

DISCUSSION PAPER SERIES

DP12470

IDENTIFICATION OF COUNTERFACTUALS IN DYNAMIC DISCRETE CHOICE MODELS

Myrto Kalouptsi, Paul Scott and Eduardo Souza-
Rodrigues

**INDUSTRIAL ORGANIZATION and
LABOUR ECONOMICS**



IDENTIFICATION OF COUNTERFACTUALS IN DYNAMIC DISCRETE CHOICE MODELS

Myrto Kalouptsi, Paul Scott and Eduardo Souza-Rodrigues

Discussion Paper DP12470
Published 28 November 2017
Submitted 28 November 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION and LABOUR ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Myrto Kalouptsi, Paul Scott and Eduardo Souza-Rodrigues

IDENTIFICATION OF COUNTERFACTUALS IN DYNAMIC DISCRETE CHOICE MODELS

Abstract

Dynamic discrete choice models (DDC) are not identified nonparametrically. However, the non-identification of DDC models does not necessarily imply non-identification of counterfactuals of interest. Using a novel approach that can accommodate both nonparametric and restricted payoff functions, we provide necessary and sufficient conditions for the identification of counterfactual behavior and welfare for a broad class of counterfactuals. The conditions are simple to check and can be applied to virtually all counterfactuals in the DDC literature. To explore the robustness of counterfactual results to model restrictions in practice, we consider a numerical example of a monopolist's entry problem, as well as an empirical model of agricultural land use. In each case, we provide examples of both identified and non-identified counterfactuals of interest.

JEL Classification: N/A

Keywords: identification, dynamic discrete choice, counterfactual, welfare

Myrto Kalouptsi - myrto@fas.harvard.edu
Harvard University and CEPR

Paul Scott - ptscott@stern.nyu.edu
New York University

Eduardo Souza-Rodrigues - e.souzarodrigues@utoronto.ca
University of Toronto

Identification of Counterfactuals in Dynamic Discrete Choice Models*

Myrto Kalouptsi[†], Paul T. Scott[‡], Eduardo Souza-Rodrigues[§]

Harvard University, CEPR and NBER, NYU Stern School of Business, University of Toronto

November 2017

Abstract

Dynamic discrete choice models (DDC) are not identified nonparametrically. However, the non-identification of DDC models does not necessarily imply non-identification of counterfactuals of interest. Using a novel approach that can accommodate both nonparametric and restricted payoff functions, we provide necessary and sufficient conditions for the identification of counterfactual behavior and welfare for a broad class of counterfactuals. The conditions are simple to check and can be applied to virtually all counterfactuals in the DDC literature. To explore the robustness of counterfactual results to model restrictions in practice, we consider a numerical example of a monopolist's entry problem, as well as an empirical model of agricultural land use. In each case, we provide examples of both identified and non-identified counterfactuals of interest.

KEYWORDS: Identification, Dynamic Discrete Choice, Counterfactual, Welfare.

*We are grateful to Victor Aguirregabiria, Ivan Canay, Phil Haile, Bo Honore, Jakub Kastl, Yuichi Kitamura, Arthur Lewbel, Yao Luo, Rob McMillan, Ariel Pakes, Rodrigo Pinto, John Rust, Junichi Suzuki, Elie Tamer, Ali Yurukoglu, and seminar participants at various universities for many helpful comments. Charis Katsiardi and Adrian Torchiana provided excellent research assistance. The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme Grant Agreement no. 230589 and Agence Nationale de la Recherche Projet ANR-12-CHEX-0012-01.

[†]Department of Economics, Harvard University, Littauer Center, Cambridge, MA 02138, myrto@fas.harvard.edu

[‡]Stern School of Business, New York University, Kaufman Management Center, 44 W. 4th St., New York, NY 10012, ptscott@stern.nyu.edu.

[§]Department of Economics, University of Toronto, Max Gluskin House, 150 St. George St., Toronto, Ontario M5S 3G7, Canada, e.souzarodrigues@utoronto.ca

1 Introduction

Since the seminal contributions of Rust (1994) and Magnac and Thesmar (2002), it is well known that dynamic discrete choice (DDC) models are not identified nonparametrically: several payoff functions can rationalize observed choice behavior. As a result, researchers must impose additional restrictions to identify and estimate DDC models, often with the goal of performing counterfactuals, such as policy interventions affecting the environment under study. Counterfactual experiments alter the primitives of DDC models, including changes to the set of actions and states available to agents, the payoff function, the process governing state transitions, and/or the discount factor. When all models consistent with observed data generate the same behavioral response in a given counterfactual environment, then the counterfactual can be said to be identified. In some cases, however, different models that are equally compatible with observed data generate different behavior responses in a counterfactual environment; in such cases, the counterfactual is not identified. While restrictions that select among models are sometimes well-informed by the economics of the problem, often there is little guidance as to reasonable restrictions that are necessary to identify the model, generating concern about the robustness of the empirical findings.

A recent body of innovative work – Aguirregabiria (2010), Aguirregabiria and Suzuki (2014), Norets and Tang (2014), and Arcidiacono and Miller (2017) – has made valuable progress in this area. These papers 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.

This paper builds on that body of research in three respects. First, we propose a general framework that allows us to investigate the identification of virtually all counterfactuals in the applied DDC literature.¹ Examples include altering the choice set available to agents (e.g., removing social security or welfare programs, offering new insurance plans to farmers), assigning the primitives of one group of agents to those of another (e.g., assuming preferences of labor market cohorts are equal, or firm entry costs are identical across markets), and changing payoffs proportionally (e.g., subsidies that reduce firms’ entry/sunk costs by some percentage), among others.² Second, we investigate whether (and under what conditions) there are additional restrictions that suffice to identify the counterfactual, even if they are insufficient to identify the full model. Third, we add to existing results that focus on counterfactual behavior by considering the identification of

¹We consider any (simultaneous) change in the primitives, with the exception of non-differentiable changes in the payoff function and changes in the distribution of the idiosyncratic shocks. We have not seen such counterfactuals implemented in practice.

²All of these examples go beyond the case of pre-specified additive changes in payoffs, i.e., changes that must be specified before estimating the payoff function. For instance, if we wish to, say, assign the payoff of minorities to those of white women to investigate the evolution of racial-gap in labor markets, in the spirit of Keane and Wolpin (2010), we would first have to estimate these primitives. A taxonomy of counterfactuals, as well as several examples encountered in applied work, are discussed in Section 3.1.

counterfactual welfare, which is often the ultimate object of interest to policy makers.

To that end, we develop a novel approach that allows us to derive the set of necessary and sufficient conditions to identify counterfactual behavior and welfare for a broad class of counterfactuals. We consider counterfactuals that involve almost any change in the primitives, so our results can be used on a case-by-case basis to investigate the identification of particular policy interventions of interest. We first note that Magnac and Thesmar’s (2002) underidentification result implies that there is a subset of the model parameters that can be expressed as a known function of another subset of parameters (denoted “free parameters”). We then find a convenient representation that directly relates counterfactual choice probabilities and welfare to the free parameters. Based on this representation, we can determine the conditions under which counterfactual behavior and welfare are identified. When such conditions are satisfied, it is not necessary to identify all the individual structural parameters of the model to identify the effects of policy interventions. This is consistent with Marschak’s (1953) view of solving well-posed economic problems with minimal assumptions; see Heckman (2000, 2010) for excellent discussion of Marschak’s approach to identification in structural models.

Our results imply that the identification of counterfactuals can be verified directly from data on the state transition process.³ In some cases, identification can be determined without even examining the data. For example, counterfactuals eliminating an action from the choice set result in identified counterfactual behavior; counterfactuals assigning the payoff parameters from one group of agents to another are not identified, except in special cases.⁴ While prior studies show that pre-specified additive shifts in flow payoffs result in identified counterfactual behavior, we show that all other counterfactual transformations of payoff functions are identified only under restrictive conditions that require verification in the data.

Given that some counterfactuals are not identified, a natural question concerns what sort of (well-informed) restrictions are needed to obtain identification. We consider parametric assumptions, which are prevalent in applied work, as well as linear restrictions on payoffs (e.g., exclusion restrictions). Even when parametric assumptions do not suffice to identify the full model, they help identify specific sets of counterfactuals. For instance, a number of papers have implemented a counterfactual that changes the volatility or long-run mean of market states (i.e., a change in the state transition process); while such counterfactuals *are not* identified in a nonparametric setting, we show that most examples of these counterfactuals in the literature *are* identified because of the parametric setting. The parametric model we consider is a special case of a model imposing linear restrictions on payoffs. For linear restrictions, we find another set of necessary conditions

³The conditions imply, for example, that the counterfactual from Rust’s (1987) seminal bus-engine replacement model is not identified.

⁴Sometimes eliminating an action also eliminates states (e.g. that may reflect past actions). We also characterize this case, providing conditions that the counterfactual state transition process (across the remaining states) must satisfy for identification. Similarly, a counterfactual that adds an action is identified provided that the counterfactual payoffs of the new action are a convex combination of the baseline payoffs.

for identification that are simple to verify for any counterfactual.

Finally, we consider the identification of welfare, which is often the ultimate object of interest to policy makers (in terms of both sign and magnitude). We find that the identification of counterfactual behavior is necessary but not sufficient for the identification of welfare. We also provide sufficient conditions for welfare identification.

To gain intuition and explore how sensitive counterfactuals can be to model restrictions in practice, we turn to two applications. The first is a numerical exercise featuring a monopoly entry model. To identify this model, the researcher must restrict scrap values, entry costs or fixed costs; this is usually accomplished by fixing one of them to zero. Such an assumption is difficult to justify however as cost or scrap value data are extremely rare.⁵ The restrictions can affect the parameter estimates and, for non-identified counterfactuals, alter the counterfactual predictions as well. For instance, when fixed costs are set to zero, the monopolist obtains high profits, which provide incentives to enter and stay in the market. To match the observed choice behavior, estimated entry costs and scrap values must thus be high as well. The estimated model with zero fixed costs is observationally equivalent to the true model (they both rationalize the data equally well). However, when we implement a counterfactual subsidy that reduces entry costs by 10 percent, for example, the predicted impacts on turnover and welfare are exaggerated in the estimated model when compared to the impacts based on the true model. When the scrap value is set to zero instead, the estimated entry costs must be low (or even negative) to rationalize the observed choice behavior. Applying the same subsidy again results in incorrect counterfactual predictions (which may even go in the wrong direction).⁶

Next, we consider the empirical relevance of our results in the context of US agricultural land use. Following Scott (2013), field owners decide whether to plant crops or not and face uncertainty regarding commodity prices, weather shocks, and government interventions. Costs of switching between land uses create dynamic incentives for landowners. To estimate the model, Scott restricts the value of a subset of the switching cost parameters. However, as there is little guidance in the literature concerning how to specify the particular values, he sets them to zero. To evaluate the impact of these restrictions on counterfactual analysis, we bring in additional data and augment Scott's estimation strategy using land resale price data. Similar to Kalouptside (2014), we treat farmland resale transaction prices as a measure of agents' value functions. The augmented estimator allows us to test Scott's identifying restrictions and reject them.⁷

We then implement two counterfactuals. First, we consider a long-run land use elasticity, which measures the sensitivity of land use to a persistent change in crop returns. The second features

⁵Using external information on entry costs and scrap values (specifically, new ship prices and demolition prices), Kalouptside (2014) shows that the latter vary dramatically over states in the shipping industry.

⁶This entry model is similar to the model in Aguirregabiria and Suzuki (2014), but we investigate a broader set of counterfactuals.

⁷Relating land resale price data to the model requires another set of assumptions about land markets. See Kalouptside (2014) for a full discussion of these restrictions.

an increase in the cost of replanting crops and resembles a fertilizer tax. We show that while the long-run elasticity is identified, the fertilizer tax is not. Thus, a model estimated with our augmented estimator and a model imposing Scott’s restrictions both predict the same long-run elasticity, but they predict different responses to the increase in fertilizer taxes (and even responses in different directions).

Related Literature. Our paper relates to several important prior studies. In addition to Rust’s (1994) and Magnac and Thesmar’s (2002) seminal contributions, Heckman and Navarro (2007) consider the identification of a semiparametric finite horizon optimal stopping time model that allows for a rich time series dependence on the unobservables. Heckman, Humphries, and Veramendi (2016) extend the work of Heckman and Navarro (2007) by incorporating both ordered and unordered choice sets, and by decomposing the identification of dynamic treatment effects into direct effects and continuation values. Under a conditional independence assumption on the unobservables, Bajari, Chu, Nekipelov, and Park (2016) study the identification of finite-horizon models with a terminal action, while Abbring and Daljord (2017) investigate the conditions needed to identify the discount factor. Pesendorfer and Schmidt-Dengler (2007) extend Magnac and Thesmar’s results to dynamic games.⁸

Regarding the identification of counterfactuals in DDC models, Aguirregabiria (2010) shows identification of counterfactual choice probabilities when the experiment consists of adding a pre-specified amount to payoffs in a finite-horizon binary choice model. Aguirregabiria and Suzuki (2014) and Norets and Tang (2014) extend Aguirregabiria’s (2010) result to infinite horizon models. They both provide another important extension by showing nonidentification of behavior under changes in transition probabilities. Arcidiacono and Miller (2017) extend these results even further to multinomial choice models for both stationary and nonstationary environments in the presence of long and short panel data.

We focus on infinite horizon multinomial choice models and complement the literature by (a) providing the first full set of necessary and sufficient conditions for identification of counterfactuals involving almost any change in the model primitives, (b) investigating the identification power of additional restrictions to the basic framework, and (c) providing identification results for counterfactual welfare. In a companion paper (Kalouptsi, Scott, and Souza-Rodrigues (2017a)), we consider the identification of counterfactual behavior in dynamic games.⁹

The paper is organized as follows: Section 2 presents the dynamic discrete choice framework and

⁸Recent work considers the identification of the distribution of the idiosyncratic shocks when agents can make both discrete and continuous choices (Blevins (2014)), or in the presence of continuous states and exclusion restrictions (Chen (2017)), or under linear-in-parameters payoff functions (Buchholz, Shum, and Xu (2017)).

⁹Aguirregabiria and Suzuki (2014) provide results in the context of a monopolist entry/exit problem. In addition to point identification, Norets and Tang (2014) relax the assumption that the distribution of the idiosyncratic shocks is known by the econometrician, and, as a consequence, obtain some partial identification results. We do not cover partial identification as in Norets and Tang (2014), and we do not consider nonstationary settings as in Aguirregabiria (2010) and Arcidiacono and Miller (2017).

reconstructs the known results on the nonparametric underidentification of standard DDC models. Section 3 contains our main results relating to the identification of counterfactual behavior. Section 4 considers the identification of counterfactual behavior in parametric models and in models with linear restrictions on payoffs. Section 5 considers the identification of counterfactual welfare effects. Section 6 discusses our two applications: a numerical monopoly entry model and an empirical model of agricultural land use. Section 7 concludes. All proofs are presented in Appendix A. The details of the dataset and the implementation of the empirical application are discussed in Appendix B and C.

2 Modeling Framework

In each period $t \in \{1, 2, \dots\}$, agent i chooses one action a_{it} from the finite set $\mathcal{A} = \{1, \dots, A\}$. The per period payoff depends on the state variables $(x_{it}, \varepsilon_{it})$, where x_{it} is observed by the econometrician and ε_{it} is not. We assume $x_{it} \in \mathcal{X} = \{x_1, \dots, X\}$, $X < \infty$, while $\varepsilon_{it} = (\varepsilon_{it}(1), \dots, \varepsilon_{it}(A))$ is i.i.d. across agents and time and has joint distribution G with continuous support on \mathbb{R}^A . The transition distribution function for $(x_{it}, \varepsilon_{it})$ factors as follows:

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

and the per period payoff function is given by

$$\pi(a, x_{it}, \varepsilon_{it}) = \pi(a, x_{it}) + \varepsilon_{it}(a).$$

Agent i chooses a sequence of actions to maximize the expected discounted sum of current and future payoffs. Let $V(x_{it}, \varepsilon_{it})$ denote the agent's value function. By Bellman's principle of optimality,

$$V(x_{it}, \varepsilon_{it}) = \max_{a \in \mathcal{A}} \{ \pi(a, x_{it}) + \varepsilon_{it}(a) + \beta E[V(x_{it+1}, \varepsilon_{it+1}) | a, x_{it}] \},$$

where $\beta \in (0, 1)$ is the discount factor. Following the literature, we define the *ex ante value function* $V(x_{it}) \equiv \int V(x_{it}, \varepsilon_{it}) dG(\varepsilon_{it})$, as well as the *conditional value function*

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

The agent's optimal policy is given by the conditional choice probabilities (CCPs):

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

where $\mathbf{1}\{\cdot\}$ is the indicator function. We define the vectors $p(x) = [p_1(x), \dots, p_{A-1}(x)]$ and $p = [p(x_1), \dots, p(X)]$.

The following results provide relations between key objects of the model and are widely used in the literature. The first statement was proved by Hotz and Miller (1993), while the second by Arcidiacono and Miller (2011).

Lemma 1. (i) For all $(a, x) \in \mathcal{A} \times \mathcal{X}$ and a given reference action j ,

$$v_a(x) - v_j(x) = \phi_{aj}(p(x)),$$

where $\phi_{aj}(\cdot)$ are functions mapping the simplex in \mathbb{R}^A onto \mathbb{R} and are derived only from G .

(ii) For any $(a, x) \in \mathcal{A} \times \mathcal{X}$, there exists a real-valued function $\psi_a(p(x))$ such that

$$V(x) - v_a(x) = \psi_a(p(x)),$$

where the functions ψ_a are derived only from G . Moreover, $\psi_a(p(x)) = \psi_j(p(x)) - \phi_{aj}(p(x))$, all $a \neq j$.¹⁰

2.1 Nonparametric Identification of Payoffs

A DDC model consists of the primitives $(\mathcal{A}, \mathcal{X}, \beta, \pi, G, F)$ that generate the endogenous objects $\{p_a, v_a, V, a \in \mathcal{A}\}$. Typically, the available data consist of agents' actions at different states, $y_{it} \equiv (a_{it}, x_{it})$. We assume the joint distribution of y_{it} , $\Pr(y)$, is known, which implies the CCPs $p_a(x)$ and the transition F are also known. Further, we follow the literature and assume that (β, G) is known as well.¹¹

The objective is to identify the payoff function π . Intuitively, π has AX parameters, and there are only $(A - 1)X$ observed CCPs; thus there are X free payoff parameters and X restrictions will need to be imposed. Formally, the analysis is based on the following relationships between the primitives and the endogenous objects:

$$\pi_a = v_a - \beta F_a V, \quad \text{for } a = 1, \dots, A \quad (1)$$

¹⁰To derive the Arcidiacono-Miller Lemma from the Hotz-Miller inversion, note that

$$\begin{aligned} V(x) &= \int \max_{j \in \mathcal{A}} \{v_j(x) + \varepsilon(j)\} dG(\varepsilon) \\ &= \int \max_{j \in \mathcal{A}} \{v_j(x) - v_a(x) + \varepsilon(j)\} dG(\varepsilon) + v_a(x) \\ &= \int \max_{j \in \mathcal{A}} \{\phi_{ja}(p(x)) + \varepsilon(j)\} dG(\varepsilon) + v_a(x), \end{aligned}$$

and take $\psi_a(p(x)) = \int \max_{j \in \mathcal{A}} \{\phi_{ja}(p(x)) + \varepsilon(j)\} dG(\varepsilon)$. When ε_{it} follows the type 1 extreme value distribution, then $\psi_a(p(x)) = \gamma - \log p_a(x)$, where γ is the Euler constant; $\phi_{aj}(p(x)) = \log p_a(x) - \log p_j(x)$; and $p_a(x) = \exp v_a(x) / \sum_{j \in \mathcal{A}} \exp v_j(x)$.

¹¹In that sense, identification is not entirely non-parametric. Norets and Tang (2014) have considered the problem of identifying π when G is unknown. Blevins (2014), Chen (2017), and Buchholz, Shum, and Xu (2017) consider identification of G under different model assumptions. Magnac and Thesmar (2002, Proposition 4), and Abbring and Daljord (2017) investigate sufficient conditions for identification of the discount factor.

$$v_a - v_j = \phi_{aj}, \text{ for } a = 1, \dots, A, a \neq j \quad (2)$$

$$V = v_a + \psi_a, \text{ for } a = 1, \dots, A, \quad (3)$$

where $\pi_a, v_a, V, \phi_{aj}, \psi_a \in \mathbb{R}^X$, with $\pi_a(x) = \pi(a, x)$; and F_a is the transition matrix with (m, n) element equal to $\Pr(x_{it+1} = x_n | a, x_{it} = x_m)$. Note that using the observed choice probabilities, p , we can compute ϕ_{aj} , as well as ψ_a , for all a . Proposition 1 below formalizes the underidentification problem and can be considered a rephrasing of Proposition 2 in Magnac and Thesmar (2002). All proofs are in Appendix A.

