters,” *Econometrica*, 87(4), 1397–1432.
- KALOUPTSIDI, M. (2014): “Time to build and fluctuations in bulk shipping,” *The American Economic Review*, 104(2), 564–608.
- KALOUPTSIDI, M., L. LIMA, AND E. SOUZA-RODRIGUES (2021): “On Estimating Counterfactuals Directly in Dynamic Models,” Discussion paper, University of Toronto.
- KALOUPTSIDI, M., P. T. SCOTT, AND E. SOUZA-RODRIGUES (2017): “On the Non-identification of Counterfactuals in Dynamic Discrete Games,” *International Journal of Industrial Organization*, 50, 362–371.
- KALOUPTSIDI, M., P. T. SCOTT, AND E. A. SOUZA-RODRIGUES (2021): “Identification of Counterfactuals in Dynamic Discrete Choice Models,” *Quantitative Economics*, 12(2), 351–403.
- KASAHARA, H., AND K. SHIMOTSU (2009): “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices,” *Econometrica*, 77(1), pp. 135–175.
- KEANE, M. P., AND A. MERLO (2010): “Money, Political Ambition, and the Career Decisions of Politicians,” *American Economic Journal: Microeconomics*, 2(3), 186–215.
- KEANE, M. P., AND K. I. WOLPIN (1997): “The Career Decisions of Young Men,” *Journal of Political Economy*, 105(3), 473–522.
- (2010): “The Role of Labor and Marriage Markets, Preference Heterogeneity and the Welfare System in the Life Cycle Decisions of Black, Hispanic and White Women,” *International Economic Review*, 51(3), 851–892.

- KITAMURA, Y., AND J. STOYE (2018): “Nonparametric Analysis of Random Utility Models,” *Econometrica*, 86(6), 1883–1909.
- KITAMURA, Y., AND J. STOYE (2019): “Nonparametric Counterfactuals in Random Utility Models,” Discussion paper, Yale University.
- KRANTZ, S. G., AND H. R. PARKS (2003): *The Implicit Function Theorem: History, Theory, and Applications*. Birkäuser Basel, 1 edn.
- LIN, H. (2015): “Quality Choice and Market Structure: A Dynamic Analysis of Nursing Home Oligopolies,” *International Economic Review*, 56(4), 1261–1290.
- MAGNAC, T., AND D. THESMAR (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70(2), 801–816.
- MANSKI, C. F. (2007): “Partial Identification of Counterfactual Choice Probabilities,” *International Economic Review*, 48, 1393–1410.
- MARSCHAK, J. (1953): “Economic Measurements for Policy and Prediction,” in *Cowles Commission Monograph 14: Studies in Econometric Methods*, ed. by W. C. Hood, and T. Koopmans. New York: Wiley.
- McFADDEN, D. (1974): *Conditional logit analysis of qualitative choice behavior* chap. 4, pp. 105–142. Academic Press, New York.
- MILLER, R. (1984): “Job matching and occupational choice,” *Journal of Political Economy*, 92(6), 1086–1120.
- MOGSTAD, M., A. SANTOS, AND A. TORGOVITSKY (2018): “Using Instrumental Variables for Inference about Policy Relevant Treatment Parameters,” *Econometrica*, 86(5), 1589–1619.
- MORALES, E., G. SHEU, AND A. ZAhLER (2019): “Extended gravity,” *The Review of Economic Studies*, 86(6), 2668–2712.
- NORETS, A. (2010): “Continuity and Differentiability of Expected Value Functions in Dynamic Discrete Choice Models,” *Quantitative Economics*, 1(2), 305–322.
- (2011): “Semiparametric Identification of Dynamic Multinomial Choice Models,” Discussion paper, Brown University.
- NORETS, A., AND X. TANG (2014): “Semiparametric Inference in dynamic binary choice models,” *The Review of Economic Studies*, 81(3), 1229–1262.

- PAKES, A. (1986): “Patents as options: Some estimates of the value of holding European patent stocks,” *Econometrica*, 54(4), 755–784.
- PAKES, A., M. OSTROVSKY, AND S. BERRY (2007): “Simple estimators for the parameters of discrete dynamic games (with entry/exit examples),” *The RAND Journal of Economics*, 38(2), 373–399.
- PESENDORFER, M., AND P. SCHMIDT-DENGLER (2008): “Asymptotic Least Squares Estimators for Dynamic Games,” *The Review of Economic Studies*, 75(3), 901–928.
- ROCKAFELLAR, R. T., AND R. J.-B. WETS (2009): *Variational analysis*, vol. 317. Springer Science & Business Media.
- ROMANO, J. P., AND A. M. SHAIKH (2008): “Inference for identifiable parameters in partially identified econometric models,” *Journal of Statistical Planning and Inference*, 138(9), 2786–2807.
- (2012): “On the uniform asymptotic validity of subsampling and the bootstrap,” *The Annals of Statistics*, 40(6), 2798–2822.
- ROSENZWEIG, M. R., AND K. I. WOLPIN (1993): “Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India,” *Journal of Political Economy*, 101(2), 223–244.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: an Empirical Model of Harold Zurcher,” *Econometrica*, 55(5), 999–1033.
- (1988): “Maximum likelihood Estimation of Discrete Control Processes,” *SIAM Journal on Control and Optimization*, 26(5), 1006–1024.
- (1994): “Structural Estimation of Markov Decision Processes,” *Handbook of Econometrics* 4, 4, 3081–3143.
- RYAN, S. P. (2012): “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 80(3), 1019–1061.
- SCHIRALDI, P. (2011): “Automobile replacement: a dynamic structural approach,” *The RAND Journal of Economics*, 42(2), 266–291.
- TRAIN, K. (2009): *Discrete Choice Methods with Simulation*. Cambridge UP.
- VARELA, M. J. (2018): “The costs of growth: Accelerated growth and crowd-out in the Mexican supermarket industry,” *International Journal of Industrial Organization*, 61, 1–52.

WEI, C., AND S. LI (2014): “The Cost of Greening Stimulus: A Dynamic Discrete Choice Analysis of Vehicle Scrappage Programs,” Working Papers 2014-12, The George Washington University, Institute for International Economic Policy.

WOLPIN, K. (1984): “An estimable dynamic stochastic model of fertility and child mortality,” *Journal of Political Economy*, 92(5), 852–874.

ZIEGLER, G. M. (2012): *Lectures on Polytopes*, vol. 152. Springer Science & Business Media.

Appendix

A Proofs

A.1 Proof of Proposition 1

The identified set $\tilde{\mathbf{P}}^I$ defined in (17) is sharp by construction because equations (8), (9), and (10) contain all model restrictions, and equation (16) fully characterizes \tilde{p} as an (implicit) function of π (see the arguments in footnotes 13 and 18 in the main text, and the characterization of $\tilde{\mathbf{P}}^I$ below).

Clearly, $\tilde{\mathbf{P}}^I$ is empty whenever Π^I is empty. Assume hereafter that Π^I is non-empty. Recall that the identified set is characterized by the equations (8), (9), (10), and (16). By combining (8) and (9), we get

$$(R_{-J}^{eq}M_{-J} + R_J^{eq})\pi_J = r^{eq} - R_{-J}^{eq}b_{-J}(p),$$

which is of the form:

$$Q^{eq}\pi_J = q^{eq}, \tag{A1}$$

where

$$Q^{eq} = R_{-J}^{eq}M_{-J} + R_J^{eq} \tag{A2}$$

is a $d \times X$ matrix, and $q^{eq} = r^{eq} - R_{-J}^{eq}b_{-J}(p) \in \mathbb{R}^d$. Equation (A1) incorporates all equality restrictions on π , and expresses them in terms of the “free parameter” $\pi_J \in \mathbb{R}^X$.

We assume that $\text{rank}(Q^{eq}) = d$ and that the first d columns of Q^{eq} are independent.³² We write (A1) as

$$Q^{eq}\pi_J = [Q_1 \ Q_2] \begin{bmatrix} \pi_J^1 \\ z \end{bmatrix} = Q_1\pi_J^1 + Q_2z = q^{eq},$$

³²In the more general case, the independent columns of Q^{eq} must be permuted to the front.

where Q_1 is $d \times d$ and non-singular. Then $\pi_J^1 = Q_1^{-1}q^{eq} - Q_1^{-1}Q_2z$. Therefore,

$$\pi_J = \begin{bmatrix} \pi_J^1 \\ z \end{bmatrix} = \begin{bmatrix} -Q_1^{-1}Q_2 \\ I \end{bmatrix} z + \begin{bmatrix} Q_1^{-1}q^{eq} \\ 0 \end{bmatrix} = P_Q z + r_Q. \quad (\text{A3})$$

This gives a complete parameterization of all π_J in terms of the “free” $X - d$ parameters in the vector z . Represent the elements of this set by $\pi_J(z)$. Note that in the absence of the equality restrictions (9), we can just take $\pi_J = z$.

Similarly, combine (8) and (10), to get

$$(R_{-J}^{iq}M_{-J} + R_J^{iq})\pi_J \leq r^{iq} - R_{-J}^{iq}b_{-J}(p),$$

which is of the form:

$$Q^{iq}\pi_J \leq q^{iq},$$

where $Q^{iq} = R_{-J}^{iq}M_{-J} + R_J^{iq}$ is an $m \times X$ matrix, and $q^{iq} = r^{iq} - R_{-J}^{iq}b_{-J}(p) \in \mathbb{R}^m$. Substituting π_J in the inequality above by $\pi_J(z)$ defined in (A3) and rearranging, we get the m inequalities defined in terms of $z \in \mathbb{R}^X$:

$$Q^{iq}P_Q z \leq q^{iq} - Q^{iq}r_Q. \quad (\text{A4})$$

Define the set

$$\mathcal{Z} = \left\{ z \in \mathbb{R}^{X-d} : Q^{iq}P_Q z \leq q^{iq} - Q^{iq}r_Q \right\}. \quad (\text{A5})$$

Clearly, \mathcal{Z} is a convex polyhedron. By construction, any vector $\pi = [\pi'_{-J}, \pi'_J]'$ such that $\pi_{-J} = M_{-J}\pi_J(z) + b_{-J}$, with $\pi_J(z)$ defined by (A3) for some $z \in \mathcal{Z}$ satisfies (8), (9), and (10). I.e., for any given $z \in \mathcal{Z}$, we can find one π satisfying all model restrictions.

Next, combine (8) and (16) to obtain

$$\underbrace{\begin{bmatrix} I & -\widetilde{M}_{-J} \end{bmatrix}}_{=\widetilde{\mathbf{M}}} \underbrace{\begin{bmatrix} \mathcal{H}_{1,-J} & \mathcal{H}_{1J} \\ \mathcal{H}_{2,-J} & \mathcal{H}_{2J} \end{bmatrix}}_{=\mathcal{H}} \underbrace{\begin{bmatrix} M_{-J}\pi_J + b_{-J}(p) \\ \pi_J \end{bmatrix}}_{=\pi} = \widetilde{b}_{-J}(\widetilde{p}) - \widetilde{\mathbf{M}}g,$$

or,

$$C\pi_J + (\mathcal{H}_{1,-J} - M_{-J}\mathcal{H}_{2,-J})b_{-J}(p) = \widetilde{b}_{-J}(\widetilde{p}) - g_{-J} + \widetilde{M}_{-J}g_J, \quad (\text{A6})$$

where C is the $\widetilde{A}\widetilde{X} \times X$ matrix defined in equation (18).