Proposition 1. *Let $J \in \mathcal{A}$ be some reference action. For each $a \neq J$, the payoff function π_a can be represented as an affine transformation of π_J :*

$$\pi_a = A_a \pi_J + b_a, \quad (4)$$

where $A_a = (I - \beta F_a)(I - \beta F_J)^{-1}$ and $b_a = A_a \psi_J - \psi_a$.

Given $\Pr(y)$, one can compute A_a and b_a directly from the data for all $a \neq J$. Proposition 1 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 long as the extra assumptions add X linearly independent restrictions to the $(A - 1)X$ restrictions expressed by (4), π will be uniquely determined. Further, whichever extra restrictions are imposed, they are equivalent to stipulating the payoffs of a reference action; i.e. if π_J^* is the vector of payoffs for the reference action identified by some set of restrictions and (4), then that set of restrictions is equivalent to stipulating $\pi_J = \pi_J^*$ *a priori*.

In the remainder of the paper, it will be useful to represent (4) for all actions $a \neq J$ at once using the compact notation

$$\pi_{-J} = A_{-J} \pi_J + b_{-J} \quad (5)$$

where π_{-J} stacks π_a for all $a \neq J$; and the matrix A_{-J} and the vector b_{-J} are defined similarly. The underidentification problem is therefore represented by the free parameter π_J .

3 Identification of Counterfactual Behavior

This section presents our main results on the identification of counterfactual behavior. We begin with a taxonomy of counterfactuals; we then provide the necessary and sufficient conditions for identification; next, we investigate several special cases of practical interest; and finally we provide

some intuition for the results.¹²

3.1 Taxonomy of Counterfactuals

A counterfactual is defined by the tuple $\{\tilde{\mathcal{A}}, \tilde{\mathcal{X}}, \tilde{\beta}, h, h^s\}$. The sets $\tilde{\mathcal{A}} = \{1, \dots, \tilde{A}\}$ and $\tilde{\mathcal{X}} = \{x_1, \dots, \tilde{X}\}$ denote the new set of actions and states respectively. The new discount factor is $\tilde{\beta}$. The function $h : \mathbb{R}^{A \times X} \rightarrow \mathbb{R}^{\tilde{A} \times \tilde{X}}$ transforms the payoff function π into the counterfactual payoff $\tilde{\pi}$, so that $\tilde{\pi} = h(\pi)$, where $h(\pi) \equiv [h_1(\pi), \dots, h_{\tilde{A}}(\pi)]$, with $h_a(\pi) = h_a(\pi_1, \dots, \pi_A)$ for each $a \in \tilde{\mathcal{A}}$. Finally, 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} . Below, we discuss a number of special cases encountered in applied work.¹³

Affine Payoff Counterfactuals. In affine payoff counterfactuals, 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, so that:

$$\tilde{\pi} = \mathcal{H}\pi + g, \quad (6)$$

where $\mathcal{H} \in \mathbb{R}^{\tilde{A}\tilde{X} \times AX}$ and g is a $\tilde{A}\tilde{X} \times 1$ vector. It is helpful to write this in a block-matrix equivalent form:

$$\tilde{\pi} = \begin{bmatrix} H_{11} & H_{12} & \cdots & H_{1A} \\ \vdots & \vdots & \vdots & \vdots \\ H_{\tilde{A}1} & H_{\tilde{A}2} & \cdots & H_{\tilde{A}A} \end{bmatrix} \pi + g \quad (7)$$

where the submatrices H_{aj} have dimension $\tilde{X} \times X$ for each pair $a \in \tilde{\mathcal{A}}$ and $j \in \mathcal{A}$.

When the counterfactual does not change the set of actions and states (i.e. $\tilde{\mathcal{A}} = \mathcal{A}$ and $\tilde{\mathcal{X}} = \mathcal{X}$), \mathcal{H} is a square matrix. When, further, $\tilde{\pi}_a$ depends solely on π_a , \mathcal{H} is block-diagonal and for all $a \in \mathcal{A}$,

$$\tilde{\pi}_a = H_a \pi_a + g_a. \quad (8)$$

We call these ‘‘action diagonal counterfactuals.’’ Below, we contrast three simple special cases of (8) that are common in applications.

¹²Note that for dynamic games, we can always treat the problem of solving for an individual player’s best response (holding the opponent’s strategy fixed) as a single-agent problem. Our identification results can therefore be applied to investigate identification of counterfactual best responses in dynamic games. A full analysis naturally requires strategic considerations and the possibility of multiple equilibria. See Kalouptsi, Scott, and Souza-Rodrigues (2017a) for discussion of how strategic interactions makes the identification of counterfactuals in dynamic games particularly difficult.

¹³As previously mentioned, the discount factor is typically assumed known. We however allow for changes in β for completeness. When β is identified under further restrictions (Magnac and Thesmar (2002), Abbring and Daljord (2017)), one may be interested in investigating behavior when the discount factor takes different values. For instance, Conlon (2012) studies the evolution of the LCD TV industry when consumers are myopic in the counterfactual experiment (i.e., $\tilde{\beta} = 0$).

Pre-Specified Additive Changes. This counterfactual takes H_a as the identity matrix for all a (i.e. $\mathcal{H} = I$), so that $\tilde{\pi}_a = \pi_a + g_a$. For instance, Keane and Wolpin (1997) investigate a hypothetical college tuition subsidy that pays students \$2,000. Schiraldi (2011) and Li and Wei (2014) study automobile scrappage subsidies that depend on the car’s model and age. Duflo, Hanna and Ryan (2012) implement optimal bonus incentives for teachers in rural India, where the bonus depends on the number of classes the teachers attend.

“Pre-specified additive changes” have been considered by the previous theoretical literature (Aguirregabiria (2010), Aguirregabiria and Suzuki (2014), Norets and Tang (2014), and Arcidiacono and Miller (2017)). Note that “additive changes” are flexible to the extent that g can vary across actions and states. But g is not allowed to depend on π , so the researcher must be able to specify g before estimating the model. Therefore, it is not possible to represent an arbitrary counterfactual $\tilde{\pi} = h(\pi)$ by an “additive changes” in practice: this would require setting $g = h(\pi) - \pi$. In other words, payoffs must be changed by amounts that can be specified *without estimating the model*.

Proportional Changes. This counterfactual sets \mathcal{H} to be diagonal and $g = 0$. It imposes percentage changes on original payoffs, i.e. $\tilde{\pi}_a(x) = \lambda_a(x) \pi_a(x)$. A common example involves entry subsidies represented by percentage changes on entry/sunk costs: for instance, Das, Roberts and Tybout (2007) study firms’ exporting decisions; Varela (2013) studies supermarket entry; Lin (2015) investigates entry and quality investment in the nursing home industry; and Igami (2017) studies innovation in the hard drive industry.

Changes in Types. Another common counterfactual encountered in practice replaces the primitives of one type of agents by those of another, where types can be broadly defined to include markets or regions. For instance, Keane and Wolpin (2010) replace the primitives of minorities by those of white women to investigate the racial-gap in labor markets. Eckstein and Lifshitz (2011) substitute the preference/costs parameters of the 1955’s cohort by those of other cohorts to study the evolution of labor market conditions. Ryan (2012) replaces the entry costs post the Clean Air Act Amendment (CAA) by those before the CAA in the cement industry. Dunne, Klimek, Roberts and Xu (2013) substitute entry costs in Health Professional Shortage Areas (HPSA) by those in the non-HPSA for dentists and chiropractors.

To represent such a counterfactual we can explicitly add a time-invariant state, denoted by s , the type, so that the payoff is written $\pi_a(x, s)$. For example, if there are two types, $s \in \{s_1, s_2\}$, a counterfactual in which the payoff of type s_1 is replaced by that of type s_2 is represented by

$$\begin{bmatrix} \tilde{\pi}_a(s_1) \\ \tilde{\pi}_a(s_2) \end{bmatrix} = \begin{bmatrix} 0 & I \\ 0 & I \end{bmatrix} \begin{bmatrix} \pi_a(s_1) \\ \pi_a(s_2) \end{bmatrix}. \quad (9)$$

Note that H_a is not diagonal in this case.

Changes in Choice Sets and State Space. Eliminating an action leads to $\tilde{\mathcal{A}} = \mathcal{A} - \{j\}$, where j is the action to be eliminated. In this case, $\tilde{\pi}$ satisfies (7) with $H_{aa} = I$ and $H_{ak} = 0$ for $a \in \tilde{\mathcal{A}}$ and $k \in \mathcal{A}$, $a \neq k$. For instance, if $A = 3$ and we drop action $j = 3$, (7) becomes

$$\begin{bmatrix} \tilde{\pi}_1 \\ \tilde{\pi}_2 \end{bmatrix} = \begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \end{bmatrix} \begin{bmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \end{bmatrix}.$$

Note that in some cases, changing the set of actions also changes the set of states (e.g. when $x_t = a_{t-1}$ as in many entry models). Rust and Phelan (1997) eliminate social security in a retirement decision model. Gilleskie (1998) restricts access to medical care in the first days of illness. Crawford and Shum (2005) do not allow patients to switch medications to study the impact of experimentation. Keane and Wolpin (2010) eliminate a welfare program. Keane and Merlo (2010) eliminate the option of private jobs for politicians who leave congress.

A counterfactual that adds a new action can also be represented by (7): take $\tilde{\mathcal{A}} = \mathcal{A} \cup \{j\}$, where j is the new action and let $H_{aa} = I$ and $H_{ak} = 0$ for $a \neq k, j$. Note that, adding an action also requires specifying its payoff $\tilde{\pi}_j$, the new transition matrix \tilde{F}_j , and possibly new states. Rosenzweig and Wolpin (1993), for example, add an insurance option for farmers in rural India.

Changes in Transitions. Finally, a counterfactual that changes the state transition process is represented by a function h_s that transforms F to \tilde{F} . This is the second type of counterfactual that has been considered by the previous theoretical literature (Aguirregabiria and Suzuki (2014), Norets and Tang (2014), and Arcidiacono and Miller (2017)). Such counterfactuals may involve changes in the long-run mean or volatility of market-level variables. For example, Hendel and Nevo (2007) study consumers' long-run responsiveness to prices using supermarket scanner data. Collard-Wexler (2013) explores the effects of demand volatility in the ready-mix concrete industry. Kalouptsi (2014) investigates the impact of time to build on industry fluctuations for the case of the shipping industry. Chan, Hamilton and Papageorge (2016) evaluate the value, and the impact on risky behavior, of an HIV treatment breakthrough (known as HAART) that affects the likelihood of HIV infection.

3.2 Identification of Counterfactual CCP: The General Case

We now present our main theorem, which provides a general framework to investigate identification of counterfactual behavior; then, we turn to the special cases and provide some intuition for the results. The starting point is equation (4). This relationship is convenient for two reasons. First, it does not depend on continuation values. Second, the CCP vector generated by the model

primitives is the unique vector that satisfies (4).¹⁴

The counterfactual $\{\tilde{\mathcal{A}}, \tilde{\mathcal{X}}, \tilde{\beta}, h, h^s\}$ determines a new set of primitives $(\tilde{\mathcal{A}}, \tilde{\mathcal{X}}, \tilde{\beta}, \tilde{\pi}, G, \tilde{F})$, with $\tilde{\pi} = h(\pi)$ and $\tilde{F} = h^s(F)$, which in turn leads to a new optimal behavior: the counterfactual CCP, denoted by \tilde{p} . The counterfactual counterpart to (4) for any $a \in \tilde{\mathcal{A}}$, with $a \neq J$, is

$$\tilde{\pi}_a = \tilde{A}_a \tilde{\pi}_J + \tilde{b}_a(\tilde{p}), \quad (10)$$

where $\tilde{A}_a = (I - \tilde{\beta}\tilde{F}_a)(I - \tilde{\beta}\tilde{F}_J)^{-1}$, $\tilde{b}_a(\tilde{p}) = \tilde{A}_a \psi_J(\tilde{p}) - \psi_a(\tilde{p})$, and we take without loss of generality a reference action J that belongs to both \mathcal{A} and $\tilde{\mathcal{A}}$.

It is clear from (4) and (10) that \tilde{p} is a function of the free parameter π_J . Because the lack of identification of the model is represented by this free parameter, the counterfactual CCP \tilde{p} is identified if and only if it does not depend on π_J . To determine whether or not this is the case, we apply the implicit function theorem to (10).

Before presenting the general case, we consider a binary choice example to fix ideas. Take $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and assume $\tilde{\pi}_a$ is action diagonal so that $\tilde{\pi}_a = h_a(\pi_a)$. Take $J = 2$, and rewrite (10) as

$$h_1(\pi_1) = \tilde{A}_1 h_2(\pi_2) + \tilde{b}_1(\tilde{p}). \quad (11)$$

The implicit function theorem allows us to locally solve (11) with respect to \tilde{p} provided the matrix

$$\frac{\partial}{\partial \tilde{p}} \left[h_1(\pi_1) - \tilde{A}_1 h_2(\pi_2) - \tilde{b}_1(\tilde{p}) \right]$$

is invertible. We prove this matrix is indeed invertible in the general case (see Lemma 1 below). Then, it follows from the implicit function theorem that \tilde{p} does not depend on the free parameter π_2 if and only if

$$\frac{\partial}{\partial \pi_2} \left[h_1(\pi_1) - \tilde{A}_1 h_2(\pi_2) - \tilde{b}_1(\tilde{p}) \right] = 0.$$

Because $\pi_1 = A_1 \pi_2 + b_1(p)$ from (4), the above equation simplifies to

$$\frac{\partial h_1(\pi_1)}{\partial \pi_1} A_1 - \tilde{A}_1 \frac{\partial h_2(\pi_2)}{\partial \pi_2} = 0. \quad (12)$$

The equality depends on the (known) counterfactual transformation $\{h, h^s\}$ and on the data, F , through A_1 and \tilde{A}_1 . So, in practice, one only needs to verify whether (12) holds for the particular combination $\{h, h^s\}$ of interest.

To facilitate the passage to the general case, rearrange the equality above in matrix form as

¹⁴Note that a unique CCP vector p is indeed guaranteed from (4): since the Bellman is a contraction mapping, V is unique; from (1) so are v_a and thus so is p .

follows:

$$\left[\frac{\partial h_1(\pi_1)}{\partial \pi_1} - \tilde{A}_1 \frac{\partial h_2(\pi_2)}{\partial \pi_2} \right] \begin{bmatrix} A_1 \\ I \end{bmatrix} = 0$$

or

$$\begin{bmatrix} I & -\tilde{A}_1 \end{bmatrix} \begin{bmatrix} \frac{\partial h_1(\pi)}{\partial \pi_1} & \frac{\partial h_1(\pi)}{\partial \pi_2} \\ \frac{\partial h_2(\pi)}{\partial \pi_1} & \frac{\partial h_2(\pi)}{\partial \pi_2} \end{bmatrix} \begin{bmatrix} A_1 \\ I \end{bmatrix} = 0,$$

where, in this example, $\frac{\partial h_1(\pi)}{\partial \pi_2} = \frac{\partial h_2(\pi)}{\partial \pi_1} = 0$. Using the property of the Kronecker product $vec(ABC) = (C' \otimes A)vec(B)$, our condition becomes:

$$\left(\begin{bmatrix} A'_1 & I \end{bmatrix} \otimes \begin{bmatrix} I & -\tilde{A}_1 \end{bmatrix} \right) vec(\nabla h(\pi)) = 0,$$

where $\nabla h(\pi)$ is the matrix with elements $\frac{\partial h_a(\pi)}{\partial \pi_j}$ for $a, j = 1, 2$. So, to identify the counterfactual CCPs, $vec(\nabla h(\pi))$ must lie in the nullspace of a matrix determined by A_1 and \tilde{A}_1 .

Moving from the binary to the general model, take (10) together with $\tilde{\pi}_a = h_a(\pi)$ and stack all π_a for $a \neq J$ to obtain:

$$h_{-J}(\pi) = \tilde{A}_{-J}h_J(\pi) + \tilde{b}_{-J}(\tilde{p}), \quad (13)$$

where $h_{-J}(\pi)$ stacks $h_a(\pi)$ for all $a \in \tilde{\mathcal{A}}$ except for J , and the matrix \tilde{A}_{-J} and vector $\tilde{b}_{-J}(\tilde{p})$ are defined similarly. The next lemma guarantees that the implicit function theorem can be applied to (13).

Lemma 1. *The function $\tilde{b}_{-J}(\cdot)$ is continuously differentiable and its Jacobian is everywhere invertible.*

We state our main theorem below. $vecbr(C)$ rearranges the blocks of matrix C into a block column by stacking the block rows of C ; the symbol \boxtimes denotes the block Kronecker product.¹⁵

Theorem 1. *Consider the counterfactual transformation $\{\tilde{\mathcal{A}}, \tilde{\mathcal{X}}, \tilde{\beta}, h, h^s\}$ and suppose h is differentiable. The counterfactual conditional choice probabilities \tilde{p} are identified if and only if for all π satisfying (5),*

$$Q(A, \tilde{A})vecbr(\nabla h(\pi)) = 0, \quad (14)$$

where

$$Q(A, \tilde{A}) = \left[\begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes I, - \begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes \tilde{A}_{-J} \right].$$

¹⁵The block Kronecker product, \boxtimes , of two partitioned matrices B and C is defined by (Koning, Neudecker and Wansbeek (1991)):

$$B \boxtimes C = \begin{bmatrix} B \otimes C_{11} & \dots & B \otimes C_{1b} \\ \vdots & \ddots & \vdots \\ B \otimes C_{c1} & \dots & B \otimes C_{cb} \end{bmatrix}.$$

Note that at the entry level, Kronecker rather than ordinary products are employed.

The matrix $Q(A, \tilde{A})$ has dimension $(\tilde{A} - 1)\tilde{X}X \times (\tilde{A}\tilde{X})(AX)$, while $\text{vecbr}(\nabla h(\pi))$ has dimension $(\tilde{A}\tilde{X})(AX) \times 1$.

Theorem 1 holds that counterfactual choice probabilities \tilde{p} are identified if and only if the Jacobian matrix of h is restricted to lie in the nullspace of a matrix defined by A_{-J} and \tilde{A}_{-J} , which in turn are determined by the transition probabilities F and \tilde{F} . So model primitives, data and counterfactual transformations have to interact with each other in a specific way to obtain identification of counterfactual CCPs. The only requirement is that h is differentiable – this is a mild restriction typically satisfied in practice.¹⁶

Equation (14) is the minimal set of sufficient conditions that applied researchers need to verify to secure identification of counterfactual behavior. For instance, for “action diagonal” counterfactuals, i.e. $\tilde{\pi}_a = h_a(\pi_a)$, equation (14) is substantially simplified: \tilde{p} is identified if and only if for all π satisfying (5) and all $a \in \tilde{\mathcal{A}}$, $a \neq J$,

$$\frac{\partial h_a}{\partial \pi_a} A_a - \tilde{A}_a \frac{\partial h_J}{\partial \pi_J} = 0. \quad (15)$$

This implies that it is particularly difficult to identify counterfactual behavior when payoffs change non-linearly, since equation (15) must be satisfied for *all* admissible payoffs π .¹⁷

Finally, it is worth noting that adding a known vector to $\tilde{\pi}$ (e.g. $\tilde{\pi} = h(\pi) + g$) does not affect the Jacobian matrix of h , and so whether \tilde{p} is identified or not does not depend on vector g .

Example: Rust’s Bus Engine Replacement Problem. Rust (1987) investigates the optimal stopping problem of replacing a bus’s engine, trading-off aging and replacement costs. The choice set is $\mathcal{A} = \{\text{replace, keep}\}$; the state variable, x , is the bus mileage which evolves stochastically and is renewed upon replacement; and the payoff function is

$$\pi(a, x) = \begin{cases} -\phi(x) - c(0), & \text{if } a = \text{replace} \\ -c(x), & \text{if } a = \text{keep} \end{cases}$$

where $\phi(x)$ is the cost of replacing an engine (net of scrap value); and $c(x)$ is the operating cost at mileage x . In the counterfactual analysis, Rust varies the level of replacement costs and obtains the corresponding (long run) replacement choice probabilities, or a demand curve for engine replacement. One way to represent his counterfactual is to consider counterfactual payoffs as $\tilde{\pi}(\text{replace}, x) = -(1 + \lambda)\phi(x) - c(0)$, for various levels of λ , where λ is a parameter capturing

¹⁶ Although the implicit function theorem involves local conditions, equation (14) must be satisfied for *all* payoffs that rationalize observed choice probabilities; i.e. for all π satisfying (5). Note also that the choice of the reference action J does not affect whether or not (14) is satisfied.

¹⁷ One family of counterfactuals that satisfies (15) is a class of periodic functions satisfying $\frac{\partial h}{\partial y}(By + c) = \frac{\partial h(y)}{\partial y}$.

the shift in replacement costs. Equivalently,

$$\tilde{\pi}(\text{replace}, x) = \pi(\text{replace}, x) + \lambda(\pi(\text{replace}, x) - \pi(\text{keep}, 0)).$$

Each value of λ corresponds to one point along the demand curve for engine replacement. The counterfactual does not affect $\pi(\text{keep}, x)$, nor β or F . In Appendix A, we show that Theorem 1 implies that Rust’s counterfactual is not identified. Showing this for a particular specification requires only a simple calculation evaluating equation (14).¹⁸

Rust’s counterfactual falls within the class of affine payoff transformations, a class of counterfactuals for which equation (14) is simplified. As we show below, we can derive more intuitive conditions for the identification of such counterfactuals.

3.3 Identification of Counterfactual CCP: Special Cases

In this section, we discuss several special cases of interest following the taxonomy presented in Section 3.1. As Theorem 1 states generally, a counterfactual is identified when the way the counterfactual changes payoffs satisfies conditions that depend on the state transition process. Corollary 1 shows how such conditions simplify when the payoff transformation $h(\cdot)$ is affine, i.e., $\tilde{\pi} = \mathcal{H}\pi + g$, or $\tilde{\pi}_a = \sum_{j \in \mathcal{A}} H_{aj}\pi_j + g_a$. The affine case is prevalent in applied work.

Corollary 1. (“Affine Payoff” Counterfactual) Assume $\tilde{\pi} = \mathcal{H}\pi + g$.

(i) The counterfactual CCP \tilde{p} is identified if and only if for all π satisfying (5) and all $a \in \tilde{\mathcal{A}}$, $a \neq J$,

$$\sum_{l \in \mathcal{A}, l \neq J} \left(H_{al} - \tilde{A}_a H_{Jl} \right) A_l + H_{aJ} - \tilde{A}_a H_{JJ} = 0. \quad (16)$$

(ii) Further, if the counterfactual is “action diagonal,” $\tilde{\pi}_a = H_a \pi_a + g_a$, this becomes, for all $a \in \tilde{\mathcal{A}}$,

$$H_a A_a - \tilde{A}_a H_J = 0. \quad (17)$$

Recalling that $A_a = (I - \beta F_a)(I - \beta F_J)^{-1}$, and noting that the transition matrices F_a can be

¹⁸ The representation is appropriate, for instance, if the replacement cost $\phi(x)$ represents labor costs of rebuilding an old engine (which potentially varies with the mileage on the engine), and the counterfactual of interest involves increasing wages. Formally, $\tilde{\pi} = \mathcal{H}\pi$, where \mathcal{H} is not block-diagonal:

$$\mathcal{H} = \begin{bmatrix} (1 + \lambda)I & [-\lambda \mathbf{1}, 0] \\ 0 & I \end{bmatrix},$$

where I is the identity matrix and $\mathbf{1}$ is a vector of ones. To evaluate equation (14), consider a simple version of the model in which $a = \text{replace}$, and where the state space is simply $\mathcal{X} = \{\text{new}, \text{old}\}$ with deterministic transitions. Then $Q(A, \tilde{A}) = [A_{\text{replace}} \otimes I, I \otimes I, A'_{\text{replace}} \otimes A_{\text{replace}}, I \otimes A_{\text{replace}}]$. Finally, with $\beta = .99$ and $\lambda = 0.1$ (representing a 10% increase in replacement costs), we obtain $\|Q(A, \tilde{A}) \text{vec}(\nabla h(\pi))\| = 1.89$, where $\|\cdot\|$ is the matrix 2-norm. Equation (14) is violated, implying the counterfactual is not identified.

estimated, it is clear that conditions (16) and (17) can be easily verified from the data. The next set of results make direct use of Corollary 1.

Changes in Payoffs. We now consider counterfactuals that only change agents’ payoff functions, holding fixed the remaining primitives. As already mentioned, previous work on transformations of payoffs have shown that one particular class of transformations yield identified counterfactual behavior: “pre-specified additive changes,” which are of the form $\tilde{\pi}(a, x) = \pi(a, x) + g(a, x)$ (Aguirregabiria (2010), Aguirregabiria and Suzuki (2014)). Norets and Tang (2014) also proved identification when $\tilde{\pi} = \lambda\pi + g$, where λ is a scalar. The identification of \tilde{p} for this class of counterfactuals is an immediate implication of Corollary 1. Note that in this case, we have $H_a = \lambda I$ for all a , and so equation (17) is clearly satisfied.

Affine payoff counterfactuals are substantially more general than the pre-specified additive changes addressed by earlier work, allowing for changes that depend on parameters that the econometrician must estimate. For instance, the counterfactual change in bus engine replacement costs in the example presented in Section 3.2 is affine, but it cannot be expressed as a pre-specified additive change.

Following the taxonomy, we now consider “proportional changes” counterfactuals. For this class, recall that we take $\tilde{\pi}_a(x) = \lambda_a(x)\pi_a(x)$, or, compactly, $\tilde{\pi}_a = H_a\pi_a$, with H_a diagonal.

Proposition 2. (“Proportional Changes”) Consider $\tilde{\pi}_a = H_a\pi_a$ with H_a diagonal. Assume $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and $\tilde{F} = F$. Then, to identify \tilde{p} it is necessary that $H_a = H$, for all a .

Suppose further that H has d distinct diagonal elements $\lambda_1, \dots, \lambda_d$, each occurring n_1, \dots, n_d times so that H can be written as $H = \text{diag}(\lambda_1 I_{n_1}, \dots, \lambda_d I_{n_d})$. The following statements are equivalent:

- (i) \tilde{p} is identified
- (ii) A_a is block diagonal with diagonal blocks, $(A_a)_i$, of conformable sizes n_1, \dots, n_d
- (iii) Let $(F_a)_{ij}$ be the $n_i \times n_j$ submatrix of F_a that conforms with $H = \text{diag}(\lambda_1 I_{n_1}, \dots, \lambda_d I_{n_d})$. For all $a \in \mathcal{A}$ and $i \neq j$, the block partitions of F_a and F_J satisfy

$$(F_a)_{ij} = (A_a)_i (F_J)_{ij}, \quad (18)$$

where $(A_a)_i \equiv (I - \beta(F_a)_{ii})(I - \beta(F_J)_{ii})^{-1}$.

The necessary conditions to identify \tilde{p} in the case of “proportional changes” are restrictive. For one, if we change the payoff of action a in state x by $\lambda(x)$, $\tilde{\pi}(a, x) = \lambda(x)\pi(a, x)$, then it is necessary to change the payoff of any other action a in state x by exactly the same proportion $\lambda(x)$. Furthermore, identification requires special conditions on the A_a matrices (part ii), which are equivalent to special conditions on the transition process F (part iii). In particular, an implication of the proposition is that, if all diagonal elements of H are pairwise distinct, then identification of \tilde{p} requires $F_a = F_J$ for all $a \in \mathcal{A}$. This condition however will not be satisfied in any dynamic

model of interest.¹⁹

Another set of payoff transformations involves assigning the payoffs of one type of agent to those of another type; i.e., the “changes in types” counterfactual. We prove the following proposition for two types and two actions for notational simplicity; the extension to multiple types and actions is entirely straightforward.

Proposition 3. (*“Changes in Types”*) *Suppose the payoff is $\pi_a(x, s)$, where s is a time-invariant state (type) that takes two values, $s \in \{s_1, s_2\}$ and that there are two actions $\mathcal{A} = \{a, J\}$. Suppose also that $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and $\tilde{F} = F$.*

(i) *If the counterfactual replaces the payoff of type s_1 by that of s_2 for one of the actions, then \tilde{p} is not identified.*

(ii) *If the counterfactual replaces the payoff s_1 by that of s_2 for all actions, then \tilde{p} is identified if and only if $(I - \beta F_a^{s_1})(I - \beta F_J^{s_1})^{-1} = (I - \beta F_a^{s_2})(I - \beta F_J^{s_2})^{-1}$, where F_a^s is the transition matrix corresponding to type s .*

Proposition 3 results in nonidentification if payoffs of a subset of actions are replaced; and requires strong restrictions on transition probabilities if payoffs at all actions are replaced (much like the “proportional changes” case). It is worth noting that if the types have the same transitions ($F_a^{s_1} = F_a^{s_2}$), the condition is satisfied and the counterfactual is identified.

Proposition 2 and 3 express specific and verifiable restrictions on the transition process that must hold for the identification of counterfactual behavior. A natural question then is whether some transformations of payoffs – and which – can be said to be identified for *any* transition process. In other words, when can we say a counterfactual transformation of payoffs is identified without having to estimate or specify the specific transition process? In such a case, the researcher can establish identification ex ante, regardless of the data at hand. Our next result holds that only a limited subset of payoff transformations can be said to be identified without verifying the restrictions on the transition process.

Proposition 4. *Assume $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, $\tilde{F} = F$, and $\tilde{\pi} = \mathcal{H}\pi + g$. Counterfactual behavior \tilde{p} is identified without restrictions on the transition process F (i.e., counterfactual behavior is identified for every transition F) if and only if*

- (a) $\mathcal{H} = \mathcal{H}_1 + \mathcal{H}_2$, where $\mathcal{H}_1 = \lambda I$, λ is a scalar, and \mathcal{H}_2 has identical rows;
- (b) or, equivalently, the transformation of payoffs can be expressed in the following form for all $a \in \mathcal{A}$ and $x \in \mathcal{X}$:

$$\tilde{\pi}_a(x) = \lambda\pi_a(x) + L(\pi) + g_a(x) \tag{19}$$

where $L(\cdot)$ is the scalar-valued function given by $L(\pi) = \sum_{j \in \mathcal{A}} \sum_{x \in \mathcal{X}} \rho_{jx} \pi_j(x)$, and the vector $[\rho_{11}, \dots, \rho_{AX}]$ corresponds to one row of \mathcal{H}_2 .²⁰

¹⁹In general, if a diagonal element of H is unique, i.e. $n_i = 1$ for some i , identification requires that the i -row of F_a and F_J are identical. That is, conditional on state x_i , the transition probabilities do not depend on the action.

²⁰To connect parts (a) and (b), note that $\tilde{\pi} = \mathcal{H}\pi + g = \mathcal{H}_1\pi + \mathcal{H}_2\pi + g$, and that $\mathcal{H}_2\pi$ is a constant vector

Equation (19) shows that only the following changes are identified regardless of the transition process:

1. Pre-specified additive changes $g_a(x)$ which may depend arbitrarily on actions and states but does not depend on the baseline payoff function.
2. Multiplication of baseline payoffs by a scalar λ , which does not depend on actions or states. For $\lambda > 0$, this resembles a *change in the scale of the payoff function*.²¹
3. Addition of a scalar-valued function $L(\pi)$, which does not depend on actions or states. This corresponds to a *change in the level of the payoff function*.²²

Proposition 4 states that the only meaningful counterfactual transformations of payoffs that can be said to be identified with no restrictions on the state transition process are (1) pre-specified additive changes and (2) changes in the level and scale of the payoff function. As described above, different versions of the “if” direction of Proposition 4 have been proved in the literature. Our result shows that *any other* counterfactual transformations of payoffs are identified only under restrictions on the state transition processes that require verification in the data. Theorem 1 provides the conditions the transition process must satisfy most generally; Corollary 1 and Propositions 2 and 3 provide simpler conditions for certain classes of payoff transformations.

Changes in Transitions. For completeness, we include a corollary of Theorem 1, first proven by Aguirregabiria and Suzuki (2014) and Norets and Tang (2014), regarding counterfactuals that only change the state transitions.

Corollary 2. (“Change in Transition”) Assume $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and $\tilde{\pi} = \pi$, but $\tilde{F} \neq F$. Then \tilde{p} is identified if and only if $A_a = \tilde{A}_a$, for all $a \in \tilde{\mathcal{A}}$, $a \neq J$.

Corollary 2 is an immediate implication of equation (17). Identification again requires restrictions on state transitions. However, in Section 4.1, we explain how parametric assumptions on the payoff function can lead to less restrictive requirements for the identification of this class of counterfactuals.²³

because \mathcal{H}_2 has identical rows. So, $\mathcal{H}_2\pi$ corresponds to $L(\pi)1$, where 1 is a vector of ones.

²¹Strictly speaking, multiplication of π by a positive scalar is *not* equivalent to a scale normalization of the whole utility function, $\pi + \varepsilon$, because the distribution of the idiosyncratic shocks ε is fixed. More formally, $\tilde{\pi} = \lambda\pi$ for $\lambda > 0$ is equivalent to multiplying the variance of the idiosyncratic shocks by λ^{-2} .

²²Note that a shift in the level of payoffs by the same number L for every action and state does not affect agents’ incentives.

²³Note that identification of a counterfactual that only changes the discount factor, $\tilde{\beta} \neq \beta$, requires the same condition: $A_a = \tilde{A}_a$, for all $a \neq J$, but with $\tilde{A}_a = (I - \tilde{\beta}F_a)(I - \tilde{\beta}F_J)^{-1}$.

Changes in Choice Sets and State Space. Consider now a counterfactual that adds an option to agent’s choice set. This counterfactual naturally requires pre-specifying $\tilde{\pi}_j$ and \tilde{F}_j for the new choice j .

Proposition 5. (“Add an Action” Counterfactual) Suppose $\tilde{\mathcal{A}} = \mathcal{A} \cup \{j\}$, where j is the new action. Assume $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, $\tilde{F} = F$, $\tilde{\pi}_a = \pi_a$ for all $a \in \mathcal{A}$, and

$$\tilde{\pi}_j = \sum_{a \in \mathcal{A}} H_{ja} \pi_a + g_j.$$

Let $\mathbf{1}$ be an $X \times 1$ vector of ones. Then \tilde{p} is identified if and only if $\sum_{a \in \mathcal{A}} H_{ja} \mathbf{1} = 1$, and

$$\tilde{F}_j = \sum_{a \in \mathcal{A}} H_{ja} F_a + \beta^{-1} \left(I - \sum_{a \in \mathcal{A}} H_{ja} \right).$$

In words, to obtain identification it is necessary that the payoff of the new action j is a “convex combination” of existing payoffs, and the new transition matrix is an “affine” combination of existing transitions. This is reminiscent of predicting consumers’ choices when a new good is introduced in a static differentiated product demand framework. Predicting the demand for the new good requires that the attributes of the new good are a combination of the attributes of existing goods in the market. The same applies in the dynamic context; here we additionally need restrictions on the transitions in order to predict behavior when a new choice is available.

Consider next a counterfactual that eliminates one action (extensions to eliminating more actions are straightforward).

Proposition 6. (“Eliminate an Action” Counterfactual) Suppose $\tilde{\mathcal{A}} = \mathcal{A} - \{j\}$, where j is the action to be eliminated. If $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, $\tilde{F}_a = F_a$, and $\tilde{\pi}_a = \pi_a$ for all $a \in \tilde{\mathcal{A}}$, then \tilde{p} is identified.

Here, the key to identification is that transitions do not change. However, elimination of an action often implies elimination of some states as well (e.g. when $x_t = a_{t-1}$ as in many entry models), which necessarily changes transitions. In that case, identification depends on how the probability mass is reallocated from \mathcal{X} into the remainder set of states $\tilde{\mathcal{X}}$. Lemma 2 in Appendix A provides the necessary and sufficient conditions for identification in this case. Below, we consider a special case that is common in applied work. Decompose the state variables as $x = (k, w)$, where $k_t = a_{t-1}$, and w is an exogenous state (i.e., its evolution does not depend on choices a). Formally, $F_a = F^w \otimes F_a^k$, where F^w is the transition matrix for w , F_a^k is the transition for k , and \otimes denotes the Kronecker product. The monopolist entry problem presented in Section 6 satisfies these restrictions. The counterfactual CCP is indeed identified in this case.

Proposition 7. (“Eliminate an Action and States” Counterfactual) Suppose $\tilde{\mathcal{A}} = \mathcal{A} - \{j\}$, where j is the action to be eliminated. Without loss of generality, let the set of states be $\tilde{\mathcal{X}} = \{1, \dots, \bar{x}\}$

and $\mathcal{X} = \{1, \dots, \bar{x}, \bar{x} + 1, \dots, X\}$. Assume $\tilde{\pi}_a = H_a \pi_a$ with $H_a = [I_{\bar{x}}, 0]$ for all $a \in \tilde{\mathcal{A}}$, where $I_{\bar{x}}$ is the $\bar{x} \times \bar{x}$ identity matrix. Suppose $x = (w, k)$ with transition matrix $F_a = F^w \otimes F_a^k$ and $k_t = a_{t-1}$. Then, the counterfactual CCP \tilde{p} is identified.

3.4 Some Intuition

Intuitively, the reason payoffs are not identified in a DDC model is twofold. First, although we can identify the difference in continuation values $v_a - v_J$ from the observed choice probabilities (Hotz and Miller (1993)), the discrete choice nature of the data does not allow us to separate v_a from v_J . Second, we also cannot separate the two components of $v_a(x)$ nonparametrically, i.e. $\pi(a, x)$ and $E[V(x') | a, x]$, since they both depend on the same arguments.

To obtain some intuition for why some counterfactuals are identified while others are not, take a simple example of a binary choice, $\mathcal{A} = \{a, J\}$ and consider the counterfactual $\tilde{\pi}_a = H_a \pi_a + g_a$, for all a , with $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, and $\tilde{\beta} = \beta$. Next, note that by rearranging our main equation (4) in Proposition 1 (i.e. $\pi_a = A_a \pi_J + b_a(p)$), we obtain

$$(I - \beta F_a)^{-1} \pi_a - (I - \beta F_J)^{-1} \pi_J = (I - \beta F_a)^{-1} b_a(p).$$

The left-hand side is the difference of the expected discounted present value obtained by always choosing a versus always choosing J . This difference is known, as the right-hand side is known.

For the counterfactual scenario, rearrange the equivalent counterfactual equation (10) as above. Assuming the counterfactual is identified (i.e., $A_a H_a = \tilde{A}_a H_J$), it is easy to show that

$$H_a (I - \beta F_a) [(I - \beta F_a)^{-1} \pi_a - (I - \beta F_J)^{-1} \pi_J] + (g_a - \tilde{A}_a g_J) = \tilde{b}_a(\tilde{p}).$$

In words, the counterfactual CCP \tilde{p} depends on π *only through* the difference in present values: $(I - \beta F_a)^{-1} \pi_a - (I - \beta F_J)^{-1} \pi_J$. All other terms of the left-hand side, as well as the function \tilde{b}_a , are known. So, to calculate \tilde{p} , it is not necessary to identify the payoff function π .

To make the argument more transparent, consider the “additive changes” counterfactual (see Aguirregabiria (2010), and Norets and Tang (2014)). The left-hand side now becomes the sum of two terms: $(I - \beta F_a)^{-1} \pi_a - (I - \beta F_J)^{-1} \pi_J$ and $(I - \beta F_a)^{-1} g_a - (I - \beta F_J)^{-1} g_J$. Both terms are known and thus the change in the choice probabilities is also known. There is no need to re-optimize agents’ dynamic behavior in the counterfactual scenario. This is possible only because the counterfactual difference $(I - \beta \tilde{F}_a)^{-1} \tilde{\pi}_a - (I - \beta \tilde{F}_J)^{-1} \tilde{\pi}_J$ is a *known* function of the *observed* difference $(I - \beta F_a)^{-1} \pi_a - (I - \beta F_J)^{-1} \pi_J$.

When the equality $A_a H_a = \tilde{A}_a H_J$ is not satisfied, the counterfactual difference in continuation values depends *directly* on π . It is no longer sufficient to know the observed differences $(I - \beta F_a)^{-1} \pi_a - (I - \beta F_J)^{-1} \pi_J$ and therefore we cannot identify the counterfactual behavior. At least not without additional restrictions.

4 Identification of Counterfactual Behavior Under Additional Model Restrictions

In this section, we investigate when additional restrictions that may be insufficient to identify the full model may suffice to identify the counterfactual. We start with parametric models, and then we move to more general restrictions on payoffs.

4.1 Parametric Restrictions

Parametric Model. More often than not, applied work relies on parametric restrictions. We thus consider identification of a specific parametric model that is common in the literature.

We decompose the state space into two components, $x = (k, w)$, where $k \in \mathcal{K} = \{1, \dots, K\}$ are controlled states (i.e. their evolution is affected by agents' choices, and $w \in \mathcal{W} = \{1, \dots, W\}$ are exogenous (e.g. market-level states), with K, W finite. Therefore,

$$F(x'|a, x) = F^k(k'|a, k) F^w(w'|w). \quad (20)$$

and the transition matrix F_a is written as $F_a = F^w \otimes F_a^k$, where \otimes denotes the Kronecker product. In addition, the following parametric payoff is adopted:

$$\pi(a, k, w) = \theta_0(a, k) + Z(a, w)' \theta_1(a, k), \quad (21)$$

where $Z(a, w)$ is a known function of actions and states w (e.g. observed measures of variable profits or returns) and $\theta_0(a, k)$ is interpreted as a fixed cost component. For instance, in the monopolist's entry/exit problem considered in Section 6.1, $Z(a, w)' \theta_1(a, k)$ represents variable profits which may be either directly observed or a flexible function of observables such as market size and input prices, while $\theta_0(a, k)$ denotes entry/exit/fixed cost depending on the action and state.