Noting that \widetilde{p} has to satisfy \widetilde{X} restrictions as it is a collection of conditional probability vectors, let \widetilde{p}^* denote a $\widetilde{A}\widetilde{X}$ -vector of independent elements of \widetilde{p} , and denote the set of independent elements by $\widetilde{\mathbf{P}}^*$.

Substitute (A3) into (A6), rearrange it, and define the function $\mathbf{F} : \mathbb{R}^{X-d} \times \text{int}(\tilde{\mathbf{P}}^*) \rightarrow \mathbb{R}^{\tilde{A}\tilde{X}}$ given by

$$\mathbf{F}(z, \tilde{p}^*) = -CP_Q z - Cr_Q - (\mathcal{H}_{1,-J} - M_{-J}\mathcal{H}_{2,-J}) b_{-J}(p) + \tilde{b}_{-J}(\tilde{p}^*) - g_{-J} + \tilde{M}_{-J}g_J,$$

where $\text{int}(\tilde{\mathbf{P}}^*)$ is the interior of $\tilde{\mathbf{P}}^*$. Clearly, the model and counterfactual restrictions impose $\mathbf{F}(z, \tilde{p}^*) = 0$, for all $z \in \mathcal{Z}$.

The Jacobian of \mathbf{F} is given by $\nabla \mathbf{F} = \left[\frac{\partial \mathbf{F}}{\partial z}, \frac{\partial \mathbf{F}}{\partial \tilde{p}^*} \right]$, with

$$\begin{aligned} \frac{\partial \mathbf{F}}{\partial z} &= -CP_Q, \\ \frac{\partial \mathbf{F}}{\partial \tilde{p}^*} &= \frac{\partial \tilde{b}_{-J}}{\partial \tilde{p}^*}. \end{aligned}$$

Because $\frac{\partial \tilde{b}_{-J}}{\partial \tilde{p}^*}$ is everywhere invertible (see KSS), the implicit function theorem applies. Specifically, for a point $(z^0, \tilde{p}^{*0}) \in \mathbb{R}^{X-d} \times \text{int}(\tilde{\mathbf{P}}^*)$ satisfying $\mathbf{F}(z^0, \tilde{p}^{*0}) = 0$, there exist open sets $U \subseteq \mathbb{R}^{X-d}$ and $W \subseteq \text{int}(\tilde{\mathbf{P}}^*)$ such that $z^0 \in U$ and $\tilde{p}^{*0} \in W$, and there exists a continuously differentiable function $\varphi : U \rightarrow W$ satisfying $\tilde{p}^{*0} = \varphi(z^0)$ and that

$$\mathbf{F}(z, \varphi(z)) = 0,$$

for all $z \in U$. Furthermore,

$$\frac{\partial \varphi(z)}{\partial z} = - \left[\frac{\partial \mathbf{F}}{\partial \tilde{p}^*} \right]^{-1} \frac{\partial \mathbf{F}}{\partial z} = \left[\frac{\partial \tilde{b}_{-J}}{\partial \tilde{p}^*} \right]^{-1} CP_Q.$$

The rank of the matrix $\frac{\partial \varphi(z)}{\partial z}$ equals the rank of CP_Q because $\frac{\partial \tilde{b}_{-J}}{\partial \tilde{p}^*}$ is invertible everywhere. Let $\text{rank}(CP_Q) = k$. By the Rank Theorem, the image set of φ is a differentiable k -dimensional manifold in $\text{int}(\tilde{\mathbf{P}}^*)$ (see Theorem 3.5.1 in Krantz and Parks, 2003). Clearly, by restricting z to the convex polyhedron \mathcal{Z} , the image set of φ becomes a k -dimensional manifold with boundary. In the absence of the model restrictions (9), we have $\pi_J = z$ and so the image set of φ becomes a manifold with dimension that equals the rank of C .

We can construct a global function $\bar{\varphi}$ defined on the entire domain \mathcal{Z} based on the local function φ defined above. To do so, we need to show that the constructed $\bar{\varphi}$ is not a set-function on \mathcal{Z} . I.e., if for any pair of points (z^0, \tilde{p}^{*0}) and (z^0, \tilde{p}^{*1}) with $z^0 \in \mathcal{Z}$ and $\tilde{p}^{*0}, \tilde{p}^{*1} \in \text{int}(\tilde{\mathbf{P}}^*)$, if $\bar{\varphi}(z^0) = \tilde{p}^{*0}$ and $\bar{\varphi}(z^0) = \tilde{p}^{*1}$, then we must have $\tilde{p}^{*0} = \tilde{p}^{*1}$. Suppose by contradiction that there exist implicit functions φ^0 and φ^1 defined locally on the neighborhood of the points (z^0, \tilde{p}^{*0}) and (z^0, \tilde{p}^{*1}) such that $\tilde{p}^{*0} = \varphi^0(z^0)$ and $\tilde{p}^{*1} = \varphi^1(z^0)$, with $\tilde{p}^{*0} \neq \tilde{p}^{*1}$. Next, recall that for any point $z^0 \in \mathcal{Z}$, there exists only one vector of payoffs $\pi(z^0) = [\pi'_{-J}(z^0), \pi'_J(z^0)]'$ satisfying all model restrictions: This vector is given by the elements

$\pi_{-J}(z^0) = M_{-J}\pi_J(z^0) + b_{-J}$, and $\pi_J(z^0)$ defined by (A3). This leads to the counterfactual payoff $\tilde{\pi}(z^0)$, which is given by the affine function $\tilde{\pi}(z^0) = \mathcal{H}\pi(z^0) + g$. Finally, the counterfactual payoff $\tilde{\pi}(z^0)$ can generate just one conditional choice probability function in the counterfactual scenario (by the uniqueness of the solution of the Bellman equation). We therefore must have $\tilde{p}^{*0} = \tilde{p}^{*1}$ (as well as $\varphi^0 = \varphi^1 = \varphi$). The global function $\bar{\varphi}$ equals the local implicit functions everywhere.³³

We conclude that the identified set $\tilde{\mathbf{P}}^I$ is the image set of the global function $\bar{\varphi}$, defined on the domain \mathcal{Z} . Consequently, $\tilde{\mathbf{P}}^I$ is a manifold with boundary and with dimension in the interior given by the rank of CP_Q . Further, $\tilde{\mathbf{P}}^I$ is connected because $\bar{\varphi}$ is a continuous function defined on the convex domain \mathcal{Z} . In addition, when Π^I is bounded, so is the closed set \mathcal{Z} , which implies that $\bar{\varphi}(\mathcal{Z})$ is compact. Finally, we have $\text{rank}(CP_Q) \leq X - d$ because $\text{rank}(C) \leq \min\{\tilde{A}\tilde{X}, X\}$ and $\text{rank}(P_Q) = X - d$.

A.2 Proof of Proposition 2

The identified set Θ^I defined in (19) is sharp by construction. We can construct payoff vectors satisfying all model restrictions, denoted by $\pi(z)$, and obtain the counterfactual CCP from the function $\tilde{p}^* = \bar{\varphi}(z)$, where $\bar{\varphi}$ is continuously differentiable, $z \in \mathcal{Z}$, and \mathcal{Z} is defined in (A5), as explained in the proof of Proposition 1. We have therefore

$$\theta = \phi(\tilde{p}, \pi) = \phi(\bar{\varphi}(z), \pi(z)) = \bar{\phi}(z),$$

where we omit p from the notation for simplicity. When the function ϕ is continuous, so is the function $\bar{\phi}$ because $\bar{\varphi}(z)$ and $\pi(z)$ are both continuous. Clearly, Θ^I equals the image set of the function $\bar{\phi}$ defined on the domain \mathcal{Z} . The image set is connected because \mathcal{Z} is convex, and it becomes compact when \mathcal{Z} is compact (which happens when Π^I is bounded, see the proof of Proposition 1). Furthermore, when θ is a scalar, the connected set Θ^I becomes an interval.

A.3 Proof of Proposition 3 and Implementation Details

To extend the argument from Proposition 1 we need to prove two results: first, that the function $\bar{\varphi}(z|\beta, \lambda)$ is jointly continuous on z and (β, λ) , where $\bar{\varphi}(z|\beta, \lambda)$ corresponds to the function $\bar{\varphi}(z)$ defined in Proposition 1 given the parameters β and λ . Second, that the set $\bigcup_{(\beta, \lambda) \in \mathcal{B} \times \Lambda} \mathcal{Z}|\beta, \lambda$ is connected, where $\mathcal{Z}|\beta, \lambda$ denotes the set \mathcal{Z} in Proposition 1 given β and λ . If that is the case, then the image set of $\bigcup_{(\beta, \lambda) \in \mathcal{B} \times \Lambda} \mathcal{Z}|\beta, \lambda \times \{\beta, \lambda\}$ under the function $\bar{\varphi}$ is a connected set.

Joint continuity of $\bar{\varphi}(z|\beta, \lambda)$ follows because β affects the Bellman contraction mapping continuously provided it is bounded away from 1, and because of Theorem 4.1 in Rust (1988), as V is continuously

³³While different z 's can generate the same \tilde{p}^* (because the function φ is not one-to-one, which is at the heart of the identification problem in dynamic discrete choice models), a single z cannot generate more than one \tilde{p}^* .

differentiable with respect to λ and a discrete state-space implies his assumptions (A11) and (A12).³⁴

The same conditions imply that $b_{-J}(p|\beta, \lambda)$ are continuous in (β, λ) . Therefore, if $\mathcal{B} \times \Lambda$ is connected, then $\bigcup_{(\beta, \lambda) \in \mathcal{B} \times \Lambda} \mathcal{Z}_{|\beta, \lambda}$ is connected as well because $\mathcal{Z}_{|\beta, \lambda}$ depends on $b_{-J}(p|\beta, \lambda)$ continuously. It follows that the set $\bigcup_{(\beta, \lambda) \in \mathcal{B} \times \Lambda} \mathcal{Z}_{|\beta, \lambda} \times \{\beta, \lambda\}$ is connected as the cartesian product of connected sets is connected.

Implementation. We now briefly discuss some practical issues regarding the computation of the identified sets. To find the bounds for a scalar θ in practice, we need to solve the (augmented) constrained optimization problem:

$$\theta^U \equiv \max_{(\tilde{p}, \pi, \beta, \lambda) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1) \times X} \times \mathcal{B} \times \Lambda} \phi(\tilde{p}, \pi, \beta, \lambda; p) \quad (\text{A7})$$

subject to

$$\begin{aligned} \mathbf{M}(\beta) \pi &= b_{-J}(p; \beta, \lambda), \\ R^{eq} \pi &= r^{eq}, \\ R^{iq} \pi &\leq r^{iq}, \\ (\tilde{\mathbf{M}}(\beta) \mathcal{H}) \pi &= \tilde{b}_{-J}(\tilde{p}; \beta, \lambda) - \tilde{\mathbf{M}}(\beta) g. \end{aligned} \quad (\text{A8})$$

As before, the lower bound is obtained similarly, by replacing max by min. (Note that we subsumed in the notation for $\tilde{\mathbf{M}}(\beta)$ and $\tilde{b}_{-J}(\tilde{p}; \beta, \lambda)$ the fact that $\tilde{\beta}$ and \tilde{G} are continuous functions of β and G .)