Proposition 8 provides sufficient conditions for the identification of this parametric model. For notational simplicity, we focus on binary choice with $\mathcal{A} = \{a, J\}$ and assume $Z(a, w)$ is scalar. The proposition also holds in the more general case of $F^w(w'|w, a)$ and multivariate $Z(a, w)$.

Proposition 8. *Assume (20) and (21) hold. Let*

$$\begin{aligned} D_a &= (I - \beta F_a)^{-1}, \\ Z_a &= [Z_a(w_1)I_k, \dots, Z_a(W)I_k]', \end{aligned}$$

and similarly for D_J and Z_J . I_k is the identity matrix of size K and $e'_w = [0, 0, \dots, I_k, 0, \dots, 0]$ with

I_k in the w position. Suppose $W \geq 3$ and there exist w, \tilde{w}, \bar{w} such that the matrix

$$\begin{bmatrix} (e'_w - e'_{\tilde{w}}) D_a Z_a & (e'_{\tilde{w}} - e'_w) D_J Z_J \\ (e'_w - e'_{\bar{w}}) D_a Z_a & (e'_{\bar{w}} - e'_w) D_J Z_J \end{bmatrix} \quad (22)$$

is invertible. Then the parameters $[\theta_1(a, k), \theta_1(J, k)]$ are identified, but $[\theta_0(a, k), \theta_0(J, k)]$ are not identified.

Intuitively, the “slope” coefficients θ_1 are identified provided there is “sufficient variation” in w (guaranteed by the invertibility of matrix (22)). This requires w to significantly change the conditional expected values of Z_a and Z_J (evidently, it is necessary that $Z_a \neq Z_J$).²⁴ The “intercept” parameters θ_0 , however, are not identified. To identify this vector, we have to add K linearly independent restrictions to the model, much as we have to impose X linearly independent restrictions in the nonparametric setting.

Counterfactuals. In addition to the counterfactuals discussed in Section 3.1, one may be interested in changes in either θ_0 , i.e.:

$$\tilde{\pi}(a, k, w) = h_0 [\theta_0(a, k)] + Z(a, w)' \theta_1(a, k); \quad (23)$$

or in $Z'\theta_1$, i.e.:

$$\tilde{\pi}(a, k, w) = \theta_0(a, k) + h_1 [Z(a, w)' \theta_1(a, k)]. \quad (24)$$

These counterfactuals allow for changes in how the flow payoff responds to some state (k, w) , or to the outcome variable Z .

We show that transformations in $Z'\theta_1$ result in identified counterfactuals, while transformations in $\theta_0(a, k)$ may not. Indeed, since $\theta_1(a, k)$ is identified, a counterfactual that changes $Z'\theta_1$ resembles a “pre-specified additive change,” which is an identified counterfactual (see Section 3.3). In contrast, since $\theta_0(a, k)$ is not identified, one needs to follow our analysis of counterfactuals for nonparametric payoffs to establish whether a particular counterfactual is identified or not.

To illustrate, we consider affine action-diagonal counterfactuals as an example. In particular, let:

$$\tilde{\theta}_0(a) = H_0(a)\theta_0(a) \quad (25)$$

for $a = 1, \dots, J$, $\theta_0(a)$ is obtained by stacking $\theta_0(a, k)$ for all k , and $H_0(a)$ is a $K \times K$ matrix. Extending to a more general (differentiable) function of θ_0 is straightforward.

Proposition 9. (*Parametric Model*) Assume $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and that the conditions of Proposition 8 hold.

²⁴ Note that $e'_w D_a Z_a$ is the expected discounted value of Z when action a is always chosen conditional on observing state w today.

(i) \tilde{p} is identified when the counterfactual only changes the term $Z(a, w)' \theta_1(a, k)$ of $\pi(a, k, w)$ as in (24).

(ii) \tilde{p} is identified under the affine action diagonal counterfactual (25) if and only if for all $a \neq J$

$$H_0(a)A_a^k - A_a^k H_0(J) = 0$$

where $A_a^k = (I - \beta F_a^k)(I - \beta F_J^k)^{-1}$.

(iii) \tilde{p} is identified when the counterfactual only changes the transition F_a^k if and only if $A_a^k = \tilde{A}_a^k$, $a \neq J$, where $\tilde{A}_a^k = (I - \beta \tilde{F}_a^k)(I - \beta \tilde{F}_J^k)^{-1}$.

(iv) \tilde{p} is identified when the counterfactual changes the transition F^w .

In a nonparametric setting, changes in the transition process generically result in non-identified counterfactual behavior (in the sense that the necessary conditions are bound to be restrictive). However, Proposition 9 shows that the intuition from the nonparametric setting does not necessarily carry over to parametric models. When a counterfactual changes the transition process for state variables that are part of the identified component of the payoff function, counterfactual behavior is identified. For instance, the response to a change in the volatility of demand shocks in the monopolist entry/exit example is identified. Even though Aguirregabiria and Suzuki (2014) and Norets and Tang (2014) have explored changes in transitions in the nonparametric context, most implementations of these counterfactuals in practice are done in the parametric context (Hendel and Nevo (2007), Collard-Wexler (2013)) and so based on our results, are in fact identified.

4.2 Linear Restrictions on Payoffs

The parametric payoff function (21) is a special case of linear restrictions on payoffs (another example is “exclusion restrictions”). Indeed, consider a set of $d \leq X$ linearly independent payoff restrictions:

$$R\pi = r \tag{26}$$

with $R \in \mathbb{R}^{d \times AX}$, or in block-form, $R = [R_{-J}, R_J]$.²⁵ To incorporate all model restrictions in a single equation, rewrite (26) as:

$$R_{-J}\pi_{-J} + R_J\pi_J = r \tag{27}$$

and combine it with our main relationship, $\pi_{-J} = A_{-J}\pi_J + b_{-J}(p)$, to obtain:

$$(R_{-J}A_{-J} + R_J)\pi_J = r - R_{-J}b_{-J}(p).$$

²⁵To see how the parametric model (21) can be written as (26) take a simple example where $w \in \mathcal{W} = \{w_1, w_2, w_3\}$. Define the known scalar $z^a = \frac{Z(a, w_1) - Z(a, w_2)}{Z(a, w_1) - Z(a, w_3)}$. Then, $\pi(a, k, w_1) - \pi(a, k, w_2) = z^a [\pi(a, k, w_1) - \pi(a, k, w_3)]$, for any $k \in \mathcal{K}$. Rearranging and stacking all equalities for the different k 's deliver the result.

This is of the form:

$$Q\pi_J = q, \tag{28}$$

with $Q = R_{-J}A_{-J} + R_J \in \mathbb{R}^{d \times X}$ and $q = r - R_{-J}b_{-J}(p) \in \mathbb{R}^d$. Equation (28) is useful because it encompasses all restrictions imposed on the model. When $d = X$ and Q is invertible, π is identified. Conversely, if $\text{rank}(Q) < X$, the model is not identified.

Now consider, for simplicity, an affine counterfactual, $\tilde{\pi} = \mathcal{H}\pi + g$. As shown in Section 3.3, if equation (16) of Corollary 1 holds, the counterfactual CCPs are identified with no extra restrictions (R, r) . Stack the left hand side of (16) for all a and denote the resulting matrix by C . For instance, in the case of action-diagonal counterfactuals (assuming J is the last action),

$$C \equiv \begin{bmatrix} H_1 & \dots & 0 \\ \vdots & \vdots & \vdots \\ 0 & \dots & H_{J-1} \end{bmatrix} A_{-J} - \tilde{A}_{-J}H_J,$$

or, more compactly, $C = H_{-J}A_{-J} - \tilde{A}_{-J}H_J$. So, if (16) holds, $C = 0$ and \tilde{p} is identified regardless of any extra restrictions (R, r) .

In short, at one extreme we impose enough restrictions, make Q invertible, identify π and so \tilde{p} is also identified for any counterfactual. At the other extreme, we impose no restriction but consider a counterfactual such that $C = 0$, in which case \tilde{p} is also identified. The intermediate cases are investigated below. For these, the interaction between restrictions Q and counterfactuals C may result in an identified \tilde{p} . The following proposition presents the necessary condition to obtain identification.

Proposition 10. *Assume the linear restrictions (26) hold. The counterfactual CCP is identified only if there exists an $(A - 1)X \times d$ matrix M such that $C = MQ$.²⁶*

Proposition 10 establishes a direct relationship between the counterfactual transformation (embedded in C) and the restrictions (embedded in Q). An immediate implication is that, for any given counterfactual \mathcal{H} such that $\text{rank}(C) = c$, it is necessary to impose *at least* $d \geq c$ linearly independent restrictions on π to identify \tilde{p} . Evidently, when $\text{rank}(C) = X$, identification of \tilde{p} requires identification of the full model (i.e., $d = X$). Yet, it is possible to identify the counterfactual with fewer restrictions when $\text{rank}(C) < X$ (the parametric model of Section 4.1 is an example of that).²⁷

The next corollary provides a simple way to verify in practice whether a combination of restrictions and counterfactual transformation identifies \tilde{p} .

²⁶For instance, when $C = 0$, we take $M = 0$; and when $C \neq 0$ and Q is invertible, we take $M = CQ^{-1}$.

²⁷If the researcher has reasons to believe the restrictions (26) are true, Proposition 10 allows us to explore the range of possible counterfactuals \mathcal{H} that are identified. Conversely, given a counterfactual of interest \mathcal{H} , the proposition tells us the type of restrictions that must be imposed to identify \tilde{p} .

Corollary 3. *If the restrictions are linearly independent (i.e. $\text{rank}(Q) = d \leq X$) then the counterfactual CCP is identified only if*

$$C \left(I - Q' (QQ')^{-1} Q \right) = 0. \quad (29)$$

5 Identification of Counterfactual Welfare

In this section, we discuss the identification of counterfactual welfare and provide the minimal set of sufficient conditions for identification. For simplicity, we only consider affine action-diagonal counterfactuals; i.e. $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and $\tilde{\pi}_a = H_a \pi_a + g_a$, all a . Extensions to more general cases are straightforward, but at the cost of substantially more cumbersome notation. The feature of interest here is the value function difference $\Delta V = \tilde{V} - V$, where \tilde{V} is the counterfactual value function.

Proposition 11. *(Welfare) Assume $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and $\tilde{\pi}_a = H_a \pi_a + g_a$, all a . The welfare difference ΔV is identified if, for all $a \neq J$,*

$$H_a A_a - \tilde{A}_a H_J = 0,$$

and

$$H_J = (I - \beta \tilde{F}_J) (I - \beta F_J)^{-1}. \quad (30)$$

Proposition 11 shows that identification of \tilde{p} (which is implied by the proposition's first condition) is not sufficient to identify ΔV ; we also need (30). The second condition is satisfied, for instance, when the counterfactual transformation does not affect option J : $H_J = I$ and $\tilde{F}_J = F_J$. For “proportional changes” counterfactuals the two conditions are satisfied only when all matrices H_a equal the identity matrix; i.e. $\tilde{\pi} = \pi$, which is equivalent to saying that ΔV is not identified. On a positive note, an immediate implication of Proposition 11 is that the welfare impact of “additive changes” is identified. Therefore, “additive changes” are robust to nonidentification of the model primitives: both \tilde{p} and ΔV are identified.

Proposition 11 is an immediate consequence of Lemma 5 in the Appendix. Lemma 5 provides the full set of necessary and sufficient conditions to identify ΔV , and shows that identification of \tilde{p} together with (30) are “almost” necessary to identify welfare.²⁸

Finally, the next corollary considers identification of ΔV for the parametric model of Section 4.1. As expected, identification is guaranteed when counterfactuals change $Z'\theta_1$ and/or F^w .

²⁸ “Almost” here is in the following sense: when at least one of the two conditions does not hold, the nonlinear system of equations necessary to identify ΔV is generically inconsistent. So, in practice, identification of ΔV is unlikely when either \tilde{p} is not identified or (30) fails.

Corollary 4. (*Welfare, Parametric Model*) Assume the conditions of Proposition 8 hold. Suppose $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and $\tilde{\theta}_0(a) = H_0(a)\theta_0(a)$. The welfare difference ΔV is identified if, for all $a \neq J$,

$$H_0(a)A_a^k - \tilde{A}_a^k H_0(J) = 0,$$

and

$$H_0(J) = (I - \beta\tilde{F}_J^k)(I - \beta F_J^k)^{-1}.$$

Furthermore, if $H_0(a) = I$ and $\tilde{F}_a^k = F_a^k$ for all a , then ΔV is identified for any counterfactual transformation on $Z(a, w)'\theta_1(a, k)$ and F^w .

6 Applied Examples

6.1 Numerical Example: Monopolist Entry and Exit Problem

This section illustrates some of our theoretical results using a monopolist's entry/exit problem; the model adopts the parameterization of Section 4.1 and has been commonly used in the literature (Aguiregabiria and Suzuki (2014), Das, Roberts and Tybout (2007), Varela (2013), Lin (2015), Igami (2017)). Consider a monopolist deciding whether to be active or exit from a market, so that $\mathcal{A} = \{\text{active}, \text{inactive}\}$. Let $x = (k, w)$ with $k_{it} = a_{it-1}$, and w the firm's aggregate demand shocks determining variable profits, $\bar{\pi}(w)$. The flow payoff is

$$\pi(a, k, w) = \begin{cases} k\phi^s & \text{if } a = 0 \text{ (inactive)} \\ k(\bar{\pi}(w) - fc) - (1 - k)\phi^e, & \text{if } a = 1 \text{ (active)} \end{cases} \quad (31)$$

where ϕ^s is the scrap value, fc is the fixed cost, and ϕ^e is the entry cost.²⁹

We assume that w is observable and can take three values: high, medium or low, $w \in \{w^H, w^M, w^L\}$. Moreover, it follows a first-order Markov process. Variable profits $\bar{\pi}(w)$ are determined by static profit maximization. We assume the econometrician knows (or estimates): (a) the true CCP, $\Pr(\text{active}|k, w)$; (b) the transition of the demand shocks, $\Pr(w_{t+1}|w_t)$; and (c) the variable profits $\bar{\pi}(w)$, which can be recovered "offline," using price and quantity data.³⁰

²⁹The model falls within the parametric framework of Section 4.1, with $\theta_0(0, 0) = 0$, $\theta_0(1, 0) = -\phi^e$, $\theta_0(0, 1) = \phi^s$, and $\theta_0(1, 1) = -fc$. Variable profits, $\bar{\pi}(w)$, are estimated outside of the dynamic problem using price and quantity data, one may then take $Z(a, w) = \bar{\pi}(w)$ and $\theta_1(a, k) = 1$ (alternatively, one might assume a reduced form profit function $\bar{\pi}(w) = w'\gamma$, in which $Z(a, w) = w$ and $\theta_1(a, k) = \gamma$, with γ identified under sufficient variation on w).

³⁰We assume the monopolist faces the demand curve $P_t = w_t - \eta Q_t$ and has constant marginal cost c , so that $\bar{\pi}(w_t; \eta, c) = (w_t - c)^2 / 4\eta$. The idiosyncratic shocks ε_{it} follow a type 1 extreme value distribution. We ignore sampling variation for simplicity and set: $c = 11$, $\eta = 1.5$, $w = (20, 17, 12)$, $\beta = 0.95$, $fc = 5.5$, $\phi^s = 10$, $\phi^e = 9$, while the transition matrix for w is

$$F(w_{t+1}|w_t) = \begin{bmatrix} 0.4 & 0.35 & 0.25 \\ 0.3 & 0.4 & 0.3 \\ 0.2 & 0.2 & 0.6 \end{bmatrix}.$$

First, we solve the true model and obtain the baseline CCPs. Then, we recover π using Proposition 1. Typically, researchers identify the model by setting either $\phi^s = 0$ or $fc = 0$. As previously discussed, there is little guidance to justify these assumptions because cost or scrap value data are extremely rare. We estimate the model twice to compare the different sets of restrictions. When $\phi^s = 0$, so that $\pi(\text{inactive}, k, w) = 0$ for all (k, w) , identification of π follows directly from (4). When instead $fc = 0$, since $\pi(\text{active}, 1, w) = \bar{\pi}(w)$ is assumed known, we can recover the remaining elements of π combining (4) and (26).

Table 1: Numerical Example – True vs. Estimated Profits

States: (k, w)	True Profit	Estimated Profit	Estimated Profit
		<i>scrap value = 0</i>	<i>fixed cost = 0</i>
<i>a = inactive</i>			
$\pi(a, k = 0, w_H) = 0$	0	0	0
$\pi(a, k = 0, w_M) = 0$	0	0	0
$\pi(a, k = 0, w_L) = 0$	0	0	0
$\pi(a, k = 1, w_H) = \phi^s$	10	0	120
$\pi(a, k = 1, w_M) = \phi^s$	10	0	120
$\pi(a, k = 1, w_L) = \phi^s$	10	0	120
<i>a = active</i>			
$\pi(a, k = 0, w_H) = -\phi^e$	-9	0.5	-113.5
$\pi(a, k = 0, w_M) = -\phi^e$	-9	0.5	-113.5
$\pi(a, k = 0, w_L) = -\phi^e$	-9	0.5	-113.5
$\pi(a, k = 1, w_H) = \bar{\pi}(w_H; \eta, c) - fc$	8	7.5	13.5
$\pi(a, k = 1, w_M) = \bar{\pi}(w_M; \eta, c) - fc$	0.5	0	6
$\pi(a, k = 1, w_L) = \bar{\pi}(w_L; \eta, c) - fc$	-5.33	-5.83	0.167

Table 1 presents the true and the two estimated payoff functions. Note that under the first restriction ($\phi^e = 0$) entry costs change sign. This is because if there is no scrap value, entering the market becomes less attractive and entry costs must become low (in fact, negative) to capture the observed entry patterns. Under the second restriction ($fc = 0$), both entry costs and scrap values are considerably larger than their true values. To see why, consider an active firm ($k = 1$). Fixing $fc = 0$ implies higher profits when active, which gives incentives to stay more often in the market. To match the observed CCP, scrap values must increase to provide incentives to exit and match the observed exit rate. Similarly, when the firm is out ($k = 0$), increasing profits when active provides incentives to enter. Entry costs must then increase to compensate for this incentive.

Given the recovered payoffs, we implement four counterfactuals and compare the true and the inferred counterfactual CCPs and welfare. In the first two, the government provides subsidies to encourage entry. Counterfactual 1 is an additive subsidy that reduces entry costs: $\tilde{\pi}(\text{active}, 0, w) = \pi(\text{active}, 0, w) + g$. Counterfactual 2 is a proportional subsidy: $\tilde{\pi}(\text{active}, 0, w) = H\pi(\text{active}, 0, w)$.

As shown in Section 3, while the counterfactual CCPs and welfare are identified in the first case, they are not identified in the second scenario.³¹

Table 2 presents the results from counterfactuals 1 and 2 for the true model and the two estimated models. In both counterfactuals, the true counterfactual probability of entering increases compared to the baseline because of the subsidy; and the probability of staying in the market decreases because it is cheaper to re-enter in the future. So, the monopolist enters and exits more often than in the baseline case.

In counterfactual 1 (additive subsidy), as expected, the counterfactual CCPs and welfare are identical in the true model and under both estimated models. In contrast, counterfactual 2 (proportional subsidy) results in very different outcomes under the two restrictions. When $\phi^s = 0$, the changes in the CCPs are all in the wrong direction: while the true entry probability increases relative to the baseline, the predicted counterfactual entry probability decreases. Similarly, the counterfactual exit probability decreases in the true model, while it increases in the estimated model. Welfare also has the wrong sign in all states. This is a direct consequence of the fact that the identified entry cost under this restriction has the wrong sign: in the true model, multiplying $\pi(\text{active}, 0, w)$ by H represents a subsidy, but in the estimated model, it becomes a tax. This illustrates the importance of the identifying restrictions in driving conclusions, especially when the researcher does not know the sign of the true parameter. When instead we restrict $fc = 0$, since the estimated entry costs and scrap values are magnified, it is profitable to enter and exit the market repeatedly when the entry cost is reduced by 10 percent in the counterfactual. Predicted turnover and welfare are therefore excessive.³²

Counterfactual 3 changes the transition process $\Pr(w_{t+1}|w_t)$. Because $\bar{\pi}(w)$ is known, the counterfactual behavior and welfare are identified (Proposition 9 and Corollary 4)). Top panel of Table 3 confirms the results.³³

Finally, counterfactual 4 implements a “change in types” experiment. We add a second market with different parameter values: market 2 is more profitable than market 1 both through lower entry costs and higher variable profits. We identify the parameters for market 2 as before and perform a counterfactual that substitutes the entry cost of market 1 by the estimated entry cost

³¹We choose the additive and proportional subsidies so that the true counterfactual CCP and welfare are the same. As $\pi(\text{active}, 0, w) = -\phi^e$, and the true $\phi^e = 9$, we set $g = 0.9$ and $H = 0.9$, so that in both cases the true counterfactual entry cost becomes $\tilde{\pi}(\text{active}, 0, w) = -8.1$.

³²Most entry model studies involve dynamic games rather than single-agent models, but the concerns we raise about the identification of proportional entry subsidies still apply to dynamic games, as discussed in footnote 12. If the best responses are not identified, there seems to be little hope that counterfactual equilibrium might be identified (see Kalouptsi, Scott, and Souza-Rodrigues (2017a)).

³³We set $\tilde{\Pr}(w'|w) = 1/3$, for all (w', w) . Aguirregabiria and Suzuki (2014) also implement a change in transitions in a similar model. But they consider a change in F_a^k , i.e., a change in the transition of states that enter the nonidentified part of payoffs. As expected, their counterfactual is not identified. Similar to our counterfactual 2 under the restriction $\phi^s = 0$, they obtained counterfactual predictions in the wrong direction.

Table 2: Counterfactuals 1 and 2 – Additive and Prop Entry Subsidies

States: (k, w)	Baseline	True CF	Estimated CF <i>scrap value = 0</i>	Estimated CF <i>fixed cost = 0</i>
<i>CF1: $\tilde{\pi}_0 = \pi_0, \tilde{\pi}_1 = \pi_1 + g$</i>				
CCP: Pr (active x)				
$(k = 0, w_H)$	93.61%	94.95%	94.95%	94.95%
$(k = 0, w_M)$	87.48%	90.27%	90.27%	90.27%
$(k = 0, w_L)$	72.99%	80.33%	80.33%	80.33%
$(k = 1, w_H)$	99.99%	99.99%	99.99%	99.99%
$(k = 1, w_M)$	80.91%	69.59%	69.59%	69.59%
$(k = 1, w_L)$	0.48%	0.29%	0.29%	0.29%
Welfare: $\tilde{V} - V$				
$(k = 0, w_H)$	-	5.420	5.420	5.420
$(k = 0, w_M)$	-	5.445	5.445	5.445
$(k = 0, w_L)$	-	5.539	5.539	5.539
$(k = 1, w_H)$	-	4.535	4.535	4.535
$(k = 1, w_M)$	-	4.727	4.727	4.727
$(k = 1, w_L)$	-	5.219	5.219	5.219
<i>CF2: $\tilde{\pi}_0 = \pi_0, \tilde{\pi}_1 = H\pi_1$</i>				
CCP: Pr (active x)				
$(k = 0, w_H)$	93.61%	94.95%	93.53%	99.87%
$(k = 0, d_M)$	87.48%	90.27%	87.31%	99.84%
$(k = 0, d_L)$	72.99%	80.33%	72.53%	99.81%
$(k = 1, d_H)$	99.99%	99.99%	99.99%	90.59%
$(k = 1, w_M)$	80.91%	69.59%	81.44%	0.44%
$(k = 1, w_L)$	0.48%	0.29%	0.49%	0.00%
Welfare: $\tilde{V} - V$				
$(k = 0, w_H)$	-	5.420	-0.289	88.255
$(k = 0, w_M)$	-	5.445	-0.290	88.829
$(k = 0, w_L)$	-	5.539	-0.295	89.756
$(k = 1, w_H)$	-	4.535	-0.239	77.068
$(k = 1, w_M)$	-	4.727	-0.248	82.836
$(k = 1, w_L)$	-	5.219	-0.278	84.802

Table 3: Counterfactuals 3 and 4 – Change in $F(w_{t+1}/w_t)$ and Change Markets' Entry Costs

States: (k, w)	Baseline	True CF	Estimated CF <i>scrap value = 0</i>	Estimated CF <i>fixed cost = 0</i>
<i>CF3: $\tilde{\pi}_0 = \pi_0, \tilde{\pi}_1 = \pi_1, \tilde{F}^w \neq F^w$</i>				
CCP: Pr (active x)				
$(k = 0, w_H)$	93.61%	86.97%	86.97%	86.97%
$(k = 0, w_M)$	87.48%	86.97%	86.97%	86.97%
$(k = 0, w_L)$	72.99%	86.97%	86.97%	86.97%
$(k = 1, w_H)$	99.99%	99.99%	99.99%	99.99%
$(k = 1, w_M)$	80.91%	80.19%	80.19%	80.19%
$(k = 1, w_L)$	0.48%	1.17%	1.17%	1.17%
Welfare: $\tilde{V} - V$				
$(k = 0, w_H)$	-	0.542	0.542	0.542
$(k = 0, w_M)$	-	1.347	1.347	1.347
$(k = 0, w_L)$	-	2.530	2.530	2.530
$(k = 1, w_H)$	-	0.468	0.468	0.468
$(k = 1, w_M)$	-	1.350	1.350	1.350
$(k = 1, w_L)$	-	1.808	1.808	1.808
<i>CF4: $\tilde{\pi}_0^1 = \pi_0^2, \tilde{\pi}_1^1 = \pi_1^1$</i>				
CCP: Pr (active x)				
$(k = 0, w_H)$	93.61%	97.28%	95.22%	100.00%
$(k = 0, w_M)$	87.48%	95.08%	90.83%	100.00%
$(k = 0, w_L)$	72.99%	91.44%	81.74%	100.00%
$(k = 1, w_H)$	99.99%	99.95%	99.99%	0%
$(k = 1, w_M)$	80.91%	36.86%	66.67%	0%
$(k = 1, w_L)$	0.48%	0.09%	0.02%	0%
Welfare: $\tilde{V} - V$				
$(k = 0, w_H)$	-	19.778	6.684	482.861
$(k = 0, w_M)$	-	19.883	6.715	483.667
$(k = 0, w_L)$	-	20.198	6.831	484.849
$(k = 1, w_H)$	-	16.816	5.602	449.773
$(k = 1, w_M)$	-	17.752	5.846	457.934
$(k = 1, w_L)$	-	19.044	6.438	459.999

of market 2.³⁴

The bottom panel of Table 3 presents the results. Similar to counterfactual 2, turnover increases in the true counterfactual compared to the baseline; and again, the two identifying restrictions generate very different outcomes. This is expected given Proposition 3. Under the first restriction ($\phi^s = 0$), counterfactual CCPs and welfare are all in the right direction, even though entry costs have the wrong sign in both markets. This happens because replacing the market 1 entry cost by the market 2 entry cost amounts to an increase in entry costs in the restricted model. Even though the CCP moves in the right direction, the magnitude is bound to be wrong and turnover under this restriction is not as large as the true counterfactual turnover. Under the second identifying restrictions ($fc = 0$), turnover and welfare are again exaggerated, to the point that counterfactual choice probabilities are (numerically close to) either zero or one.

6.2 Empirical Example: Agricultural Land Use Model

In this section, we explore the impact of identifying restrictions on counterfactuals using actual data on agricultural land use. We estimate a dynamic model of farmers’ planting choices and perform two counterfactuals of interest: the long-run land use elasticity and a fertilizer tax. We emphasize the impact of identifying restrictions on counterfactuals and relegate the details of the estimation methodology to Appendix C.

Empirical Model. Each year, field owners decide whether to plant crops or not; i.e. $\mathcal{A} = \{c, nc\}$, where c stands for “crops” and nc stands for “no crops” (e.g. pasture, hay, non-managed land). Fields are indexed by i and counties are indexed by m . We partition the state x_{imt} into:

1. time-invariant field and county characteristics, s_{im} , e.g. slope, soil composition;
2. number of years since field was last in crops, $k_{imt} \in \mathcal{K} = \{0, 1, \dots, K\}$; and³⁵
3. aggregate state, w_{mt} (e.g. input and output prices, government policies).

Per period payoffs are specified as in (21) so that $\pi(a, k, s, w) = \theta_0(a, k, s) + \theta_1 Z(a, w)$; here, $\theta_0(a, k, s)$ captures switching costs between land uses and $Z(a, w)$ are observable measures of returns. The dependence of θ_0 on k is what creates dynamic incentives for landowners. The action of “no crops” leaves the land idle, slowly reverting it to natural vegetation, rough terrain, etc. The farmer needs to clear the land in order to convert to crop and start planting. The costs of switching to crop may be rising as the terrain gets rougher. At the same time, however, there may

³⁴For market 2, we set: $c_2 = 9$, $\eta_2 = 1.7$, $w_2 = (18, 15, 11)$, $fc_2 = 3$, $\phi_2^s = 8$, $\phi_2^e = 6$. The discount factor and transition matrix in market 2 is the same as in market 1. The estimated profit under the first restriction ($\phi_2^s = 0$) is $\phi_2^e = 1.6$, $\pi^2(\text{active}, 1, w) = (8.52, 1.89, -2.82)$; and under the second restriction ($fc_2 = 0$) is: $\phi_2^e = -63$, $\phi_2^s = 68$, $\pi^2(\text{active}, 1, w) = (11.91, 5.29, 0.59)$.

³⁵Ideally, k_{imt} would include detailed information on past land use. We consider the years since the field was in crop (bounded by K) for computational tractability and due to data limitations.

be benefits to switching, e.g. planting crops may be more profitable after the land is left fallow for a year. In summary, we expect $\theta_0(a, k, s)$ to differ across actions and states.

The transition of state variables follows the decomposition (20), which implies that farmers are small (price takers) and that there are no externalities across fields. The transition rule of k is: $k'(a, k) = 0$ if $a = c$, and $k'(a, k) = \min\{k + 1, K\}$ if $a = nc$. I.e. if “no crop” is chosen, the field state since last crop increases by one, up to K . If “crop” is chosen, the field state is reset to zero. Planting crops is therefore a “renewal” action. We return to the market state w below.

Data. First, we collected high-resolution annual land use data in the United States obtained from the Cropland Data Layer (CDL) database. The CDL was merged with an extensive dataset of land transactions obtained from DataQuick (which includes information on price, acreage, field address and other characteristics). Then, we incorporated detailed data from NASA’s Shuttle Radar Topography Mission database (with fine topographical information on altitude, slope and aspect); the Global Agro-Ecological Zones dataset (with information on soil categories and on protected land); and various public databases on agricultural production and costs from the USDA. The final dataset goes from 2010 to 2013 for 515 counties and from 2008 to 2013 for 132 counties.

Further details about the construction of the dataset, as well as some summary statistics, are presented in Appendix B. Here we only emphasize that land use exhibits substantial persistence. The average proportion of cropland in the sample is 15%; the probability of keeping the land in crop is about 85%, while the probability of switching to crops after two years as non-crop is quite small: 1.6%. Finally, the proportion of fields that switch back to crops after one year as “no crop” ranges from 27% to 43% on average depending on the year, which suggests some farmers enjoy benefits from leaving land fallow for a year.

Estimation. Following Scott (2013), we augment this land use model to allow for unobserved market states. This may be important as the econometrician may not be able to capture the entire information set of the agent (commodity prices, government policy, etc.). See Appendix C for details.³⁶

The parameters of interest are θ_1 and $\theta_0(a, k, s)$, all a, k, s . The slope θ_1 is identified provided there is sufficient variation on $Z(a, w)$. Switching costs between land uses, $\theta_0(a, k, s)$, on the other hand are not identified. The sensitivity of certain counterfactuals to identifying restrictions on payoffs calls out for some means to assess the accuracy of these restrictions. To do so, we present and compare two estimators.

First, we estimate the model using the observed data on farmers’ actions and states, following

³⁶Unobserved market states are relevant for many applications, such as when the available observed market states are insufficient to capture the agents’ full information set, or where it is difficult to capture the evolution of market states adequately, or when the dimensionality of the state space is large. In older versions of this paper, and in Kalouptsi, Scott and Souza-Rodrigues (2017b), we provide the details of a general setup with unobserved market states and characterize the identification of payoffs.

Scott (2013). We call this the “CCP estimator.” Naturally, in this case we need some identifying restrictions, and as in Scott (2013), we impose $\theta_0(\text{nocrop}, k, s) = 0$ for all k and s . As is common in applied work, there is little guidance to specify the particular values that $\theta_0(a, k, s)$ should take. To evaluate the impact of these restrictions in this real-data setup, we bring in additional data, namely, the land resale prices.

Our second estimator makes use of resale prices to avoid the restrictions on $\theta_0(a, k, s)$. We call it the “V-CCP estimator.” Recall that we need K restrictions to identify the model. Returning to our base equations (1)-(3) it is clear that if V is known, so are the payoffs (recover v_a from (3) and π_a from (1)); we use resale prices to make measurements of V . Our estimator is designed so that the *only* role of the resale price data is to avoid the identifying restrictions $\theta_0(\text{nocrop}, k, s) = 0$ for all k, s . By construction of the estimator, θ_1 is the same as that of the CCP estimator.³⁷

Appendix C explains both estimators in detail. Here, we only emphasize that the CCP estimator imposes restrictions on $\theta_0(a, k, s)$ for identification while the “V-CCP” estimator replaces these *a priori* restrictions with more data-driven restrictions.

Results. We now turn to our results. Table 4 presents the estimated parameters using the CCP and V-CCP estimators. For brevity we only present the average of $\theta_0(a, k, s) / \theta_1$ across field types s (we divide by θ_1 so that the parameters can be interpreted in dollars per acre). We set $K = 2$ due to data limitations and because after 2 years out of crops there are very few conversions back to crops in the data.³⁸

The mean switching cost parameters from the CCP estimator are all negative and increase in magnitude with k . One may interpret this as follows: when $k = 0$, crop was planted in the previous year. According to the estimates, preparing the land to replant crops costs on average \$722/acre. When $k = 1$, the land was not used to produce crop in the previous year. In this case, it costs more to plant crops than when $k = 0$. Conversion costs when $k = 2$ are even larger. Of course such interpretation hinges on the assumption that $\theta_0(\text{nocrop}, k, s) = 0$ for all k, s . As is typical in switching cost models, estimated switching costs are somewhat large in order to explain the observed persistence in choices; unobserved heterogeneity – which is beyond the scope of this paper – can alleviate this (see Scott (2013)).

The estimated parameters of the V-CCP estimator do not impose $\theta_0(\text{nocrop}, k, s) = 0$. When $k = 0$, switching out of crops is now expensive on average (not zero anymore). In fact, we reject

³⁷There are numerous ways to model resale markets, and different models may imply different mappings between transaction prices and agents’ value function. Here, we essentially consider the simplest possible setting: in a world with a large number of homogeneous agents, a resale transaction price must equal the value of the asset. A similar approach is adopted in Kalouptsi (2014, 2017). To address concerns that transacted fields may be selected, we compare the transacted fields (in DataQuick) to all US fields (in the CDL) in Table 8 of Appendix B. Overall, the two sets of fields look similar. We also explore whether land use changes upon resale and find no such evidence (see Table 10 in Appendix C).

³⁸We weight observations as in Scott (2013) and cluster standard errors by year. We construct the confidence intervals for $\bar{\theta}_0(a, k) / \theta_1$ by sampling from the estimated asymptotic distribution of $(\hat{\theta}_0, \hat{\theta}_1)$. The details of the first stage estimator are in Appendix C.

Table 4: Empirical Results

Estimator:	CCP	V-CCP
$\theta_0(crop, 0) / \theta_1$	-721.93 (-1150,-562)	-1228.9 (-2250,-852)
$\bar{\theta}_0(crop, 1) / \theta_1$	-2584.4 (-4610,-1840)	-1119.4 (-3130,-379)
$\bar{\theta}_0(crop, 2) / \theta_1$	-5070.8 (-9220,-3541)	-4530.4 (-8340,-3120)
$\bar{\theta}_0(nocrop, 0) / \theta_1$	0	-2380.3 (-3540,-1950)
$\bar{\theta}_0(nocrop, 1) / \theta_1$	0	470.05 (-393,788)
$\bar{\theta}_0(nocrop, 2) / \theta_1$	0	-454.58 (-996,-255)
θ_1^{-1}	734.08 (418,1050)	734.08 (418,1050)

90% confidence intervals in parentheses. θ_0 values are means across all fields in the sample, divided by θ_1 so that they can be interpreted in dollars. Note that θ_1^{-1} is proportional to standard derivation of idiosyncratic shocks, when payoff function is measured in dollars.

the null hypothesis of $\bar{\theta}(nocrop, k) = 0$, for all k . This is reasonable because the “no crop” option incorporates, in addition to fallow land, pasture, hay, and other land uses. While staying out of crops for one year may be the result of the decision to leave land fallow, staying out of crops for longer periods reflects other land usages (since land will likely not stay idle forever) with their associated preparation costs. Furthermore, the estimated value of $\theta_0(crop, k, s)$ also is affected when we drop the restriction. The absolute value of the estimated $\bar{\theta}_0(crop, 0)$ is now larger than the absolute value of $\bar{\theta}_0(crop, 1)$. This reflects the benefits of leaving land fallow for one year (i.e. smaller replanting costs). This potential benefit is not apparent when we restrict $\theta_0(nocrop, k, s)$. Given that the probability of planting crops after one year of fallow is lower than the probability of planting crops after crops in the data (in most counties), in order to rationalize the choice probabilities, the restricted model (imposing $\theta_0(nocrop, k, s) = 0$) must assign higher costs to crops after fallow than after crops. We view this as an appealing feature of the V-CCP model – it is arguably not plausible that leaving land out of crops for one year would increase the costs of planting crops in the following year dramatically.³⁹

Counterfactuals. We implement two counterfactuals: the long-run elasticity (LRE) of land use and an increase in the costs of replanting crops.

The LRE measures the long-run sensitivity of land use to an (exogenous) change in crop returns, $Z(c, w)$. To calculate it, we compare the steady-state acreage distribution in the data obtained when Z^c is held fixed at their average recent levels and when Z^c is held fixed at 10% higher levels. The LRE is defined as the arc elasticity between the total acreage in the two steady states.⁴⁰

As shown in Table 5, the CCP and V-CCP estimators give *exactly* the same LRE. This is no coincidence. By Proposition 8, θ_1 is identified and by Proposition 9(i), a counterfactual that changes only the identified part of payoffs is also identified. Therefore, the LRE is not affected by identifying restrictions on θ_0 , and the *only* difference between the CCP estimator and the V-CCP estimator is that the latter relies on land values to identify the profit function while the former relies on *a priori* restrictions.

The second counterfactual increases the crop replanting costs as

$$\tilde{\theta}(crop, 0) = \theta(crop, 0) + \lambda(\theta(crop, 1) - \theta(crop, 0)).$$

The difference $\theta(crop, 1) - \theta(crop, 0)$ captures the benefits of leaving land out of crops for a year.

³⁹One could also argue that it is not plausible that staying out of crops for only two years would lead to dramatically higher costs of planting crops. However, as mentioned previously, we observe very few fields in the data with field state $k = 2$ which have not been out of crops for longer than two years; i.e., fields which have been out of crops for at least two years have typically been out of crops for a long time.

⁴⁰See Scott (2013) for a formal definition and further discussion. The LREs estimated here are somewhat higher than those found in Scott (2013) (although not significantly so). We find that this is largely due to our different sample combined with the absence of unobserved heterogeneity: when Scott’s estimation strategy is applied to our sample of counties ignoring unobserved heterogeneity, LREs are very similar to those presented here.

Table 5: Policy Counterfactuals

Estimator:	CCP	V-CCP
Long-run elasticity	0.57	0.57
Fertilizer tax	0.32	-0.16

Fertilizer tax statistic is percentage change in long-run cropland.

Long-run elasticity is a 10 percent arc elasticity.

One such benefit is to allow soil nutrient levels to recover, reducing the need for fertilizer inputs. When it is difficult to measure the fertilizer saved by leaving land fallow, one can use the switching cost parameters to implement a counterfactual that resembles a fertilizer tax. A motivation for this type of counterfactual is that higher fertilizer prices would be a likely consequence of pricing greenhouse gas emissions, as fertilizer production is very fossil-fuel intensive. Here we impose $\lambda = 0.1$. So, this exercise changes the costs of replanting crops in a way that reflects 10% of the benefits of leaving land out of crops for one year.⁴¹ Formally, $\theta_0(a)$ is a 3×1 vector, and we take $\tilde{\theta}_0(a) = H_0(a)\theta_0(a)$, with $H_0(nocrop) = I$, and

$$H_0(crop) = \begin{bmatrix} 1 - \lambda & \lambda & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Because $H_0(nocrop)$ is diagonal but $H_0(crop)$ is not, these are not similar matrices. By Proposition 9(ii), the counterfactual choice probability is not identified.

Indeed, as shown in Table 5, the identifying restrictions do matter when it comes to this counterfactual. The CCP estimator leads to a 32% increase in cropland, while the V-CCP estimator predicts a decrease in cropland, as expected. In other words, the CCP estimator errs in predicting not just the magnitude, but also the sign of the change in crop acreage. The reason behind this is that the CCP estimator cannot capture the benefits from leaving land fallow (on average) and thus interprets this counterfactual as a subsidy rather than a tax.

To summarize, when we only change the identifying restrictions (i.e. moving from the CCP to the V-CCP estimator), the LRE does not change, as it involves only a transformation of the identified component of the profit function. However, the land use pattern in the second counterfactual, which involves a transformation of the non-identified part of payoffs, is substantially altered when we modify the identifying restrictions.

⁴¹As with the LRE, we fix $Z(c, w)$ and $Z(nc, w)$ at their mean level for each county.