The optimization (A7)–(A8) is well-behaved and can be solved using standard software, as before. Yet, two additional computational challenges arise, besides the fact that the search now must be performed on a larger space. First, the discount factor β enters the constraints through the matrix $\mathbf{M}(\beta)$, which stacks M_a for all $a \neq J$, which in turn require calculating inverse matrices for any given β ; see equation (3). This challenge can be avoided when action J is a renewal or terminal action (or under general forms of finite dependence), because M_a simplifies and does not involve an inverse matrix.³⁵ (The same argument holds for the matrix $\tilde{\mathbf{M}}(\beta)$.) Second, for flexible distributions $G(\varepsilon; \lambda)$ for which there is no closed-form expression for ψ_a , one may need to use the linear program algorithm proposed by Chiong, Galichon, and Shum (2016) in order to calculate ψ_a (and obtain the corresponding b_{-J} and \tilde{b}_{-J}) for each given λ .

A.4 Proof of Theorem 1

First we introduce conditions used in the statement of Theorem 1. Formally, the (populational) minimization problem (24) projects $b_{-J}(\mathbf{p})$ on the manifold $S(\mathbf{p}, \theta)$ under the weighted norm $\|x\|_{\Omega} = x' \Omega x$,

³⁴The other assumptions in Rust's (1988) Theorem 4.1 are already satisfied by our model or hold vacuously. See the discussion in Rust (1988), especially page 1015, and Norets (2010) for a more general result.

³⁵When action J is either a renewal or a terminal action, then for all $a, j \in \mathcal{A}$, $F_a F_j = F_j F_a$, which implies that $M_a = I + \beta(F_j - F_a)$, for all $a \in \mathcal{A}$. The matrix M_a is linear in β under more general forms of finite dependence.

for $x \in \mathbb{R}^{AX}$, where

$$S(\mathbf{p}, \theta) := \{\mathbf{M}(\mathbf{p})\pi, \pi \in \mathbb{R}^{(A+1)X} : \mathcal{R}(\theta, \pi, \tilde{p}; \mathbf{p}) = 0, \text{ and } R^{iq}\pi \leq r^{iq} \text{ hold for some } \tilde{p} \in \tilde{\mathbf{P}}\}.$$

The value of $J(\theta)$ is the squared length of the projection residual vector.

It is useful to impose a mild requirement on $S(\mathbf{p}, \theta)$ in terms of its local geometric property. To this end, we introduce the notion of tangent cone:

Definition 1. For a (possibly non-convex) set $A \subset \mathbb{R}^d$, the tangent cone of A at $x \in A$, henceforth denoted by $T_A(x)$, is given by

$$T_A(x) := \limsup_{\tau \downarrow 0} \tau^{-1}(A \ominus x),$$

where \ominus denotes the usual Minkowski difference.

See, e.g., Section 6A of Rockafellar and Wets (2009) for a discussion on the role of a tangent cone and other related concepts.

Condition 1. (i) For every $(\mathbf{p}, \theta_0) \in \mathcal{P}$, the tangent cone $T_{S(\mathbf{p}, \theta_0)}(x)$ of $S(\mathbf{p}, \theta_0)$ is convex at each $x \in \mathbb{R}^{AX} \in S(\mathbf{p}, \theta_0)$. (ii) The tuning parameter sequence h_N is chosen such that $h_N \rightarrow \infty$ and $h_N/N \rightarrow 0$ as $N \rightarrow \infty$. (iii) Both ϕ and h^s are continuously differentiable functions.

Now we prove Theorem 1. Consider a sequence $\{(\mathbf{p}_N, \theta_N) \in \mathcal{P}, N \in \mathbb{N}\}$. Recall that p and F are determined by \mathbf{p} . Let $(p_N, F_N) := (p(\mathbf{p}_N), F(\mathbf{p}_N))$. In what follows we use symbols such as $S_N, \widehat{S}_N, \bar{\Pi}_N, (\bar{V}_N, V_N, v), (\bar{W}_N, W_N, w), B, \mu_N, \binom{\eta}{\zeta}$ and Σ (and their appropriate subsample counterparts with an asterisk symbol $*$ in superscript) while omitting their dependence on θ_N to ease the notational burden in the proof.

Let $S_N := S(\mathbf{p}_N, \theta_N)$ and $\widehat{S}_N := S(\widehat{\mathbf{p}}_N, \theta_N)$, where

$$S(\mathbf{p}, \theta) := \{\mathbf{M}(\mathbf{p})\pi, \pi \in \mathbb{R}^{(A+1)X} : \mathcal{R}(\theta, \pi, \tilde{p}; \mathbf{p}) = 0, \text{ and } R^{iq}\pi \leq r^{iq} \text{ hold for some } \tilde{p} \in \tilde{\mathbf{P}}\}.$$

Then writing $\|x\|_{\Omega}^2 := x'\Omega x$ for $x \in \mathbb{R}^{AX}$,

$$\begin{aligned} N\widehat{J}_N(\theta_N) &= \min_{x \in \widehat{S}_N} N \|b_{-J}(\widehat{\mathbf{p}}_N) - x\|_{\widehat{\Omega}_N}^2 \\ &= \min_{x \in \widehat{S}_N} \|\sqrt{N}[b_{-J}(\widehat{\mathbf{p}}_N) - b_{-J}(\mathbf{p}_N)] - \sqrt{N}[x - b_{-J}(\mathbf{p}_N)]\|_{\widehat{\Omega}_N}^2 \\ &= \min_{\xi \in \sqrt{N}(\widehat{S}_N \ominus b_{-J}(\mathbf{p}_N))} \|\sqrt{N}[b_{-J}(\widehat{\mathbf{p}}_N) - b_{-J}(\mathbf{p}_N)] - \xi\|_{\widehat{\Omega}_N}^2, \end{aligned} \tag{A9}$$

where \ominus denotes the usual Minkowski difference, and for $c \in \mathbb{R}_{++}$ and a set $A \in \mathbb{R}^d$, we let cA denote the set A dilated by the factor c , that is, $\{cx : x \in A\}$.

To show the theorem, it suffices to consider sequences $\mathbf{p}_N, N \in \mathbb{N}$ such that

- (i) $\inf_{x \in \text{bdy}(S_N)} \|b_{-J}(\mathbf{p}_N) - x\|_\Omega = O(1/\sqrt{N})$, where $\text{bdy}(S_N)$ is the boundary of S_N , and
- (ii) Each sequence $\{\mathbf{p}_N, N = 1, 2, \dots\}$ converges.

Suppose $\mathbf{p}_N, N \in \mathbb{N}$ satisfies (i) and (ii). The restrictions imposed on \mathcal{P} guarantee that along the sequence \mathbf{p}_N it holds that

$$\sqrt{N}[b_{-J}(\widehat{\mathbf{p}}_N) - b_{-J}(\mathbf{p}_N)] \xrightarrow{d} \eta,$$

where η is a zero mean Gaussian vector. In what follows we also use the following notation: for finite sets $V, W \subseteq \mathbb{R}^d$ we let $\text{conv}(V)$ and $\text{cone}(W)$ denote the convex hull of V and the cone spanned by W , respectively; then the Minkowski sum $\text{conv}(V) \oplus \text{cone}(W)$ is a polyhedron. We approximate the last term in equation (A9) following Chernoff (1954). Under Condition 1 we have:

$$N\widehat{J}_N(\theta_N) \stackrel{d}{=} \min_{\xi \in \bar{\Pi}_N} \|\eta - \xi\|_\Omega^2 + o_p(1), \quad (\text{A10})$$

where $\bar{\Pi}_N = \text{conv}(\bar{V}_N) \oplus \text{cone}(\bar{W}_N)$ is a random polyhedron, with $\bar{V}_N = V_N + v$, $V_N \in \mathbb{R}^{AX \times m}$, $\bar{W}_N = W_N + w$, $W_N \in \mathbb{R}^{AX \times n}$, and v and w are $\mathbb{R}^{AX \times m}$ -valued and $\mathbb{R}^{AX \times n}$ -valued zero-mean Gaussian random matrices, respectively, for some $m, n \in \mathbb{N}$. Note that the estimation uncertainty in \widehat{S}_N makes the polyhedron $\bar{\Pi}_N$ that appears in the asymptotic approximation (A10) random. Also define a (deterministic) sequence of polyhedra $\Pi_N = \text{conv}(V_N) \oplus \text{cone}(W_N)$. By the representation theorem for polyhedra (see, for example, Theorem 1.2 in Ziegler (2012)) we can write

$$\Pi_N = \{\xi : B\xi \leq \mu_N\} \text{ for some } B \in \mathbb{R}^{\ell \times AX},$$

where $\mu_N \geq 0$ for all N and $\mu_N = O(1)$.

Recalling that each transition matrix $F_a, a \in \mathcal{A}$, depends on \mathbf{p}_N (as so does F), write

$$\det(M_a(\mathbf{p}_N)) = \det\left((1 - \beta F_a(\mathbf{p}_N))(1 - \beta F_J(\mathbf{p}_N))^{-1}\right) = \frac{\det(I - \beta F_a(\mathbf{p}_N))}{\det(I - \beta F_J(\mathbf{p}_N))}.$$

Let $\{\lambda_a^i(\mathbf{p}_N)\}$ and $\{\lambda_J^i(\mathbf{p}_N)\}$ be the eigenvalues of $F_a(\mathbf{p}_N)$ and $F_J(\mathbf{p}_N)$, then

$$\begin{aligned} \det(M_a(\mathbf{p}_N)) &= \frac{\det(\beta^{-1}I - F_a(\mathbf{p}_N))}{\det(\beta^{-1}I - F_J(\mathbf{p}_N))} \\ &= \frac{\prod_{i=1}^X (\beta^{-1} - \lambda_a^i(\mathbf{p}_N))}{\prod_{i=1}^X (\beta^{-1} - \lambda_J^i(\mathbf{p}_N))} \\ &> c, \quad \text{for every } a \in \mathcal{A} \text{ and every } N \in \mathbb{N} \end{aligned} \quad (\text{A11})$$

holds for some $c > 0$ that does not depend on N as β is fixed in the unit interval $(0, 1)$ and $\{\lambda_a^i(\mathbf{p}_N)\}$ and $\{\lambda_J^i(\mathbf{p}_N)\}$ are inside the unit circle for every N .

Note that the approximation (A10) holds for any sequence $\{V'_N, W'_N\}_{N \in \mathbb{N}}$ such that $V'_N = V_N + o(1)$

and $W'_N = W_N + o(1)$, and with Condition 1 and (A11) we can choose $\{V_N, W_N\}_{N \in \mathbb{N}}$ such that the matrix B above does not depend on N . Then we have an alternative representation for the random polyhedron $\bar{\Pi}_N$ as well: for some positive definite matrix Σ it holds that

$$\bar{\Pi}_N = \{\xi : B\xi \leq \mu_N + \zeta\},$$

where the vector $\begin{pmatrix} \eta \\ \zeta \end{pmatrix} \sim N(0, \Sigma)$. In sum, we have

$$N\widehat{J}_N(\theta_N) \stackrel{d}{=} \min_{\xi: B\xi \leq \mu_N + \zeta} \|\eta - \xi\|_{\Omega}^2 + o_p(1). \quad (\text{A12})$$

Next we turn to the subsample statistic $\widehat{J}_{h_N}^*(\theta_N)$. To show the uniform validity of subsampling we can instead analyze the asymptotic behavior of the statistic \widehat{J}_{h_N} , the \widehat{J} -statistic calculated from a random sample of size h_N , drawn according to \mathbf{p}_N (Romano and Shaikh, 2012). That is, we now study the limiting behavior of the CDF $G_{h_N}(x, \mathbf{p}_N)$, $N = 1, 2, \dots$, where $G_{\ell}(x, \mathbf{p}) := \Pr_{\mathbf{p}}\{\ell\widehat{J}_{\ell}(\theta_N) \leq x\}$ for $\ell \in \mathbb{N}$. Then proceeding as before, along the sequence \mathbf{p}_N we have

$$h_N\widehat{J}_{h_N}(\theta_N) \stackrel{d}{=} \min_{\xi \in \bar{\Pi}_{h_N, N}^*} \|\eta^* - \xi\|_{\Omega}^2 + o_p(1), \quad (\text{A13})$$

where $\bar{\Pi}_{h_N, N}^* = \text{conv}(\bar{V}_{h_N, N}^*) \oplus \text{cone}(\bar{W}_{h_N, N}^*)$, with $\bar{V}_{h_N, N}^* = V_{h_N, N}^* + v^*$, $V_{h_N, N}^* \in \mathbb{R}^{AX \times m}$, $\bar{W}_{h_N, N}^* = W_{h_N, N}^* + w^*$, $W_{h_N, N}^* \in \mathbb{R}^{AX \times m}$, and η^* , v^* and w^* are zero-mean Gaussian random elements taking values in \mathbb{R}^{AX} , $\mathbb{R}^{AX \times m}$ and $\mathbb{R}^{AX \times n}$ with $(\eta^*, v^*, w^*) \stackrel{d}{=} (\eta, v, w)$. Define $\Pi_{h_N, N}^* = \text{conv}(V_{h_N, N}^*) \oplus \text{cone}(W_{h_N, N}^*)$ and observe that it has a half-space based representation $\Pi_{h_N, N}^* = \{\xi : B\xi \leq \sqrt{\frac{h_N}{N}}\mu_N\}$. We now have

$$\bar{\Pi}_N^* = \left\{ \xi : B\xi \leq \sqrt{\frac{h_N}{N}}\mu_N + \zeta^* \right\}.$$

Recall that $\mu_N = O(1)$, and moreover, we have $\begin{pmatrix} \eta^* \\ \zeta^* \end{pmatrix} \sim N(0, \Sigma)$. Therefore

$$h_N\widehat{J}_{h_N}(\theta_N) \xrightarrow{d} \min_{\xi: B\xi \leq \zeta} \|\eta - \xi\|_{\Omega}^2. \quad (\text{A14})$$

In sum, for every sequence \mathbf{p}_N , $N \in \mathbb{N}$ satisfying conditions (i) and (ii) above, by (A12) and (A14) and noting $\mu_N \geq 0$ for every N , we have

$$\limsup_{N \rightarrow \infty} \sup_x (G_{h_N}(x, \mathbf{p}_N) - G_N(x, \mathbf{p}_N)) \leq 0.$$

We can now invoke Theorem 2.1 in Romano and Shaikh (2012) to conclude.

B Gradient of ϕ involving Ergodic Distribution

In this section, we show how to calculate the gradient of the function ϕ when it involves counterfactual average effects based on ergodic distributions of the state variables.

Assume the function ϕ is given by

$$\phi(\tilde{p}, \pi; p) = \sum_{x \in \tilde{\mathcal{X}}} \tilde{Y}(x; \pi) \tilde{f}^*(x) - \sum_{x \in \mathcal{X}} Y(x; \pi) f^*(x),$$

where $\tilde{Y}(x; \pi)$ and $Y(x; \pi)$ are outcome variables of interest in the counterfactual and baseline scenarios and that may depend on baseline payoffs π (e.g., consumer surplus, or the firm value); and $\tilde{f}^*(x)$ and $f^*(x)$ are the counterfactual and baseline steady-state distributions of the state variables (which in turn depend on state transitions and choice probabilities, as shown below). Note that all target parameters θ presented in our firm entry/exit example (discussed in details in Online Appendix G) and in our Monte Carlo study (Section E of this appendix) are of this type; in the empirical application of Section 6 in the main text, we take the ratio of objects of this type.

The term $\tilde{f}^*(x)$ is the ergodic distribution of the (endogenous) Markovian process for the state variables

$$\tilde{F}(x'|x) = \sum_{a \in \tilde{\mathcal{A}}} \tilde{p}(a|x) \tilde{F}(x'|a, x).$$

(A similar expression holds for $f^*(x)$.) In matrix notation, we have

$$\phi(\tilde{p}, \pi; p) = \tilde{\mathbf{Y}}' \tilde{f}^* - \mathbf{Y}' f^*,$$

where $\tilde{\mathbf{Y}}$ and \mathbf{Y} are vectors of the outcome variables in the counterfactual and baseline; and \tilde{f}^* is the vector of the ergodic distribution satisfying the steady-state condition

$$\tilde{f}^{*'} = \tilde{f}^{*'} \tilde{\mathbf{F}}, \tag{B15}$$

where

$$\tilde{\mathbf{F}} = \sum_{a \in \tilde{\mathcal{A}}} \tilde{P}_a \tilde{F}_a, \tag{B16}$$

and \tilde{P}_a is a diagonal matrix with \tilde{p}_a in its diagonal, and \tilde{F}_a is the counterfactual transition matrix conditional on the choice a . (Again, a similar expression holds for f^* .) Importantly, the ergodic distribution \tilde{f}^* depends directly on \tilde{p} (through equations (B15)–(B16)), and indirectly on the baseline payoff π , since \tilde{p} depends on π through equation (16) presented in the main text.

We want to know the derivative of ϕ with respect to π and \tilde{p} , holding all other arguments of ϕ constant (e.g., the baseline CCP p and the state transitions F). We focus on the derivative with respect to π ;

obtaining the derivative with respect to \tilde{p} is similar but simpler, and is therefore omitted. Clearly, we have

$$\frac{\partial \phi}{\partial \pi'} = \left(\tilde{f}^{*'} \frac{\partial \tilde{\mathbf{Y}}}{\partial \pi'} \right) - \left(f^{*'} \frac{\partial \mathbf{Y}}{\partial \pi'} \right) + \left(\tilde{\mathbf{Y}}' \frac{\partial \tilde{f}^*}{\partial \pi'} \right).$$

The derivatives $\frac{\partial \tilde{\mathbf{Y}}}{\partial \pi'}$ and $\frac{\partial \mathbf{Y}}{\partial \pi'}$ depend on the specific outcome of interest. Here, we focus on the third term of the right-hand-side, $\frac{\partial \tilde{f}^*}{\partial \pi'}$. By the chain rule, we have

$$\frac{\partial \tilde{f}^*}{\partial \pi'} = \frac{\partial \tilde{f}^*}{\partial \tilde{p}'} \frac{\partial \tilde{p}}{\partial \pi'}.$$

By equation (16), we know that

$$\frac{\partial \tilde{p}}{\partial \pi'} = \left(\frac{\partial \tilde{b}_{-J}}{\partial \tilde{p}'} \right)^{-1} \tilde{\mathbf{M}} \mathcal{H}.$$

See Section C.2 of this appendix for a discussion of how one can take advantage of the structure of \tilde{b}_{-J} to reduce the cost of inverting $\left(\frac{\partial \tilde{b}_{-J}}{\partial \tilde{p}'} \right)$ in practice.

We now derive the remaining term $\frac{\partial \tilde{f}^*}{\partial \tilde{p}'}$. Let x, x', \tilde{x} be arbitrary states and $\tilde{a} \neq J$. Then (B15) pointwise becomes

$$\tilde{f}^*(x') = \sum_x \tilde{f}^*(x) \sum_a \tilde{p}_a(x) \tilde{F}(x'|x, a).$$

Therefore,

$$\frac{\partial \tilde{f}^*(x')}{\partial \tilde{p}_{\tilde{a}}'} = \sum_x \frac{\partial \tilde{f}^*(x)}{\partial \tilde{p}_{\tilde{a}}'} \tilde{F}(x'|x) + \tilde{f}^*(\tilde{x}) \left[\tilde{F}(x'|\tilde{x}, \tilde{a}) - \tilde{F}(x'|\tilde{x}, J) \right].$$

This is written compactly in matrix as,

$$\frac{\partial \tilde{f}^*}{\partial \tilde{p}'_a} = -(\tilde{\mathbf{F}}' - I)^+ (\tilde{F}'_a - \tilde{F}'_J) \tilde{\mathbf{f}}^*, \quad (\text{B17})$$

where $(\tilde{\mathbf{F}}' - I)^+$ is the pseudo-inverse of $(\tilde{\mathbf{F}}' - I)$, and $\tilde{\mathbf{f}}^*$ is a diagonal matrix with \tilde{f}^* in its diagonal.

C A Proposed Algorithm Based on Stochastic Search

In this section, we discuss practical and computational aspects of calculating the identified set of low-dimensional outcomes of interest θ . As explained in the main text, in our experience, standard solvers are highly efficient in solving (20)–(21) when the researcher can provide the gradient of ϕ . However, when numerical (or analytical) gradients are costly to evaluate in practice, standard solvers can be slow in converging to the maximum (again, in our experience). For such cases, we propose a stochastic algorithm that exploits the structure of the problem and combines the strengths of alternative stochastic search procedures, as we explain below.

Our proposed algorithm builds upon a couple of observations. First, while a search over π to maximize ϕ is feasible, it is computationally costly and takes a long time to converge when it is expensive to calculate the gradient of ϕ numerically. In high-dimensional problems, this may become impractical. This procedure searches over the admissible values that π can take, and, for each candidate, it finds the corresponding counterfactual CCP by solving the nonlinear equation (16), and then it evaluates ϕ – and its (numerical) derivative, to obtain updated directions for π – until reaching the maximum value for θ . Although finding admissible values for π is not difficult in high-dimensional problems (as it only depends on linear constraints), and solving the nonlinear equation (16) once is not computationally costly (as standard quasi-Newton methods can be used to find \tilde{p}), solving (16) too many times and calculating the gradient of ϕ numerically can be demanding. Unless the econometrician imposes a sufficient number of assumptions to make π effectively a low dimensional vector (e.g., 3-dimensional or smaller), this method takes a long time to converge, as it requires too many evaluations before we can increase θ substantially in the direction of its maximum.

Second, it is possible to perform a search over \tilde{p} , instead of over π , to calculate θ^U . For any given \tilde{p} , existence of a π satisfying linear constraints is computationally cheap (for example, existence can be easily checked as a solution to a linear programming problem). If there is no such π satisfying all restrictions, we discard \tilde{p} , since it does not belong to the identified set $\tilde{\mathbf{P}}^I$. If there exists *some* π satisfying the restrictions, we keep \tilde{p} , and compute the corresponding θ . This approach may be particularly useful when ϕ is not a direct function of π , in which case it is not necessary to find a particular π to calculate θ – existence of *some* π suffices. The difficulty here is that, while an exhaustive grid search over \tilde{p} can be used to find the maximum θ^U , grid search is unfeasible for empirically-relevant high-dimensional problems. An alternative would be to perform a stochastic search (to find good directions for \tilde{p}).³⁶ Yet, and more importantly, the random search must be performed on the $\tilde{A}\tilde{X}$ -space $\tilde{\mathbf{P}}$, while the identified set $\tilde{\mathbf{P}}^I$ can be of much smaller dimension: $X - d$, or smaller (depending on the rank of CP_Q ; see Proposition 1). In other words, $\tilde{\mathbf{P}}^I$ may be a “thin” set in $\tilde{\mathbf{P}}$. The combination of a “thin” set with an unknown shape (recall that \tilde{p} is a nonlinear function of π – see equation (16)) makes it difficult to find points within that set randomly. Further, it is easy for perturbation methods to “exit” the set, increasing the costs of finding the maximum θ . Note that searching over π to maximize θ does not suffer from this problem because finding admissible values (and updated directions) for π are computationally easier.