7 Conclusions

This paper studies the identification of counterfactuals in dynamic discrete choice models. We provide the set of necessary and sufficient conditions that determine whether counterfactual behavior and welfare are identified for a broad class of counterfactuals of interest. We also investigate the identification power of additional restrictions that may not suffice to identify the full model but may suffice to identify the counterfactual. For a large class of interventions (involving affine changes in payoffs), the conditions are straightforward to verify in practice.

We investigate relevant counterfactuals in two applied examples (a monopolist’s entry/exit decisions and a farmer’s land use decisions). The results call for caution while leaving room for optimism: although counterfactual behavior and welfare can be sensitive to identifying restrictions imposed on the model, there exists important classes of counterfactuals that are robust to such restrictions.

References

- [1] Abbring, J. H, and O. Daljord (2017). “Identifying the Discount Factor in Dynamic Discrete Choice Models,” Becker Friedman Institute for Research in Economics Working Paper No. 2017-17
- [2] 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.
- [3] 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.
- [4] Arcidiacono, P. and R. Miller (2011). “CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity,” *Econometrica*, 79, 1823-1867.
- [5] Arcidiacono, P. and R. Miller (2017). “Identifying Dynamic Discrete Choice Models off Short Panels,” Mimeo.
- [6] Bajari, Chu, Nekipelov, and Park (2016). “Identification and Semiparametric Estimation of a Finite Horizon Dynamic Discrete Choice Model with a Terminating Action,” *Quantitative Marketing and Economics*, vol. 14, issue 4, pages 271-323.
- [7] Blevins, J. (2014). “Nonparametric Identification of Dynamic Decision Processes with Discrete and Continuous Choices,” *Quantitative Economics*, 5, 531-554.
- [8] Buchholz, N., H. Xu, and M. Shum (2016). “Semiparametric Estimation of Dynamic Discrete-Choice Models,” Mimeo.

- [9] Chan, T. Y., B. H. Hamilton and N. Papageorge (2016). “Health, Risky Behavior and the Value of Medical Innovation for Infectious Diseases,” *Review of Economic Studies*, 83(4), 1465–1510
- [10] Chen, L. (2017). “Identification of Discrete Choice Dynamic Programming Models with Non-parametric Distribution of Unobservables,” *Econometric Theory*, Vol 33(3), 551-577.
- [11] Collard-Wexler, A. (2013). “Demand Fluctuations in the Ready-mix Concrete Industry,” *Econometrica*, 81(3), 1003–1007.
- [12] Conlon, C. T. (2012). “A Dynamic Model of Prices and Margins in the LCD TV Industry,” Working Paper, New York University.
- [13] Crawford, G. S. and M. Shum (2005). “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 73, 4, 1137-1173.
- [14] Das, S., M.J. Roberts, and J.R Tybout. (2007). “Market Entry Costs, Producer Heterogeneity, and Export Dynamics,” *Econometrica*, 75(3), 837–873.
- [15] Duflo, E., R. Hanna, and S. P. Ryan (2012). “Incentives Work: Getting Teachers to Come to School,” *The American Economic Review*, 102, 1241–1278.
- [16] Dunne, T., S. Klimek, M. Roberts, and Y. Xu (2013). “Entry and Exit and the Determinants of Market Structure,” *The RAND Journal of Economics*, 44 (3), 462–487.
- [17] Eckstein, Z. and O. Lifshitz (2011). “Dynamic Female Labor Supply,” *Econometrica*, 79(6), 1675–1726.
- [18] Gilleskie, D. (1998). “A Dynamic Stochastic Model of Medical Care Use and Work Absence,” *Econometrica*, 66, 1-45.
- [19] Heckman, J. J. (2000). “Causal Parameters and Policy Analysis in Economics: A Twentieth Century Retrospective.” *Quarterly Journal of Economics*, 115(1), 45-97.
- [20] Heckman, J. J. (2010). “Building Bridges between Structural and Program Evaluation Approaches to Evaluating Policy.” *Journal of Economic Literature*, 48(2), 356-98.
- [21] Heckman, J., J. E. Humphries, and G. Veramendi (2016). “Dynamic Treatment Effects,” *Journal of Econometrics*, 191, Issue 2, 276-292.
- [22] Heckman, J. J., and S. Navarro (2007). “Dynamic Discrete Choice and Dynamic Treatment Effects,” *Journal of Econometrics*, 136, 341–396.
- [23] Hendel, I., and A. Nevo (2006). “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74(6), 1637-1674.

- [24] 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.
- [25] Igami, M. (2017) “Estimating the Innovator’s Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981–1998,” *Journal of Political Economy*, 125(3), 798–847
- [26] Lin, H. (2015). “Quality Choice and Market Structure: A Dynamic Analysis of Nursing Home Oligopolies,” *International Economic Review*, Vol. 56, Issue 4, 1261–1290.
- [27] Li, S. and C. Wei (2014). “The Cost of Green Stimulus: A Dynamic Discrete Analysis of Vehicle Scrappage Programs,” Working Paper, The George Washington University.
- [28] Kalouptside, M. (2014). “Time to Build and Fluctuations in Bulk Shipping,” *American Economic Review*, 104(2): 564-608.
- [29] Kalouptside, M. (2017). “Detection and Impact of Industrial Subsidies: The Case of Chinese Shipbuilding,” Forthcoming, *Review of Economic Studies*.
- [30] Kalouptside, M., P. Scott, and E. Souza-Rodrigues (2017a). “On the Non-Identification of Counterfactuals in Dynamic Discrete Games.” *International Journal of Industrial Organization*, Volume 50, 362-371.
- [31] Kalouptside, M., P. Scott, and E. Souza-Rodrigues (2017b). “Linear IV Regression Estimators for Single-Agent Dynamic Discrete Choice Models,” Mimeo.
- [32] Keane, M. P. and A. Merlo (2010). “Money, Political Ambition, and the Career Decisions of Politicians.” *American Economic Journal: Microeconomics*, 2(3): 186-215.
- [33] Keane, M. P., and Wolpin, K. I. (1997). “The Career Decisions of Young Men,” *Journal of Political Economy*, 105 (3), 473–522.
- [34] Keane, M. P., and Wolpin, K. I. (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*, Vol. 51, No. 3, 851-892.
- [35] Koning, R. H., H. Neudecker and T. Wansbeek. (1991). “Block Kronecker Products and the vecb Operator”. *Linear Algebra and Its Applications*. 165-184.
- [36] Magnac, T. and D. Thesmar (2002). “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70(2), 801–816.
- [37] Marschak, J. (1953). “Economic Measurements for Policy and Prediction.” In *Studies in Econometric Method*, ed. William C. Hood and Tjalling C. Koopmans, 1-26. New York: Wiley.

- [38] Norets, A. and X. Tang (2014). “Semiparametric Inference in Dynamic Binary Choice Models,” *Review of Economic Studies*, 81(3), 1229-1262.
- [39] Pesendorfer, M. and Schmidt-Dengler (2008). “Asymptotic Least Square Estimators for Dynamic Games,” *Review of Economic Studies*, 75, 901–908.
- [40] 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.
- [41] Rust, J. (1987). “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999-1033.
- [42] Rust, J. (1994). “Structural Estimation of Markov Decision Processes,” In: Engle, R., McFadden, D. (Eds.), *Handbook of Econometrics*. North-Holland, New York, 3081-3143.
- [43] Rust, J., and C. Phelan (1997). “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets,” *Econometrica*, 65, 781-832.
- [44] Ryan, S. P. (2012). “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 80. 1019-1061.
- [45] Schiraldi, P. (2011). “Automobile Replacement: a Dynamic Structural Approach,” *The RAND Journal of Economics*, Volume 42, Issue 2, pages 266–291.
- [46] Scott, P. T. (2013). “Dynamic Discrete Choice Estimation of Agricultural Land Use,” Working Paper, Toulouse School of Economics.
- [47] Varela, M. (2013). “The costs of Growth: Estimating Entry Costs During Rollouts,” Working Paper, University of Arizona.

A Appendix: Proofs

Appendix A summarizes the proofs of the claims given in the main body of the paper. The subsections here correspond to the main paper's sections.

A.1 Nonparametric Identification of Payoffs

A.1.1 Proof of Proposition 1

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 $a = J$

$$V = (I - \beta F_J)^{-1} (\pi_J + \psi_J).$$

After substituting for V , we obtain the result.⁴²

A.2 Identification of Counterfactual Behavior

A.2.1 Proof of Lemma 1

To prove Lemma 1, we make use of Lemma 2 below. Assume without loss of generality that $J = A$. Define

$$\frac{\partial \phi_{-J}}{\partial p} = \begin{bmatrix} \Phi_{11} & \Phi_{12} & \cdots & \Phi_{1,A-1} \\ \Phi_{21} & \Phi_{22} & \cdots & \Phi_{2,A-1} \\ \vdots & \vdots & \vdots & \vdots \\ \Phi_{A-1,1} & \Phi_{A-1,2} & \cdots & \Phi_{A-1,A-1} \end{bmatrix} \equiv \Phi,$$

where Φ_{ij} are the $X \times X$ matrices with elements $\frac{\partial \phi_{iJ}(p(x))}{\partial p_j(x')}$, with $x, x' \in \mathcal{X}$ for each $i, j = 1, \dots, A-1$. Note that each Φ_{ij} is diagonal because $\frac{\partial \phi_{iJ}(p(x))}{\partial p_j(x')} = 0$ when $x \neq x'$.

Next, define the diagonal matrices, P_a , with diagonal element i , $p_a(x_i)$ for $a = 1, \dots, A-1$; and let $P = [P_1, P_2, \dots, P_{A-1}]$.

Lemma 2. *The Arcidiacono-Miller function $\psi_J(p)$ is continuously differentiable with derivative:*

$$\frac{\partial \psi_J}{\partial p} = P\Phi.$$

⁴² $(I - \beta F_J)$ is invertible because F_J is a stochastic matrix and hence the largest eigenvalue is equal or smaller than one. The eigenvalues of $(I - \beta F_J)$ are given by $1 - \beta\gamma$, where γ are the eigenvalues of F_J . Because $\beta < 1$ and $\gamma \leq 1$, we have $1 - \beta\gamma > 0$.

Proof. Recall that

$$\psi_J(p(x)) = \int \max_{k \in \mathcal{A}} \{\phi_{kJ}(p(x)) + \varepsilon_k\} dG(\varepsilon).$$

Because $\phi_{jJ}(p(x))$ is a continuously differentiable function, as shown by Hotz and Miller (1993), so is $\psi_J(p(x))$. For $x \neq x'$, $\frac{\partial \psi_J(p(x))}{\partial p_a(x')} = 0$ for all a , because $\frac{\partial \phi_{kJ}(p(x))}{\partial p_a(x')} = 0$ for all k . For $x = x'$, apply the Chain Rule and obtain

$$\begin{aligned} \frac{\partial \psi_J(p(x))}{\partial p_a(x)} &= \int \frac{\partial}{\partial p_a(x)} \left[\max_{k \in \mathcal{A}} \{\phi_{kJ}(p(x)) + \varepsilon_k\} \right] dG(\varepsilon) \\ &= \sum_{j=1}^{J-1} \int 1 \left\{ j = \arg \max_{k \in \mathcal{A}} \{\phi_{kJ}(p(x)) + \varepsilon_k\} \right\} dG(\varepsilon) \frac{\partial \phi_{jJ}(p(x))}{\partial p_a(x)} \\ &= \sum_{j=1}^{J-1} p_j(x) \frac{\partial \phi_{jJ}(p(x))}{\partial p_a(x)} \end{aligned}$$

Note that

$$\frac{\partial \psi_J}{\partial p} = [\Psi_1, \dots, \Psi_{J-1}]$$

where Ψ_a is the $X \times X$ diagonal matrix with elements $\frac{\partial \psi_J(p(x))}{\partial p_a(x)}$, $x \in \mathcal{X}$, for $a = 1, \dots, J-1$. Hence,

$$\frac{\partial \psi_J}{\partial p} = [P_1, P_2, \dots, P_{J-1}] \Phi.$$

■

To simplify notation, consider the function b_{-J} , instead of \tilde{b}_{-J} . Recall the definition of $b_{-J}(p) : \mathbb{R}^{(A-1)X} \rightarrow \mathbb{R}^{(A-1)X}$ in Section 2. Because $\psi_a = \psi_J - \phi_{aJ}$, we have

$$b_{-J}(p) = \begin{bmatrix} A_1 - I \\ \vdots \\ A_{J-1} - I \end{bmatrix} \psi_J(p) + \phi_{-J}(p) = \mathbf{A} \psi_J(p) + \phi_{-J}(p),$$

where \mathbf{A} has dimension $(A-1)X \times X$ and $\psi_J(p)$ is a column vector with entries $\psi_J(p(x))$, $x \in \mathcal{X}$, and $\phi_{-J}(p)$ is an $(A-1)X$ -valued function with elements $\phi_{aJ}(p(x))$. Because both functions $\psi_J(p)$ and $\phi_{-J}(p)$ are differentiable, by Lemma 2 we have

$$\frac{\partial b_{-J}}{\partial p} = \mathbf{A} \frac{\partial \psi_J}{\partial p} + \frac{\partial \phi_{-J}}{\partial p} = [\mathbf{A}P + I] \Phi$$

Note that, by the Hotz-Miller inversion (Hotz and Miller (1993)), all block-matrices Φ_{ij} of Φ are invertible. Further, the blocks are all linearly independent, so Φ is invertible as well. Thus $\left[\frac{\partial b_{-J}(p)}{\partial p} \right]$ will be invertible if $[\mathbf{A}P + I]$ is. Using the identity $\det(I + AB) = \det(I + BA)$ and the

property $\sum_a P_a = I$, we obtain

$$\det(\mathbf{A}P + I) = \det\left(I + \sum_{a=1}^{J-1} P_a (A_a - I)\right) = \det\left(P_J + \sum_{a=1}^{J-1} P_a A_a\right)$$

But $A_a = (I - \beta F_a)(I - \beta F_J)^{-1}$ and therefore

$$\begin{aligned} \det(\mathbf{A}P + I) &= \det\left(P_J + \sum_{a=1}^{J-1} P_a (I - \beta F_a)(I - \beta F_J)^{-1}\right) \\ &= \det\left(P_J(I - \beta F_J) + \sum_{a=1}^{J-1} P_a (I - \beta F_a)\right) \det((I - \beta F_J)^{-1}) \\ &= \det\left(\sum_{a=1}^J P_a (I - \beta F_a)\right) \det((I - \beta F_J)^{-1}) \\ &= \det\left(I - \beta \sum_{a=1}^J P_a F_a\right) \det((I - \beta F_J)^{-1}) \end{aligned}$$

Note that $\sum_{a=1}^J P_a F_a$ is a stochastic matrix, since all its elements are non-negative and

$$\left(\sum_{a=1}^J P_a F_a\right) \mathbf{1} = \sum_{a=1}^J P_a \mathbf{1} = \left(\sum_{a=1}^J P_a\right) \mathbf{1} = \mathbf{1},$$

where $\mathbf{1}$ is a $X \times 1$ vector of ones. Thus, $\det\left(I - \beta \sum_{a=1}^J P_a F_a\right)$ is nonzero and $\det(\mathbf{A}P + I) \neq 0$.

A.2.2 Proof of Theorem 1

Assume without loss of generality that action $J = A$ belongs to both sets \mathcal{A} and $\tilde{\mathcal{A}}$. The implicit function theorem allows us to locally solve (13) with respect to \tilde{p} provided the matrix

$$\frac{\partial}{\partial \tilde{p}} \left[h_{-J}(\pi) - \tilde{A}_{-J} \tilde{\pi}_J - \tilde{b}_{-J}(\tilde{p}) \right] = -\frac{\partial}{\partial \tilde{p}} \tilde{b}_{-J}(\tilde{p})$$

is invertible; this is proved in Lemma 1.⁴³

The vector \tilde{p} does not depend on the free parameter π_J if and only if

$$\frac{\partial}{\partial \pi_J} \left[h_a(\pi_1, \pi_2, \dots, \pi_J) - \tilde{A}_a h_J(\pi_1, \pi_2, \dots, \pi_J) - \tilde{b}_a(\tilde{p}) \right] = 0$$

⁴³ Because $\pi \in \mathbb{R}^{A \times X}$, the set of payoffs that satisfies (5) is an open linear manifold. Therefore, for any point π in this manifold, there exists a neighborhood for which the implicit function theorem is valid.

for all $a \in \tilde{\mathcal{A}}$, with $a \neq J$, and all π satisfying (5). But, the above yields

$$\sum_{l \in \mathcal{A}, l \neq J} \frac{\partial h_a}{\partial \pi_l} \frac{\partial \pi_l}{\partial \pi_J} + \frac{\partial h_a}{\partial \pi_J} = \tilde{A}_a \left(\sum_{l \in \mathcal{A}, l \neq J} \frac{\partial h_J}{\partial \pi_l} \frac{\partial \pi_l}{\partial \pi_J} + \frac{\partial h_J}{\partial \pi_J} \right)$$

where, for each $a \in \tilde{\mathcal{A}}$ and $l \in \mathcal{A}$, the matrix $\left[\frac{\partial h_a}{\partial \pi_l} \right]$ has dimension $\tilde{X} \times X$; while \tilde{A}_a is an $\tilde{X} \times \tilde{X}$ matrix. Using (4),

$$\sum_{l \in \mathcal{A}, l \neq J} \frac{\partial h_a}{\partial \pi_l} A_l + \frac{\partial h_a}{\partial \pi_J} = \tilde{A}_a \left(\sum_{l \in \mathcal{A}, l \neq J} \frac{\partial h_J}{\partial \pi_l} A_l + \frac{\partial h_J}{\partial \pi_J} \right) \quad (32)$$

or,

$$\begin{bmatrix} \frac{\partial h_a}{\partial \pi_1} & \frac{\partial h_a}{\partial \pi_2} & \cdots & \frac{\partial h_a}{\partial \pi_J} \end{bmatrix} \begin{bmatrix} A_{-J} \\ I \end{bmatrix} = \tilde{A}_a \begin{bmatrix} \frac{\partial h_J}{\partial \pi_1} & \frac{\partial h_J}{\partial \pi_2} & \cdots & \frac{\partial h_J}{\partial \pi_J} \end{bmatrix} \begin{bmatrix} A_{-J} \\ I \end{bmatrix}$$

For $a \in \tilde{\mathcal{A}}$, define the $\tilde{X} \times AX$ matrix (recall $J = A$)

$$\nabla h_a(\pi) = \begin{bmatrix} \frac{\partial h_a}{\partial \pi_1} & \frac{\partial h_a}{\partial \pi_2} & \cdots & \frac{\partial h_a}{\partial \pi_J} \end{bmatrix}.$$

Then, stacking the above expressions for all $a \in \tilde{\mathcal{A}}$, with $a \neq J$, we obtain

$$\nabla h_{-J}(\pi) \begin{bmatrix} A_{-J} \\ I \end{bmatrix} = \tilde{A}_{-J} \nabla h_J(\pi) \begin{bmatrix} A_{-J} \\ I \end{bmatrix}.$$

Now apply the property $vecbr(BCA') = (A \boxtimes B) vecbr(C)$ to obtain:

$$\begin{aligned} & \left(\begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes I \right) vecbr(\nabla h_{-J}(\pi)) - \left(\begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes \tilde{A}_{-J} \right) vecbr(\nabla h_J(\pi)) = 0 \\ & \underbrace{\left[\begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes I, - \begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes \tilde{A}_{-J} \right]}_{(\tilde{A}-1)\tilde{X}X \times (\tilde{A}\tilde{X})(AX)} \underbrace{\begin{bmatrix} vecbr(\nabla h_{-J}(\pi)) \\ vecbr(\nabla h_J(\pi)) \end{bmatrix}}_{(\tilde{A}\tilde{X})(AX) \times 1} = 0, \end{aligned}$$

which is (14). Note that $\begin{bmatrix} A'_{-J} & I \end{bmatrix}$ is an $X \times AX$ matrix, while $\left(\begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes I \right)$ is an $(\tilde{A}-1)\tilde{X}X \times (\tilde{A}-1)\tilde{A}\tilde{X}X$ matrix. Similarly, $\left(\begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes \tilde{A}_{-J} \right)$ is an $(\tilde{A}-1)\tilde{X}X \times A\tilde{X}X$ matrix, and \tilde{A}_{-J} is an $(\tilde{A}-1)\tilde{X} \times \tilde{X}$ matrix.⁴⁴

⁴⁴ With abuse of notation, the identity matrix in $\begin{bmatrix} A'_{-J} & I \end{bmatrix}$ is an $X \times X$ matrix, while the identity matrix after \boxtimes in $\left(\begin{bmatrix} A'_{-J} & I \end{bmatrix} \boxtimes I \right)$ is $(\tilde{A}-1)\tilde{X} \times (\tilde{A}-1)\tilde{X}$.