These trade-offs led us to consider an algorithm that exploits the structure of the problem and combines the strengths of these alternative search procedures. Intuitively, we move in the “ \tilde{p} -world” (to avoid solving the nonlinear equation (16) repeatedly), but we keep a close eye on the “ π -world” (to keep track of the

³⁶For instance, one possibility is to perturb \tilde{p} completely randomly ($\tilde{p} + \varepsilon$) and check whether the perturbed vector lies in the identified set $\tilde{\mathbf{P}}^I$ (or within a tolerance level) – where checking this amounts to checking existence of π satisfying linear restrictions, as mentioned above. We then keep the perturbed \tilde{p} 's that deliver large values for θ (and perturb them further), and discard those with small values of θ . We iterate until θ cannot be increased any longer. (This is similar to genetic algorithm, or to stochastic search methods more generally.)

model restrictions and search in relevant directions). Searching in relevant directions without solving (16) and computing the numerical gradient of ϕ in every step improves substantially how fast θ moves on each iteration to the maximum.

C.1 The Proposed Stochastic Search Algorithm

We now present our proposed algorithm. To guarantee that \tilde{p} are positive and add up to one, we work with the transformation

$$\tilde{\delta} = \ln \tilde{p}_{-J} - \ln \tilde{p}_J,$$

where $\ln \tilde{p}_a$ is the $\tilde{X} \times 1$ vector with elements $\ln \tilde{p}_a(x)$, for all $x \in \tilde{\mathcal{X}}$, and $\ln \tilde{p}_{-J}$ stacks $\ln \tilde{p}_a$ for all $a \neq J$. The functions of \tilde{p} , namely \tilde{b}_{-J} and ϕ , are adjusted accordingly. To simplify the notation, we drop the subscript of the function \tilde{b}_{-J} , as well as the argument p of the function ϕ . The algorithm proceeds as follows (we discuss the most important steps below):

The Proposed Stochastic Search Algorithm:

1. Initialize $k = 0 \in \mathbb{N}$.

Set π^k satisfying $R^{eq}\pi^k = r^{eq}$ and $R^{iq}\pi^k \leq r^{iq}$. Find $\tilde{\delta}^k$ by solving (16) with π^k .

Calculate $\theta^k = \phi(\tilde{\delta}^k, \pi^k)$.

2. Increment k .

3. Set (perturbed) direction $\Delta\pi^k$. Given $\Delta\pi^k$, set direction for $\tilde{\delta}^k$,

$$\Delta\tilde{\delta}^k = \left(\frac{\partial \tilde{b}}{\partial \tilde{\delta}} \right)^{-1} \tilde{\mathbf{M}} \mathcal{H} \Delta\pi^k,$$

where $\left(\frac{\partial \tilde{b}}{\partial \tilde{\delta}} \right)$ is the derivative of \tilde{b} with respect to $\tilde{\delta}$ evaluated at $\tilde{\delta}^k$.

4. Solve for $\alpha \in \mathbb{R}$:

$$\alpha^* = \operatorname{argmax}_{\alpha} \phi(\tilde{\delta}^k + \alpha \Delta\tilde{\delta}^k, \pi^k + \alpha \Delta\pi^k),$$

subject to the constraints (21), allowing (16) to be violated at most by a tolerance level, $\text{tol} > 0$.

5. Set $\tilde{\delta}^* = \tilde{\delta}^k + \alpha^* \Delta\tilde{\delta}^k$ and $\pi^* = \pi^k + \alpha^* \Delta\pi^k$.

6. Update $\tilde{\delta}^k$:

$$\tilde{\delta}^{k+1} = \tilde{\delta}^* - \left(\frac{\partial \tilde{b}^*}{\partial \tilde{\delta}} \right)^{-1} \left(\tilde{b}^* - \tilde{\mathbf{M}}g - \tilde{\mathbf{M}}\mathcal{H}\pi^* \right),$$

where $\left(\frac{\partial \tilde{b}^*}{\partial \tilde{\delta}}\right)$ is the derivative of \tilde{b} with respect to $\tilde{\delta}$ evaluated at $\tilde{\delta}^*$, and $\tilde{b}^* = \tilde{b}(\tilde{\delta}^*)$.
Set $\pi^{k+1} = \pi^*$.

7. Calculate $\theta^{k+1} = \phi(\tilde{\delta}^{k+1}, \pi^{k+1})$.

If $\|\theta^{k+1} - \theta^k\| \leq \epsilon$ go to 8; otherwise go to 2.

8. Set $\pi = \pi^{k+1}$. Solve (16) exactly for π , and get $\tilde{\delta}$. Return $\theta^U = \phi(\tilde{\delta}, \pi)$.

We now discuss the rationale for each step. In Subsection C.2, we provide further details for the implementation of each step, as well as a discussion about the overall cost of the algorithm.

Step 1. The first step requires finding a π that satisfies the model restrictions (9) and (10) so that we obtain an initial \tilde{p} (or $\tilde{\delta}$) that lies inside the (potentially “thin”) set $\tilde{\mathbf{P}}^I$ by construction, and a corresponding θ in the identified set Θ^I . Such initial π can be obtained as any solution to the following quadratic programming problem

$$\min_{\pi} (R^{eq}\pi - r^{eq})' (R^{eq}\pi - r^{eq}) + (R^{iq}\pi - r^{iq})'_+ (R^{iq}\pi - r^{iq})_+, \quad (\text{C1})$$

where $(x)_+ = \max\{x, 0\}$. Another option is to start with a few points and project them into the identified set for π , which can also be done easily. Of note, if the minimum of (C1) is strictly greater than zero, then there is no π that satisfies all the constraints. Given π , we can solve (16) numerically using some quasi-Newton method.

Step 3. After we have our starting point π (and corresponding $\tilde{\delta}$), we need to obtain an updated direction $\Delta\pi$ (and $\tilde{\Delta}\tilde{\delta}$). Overall, the idea of first providing a direction and only then optimizing (as we do here) is a standard way to solve complex optimization problems. Ideally, we would use the gradient of ϕ , but calculating this gradient can be expensive in some cases, as mentioned previously. An alternative is to either get a completely random direction for $\Delta\pi$ (e.g., $\Delta\pi = \eta$, where η is a random vector drawn from, say, a multivariate standard normal distribution), or a random direction weighted by states that are more important (e.g., in terms of the ergodic distribution of the state variables).³⁷

It is also important to not let an updated point get too close to the boundary of the inequality constraints (10). We follow the insights of interior-point methods to help the algorithm not get stuck early on a boundary. Specifically, we add a term to $\Delta\pi$ that moves it way from the most binding ones.

³⁷In practice, to weight the random direction η by states that are more important in terms of the steady-state distribution, we draw η from a normal distribution with zero mean and a diagonal variance-covariance matrix with a diagonal that equals the probabilities of the state variables under the ergodic distribution. The ergodic distribution is based on the latest updated \tilde{p} .

Formally,

$$\Delta\pi = \eta - \lambda \left(\frac{1}{r^{iq} - R^{iq}\pi} \right)' R^{iq},$$

where $\lambda = \frac{\lambda_0}{N}$, with $\lambda_0 > 0$ and $N =$ the number of iterations; and $\left(\frac{1}{r^{iq} - R^{iq}\pi} \right)$ denotes the $m \times 1$ vector with the reciprocal elements of the vector $r^{iq} - R^{iq}\pi$ (recall that m is the number of inequality restrictions, so that R^{iq} is $m \times (A + 1)X$ and r^{iq} is $m \times 1$). The adjustment term $\lambda \left(\frac{1}{r^{iq} - R^{iq}\pi} \right)' R^{iq}$ is a common way to handle inequality constraints. This is a simple implementation of an interior-point method.³⁸

We link the direction $\Delta\tilde{\delta}$ with $\Delta\pi$ based on equation (16). We do so because completely random directions on \tilde{p} (or more precisely, on $\tilde{\delta}$) will likely push \tilde{p} outside of the “thin” set $\tilde{\mathbf{P}}^I$. The direction $\Delta\tilde{\delta}$ is obtained by differentiating the inverse function \tilde{b}^{-1} with respect to π in the direction $\Delta\pi$.

Step 4. Given $\Delta\tilde{\delta}$, we now find how far in that direction we should go without moving away too much from the identified set $\tilde{\mathbf{P}}^I$. To that end, we allow for small violations in equation (16) when searching for α^* . Specifically, we replace the restriction (16) by $\left\| \tilde{b} - \tilde{\mathbf{M}}g - \tilde{\mathbf{M}}\mathcal{H}\pi \right\| \leq \text{tol}$, where $\|\cdot\|$ is some matrix norm and $\text{tol} > 0$ is a tolerance level. Here, the optimization is one-dimensional (line-search). We use a simple golden rule search, but even more crude approaches work.

Step 5. We now update both $\tilde{\delta}$ and π in their respective directions $\alpha^*\Delta\tilde{\delta}$ and $\alpha^*\Delta\pi_J$, where α^* is obtained in step 4.

Step 6. This step is important because at the end of step 5 it is common that the intermediary $\tilde{\delta}^*$ violates the nonlinear system (16) by the maximum tolerance tol . So this step insures that we move $\tilde{\delta}$ back to the set that violates (16) by strictly less than tol . Not doing so would constrain the directions that $\Delta\tilde{\delta}$ can move in the next iteration and slow down the algorithm considerably.

Step 7. The $\epsilon > 0$ in step 7 specifies the tolerance for convergence. We focus on convergence on θ because verifying a “derivative equals zero” condition for convergence is difficult given the high-dimensionality of the problem and the complexity of computing derivatives of ϕ (analytically or numerically).

Step 8. After convergence, we solve the nonlinear system (16) exactly to guarantee that \tilde{p} lies in the identified set $\tilde{\mathbf{P}}^I$, and so that the computed θ^U belongs to Θ^I .

³⁸Intuitively, to maximize $f(x)$ subject to $g(x) \leq 0$, an interior-point method can make use of the logarithmic “barrier function” $B(x, \lambda) = f(x) - \lambda \sum_{i=1}^n \log(g_i(x))$, where n is the dimension of g . The gradient of B is $\frac{\partial f}{\partial x} - \lambda \sum_{i=1}^n \frac{1}{g(x)} \frac{\partial g}{\partial x}$. The idea is that when some element $g_i(x)$ is close to zero for some trial x , the barrier function “explodes” to minus infinity, so that the algorithm does not get stuck on a boundary. However, because the solution may indeed lie on the boundary, it is necessary to allow for the possibility that $g_i(x) = 0$ at the optimum. To do so, λ must converge to zero as the number of iterations grows larger. In the present case, we take $\lambda = \frac{\lambda_0}{N} \rightarrow 0$ (as $N \rightarrow \infty$). The term $\left(\frac{1}{r^{iq} - R^{iq}\pi} \right)' R^{iq}$ is the derivative of the sum of the logs of $(r^{iq} - R^{iq}\pi)$ with respect to π (i.e., the derivative of $\lambda \sum \log(r^{iq} - R^{iq}\pi)$, where the summation runs from 1 to m).

One of the main computational cost of this algorithm is to calculate the inverse matrix $\left(\frac{\partial \tilde{b}}{\partial \tilde{\delta}}\right)^{-1}$, used in steps 3 and 6. In the next subsection we discuss conditions under which calculating $\left(\frac{\partial \tilde{b}}{\partial \tilde{\delta}}\right)^{-1}$ is not as costly.

C.2 Further Comments on Implementation

We now comment on the computational costs of the algorithm.

1. The matrix M_a equals $(I - \beta F_a)(I - \beta F_J)^{-1}$, which involves the inversion of a $X \times X$ matrix. The computational cost of inverting a matrix is of the order of $O(X^3)$ in general, but there are ways to reduce this cost. When action J is renewal or terminal, the matrix simplifies to $M_a = I + \beta(F_J - F_a)$, for all $a \in \mathcal{A}$, which can be calculated fast since it involves no matrix inversion (see footnote 35). When there are no terminal or renewal actions, computing M_a requires calculating the inverse of $(I - \beta F_J)$. As F_J is a transition matrix, we can approximate that inverse based on the geometric series:

$$(I - \beta F_J)^{-1} = \sum_{\tau=0}^{\infty} \beta^\tau F_J^\tau.$$

By truncating the series, we can reduce the computational cost and obtain a reasonable approximation (more on that below). Note that both M_a and \tilde{M}_a can be precomputed, so they do not add costs to the iterative procedure.

2. When we find the direction $\Delta \tilde{\delta}$ implied by $\Delta \pi$ we need to solve the linear system

$$\Delta \tilde{\delta} = \left(\frac{\partial \tilde{b}}{\partial \tilde{\delta}}\right)^{-1} \tilde{\mathbf{M}} \mathcal{H} \Delta \pi$$

Usually this would cost $O(A^3 X^3)$. However, we can take advantage of the structure of the function \tilde{b} . Recall that $\tilde{b}_a(\tilde{p}) = \tilde{M}_a \psi_J(\tilde{p}) - \psi_a(\tilde{p})$. To simplify notation, take $\tilde{\psi}_J = \psi_J(\tilde{p})$ and $\tilde{\psi}_a = \psi_a(\tilde{p})$, let $\tilde{\psi}_{-J}$ stack $\tilde{\psi}_a$ for all actions $a \neq J$. For expositional convenience, consider the case of three actions with reference action $J = 3$:

$$\tilde{b} = \tilde{M}_{-J} \tilde{\psi}_J - \tilde{\psi}_{-J} = \tilde{\delta} - \left(\begin{bmatrix} \mathbf{I} \\ \mathbf{I} \end{bmatrix} - \tilde{M}_{-J} \right) \log \left(1 + \sum_{j=1}^{J-1} \exp(\tilde{\delta}_j) \right),$$

where \mathbf{I} is the identity matrix. So

$$\frac{\partial \tilde{b}}{\partial \tilde{\delta}} = \mathbf{I} - \left(\begin{bmatrix} \mathbf{I} \\ \mathbf{I} \end{bmatrix} - \tilde{M}_{-J} \right) \begin{bmatrix} \tilde{P}_1 & \tilde{P}_2 \end{bmatrix},$$

where \tilde{P}_j is an $X \times X$ diagonal matrix with \tilde{p}_j as its entries.

Now we need its inverse

$$\begin{aligned} \left(\frac{\partial \tilde{b}}{\partial \tilde{\delta}} \right)^{-1} &= \mathbf{I} + \left(\begin{bmatrix} \mathbf{I} \\ \mathbf{I} \end{bmatrix} - \tilde{M}_{-J} \right) \left(\mathbf{I} - \begin{bmatrix} \tilde{P}_1 & \tilde{P}_2 \end{bmatrix} \left(\begin{bmatrix} \mathbf{I} \\ \mathbf{I} \end{bmatrix} - \tilde{M}_{-J} \right) \right)^{-1} \begin{bmatrix} \tilde{P}_1 & \tilde{P}_2 \end{bmatrix} \\ &= \mathbf{I} + \left(\begin{bmatrix} \mathbf{I} \\ \mathbf{I} \end{bmatrix} - \tilde{M}_{-J} \right) \left(\mathbf{I} - \tilde{P}_1 - \tilde{P}_2 + \tilde{P}_1 \tilde{M}_1 + \tilde{P}_2 \tilde{M}_2 \right)^{-1} \begin{bmatrix} \tilde{P}_1 & \tilde{P}_2 \end{bmatrix}, \end{aligned}$$

where the first equality follows from the Woodbury formula $(\mathbf{I} - DB)^{-1} = \mathbf{I} + D(\mathbf{I} - BD)^{-1}B$.

Next, notice that $\tilde{P}_J = \mathbf{I} - \tilde{P}_1 - \tilde{P}_2$ and that $\tilde{M}_j = (\mathbf{I} - \beta \tilde{F}_j)(\mathbf{I} - \beta \tilde{F}_j)^{-1}$. Therefore,

$$\begin{aligned} \left(\frac{\partial \tilde{b}}{\partial \tilde{\delta}} \right)^{-1} &= \mathbf{I} + \left(\begin{bmatrix} \mathbf{I} \\ \mathbf{I} \end{bmatrix} - \tilde{M}_{-J} \right) (\mathbf{I} - \beta \tilde{F}_J) \left(\tilde{P}_J (\mathbf{I} - \beta \tilde{F}_J) + \tilde{P}_1 (\mathbf{I} - \beta \tilde{F}_1) + \tilde{P}_2 (\mathbf{I} - \beta \tilde{F}_2) \right)^{-1} \begin{bmatrix} \tilde{P}_1 & \tilde{P}_2 \end{bmatrix} \\ &= \mathbf{I} + \beta \begin{bmatrix} \tilde{F}_1 - \tilde{F}_J \\ \tilde{F}_2 - \tilde{F}_J \end{bmatrix} \left(\mathbf{I} - \beta \left(\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2 \right) \right)^{-1} \begin{bmatrix} \tilde{P}_1 & \tilde{P}_2 \end{bmatrix}. \end{aligned}$$

This reduces the cost to $O(X^3)$ because the matrix

$$\left(\mathbf{I} - \beta \left(\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2 \right) \right)^{-1}$$

has dimension $X \times X$.

But we can improve on that by noticing that for a given vector v ,

$$\left(\mathbf{I} - \beta \left(\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2 \right) \right)^{-1} v = \sum_{\tau=0}^{\infty} \beta^\tau \left(\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2 \right)^\tau v.$$

Because $\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2$ is a transition matrix, we know that

$$\left(\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2 \right)^\tau v \rightarrow v^*$$

for some v^* as τ goes to infinity.³⁹ Therefore, we can approximate

$$\begin{aligned} \left(\mathbf{I} - \beta \left(\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2 \right) \right)^{-1} v &\approx \sum_{\tau=0}^{K-1} \beta^\tau \left(\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2 \right)^\tau v \\ &+ \frac{\beta^K}{1-\beta} \left(\tilde{P}_J \tilde{F}_J + \tilde{P}_1 \tilde{F}_1 + \tilde{P}_2 \tilde{F}_2 \right)^K v, \end{aligned}$$

which can be computed in $O(KX^2)$ operations. K can be taken small because we only need a reasonable approximation (and, as long as the exogenous states are not too persistent, it should mix fast).

D Inference Procedure

We now describe the implementation of the inferential procedure for the target parameter θ . It builds on the formulation developed in Kitamura and Stoye (2018), where the implications of economic models are expressed in terms of the minimum value of a quadratic form. We also emphasize ways to make the procedure computationally faster.

Recall that we construct a confidence set by inverting tests of the type $H_0 : \theta = \theta_0$, which are equivalent to testing $H'_0 : J(\theta_0) = 0$, where $J(\theta_0)$ is defined in the minimization problem (24) presented in the main text. We assume that $\hat{\mathbf{p}}_N$ is a frequency estimator for \mathbf{p} (collecting the frequency estimators for p and F), and that the researcher has R replications $\hat{\mathbf{p}}^{*r}$, $r = 1, \dots, R$. Recall that, for a fixed θ_0 , our test statistic is the empirical counterpart of $J(\theta_0)$, given by (25) in the main text, rewritten below for convenience,

$$\hat{J}_N(\theta_0) := \min_{\substack{(\tilde{p}, \pi) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X} : R^{iq} \pi \leq r^{iq}, \\ \mathcal{R}(\theta_0, \pi, \tilde{p}; \hat{\mathbf{p}}_N) = 0}} [b_{-J}(\hat{\mathbf{p}}_N) - \widehat{\mathbf{M}}_N \pi]' \Omega [b_{-J}(\hat{\mathbf{p}}_N) - \widehat{\mathbf{M}}_N \pi]. \quad (\text{D1})$$

The corresponding test statistic being simulated is

$$\hat{J}_{h_N}^{*r}(\theta_0) := \min_{\substack{(\tilde{p}, \pi) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X} : R^{iq} \pi \leq r^{iq}, \\ \mathcal{R}(\theta_0, \pi, \tilde{p}; \hat{\mathbf{p}}_{h_N}^{*r}) = 0}} [\hat{b}_{-J}^{*r} - \widehat{\mathbf{M}}_{h_N}^{*r} \pi]' \Omega [\hat{b}_{-J}^{*r} - \widehat{\mathbf{M}}_{h_N}^{*r} \pi], \quad (\text{D2})$$

for $r = 1, \dots, R$, where $\hat{\mathbf{p}}_{h_N}^{*r}$, \hat{b}_{-J}^{*r} and $\widehat{\mathbf{M}}_{h_N}^{*r}$ are the r -th subsampling replication of the estimator for \mathbf{p} , b_{-J} and \mathbf{M} ; see Section 5 in the main paper for the definitions of these elements.

Recall the testing procedure: We use the empirical distribution of $h_N \hat{J}_{h_N}^{*r}(\theta_0)$ to obtain the critical value $\hat{c}_{1-\alpha}(\theta_0)$. When the value of the test statistic is smaller than the critical value, $N \hat{J}_N(\theta_0) \leq \hat{c}_{1-\alpha}(\theta_0)$, we do not reject the null $H'_0 : J(\theta_0) = 0$, otherwise we reject it. The $1 - \alpha$ confidence set is the collection

³⁹Under the ℓ_1 norm, this convergence is a contraction and the contraction coefficient is known as Dobrushin ergodic coefficient.

of θ_0 's for which these tests do not reject the null.

D.1 Computational Algorithm for Inference

Next, we briefly discuss some practical issues when implementing the subsampling procedure. To simplify, we focus on the scalar case, $\theta \in \mathbb{R}$. We first argue that it is easier to solve relaxed versions of (20)–(21) to obtain good approximations for $\widehat{J}_N(\theta_0)$ than solving (25) directly. We then describe how to exploit continuity in a sequence of optimizations. As an aside, we also note that subsampling is amenable to parallelization, which speeds up the procedure.