A.2.3 Proof of the Bus Engine Replacement Example in Section 3.2

Let $a = 1$ if *replace*, and $a = 2$ if *keep*. Then

$$\tilde{\pi} = \begin{bmatrix} (1 + \lambda)I & -\lambda[1, 0] \\ 0 & I \end{bmatrix} \pi,$$

where 1 is a vector of ones and 0 is a matrix with zeros. Let $J = 2$. By equation (16) in Corollary 1, \tilde{p} is identified if and only if

$$\left(H_{11} - \tilde{A}_1 H_{21} \right) A_1 + H_{12} - \tilde{A}_1 H_{22} = 0$$

or $(1 + \lambda)IA_1 - \lambda[1, 0] - A_1 = 0$, which implies $\lambda A_1 = \lambda[1, 0]$. This implies A_1 is non-invertible, which is a contradiction.

A.2.4 Proof of Corollary 1

Because $\frac{\partial h_a}{\partial \pi_l} = H_{al}$, equation (32) in the proof of Theorem 1 becomes

$$\sum_{l \neq J} H_{al} A_l + H_{aJ} = \tilde{A}_a \left(\sum_{l \neq J} H_{Jl} A_l + H_{JJ} \right).$$

In the ‘‘action diagonal’’ case, $H_{al} = H_{Jl} = 0$ for all $a, J \neq l$, and the condition simplifies to

$$H_a A_a = \tilde{A}_a H_J.$$

A.2.5 Proof of Proposition 2

Equation (17) implies $H_a = A_a H_J A_a^{-1}$, for all $a \neq J$. So all H_a must be similar. Diagonal similar matrices are equal to each other, which implies $H_a = H$, for all a .

Let A_a be partitioned conformably with H :

$$A_a = \begin{bmatrix} (A_a)_{11} & \cdots & (A_a)_{1d} \\ \vdots & \vdots & \vdots \\ (A_a)_{d1} & \cdots & (A_a)_{dd} \end{bmatrix}$$

Then, $HA_a - A_a H = 0$ implies that for all $i \neq j$, $(\lambda_i - \lambda_j)(A_a)_{ij} = 0$, and since $\lambda_i \neq \lambda_j$, it must be $(A_a)_{ij} = 0$ and A_a is block-diagonal. This proves the equivalence of statements (i) and (ii).

We next prove that statement (ii) implies statement (iii). Suppose A_a is block-diagonal. Then, $A_a(I - \beta F_J) = (I - \beta F_a)$, or

$$I - A_a = \beta (F_a - A_a F_J) \tag{33}$$

The left-hand side is block diagonal and its (i, j) block is equal to zero. Therefore,

$$0 = (F_a)_{ij} - \sum_k (A_a)_{ik} (F_J)_{kj}$$

Since A_a is block-diagonal, $(A_a)_{ik} = 0$ for $i \neq k$ and thus $(F_a)_{ij} = (A_a)_{ii} (F_J)_{ij}$. Moreover, if we equate the diagonal blocks in (33), we have:

$$I - (A_a)_{ii} = \beta ((F_a)_{ii} - (A_a)_{ii} (F_J)_{ii})$$

and $(A_a)_{ii} (F_J)_{ii} = (A_a F_J)_{ii}$, since A_a is block-diagonal. Rearranging, we establish the claim.

Finally, we show the reverse. Consider the first block-row of A_a , $a_1 = [(A_a)_{11} \ (A_a)_{12} \ \dots \ (A_a)_{1d}]$. Let $e_1 = [I \ 0 \ \dots 0]$. Then $a_1 = e_1 A_a = e_1 (I - \beta F_a) (I - \beta F_J)^{-1}$. But $e_1 (I - \beta F_a) = [I - \beta (F_a)_{11} \ \dots I - \beta (F_a)_{1d}]$ and statement (iii) implies that

$$e_1 (I - \beta F_a) = [(A_a)_{11} (I - \beta (F_J)_{11}) \ \dots (A_a)_{11} (I - \beta (F_J)_{1d})] = (A_a)_{11} e_1 (I - \beta F_J)$$

Therefore,

$$a_1 = e_1 A_a = (A_a)_{11} e_1 (I - \beta F_J) (I - \beta F_J)^{-1} = (A_a)_{11} e_1 = [(A_a)_{11} \ 0 \ \dots 0]$$

We conclude that $(A_a)_{1j} = 0$ for $j \neq 1$. The same argument applied to all block-rows shows that A_a is block diagonal with block entries given by $(A_a)_i$.

A.2.6 Proof of Proposition 3

Suppose the counterfactual replaces the payoff of type s_1 by that of s_2 for action J only. Then: $H_a = I$ and

$$H_J = \begin{bmatrix} 0 & I \\ 0 & I \end{bmatrix} \quad (34)$$

From Corollary 1 for affine counterfactuals, identification requires that, $H_a A_a = A_a H_J$. We partition A_a conformably with H_J , i.e.

$$A_a = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}$$

The identification condition then leads to $A_{11} = A_{21} = 0$ with A_{12}, A_{22} arbitrary.

But in the case of time-invariant types, A_a is block diagonal, as it is not possible to transit from one type to the other (i.e. $A_{ii} = (I - \beta F_a^{s_i})(I - \beta F_J^{s_i})^{-1}$, for $i = 1, 2$ and $A_{12} = A_{21} = 0$). Invertibility of A_a implies that the diagonal blocks are invertible as well, and therefore, the counterfactual is not identified.

Next, suppose the counterfactual replaces the payoff of type s_1 by that of s_2 for all actions. Then, $H_a = H_J$ and both are given by (34).

Therefore identification requires $HA = AH$, or, $A_{21} = 0$ and $A_{11} + A_{12} = A_{22}$, $A_{21} + A_{22} = A_{22}$. Given the form of A_a , $A_{21} = 0$ is automatically satisfied, while $A_{12} = 0$ implies that identification requires $A_{22} = A_{11}$ and A_{22} arbitrary.

A.2.7 Proof of Proposition 4

Assume $\tilde{\pi} = \mathcal{H}\pi + g$. We prove the statement in three steps. First we show that all off-diagonal submatrices H_{aj} , $a \neq j$, must have identical rows. Second, we show that, when $J \geq 3$, the off-diagonal blocks in column a of \mathcal{H} must be identical to each other; i.e., $H_{ja} = H_{la} = \bar{H}_a$, for any combination of $j \neq l \neq a$. Finally, we show that all diagonal blocks must be of the form $H_{aa} = \lambda I + \bar{H}_a$ for any choice a , for some scalar λ .

(a) By Corollary 1, equation (16) must hold for all $a \neq J$ and for any arbitrary process F . Take an action $a \neq J$ and post-multiply (16) by $(I - \beta F_J)$. We get

$$\sum_{l \in \mathcal{A}, l \neq J} (H_{al} - A_a H_{Jl}) (I - \beta F_l) + H_{aJ} (I - \beta F_J) - A_a H_{JJ} (I - \beta F_J) = 0.$$

Take $F_J = I$ (this is allowable), then $A_a = (I - \beta F_a) (1 - \beta)^{-1}$ for all $l \neq J$, and

$$\sum_{l \in \mathcal{A}, l \neq J} [(H_{al} - (1 - \beta)^{-1} (I - \beta F_a) H_{Jl}) (I - \beta F_l)] + (1 - \beta) H_{aJ} - (I - \beta F_a) H_{JJ} = 0.$$

Rearranging, we get

$$\begin{aligned} & \sum_{l \in \mathcal{A}, l \neq J} [H_{al} - (1 - \beta)^{-1} H_{Jl} + (1 - \beta) H_{aJ} - H_{JJ}] + [(1 - \beta)^{-1} \beta H_{Jl} - \beta H_{al}] F_l \\ & + F_a [(1 - \beta)^{-1} \beta H_{Jl} + \beta H_{JJ}] - F_a [(1 - \beta)^{-1} \beta^2 H_{Jl}] F_l = 0. \end{aligned}$$

This equals

$$\mathbf{A} + \mathbf{B}F_a + F_a\mathbf{C} + F_a\mathbf{D}F_a = 0, \tag{35}$$

where

$$\begin{aligned} \mathbf{A} &= \sum_{l \in \mathcal{A}, l \neq J} [H_{al} - (1 - \beta)^{-1} H_{Jl} + (1 - \beta) H_{aJ} - H_{JJ}] + \sum_{l \in \mathcal{A}, l \neq a, J} [(1 - \beta)^{-1} \beta H_{Jl} - \beta H_{al}] F_l \\ \mathbf{B} &= (1 - \beta)^{-1} \beta H_{Ja} - \beta H_{aa} \\ \mathbf{C} &= \sum_{l \in \mathcal{A}, l \neq J} [(1 - \beta)^{-1} \beta H_{Jl} + \beta H_{JJ}] - \sum_{l \in \mathcal{A}, l \neq a, J} [(1 - \beta)^{-1} \beta^2 H_{Jl}] F_l \\ \mathbf{D} &= -(1 - \beta)^{-1} \beta^2 H_{Ja} \end{aligned}$$

with $F_a \geq 0$ and $F_a \mathbf{1} = 1$, where $\mathbf{1}$ is a vector of ones. Fix F_J and F_l for $l \neq a$. The left hand side of equation (35) can be viewed as a quadratic function in F_a . If this identity is satisfied for all $F_a \geq 0$, then all derivatives of the quadratic function with respect to F_a must be equal to zero.

Let the columns of F_a be

$$F_a = \begin{bmatrix} f_1 & f_2 & \dots & f_{n-1} & 1 - \sum_i^{n-1} f_i \end{bmatrix}$$

where n is the number of columns (note that $n = X$). We first take the second derivative, and so we focus on the term $F_a \mathbf{D} F_a$. For $j \neq n$, the (i, j) entry is:

$$(F_a \mathbf{D} F_a)_{i,j} = \sum_{l,k} f_{ik} d_{lk} f_{kj} = \sum_{l,k} d_{lk} f_{ik} f_{kj}$$

Isolate the last entry and substitute in:

$$\begin{aligned} (F_a \mathbf{D} F_a)_{i,j} &= \sum_{l \neq n, k} d_{lk} f_{ik} f_{kj} + \sum_k d_{lk} f_{in} f_{kj} = \sum_{l \neq n, k} d_{lk} f_{ik} f_{kj} + \left(1 - \sum_m^{n-1} f_m\right) \sum_k d_{nk} f_{kj} \\ &= \sum_{l \neq n, k} f_{ik} f_{kj} (d_{lk} - d_{nk}) \end{aligned}$$

and therefore we must have

$$d_{lk} = d_{nk}, \text{ for all } k, l \neq n$$

Consider now $j = n$. Then,

$$\begin{aligned} (F_a \mathbf{D} F_a)_{i,n} &= \sum_{l,k} f_{ik} d_{lk} f_{kn} = \sum_{l \neq n, k} f_{ik} d_{lk} \left(1 - \sum_{m=1}^{n-1} f_{km}\right) + \sum_k \left(1 - \sum_{l=1}^{n-1} f_{il}\right) d_{lk} \left(1 - \sum_{m=1}^{n-1} f_{km}\right) \\ &= \sum_{k,l,m} f_{il} f_{km} (d_{nk} - d_{lk}) \end{aligned}$$

which already holds. We conclude that \mathbf{D} has identical rows, which implies H_{Ja} also has identical rows. Because this argument holds for any $a \neq J$, and because the choice of J is arbitrary, each off-diagonal submatrix H_{aj} must have identical rows for *any* pair of actions a and j .

The following facts will be useful below. First note that $A_a \mathbf{1} = 1$ for any a , since

$$\begin{aligned} A_a \mathbf{1} &= (I - \beta F_a) (I - \beta F_J)^{-1} \mathbf{1} = (I - \beta F_a) \sum_{n=0}^{\infty} \beta^n F_J^n \mathbf{1} = \frac{1}{1 - \beta} (I - \beta F_a) \mathbf{1} \\ &= \frac{1}{1 - \beta} (1 - \beta F_a \mathbf{1}) = \frac{1}{1 - \beta} (1 - \beta) \mathbf{1} = \mathbf{1}. \end{aligned}$$

Given that, take any two actions, j and l , and let $H_{jl} = [\rho_{1jl}1, \dots, \rho_{Xjl}1]$. Then,

$$A_a H_{jl} = A_a [\rho_{1jl}1, \dots, \rho_{Xjl}1] = [\rho_{1jl}A_a 1, \dots, \rho_{Xjl}A_a 1] = H_{jl}.$$

(b) Next, consider the case $J \geq 3$, and take $j \neq a, J$. Return to equation (16). Rearrange it and isolate the terms involving j :

$$\begin{aligned} (H_{aj} - A_a H_{Jj})(I - \beta F_j) &= A_a H_{JJ}(I - \beta F_J) - H_{aJ}(I - \beta F_J) \\ &\quad - \sum_{l \in \mathcal{A}, l \neq j, J} (H_{al} - A_a H_{Jl})(I - \beta F_l). \end{aligned}$$

Fix F_J and F_l , for all $l \neq j, J$, and view this as a function of F_j . The right-hand-side does not depend on F_j , and the term $(H_{aj} - A_a H_{Jj})$ on the left-hand-side is fixed. We need the equality to hold for any F_j . Because $(I - \beta F_j)$ is full rank, the only way this equality can be satisfied for all choices of F_j is for

$$H_{aj} - A_a H_{Jj} = 0.$$

We have shown that $A_a H_{jl} = H_{jl}$ for any pair of actions j and l . We therefore obtain

$$H_{aj} = H_{Jj}.$$

The argument holds for any combination of $j \neq a \neq J$. Take the block-column j of \mathcal{H} , then all off-diagonal blocks in column j are identical to each other (in addition to having identical rows each). Denote the off-diagonal matrices in the block-column j by \overline{H}_j . I.e., $\overline{H}_j = H_{aj}$, for all pairs $a \neq j$.

(c) Now, return to the case $J \geq 2$. We investigate the block-diagonal terms of \mathcal{H} . Again, take (16) for $a \neq J$,

$$(H_{aa} - A_a H_{Ja}) A_a + \sum_{l \in \mathcal{A}, l \neq a, J} (H_{al} - A_a H_{Jl}) A_l + H_{aJ} - A_a H_{JJ} = 0.$$

Given that all off-diagonal terms H_{al} , for all pairs $a \neq l$, must satisfy $H_{al} - A_a H_{Jl} = 0$, equation (16) simplifies to

$$(H_{aa} - A_a H_{Ja}) A_a + H_{aJ} - A_a H_{JJ} = 0.$$

In addition, for all pairs $a \neq l$, we have that $H_{al} = \overline{H}_l$, which implies

$$(H_{aa} - \overline{H}_a) A_a - A_a (H_{JJ} - \overline{H}_J) = 0.$$

Furthermore, if we take $F_a = F_J$ (this is allowable), we get

$$(H_{aa} - \overline{H}_a) = (H_{JJ} - \overline{H}_J).$$

If we take $F_J = I$ instead, we get

$$(H_{aa} - \bar{H}_a)(I - \beta F_a) - (I - \beta F_a)(H_{JJ} - \bar{H}_J) = 0.$$

Rearranging, we obtain

$$(H_{aa} - \bar{H}_a)F_a = F_a(H_{aa} - \bar{H}_a).$$

So, $(H_{aa} - \bar{H}_a)$ and F_a must commute, where F_a is an arbitrary (stochastic) matrix. This implies $(H_{aa} - \bar{H}_a)$ must be of the form

$$(H_{aa} - \bar{H}_a) = \lambda I,$$

for all a , where λ is a constant. Finally, note that, for \mathcal{H} with diagonal blocks $H_{aa} = \lambda I + \bar{H}_a$, and off-diagonal blocks $H_{aj} = \bar{H}_j$, we obtain

$$\tilde{\pi} = \mathcal{H}\pi + g = \lambda\pi + \left[\sum_{j \in \mathcal{A}} \bar{H}_j \pi_j \right] 1 + g$$

where all rows of $\sum_{j \in \mathcal{A}} \bar{H}_j \pi_j$ are identical, and therefore all elements of the vector $[\sum_{j \in \mathcal{A}} \bar{H}_j \pi_j]1$ are the same.

A.2.8 Proof of Corollary 2

Because $H_a = I$ for all a , equation (17) is satisfied if and only if $A_a = \tilde{A}_a$.

A.2.9 Proof of Proposition 5

Suppose $\mathcal{A} = \{1, 2, \dots, A\}$. Without loss of generality, take the reference action to be $J = 1$ and suppose action $j = A + 1$ is new, so that $\tilde{\mathcal{A}} = \{1, 2, \dots, A + 1\}$. Assume $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{F}_a = F_a$, and $\tilde{\pi} = \mathcal{H}\pi + g$, with $\tilde{\pi}_a = \pi_a$ for all $a \in \mathcal{A}$, and

$$\tilde{\pi}_j = \sum_{a=1}^A H_{ja} \pi_a + g_j.$$

The identification condition (16) becomes $A_a = \tilde{A}_a$, for $a = 2, \dots, A$, and

$$H_{j1} + \sum_{a=2}^A H_{ja} A_a = \tilde{A}_j, \tag{36}$$

for $j = A + 1$, since $H_{al} = 0$ and $H_a = I$ for all $a, l \neq j$. The first set of restrictions are satisfied, since transitions are unaffected. Now, post-multiply (36) by $(I - \beta F_1) = (I - \beta \tilde{F}_1)$ to obtain:

$$H_{j1}(I - \beta F_1) + \sum_{a=2}^A H_{ja}(I - \beta F_a) = I - \beta \tilde{F}_j$$

or

$$\tilde{F}_j = \sum_{a=1}^A H_{ja} F_a + \beta^{-1} \left(I - \sum_{a=1}^A H_{ja} \right)$$

Since transitions are stochastic matrices, we have that $\tilde{F}_j 1 = 1$, so that

$$1 = \sum_{a=1}^A H_{ja} 1 + \beta^{-1} \left(1 - \sum_{a=1}^A H_{ja} 1 \right)$$

or $\sum_{a=1}^A H_{ja} 1 = 1$.

A.2.10 Proof of Proposition 6

If $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{F}_a = F_a$, and $\tilde{\pi}_a = \pi_a$ for all $a \in \tilde{\mathcal{A}}$, then $H_{aa} = I$ and $H_{ak} = 0$ for $a \in \tilde{\mathcal{A}}$ and $k \in \mathcal{A}$, $a \neq k$, and so (16) becomes $A_a = \tilde{A}_a$ for all $a \in \tilde{\mathcal{A}}$, which is satisfied because $\tilde{F}_a = F_a$ for all a .

A.2.11 Proof of Proposition 7

The proof relies on the following lemma:

Lemma 2. *Set the reference action to be $J = 1$ and suppose action $a = A$ is eliminated. Suppose further that the first m states are maintained and the remaining $X - m$ are eliminated. The counterfactual is specified by:*

$$\tilde{\pi}_a = \begin{bmatrix} I_m & 0 \end{bmatrix} \pi_a \quad (37)$$

We partition the transition matrix as follows:

$$F_a = \begin{bmatrix} \hat{F}_a & f_a \\ g_a & q_a \end{bmatrix} \quad (38)$$

where \hat{F}_a is the $m \times m$ top left submatrix of F_a , corresponding to the maintained states; f_a has dimension $m \times (X - m)$; g_a is $(X - m) \times m$; and q_a is $(X - m) \times (X - m)$. Counterfactual transitions adjust the maintained states as follows,

$$\tilde{F}_a = \hat{F}_a + f_a r \quad (39)$$

where r is a $m \times (X - m)$ matrix such that $r1 = 1$, to secure that \tilde{F}_a is a stochastic matrix. The

counterfactual is identified if and only if

$$(I - \beta\hat{F}_a)^{-1}f_a = (I - \beta\hat{F}_1)^{-1}f_1 \quad (40)$$

or $f_a = \hat{A}_a f_1$, where $\hat{A}_a = (I - \beta\hat{F}_a)(I - \beta\hat{F}_1)^{-1}$.

Proof. The identification condition is $H_a A_a = \tilde{A}_a H_a$ or $H_a(I - \beta F_a) = \tilde{A}_a H_a(I - \beta F_1)$. Combining (37) and (38), we obtain:

$$\begin{bmatrix} I - \beta\hat{F}_a & -\beta f_a \end{bmatrix} = \begin{bmatrix} \tilde{A}_a (I - \beta\hat{F}_1) & -\beta\tilde{A}_a f_1 \end{bmatrix}$$

or

$$I - \beta\hat{F}_a = \tilde{A}_a (I - \beta\hat{F}_1) \quad (41)$$

and

$$f_a = \tilde{A}_a f_1$$

We show that (41) is redundant when (39) holds. Indeed, (41) is written $I - \tilde{A}_a = \beta(\hat{F}_a - \tilde{A}_a\hat{F}_1)$, while by definition, $I - \tilde{A}_a = \beta(\tilde{F}_a - \tilde{A}_a\tilde{F}_1)$. Thus, $\hat{F}_a - \tilde{A}_a\hat{F}_1 = \tilde{F}_a - \tilde{A}_a\tilde{F}_1$ and using (39), $\hat{F}_a - \tilde{A}_a\hat{F}_1 = \hat{F}_a + f_a r - \tilde{A}_a(\hat{F}_a + f_a r)$, or $f_a r = \tilde{A}_a f_1 r$, which holds because of (41). ■

Next, we return to Proposition 7. Assume $x = (k, w)$, then

$$F_a = F_a^k \otimes F^w = \begin{bmatrix} f_{11}^a F^w & f_{12}^a F^w & \dots & f_{1K}^a F^w \\ \vdots & \vdots & \dots & \vdots \\ f_{K1}^a F^w & f_{K2}^a F^w & \dots & f_{KK}^a F^w \end{bmatrix}$$

where $f_{ij}^a = \Pr(k' = j | k = i, a)$ are the elements of F_a^k . Because $k_t = a_{t-1}$, F_a is a matrix with zeros except in the a -th block-column. The a -th block-column is a block-vector with blocks F^w . If action $a = A$ is eliminated from $\mathcal{A} = \{1, 2, \dots, A\}$, then for all $a \neq A$, we have $f_a = 0$, where f_a is defined in (38). Because $f_J = 0$ as well, condition (40) in Lemma 2 is trivially satisfied.