Computing $\widehat{J}_N(\theta_0)$ in Practice. When ϕ is costly to evaluate, it is difficult to solve the minimization problems (D1) and (D2) in practice. The reason is that it is difficult to search over (\widetilde{p}, π) to minimize $J(\theta_0)$ when the constraint $\theta_0 = \phi(\widetilde{p}, \pi; \mathbf{p})$ must be satisfied for a fixed θ_0 . In other words, finding particular values for counterfactual CCP \widetilde{p} and baseline flow payoff π that generate the exact (fixed) value θ_0 can be computationally costly.

One way to bypass this difficulty is to take advantage of the relationship between the optimization problems (20)–(21) and (D1) (and the subsampling version (D2)). Abstracting from sampling issues, consider the relaxed version of the maximization (20)–(21):

$$\theta^U(\epsilon) \equiv \max_{(\widetilde{p}, \pi) \in \widetilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X}} \phi(\widetilde{p}, \pi; \mathbf{p}) \quad (\text{D3})$$

subject to

$$\begin{aligned} \|\mathbf{M}(\mathbf{p}) \pi - b_{-J}(\mathbf{p})\|_{\Omega} &\leq \epsilon, \\ R^{eq} \pi &= r^{eq}, \\ R^{iq} \pi &\leq r^{iq}, \\ (\widetilde{\mathbf{M}}(\mathbf{p}) \mathcal{H}) \pi &= \widetilde{b}_{-J}(\widetilde{p}, \mathbf{p}) - \widetilde{\mathbf{M}}(\mathbf{p}) g. \end{aligned} \quad (\text{D4})$$

where $\|\cdot\|_{\Omega}$ is the matrix norm defined as $\|x\|_{\Omega} = x' \Omega x$ for $x \in \mathbb{R}^{AX}$, and $\epsilon \geq 0$. We also consider the relaxed minimization problem by replacing the max operator in (D3) by the min operator. We index the solution and the maximum of the relaxed problems by ϵ , so that we have $(\widetilde{p}^{Us}(\epsilon), \pi^{Us}(\epsilon))$ and $\theta^U(\epsilon)$, respectively. Similarly, we index the solution and the minimum for the corresponding relaxed minimization problem by ϵ , so that $(\widetilde{p}^{Ls}(\epsilon), \pi^{Ls}(\epsilon))$ and $\theta^L(\epsilon)$.

The difference between the original problem (20)–(21) and its relaxed version (D3)–(D4) is the inequality constraint $\|\mathbf{M}(\mathbf{p}) \pi - b_{-J}(\mathbf{p})\|_{\Omega} \leq \epsilon$. The problems coincide when $\epsilon = 0$. Furthermore, $\theta^L(\epsilon) \leq \theta^L(0) \equiv \theta^L$ and $\theta^U(\epsilon) \geq \theta^U(0) \equiv \theta^U$. I.e., the true identified set $[\theta^L, \theta^U]$ is contained in the interval $[\theta^L(\epsilon), \theta^U(\epsilon)]$ when $\epsilon > 0$.

Importantly, while $J(\theta_0) = 0$ for all points θ_0 in the identified set $[\theta^L, \theta^U]$, we have $J(\theta_0) \leq \epsilon$ for all points θ_0 in the wider interval $[\theta^L(\epsilon), \theta^U(\epsilon)]$ by construction. This implies that $0 < J(\theta_0) \leq \epsilon$ for all points belonging to the wider interval, $\theta_0 \in [\theta^L(\epsilon), \theta^U(\epsilon)]$, but not to the smaller set, $\theta_0 \notin [\theta^L, \theta^U]$. Take such a point, and denote it by $\theta_0^*(\epsilon)$. If we solve the following problem for $\theta_0^*(\epsilon)$,

$$J(\theta_0^*(\epsilon)) := \min_{\substack{(\tilde{p}, \pi) \in \tilde{\mathbf{P}} \times \mathbb{R}^{(A+1)X} : R^{iq} \pi \leq r^{iq}, \\ \mathcal{R}(\theta_0^*(\epsilon), \pi, \tilde{p}; \mathbf{p}) = 0}} [b_{-J}(\mathbf{p}) - \mathbf{M}(\mathbf{p})\pi]' \Omega [b_{-J}(\mathbf{p}) - \mathbf{M}(\mathbf{p})\pi], \quad (\text{D5})$$

then the minimum $J(\theta_0^*(\epsilon))$ must be strictly greater than 0 and (weakly) smaller than ϵ . Note that by taking a sequence of ϵ 's, $0 = \epsilon_0 < \epsilon_1 < \dots < \epsilon_k < \dots < \epsilon_K$, and solving the corresponding relaxed problem (D3)–(D4) for all ϵ_k , we obtain the sequence of increasing intervals

$$[\theta^L(0), \theta^U(0)] \subseteq \dots \subseteq [\theta^L(\epsilon_k), \theta^U(\epsilon_k)] \subseteq \dots \subseteq [\theta^L(\epsilon_K), \theta^U(\epsilon_K)].$$

We also obtain a sequence of J 's such that: (a) $J(\theta_0) = 0$ if $\theta_0 \in [\theta^L(0), \theta^U(0)]$, and (b) $\epsilon_{k-1} < J(\theta_0) \leq \epsilon_k$ if $\theta_0 \in [\theta^L(\epsilon_k), \theta^U(\epsilon_k)]$ and $\theta_0 \notin [\theta^L(\epsilon_{k-1}), \theta^U(\epsilon_{k-1})]$. This means that, by taking a fine grid of ϵ 's and solving their corresponding relaxed problems (D3)–(D4), we obtain good approximations to the minimum value of the problem (D5) for any given θ_0 .

Exploiting Continuity. We can exploit the continuity of the relaxed problem (D3)–(D4) with respect to ϵ . Specifically, we exploit the fact that $\theta^L(\epsilon)$ and $\theta^U(\epsilon)$ are continuous in ϵ . This continuity implies that once we solve (D3)–(D4) for some ϵ_k , it is computationally cheap to solve the problem for a close $\epsilon_{k+1} > \epsilon_k$ by using the solution to the ϵ_k -relaxed problem as initial guess to the ϵ_{k+1} -relaxed problem. This can reduce the total computational costs substantially.

Summary. We want to test the null $H_0 : \theta = \theta_0$ for various θ_0 's. First, we solve the relaxed problem (D3)–(D4) for various ϵ_k , $0 = \epsilon_0 < \epsilon_1 < \dots < \epsilon_2 < \dots < \epsilon_K$, *both* in the full sample and in *all* R replicated subsamples (taking advantage of the continuity of the solutions with respect to ϵ). For the original sample, we obtain the estimated identified set $[\hat{\theta}^L, \hat{\theta}^U]$ when we take $\epsilon_0 = 0$; for any point θ_0 in that interval, we have $\hat{J}_N(\theta_0) = 0$ by construction. For $k \geq 1$, for any point θ_0 in the interval $[\hat{\theta}^L(\epsilon_k), \hat{\theta}^U(\epsilon_k)]$ but not in $[\hat{\theta}^L(\epsilon_{k-1}), \hat{\theta}^U(\epsilon_{k-1})]$, we have that $\epsilon_{k-1} < \hat{J}_N(\theta_0) \leq \epsilon_k$. Assuming the difference between ϵ_{k-1} and ϵ_k is small enough, we approximate the test statistic by $N \hat{J}_N(\theta_0) \simeq N \times \epsilon_k$.

Applying the same reasoning in each simulated sample r , we obtain approximated values for $\hat{J}_{h_N}^{*r}(\theta_0)$, $r = 1, \dots, R$, from which we can approximate $h_N \hat{J}_{h_N}^{*r}(\theta_0)$. Inverting the empirical distribution of these values we can construct the critical value $\hat{c}_{1-\alpha}(\theta_0)$ and test the null $H_0 : \theta = \theta_0$.

Finally, applying this procedure on a grid of points θ_0 's, we construct the confidence set by the set

of non-rejected θ_0 's.⁴⁰ By exploiting (i) the relationship between the optimizations (D1) and (D3)–(D4), and (ii) the continuity of the solutions with respect to ϵ , we improve the performance of the subsampling procedure and compute an asymptotically uniformly valid $1 - \alpha$ confidence set for the true outcome of interest θ in a tractable way.

E A Monte Carlo Study

In this section, we present a Monte Carlo study to illustrate the finite-sample performance of our inference procedure. We start with the setup, and then we show the results.

E.1 Setup

We extend the firm entry/exit problem presented in Sections 3 and 4 of the main text, allowing now for a larger state space. Specifically, we assume the presence of three exogenous states, $w_t = (w_{1t}, w_{2t}, w_{3t})$, reflecting demand and supply shocks. The exogenous states are independent to each other, and each follows a discrete-AR(1) process with W support points (obtained by discretizing latent normally-distributed AR(1) processes). The (residual) inverse demand function is linear, $P_t = \bar{w} + w_{1t} + w_{2t} - \eta Q_t$, where P_t is the price of the product, Q_t is the quantity demanded, \bar{w} is the intercept, w_{1t} and w_{2t} are demand shocks, and η is the slope. We assume constant marginal costs mc_t (i.e., mc_t does not depend on Q_t), and let the supply shocks w_{3t} affect marginal costs. To simplify, we just take $mc_t = w_{3t}$. Variable profits are then $vp_t = (\bar{w} + w_{1t} + w_{2t} - mc_t)^2/4\eta$. The idiosyncratic shocks ε follow the type 1 extreme value distribution. The model parameters are presented in Table E1.

Table E1: Parameters of the Monte Carlo Data Generating Process

Demand Function:	\bar{w}	6.8	$w_{1t} \sim \text{Normal AR}(1)$:	ρ_{01}	0
	η	4		ρ_{11}	0.75
				σ_1^2	0.02
Payoff Parameters:	s	4.5	$w_{2t} \sim \text{Normal AR}(1)$:	ρ_{02}	0
	ec	5		ρ_{12}	0.75
	fc	0.5		σ_2^2	0.025
Scale parameter:	σ	1	$w_{3t} \sim \text{Normal AR}(1)$:	ρ_{03}	0
				ρ_{13}	0.75
Discount Factor:	β	0.9		σ_3^2	0.03

The counterfactual we consider is the same as in the example in Section 4: a subsidy that reduces

⁴⁰Even though the identified set is always an interval, it could be that the set of non-rejected parameters is disjoint. In our experience, all the confidence sets we generated in practice are intervals.

entry costs by 20%. The target parameter θ is the long-run average probability of staying in the market given the subsidy, where the long-run average is based on the ergodic distribution of the state variables; the specific formula for θ is provided in Online Appendix G (but note that here we do not take the difference between the counterfactual and the baseline average probabilities).

In order to analyze the sensitivity of the target parameter θ to alternative model restrictions, we follow the example again and impose the three sets of restrictions:

1. $oo = 0$, $fc \geq 0$, $ec \geq 0$, and vp is known.
2. $\pi_1(1, w^h) \geq \pi_1(1, w^l)$, and $vp - fc \leq ec \leq \frac{\mathbb{E}[vp - fc]}{1 - \beta}$, where the expectation is taken over the ergodic distribution of the state variables.
3. s does not depend on w .⁴¹

We generate 1000 Monte Carlo replications for each of the following sample sizes: the small sample, with $N = 100$ firms on separated (independent) markets and $T = 5$ time periods, and the large sample, with $N = 1000$ firms and $T = 15$ time periods. For the first sample period, the value of the state variables are drawn from their steady-state distributions. Given that each exogenous state variable w_{jt} can take W values, the dimension of the state space is $X = 2 \times W^3$. We consider three sizes for the state space: $X = 16, 54$ and 250 , which correspond to $W = 2, 3$ and 5 . The choices of the state space were dictated by the sample size, not by computational constraints, given that the method makes use of a frequency, or a nonparametric estimator for the CCP in the first stage. As discussed in the main text and in Section C of this appendix, it is feasible to solve the optimization problem (20)–(21) for state spaces that are larger in size.

In each sample, we estimate the lower and upper bounds for the target parameter, θ^L and θ^U , by solving the minimization and maximization problems (20)–(21). We estimate CCPs using frequency estimators, and we use the true transition matrix F , both in calculating test statistics and critical values (The results do not change significantly when we estimate transition probabilities as well.) We solve the problem (20)–(21) using the Knitro MATLAB function. We provide the analytical gradient of ϕ as explained in Section B of this appendix. The initial values for π are obtained by solving the following quadratic programming problem (with the GUROBI solver, using its MATLAB interface):

$$\min_{\substack{\pi \in \mathbb{R}^{(A+1)X}: \\ R^{eq} \pi = r^{eq}, R^{iq} \pi \leq r^{iq}}} [b_{-J}(\hat{\mathbf{p}}_N) - \widehat{\mathbf{M}}_N \pi]' \widehat{\Omega}_N [b_{-J}(\hat{\mathbf{p}}_N) - \widehat{\mathbf{M}}_N \pi].$$

We specify $\widehat{\Omega}_N$ to be a diagonal matrix with diagonal terms given by the square-root of the ergodic distribution of the exogenous state variables, implied by the transition process F^w . We opt for this

⁴¹When we impose Restriction 3, we replace the inequalities defined in Restriction 1 by their average versions. This does not affect the identified set, but it improves the finite-sample behavior of the estimators when the sample size is small and the state space is large.

weighting matrix so that deviations on more visited states receive greater weights and are, therefore, considered more relevant. This is the weighting matrix we use to compute $\hat{J}_N(\theta_0)$.

We calculate 90% confidence sets for θ using the procedure described in Section 5 of the main text and in Section D of this appendix. Specifically, we approximate the value of $\hat{J}_N(\theta_0)$ for any fixed θ_0 , in the main data and in each subsample, by solving ϵ -relaxed versions of the optimization problem (20)–(21). (We let ϵ range from 0 to 1 in an equally spaced grid with 50 points.) For each sample, we generate 1000 subsamples with size that is approximately $h_N \approx 8 \times \sqrt{NT}$. Specifically, we implement a standard i.i.d. subsampling, resampling firms over the full time period: For the small sample we draw 36 firms randomly, and for the large sample, we draw 65 firms. The computations were run on the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University.

E.2 Monte Carlo Results

We now discuss the results of the Monte Carlo simulations. In the baseline scenario, the long-run average probability that the firm stays in the market is 90.5%, while the long-run average probability of being active reduces to 83.3% in the counterfactual scenario (so that $\theta = 0.833$). The impact of the entry subsidy is to reduce the long-run average by 7.2 percentage points. Similar to the example presented in the main paper, the entry subsidy increases the exit rate of forward-looking firms, which translates into firms staying less often in the market in the steady state. This result is invariant to the alternative discretizations of the state space, since the discretizations are performed on the same underlying AR(1) processes.

Table E2 presents the Monte Carlo results. The top, middle, and bottom panels show the results for the alternative state spaces: small ($X = 16$), medium-sized ($X = 54$), and large ($X = 250$), respectively. In each panel, the top subpanel presents the results for the small sample ($N = 100, T = 5$), and the bottom subpanel, for the large sample ($N = 1000, T = 15$). In each subpanel, we show for each alternative set of Restrictions 1–3, (i) the populational (true) identified set, (ii) the average estimates of the lower and upper bounds, θ^L and θ^U , (iii) the average bias of the estimated bounds, (iv) the average endpoints and the average length of the 90% confidence sets, (v) the coverage probability of the confidence sets, and (vi) the average time taken to estimate θ^L and θ^U (in seconds), and the average time taken to compute the confidence intervals (in minutes).

The identified sets under the alternative Restrictions 1–3 are all compact intervals containing the true θ (Proposition 2), and vary slightly with the size of the state space. Restriction 1 alone is highly informative: the counterfactual long-run average probability of being active is between 75% and 90.5%. It does however include the baseline probability (at the upper end of the interval). Adding Restriction 2 reduces the upper bound to 87.8%, which suffices to identify the sign of the impact of the subsidy. Adding Restriction 3 pushes the upper bound further down to 86.8%.

In all cases, the estimated lower and upper bounds of the identified sets appear to be consistent,

with smaller biases in larger samples. The coverage probabilities of the confidence sets converge to the nominal level 90%, as expected (Theorem 1). And the confidence sets' average lengths are wider (though not substantially) than the length of the true identified sets, for all sample sizes and state spaces. E.g., in the small state space case and large sample, the average length of the confidence set is 0.1782 under Restriction 1, while the length of the (true) identified set is just 0.1536; and in the large state space and small sample, the average length of the confidence set under the same restriction is 0.25.

Naturally, the finite sample performance of our inference procedure depends on both the state space and the sample size. In the larger state space cases, we obtain slightly greater average biases for the point estimates. These are expected: larger state spaces imply less (effective) degrees of freedom, as the number of model parameters increases with the state space (recall that π is an $(A + 1)X$ vector).

In terms of the computer time required to solve the minimization and maximization problems (20)–(21), it takes approximately 0.03 seconds to solve both optimization problems under Restrictions 1 and 1–2, and that time is reduced to just 0.01 seconds under Restrictions 1–3, in the small state space case. Subsampling is computationally intensive but feasible: for the same state space, the average time required to run it varies from two minutes under Restriction 1 to one minute under Restrictions 1–3.

As expected, it takes longer to solve (20)–(21) when the state space is larger. E.g., under Restriction 1, it takes approximately 0.3 seconds on average in the medium-sized state space case ($X = 54$), and approximately 6 seconds on average in the large state space case ($X = 250$). It also takes longer to run the subsampling procedure: between 7 and 28 minutes on average in the medium-sized state space, and between 150 and 580 minutes on average in the large state space, depending on the sample size and the restrictions imposed. It is important to stress, however, that the average computer time here is based on a sequential implementation of subsampling, which does not take advantage of parallelization.

Table E2: Monte Carlo Results

Target Parameter: θ = Long-run Average Probability of Being Active			
Small State Space: $X = 16$			
$T = 5, N = 100$	Restrictions 1	Restrictions 1-2	Restrictions 1-3
True Identified Set	[0.7500, 0.9036]	[0.7500, 0.8763]	[0.7500, 0.8662]
Average Estimated Bounds	[0.7583, 0.9036]	[0.7579, 0.8727]	[0.7580, 0.8651]
Average Bias	[0.0083, 0.0000]	[0.0079, -0.0036]	[0.0080, -0.0011]
Confidence Sets: Average Endpoints	[0.6729, 0.9214]	[0.6734, 0.8951]	[0.6757, 0.8870]
Confidence Sets: Average Length	0.2485	0.2217	0.2113
Coverage Probability (90% nominal)	0.9060	0.9010	0.9050
Time Estimation (sec)	0.04	0.05	0.01
Time Inference (min)	2	2	1
$T = 15, N = 1000$	Restrictions 1	Restrictions 1-2	Restrictions 1-3
True Identified Set	[0.7500, 0.9036]	[0.7500, 0.8763]	[0.7500, 0.8662]
Average Estimated Bounds	0.7507, 0.9036]	[0.7507, 0.8761]	[0.7507, 0.8661]
Average Bias	[0.0007, -0.0000]	[0.0007, -0.0002]	[0.0007, -0.0001]
Confidence Sets: Average Endpoints	[0.7296, 0.9079]	[0.7296, 0.8806]	[0.7297, 0.8713]
Confidence Sets: Average Length	0.1782	0.1510	0.1417
Coverage Probability (90% nominal)	0.9090	0.9010	0.9040
Time Estimation (sec)	0.04	0.04	0.01
Time Inference (min)	2	2	0.7
Medium-sized State Space: $X = 54$			
$T = 5, N = 100$	Restrictions 1	Restrictions 1-2	Restrictions 1-3
True Identified Set	[0.7503, 0.9057]	[0.7503, 0.8784]	[0.7503, 0.8682]
Average Estimated Bounds	[0.7591, 0.9036]	[0.7581, 0.8710]	[0.7586, 0.8641]
Average Bias	[0.0089, -0.0021]	[0.0078, -0.0074]	[0.0083, -0.0041]
Confidence Sets: Average Endpoints	[0.6656, 0.9235]	[0.6589, 0.9042]	[0.6628, 0.8932]
Confidence Sets: Average Length	0.2579	0.2453	0.2304
Coverage Probability (90% nominal)	0.8940	0.9050	0.8910
Time Estimation (sec)	0.34	0.41	0.03
Time Inference (min)	28	22	12
$T = 15, N = 1000$	Restrictions 1	Restrictions 1-2	Restrictions 1-3
True Identified Set	[0.7503, 0.9057]	[0.7503, 0.8784]	[0.7503, 0.8682]
Average Estimated Bounds	[0.7509, 0.9056]	[0.7509, 0.8782]	[0.7509, 0.8680]
Average Bias	[0.0006, -0.0001]	[0.0006, -0.0002]	[0.0006, -0.0002]
Confidence Sets: Average Endpoints	[0.7292, 0.9101]	[0.7290, 0.8831]	[0.7290, 0.8748]
Confidence Sets: Average Length	0.1809	0.1541	0.1459
Coverage Probability (90% nominal)	0.9020	0.9070	0.8990
Time Estimation (sec)	0.30	0.27	0.03
Time Inference (min)	24	17	7
Large State Space: $X = 250$			
$T = 5, N = 100$	Restrictions 1	Restrictions 1-2	Restrictions 1-3
True Identified Set	[0.7504, 0.9060]	[0.7504, 0.8787]	[0.7503, 0.8685]
Average Estimated Bounds	[0.7612, 0.9027]	[0.7605, 0.8701]	[0.7593, 0.8638]
Average Bias	[0.0108, -0.0033]	[0.0102, -0.0086]	[0.0090, -0.0047]
Confidence Sets: Average Endpoints	[0.6678, 0.9253]	[0.6602, 0.9096]	[0.6621, 0.8979]
Confidence Sets: Average Length	0.2575	0.2494	0.2358
Coverage Probability (90% nominal)	0.8960	0.9090	0.9080
Time Estimation (sec)	7	8	0.7
Time Inference (min)	578	477	252
$T = 15, N = 1000$	Restrictions 1	Restrictions 1-2	Restrictions 1-3
True Identified Set	[0.7504, 0.9060]	[0.7504, 0.8787]	[0.7503, 0.8685]
Average Estimated Bounds	[0.7532, 0.9064]	[0.7532, 0.8790]	[0.7510, 0.8685]
Average Bias	[0.0028, 0.0004]	[0.0028, 0.0003]	[0.0007, 0.0000]
Confidence Sets: Average Endpoints	[0.7321, 0.9106]	[0.7287, 0.8845]	[0.7288, 0.8757]
Confidence Sets: Average Length	0.1786	0.1558	0.1469
Coverage Probability (90% nominal)	0.9070	0.9000	0.9020
Time Estimation (sec)	6	6	0.6
Time Inference (min)	505	457	150

Note: T = number of periods, N = number of markets, X = number of states.