A.3 Identification of Counterfactual Behavior Under Additional Model Restrictions

A.3.1 Proof of Proposition 8

We make use of two lemmas. Lemma 3 provides sufficient conditions for the identification of parametric models with linear-in-parameters payoff functions.

Lemma 3. *If $\pi(a, x)$ satisfies*

$$\pi(a, x; \theta) = \bar{\pi}_a(x) \theta, \quad (42)$$

where θ is a finite dimensional parameter. The parameter θ is identified provided $\text{rank} [\bar{\pi}_{-J} - A_{-J}\bar{\pi}_J] = \dim(\theta)$, where $\bar{\pi}_{-J} = [\bar{\pi}_1, \dots, \bar{\pi}_{J-1}, \bar{\pi}_{J+1}, \dots, \bar{\pi}_A]$.

Proof. Equation (42) implies $\pi_{-J} = \bar{\pi}_{-J}\theta$. Then (5) becomes $[\bar{\pi}_{-J} - A_{-J}\bar{\pi}_J]\theta = b_{-J}$. So, if the rank of the matrix $[\bar{\pi}_{-J} - A_{-J}\bar{\pi}_J]$ equals $\dim(\theta)$, then θ is uniquely determined. ■

Next, Lemma 4 provides results that are used to prove Proposition 8.

Lemma 4. Let $D_a = [I - \beta(F^w \otimes F_a^k)]^{-1}$, where I is the identity matrix of size $KW \times KW$. Let

I_k be the identity matrix of size K , and $\mathbf{1}$ be the block vector $\mathbf{1} = \begin{bmatrix} I_k \\ \vdots \\ I_k \end{bmatrix}$ of size $KW \times K$. Finally,

let $A_a^k = (I_k - \beta F_a^k) (I_k - \beta F_J^k)^{-1}$. The following properties hold:

- (i) $D_a^{-1}\mathbf{1} = (I - \beta(F^w \otimes F_a^k))\mathbf{1} = \mathbf{1}(I_k - \beta F_a^k)$.
- (ii) $D_a\mathbf{1} = (I - \beta(F^w \otimes F_a^k))^{-1}\mathbf{1} = \mathbf{1}(I_k - \beta F_a^k)^{-1}$.
- (iii) $A_a\mathbf{1} = \mathbf{1}A_a^k$.

Statements (ii) and (iii) state that the sum of block entries on each block row of D_a and A_a is constant for all block rows.

Proof. (i) Since F^w is a stochastic matrix, its rows sum to 1: $\sum_j f_{ij}^w = 1$, where f_{ij}^w is the (i, j) element of F^w . By the definition of the Kronecker product,

$$(F^w \otimes F_a^k)\mathbf{1} = \begin{bmatrix} f_{11}^w F_a^k & f_{12}^w F_a^k & \dots & f_{1W}^w F_a^k \\ \vdots & \vdots & \dots & \vdots \\ f_{W1}^w F_a^k & f_{W2}^w F_a^k & \dots & f_{WW}^w F_a^k \end{bmatrix} \begin{bmatrix} I_k \\ \vdots \\ I_k \end{bmatrix} = \begin{bmatrix} (\sum_j f_{1j}^w) F_a^k \\ \vdots \\ (\sum_j f_{Wj}^w) F_a^k \end{bmatrix} = \mathbf{1}F_a^k$$

Thus, $(I - \beta(F^w \otimes F_a^k))\mathbf{1} = \mathbf{1}(I_k - \beta F_a^k)$.

(ii) Let n be a non-negative integer. Then, $(F^w)^n$ is a stochastic matrix with rows summing to 1. Therefore,

$$(F^w \otimes F_a^k)^n = (F^w)^n \otimes (F_a^k)^n$$

and following the proof of (i), we obtain $(F^w \otimes F_a^k)^n \mathbf{1} = \mathbf{1}(F_a^k)^n$. Now,

$$D_a\mathbf{1} = \sum_{n=0}^{\infty} \beta^n (F^w \otimes F_a^k)^n \mathbf{1} = \mathbf{1} \sum_{n=0}^{\infty} \beta^n (F_a^k)^n = \mathbf{1}(I_k - \beta F_a^k)^{-1}.$$

(iii) The proof is a direct consequence of (i) and (ii). Indeed,

$$A_a\mathbf{1} = (I - \beta(F^w \otimes F_a^k))D_J\mathbf{1} = (I - \beta(F^w \otimes F_a^k))\mathbf{1}(I_k - \beta F_J^k)^{-1} = \mathbf{1}(I_k - \beta F_a^k)(I_k - \beta F_J^k)^{-1} = \mathbf{1}A_a^k.$$

■

We now prove Proposition 8; we focus on the binary choice $\{a, J\}$ for notational simplicity, but the general case is obtained in the same fashion. Let θ be the vector of $4K$ unknown parameters (e.g. $\theta_0^a = [\theta_0(a, k_1), \dots, \theta_0(a, k_K)]'$),

$$\theta = \begin{bmatrix} \theta_0^a \\ \theta_0^J \\ \theta_1^a \\ \theta_1^J \end{bmatrix}.$$

The parametric form of interest is linear in the parameters; stacking the payoffs for a given w and all k we have:

$$\pi_a(w) = [I_k, 0_k, Z_a(w)I_k, 0_k] \theta$$

and

$$\pi_J(w) = [0_k, I_k, 0_k, Z_J(w)I_k] \theta$$

Collecting $\pi_a(w)$ for all w , we get $\pi_a = \bar{\pi}_a \theta$, where

$$\bar{\pi}_a = \begin{bmatrix} I_k & 0_k & Z_a(w_1)I_k & 0_k \\ \vdots & \vdots & \vdots & \vdots \\ I_k & 0_k & Z_a(W)I_k & 0_k \end{bmatrix} \quad (43)$$

and similarly for π_J . In Lemma 3, we showed that identification hinges on the matrix $(\bar{\pi}_a - A_a \bar{\pi}_J)$. This matrix equals:

$$\bar{\pi}_a - A_a \bar{\pi}_J = \begin{bmatrix} 1, & -A_a 1, & Z_a, & -A_a Z_J \end{bmatrix} \quad (44)$$

where $Z_a = [Z_a(w_1)I_k, \dots, Z_a(w_W)I_k]'$ (the same for Z_J).

It follows from Lemma 4 that the first two block columns of (44) consist of identical blocks each (the first block column has elements I_k , and the second, $-A_a^k$). As a consequence, the respective block parameters θ_0^a, θ_0^J , are not identified unless extra restrictions are imposed.⁴⁵ The remaining parameters, θ_1^a, θ_1^J , are identified as follows.

Consider $(\bar{\pi}_a - A_a \bar{\pi}_J) \theta = b_a$, or using (44):

$$1\theta_0^a - 1A_a^k \theta_0^J + Z_a \theta_1^a - [I - \beta(F^w \otimes F_a^k)] [I - \beta(F^w \otimes F_J^k)]^{-1} Z_J \theta_1^J = b_a.$$

Left-multiplying both sides by $D_a = [I - \beta(F^w \otimes F_a^k)]^{-1}$ and using Lemma 4, we obtain:

$$1(I_k - \beta F_a^k)^{-1} \theta_0^a - 1(I_k - \beta F_J^k)^{-1} \theta_0^J + D_a Z_a \theta_1^a - D_J Z_J \theta_1^J = D_a b_a.$$

⁴⁵In the multiple choice one block column is a linear combination of the remaining $(J - 1)$ corresponding to θ_0 ; therefore we need to fix θ_0^J for one action J to identify θ_0^{-J} .

Take the w block row of the above:

$$(I_k - \beta F_a^k)^{-1} \theta_0^a - (I_k - \beta F_J^k)^{-1} \theta_0^J + e'_w D_a Z_a \theta_1^a - e'_w D_J Z_J \theta_1^J = e'_w D_a b_a \quad (45)$$

where $e'_w = [0, 0, \dots, I_k, 0, \dots, 0]$ with I_k in the w position. Since $W \geq 3$, take two other distinct block rows corresponding to \tilde{w}, \bar{w} and difference both from the above to obtain the parameter θ_1^a, θ_1^J .

A.3.2 Proof of Proposition 9

(i) Consider the counterfactual payoff $\tilde{\pi}(a, k, w) = \theta_0(a, k) + h_1 [Z'(a, w)\theta_1(a, k)]$. Since the term $Z'(a, w)\theta_1(a, k)$ is known for all (a, k, w) , we can write this as an “additive changes” as follows: $\tilde{\pi}(a, k, w) = \pi(a, k, w) + g$, where $g = h_1 [Z'(a, w)\theta_1(a, k)] - Z'(a, w)\theta_1(a, k)$ is known.

(ii) Consider the counterfactual

$$\tilde{\pi}(a, w) = H_0(a) \theta_0(a) + Z'(a, w) \theta_1(a)$$

for $a = 1, \dots, J$, where we stack $\theta_0(a, k)$ and $\theta_1(a, k)$ for all k and $H_0(a)$ is a $K \times K$ matrix. From the proof of Proposition 8, equation (45), we know that for any w , the w block row of (4) is

$$(I_k - \beta F_a^k)^{-1} \theta_0^a - (I_k - \beta F_J^k)^{-1} \theta_0^J + e'_w D_a Z_a \theta_1^a - e'_w D_J Z_J \theta_1^J = e'_w D_a b_a(p).$$

The corresponding w block row for the counterfactual scenario is

$$(I_k - \beta F_a^k)^{-1} H_0(a) \theta_0^a - (I_k - \beta F_J^k)^{-1} H_0(J) \theta_0^J + e'_w D_a Z_a \theta_1^a - e'_w D_J Z_J \theta_1^J = e'_w D_a b_a(\tilde{p}).$$

Lack of identification of θ_0 is represented by the free parameter θ_0^J . Using (45), we prove the claim.

(iii) From item (ii) above, it is clear that when F_a^k changes, \tilde{p} is identified if and only if for all $a \neq J$, $A_a^k = \tilde{A}_a^k$.

(iv) When $\tilde{F}^w \neq F^w$ and $\tilde{F}_a^k = F_a^k$, the equality $A_a^k = \tilde{A}_a^k$ trivially holds.

A.3.3 Proof of Proposition 10

To prove Proposition 10, first note that for affine counterfactuals, \tilde{p} is a function of π_J through an equation of the form:

$$C\pi_J + c_0 = b(\tilde{p})$$

where C is an $(A - 1)X \times X$ matrix and c_0 an $(A - 1)X \times 1$ vector that does not depend on π_J or \tilde{p} .⁴⁶

The counterfactual CCP is identified under the set of linear restrictions (28) if for all $\pi_J \neq \hat{\pi}_J$ such that $Q\pi_J = Q\hat{\pi}_J = q$, it holds $\tilde{p}(\pi_J) = \tilde{p}(\hat{\pi}_J)$. Proposition 10 is a direct consequence of the following proposition:

Proposition 12. *Given C and the linear restrictions (26), the counterfactual CCP is identified if and only if any of the following equivalent statements hold:*

1. For all $\pi_J \neq \hat{\pi}_J$ such that $Q\pi_J = Q\hat{\pi}_J = q$, it holds $C(\pi_J - \hat{\pi}_J) = 0$.
2. $N(Q) \subseteq N(C)$, where $N(Q)$ denotes the null space of Q .
3. There is an $(A - 1)X \times d$ matrix M such that $C = MQ$.

Proof. For the first statement, take $\pi_J \neq \hat{\pi}_J$ such that $Q\pi_J = Q\hat{\pi}_J = q$. Then $\tilde{p}(\pi_J) = \tilde{p}(\hat{\pi}_J)$ implies $b^{-1}(C\pi_J + c_0) = b^{-1}(C\hat{\pi}_J + c_0)$, or $C\pi_J + c_0 = C\hat{\pi}_J + c_0$, or $C(\pi_J - \hat{\pi}_J) = 0$. It is clear that this argument holds in reverse as well.

For the second statement, take $u \in N(Q)$ and π_J^* a particular solution of $Q\pi_J = q$. Set $\hat{\pi}_J = \pi_J^* + u$ and note that $\hat{\pi}_J$ is also a solution (since $Q\hat{\pi}_J = Q\pi_J^* + Qu$ and $Qu = 0$ since u is in the null-space of Q). Thus,

$$C(\hat{\pi}_J - \pi_J^*) = C(\pi_J^* + u - \pi_J^*) = Cu = 0$$

For the converse, suppose that there exist $\pi_J \neq \hat{\pi}_J$ such that $Q\pi_J = Q\hat{\pi}_J = q$, but $C(\pi_J - \hat{\pi}_J) \neq 0$. Then, $Q(\pi_J - \hat{\pi}_J) = q - q = 0$. Thus, $\pi_J - \hat{\pi}_J \in N(Q) \subseteq N(C)$ and thus $C(\pi_J - \hat{\pi}_J) = 0$, contradiction.

We now show that 2 and 3 are equivalent. We rely on the known fact that $N(Q) = [\text{col}(Q')]^\perp$, where $[X]^\perp$ is the orthogonal complement of X . Hence, $N(Q) = [\text{row}(Q)]^\perp$. It is easy to show that $A \subseteq B$ if and only if $[B]^\perp \subseteq [A]^\perp$. Therefore,

$$N(Q) \subseteq N(C) \Leftrightarrow [\text{row}(Q)]^\perp \subseteq [\text{row}(C)]^\perp \Leftrightarrow \text{row}(C) \subseteq \text{row}(Q) \quad (46)$$

⁴⁶ To see this, following the binary example in Section 3.2, we obtain:

$$\left(\sum_{\hat{a} \neq J} (H_{a\hat{a}} - \tilde{A}_a H_{J\hat{a}}) A_{\hat{a}} + H_{aJ} - \tilde{A}_a H_{JJ} \right) \pi_J = \tilde{A}_a g_J - g_{\tilde{a}} + b_a(\tilde{p}) - \sum_{\hat{a} \neq J} (H_{a\hat{a}} - \tilde{A}_a H_{J\hat{a}}) b_a(p)$$

so that

$$C = \sum_{\hat{a} \neq J} (H_{a\hat{a}} - \tilde{A}_a H_{J\hat{a}}) A_{\hat{a}} + H_{aJ} - \tilde{A}_a H_{JJ}$$

and

$$c_0 = \tilde{A}_a g_J - g_{\tilde{a}} - \sum_{\hat{a} \neq J} (H_{a\hat{a}} - \tilde{A}_a H_{J\hat{a}}) b_a(p).$$

But (46) states that every row of C is a linear combination of the rows of Q . The coefficients in this linear combination form the matrix M . For the converse, take a row of C ; from 3, it is a linear combination of the rows of Q and thus belongs to $\text{row}(Q)$. Therefore, since all the rows of C belong to $\text{row}(Q)$, so does their linear span. ■

A.3.4 Proof of Corollary 3

Take $C = MQ$ or $CQ' = MQQ'$, or $CQ'(QQ')^{-1} = M$. Hence, $C = CQ'(QQ')^{-1}Q$, or

$$C \left(I - Q' (QQ')^{-1} Q \right) = 0.$$

For the converse, from (29) we have $C = CQ'(QQ')^{-1}Q$. Take $M = CQ'(QQ')^{-1}$ to obtain $C = MQ$. Finally, note that $\text{rank}(C) = X - \dim(N(C)) \leq X - \dim(N(Q)) = \text{rank}(Q) = d$.

A.4 Identification of Counterfactual Welfare

A.4.1 Proof of Proposition 11

Proposition 11 is a direct consequence of Lemma 5 below.

Lemma 5. *Assume $\tilde{\mathcal{A}} = \mathcal{A}$, $\tilde{\mathcal{X}} = \mathcal{X}$, $\tilde{\beta} = \beta$, and let $h_a(\pi_a) = H_a\pi_a + g_a$, all a . Let $C = H_{-J}A_{-J} - \tilde{A}_{-J}H_J$ and $D = (I - \beta\tilde{F}_J)(I - \beta F_J)^{-1} - H_J$. Then ΔV is identified if and only if*

$$\tilde{P} \left[C - \tilde{\mathbf{A}}D \right] = D. \quad (47)$$

where the matrices $\tilde{\mathbf{A}}$ and \tilde{P} are defined as in Lemma 1 (but based on \tilde{A}_a and \tilde{p}). Furthermore, the following holds:

1. If $C = 0$, $D = 0$, then ΔV is identified.
2. If $C = 0$, $D \neq 0$, then ΔV is not identified.
3. If $C \neq 0$, $D = 0$, then, provided there is no combination of actions a and j such that $\tilde{p}_a(x) = \tilde{p}_j(x)$ for some state x , ΔV is not identified.
4. If $C \neq 0$, $D \neq 0$, then ΔV is generically not identified.

Proof. We know that

$$V = (I - \beta F_J)^{-1} (\pi_J + \psi_J(p))$$

and similarly for \tilde{V}

$$\tilde{V} = \left(I - \beta \tilde{F}_J \right)^{-1} (h_J(\pi_J) + \psi_J(\tilde{p})).$$

Then,

$$\frac{\partial \Delta V}{\partial \pi_J} = \left(I - \beta \tilde{F}_J \right)^{-1} \left(H_J + \frac{\partial \psi_J(\tilde{p})}{\partial \tilde{p}} \frac{\partial \tilde{p}}{\partial \pi_J} \right) - (I - \beta F_J)^{-1}.$$

Therefore, $\frac{\partial \Delta V}{\partial \pi_J} = 0$ if and only if

$$\frac{\partial \psi_J(\tilde{p})}{\partial \tilde{p}} \frac{\partial \tilde{p}}{\partial \pi_J} = D \quad (48)$$

From Lemma 2, we know that

$$\frac{\partial \psi_J}{\partial \tilde{p}} = \tilde{P} \tilde{\Phi},$$

where \tilde{P} and $\tilde{\Phi}$ are the counterfactual counterpart of P and Φ defined in Lemma 2. By the Implicit Function Theorem, we know that

$$\frac{\partial \tilde{p}}{\partial \pi_J} = \left[\frac{\partial \tilde{b}_{-J}(\tilde{p})}{\partial \tilde{p}} \right]^{-1} \left(H_{-J} A_{-J} - \tilde{A}_{-J} H_J \right).$$

and, by Lemma 1,

$$\left[\frac{\partial \tilde{b}_{-J}(\tilde{p})}{\partial \tilde{p}} \right]^{-1} = \tilde{\Phi}^{-1} \left(\tilde{\mathbf{A}} \tilde{P} + I \right)^{-1},$$

Thus (48) becomes:

$$\tilde{P} \left(\tilde{\mathbf{A}} \tilde{P} + I \right)^{-1} C = D \quad (49)$$

Note that⁴⁷

$$\left(\tilde{\mathbf{A}} \tilde{P} + I \right)^{-1} = I - \tilde{\mathbf{A}} \left(I + \tilde{P} \tilde{\mathbf{A}} \right)^{-1} \tilde{P}.$$

Define $M = (I + \tilde{P} \tilde{\mathbf{A}})$. Then,

$$\tilde{P} \left(\tilde{\mathbf{A}} \tilde{P} + I \right)^{-1} = \tilde{P} - \tilde{P} \tilde{\mathbf{A}} M^{-1} \tilde{P} = \tilde{P} - (M - I) M^{-1} \tilde{P} = M^{-1} \tilde{P}$$

Then, (49) becomes $M^{-1} \tilde{P} C = D$, or $\tilde{P} C = M D = (I + \tilde{P} \tilde{\mathbf{A}}) D$, or $\tilde{P} (C - \tilde{\mathbf{A}} D) = D$, which is (47).

Next, we examine the four mutually exclusive possibilities. Statement 1 is obvious. Statement 2 follows from (49). Statement 3 amounts to $\tilde{P} C = 0$. Recall \tilde{P} is $X \times (A - 1) X$ and C is $(A - 1) X \times X$. If there are no two actions a and j such that $\tilde{p}_a(x) = \tilde{p}_j(x)$ for some x , then \tilde{P} is full rank (recall $\tilde{p}_a(x) > 0$ for all a, x , by the large support assumption on $G(\varepsilon)$). This implies the matrix $(\tilde{P}' \tilde{P})^{-1}$ exists, and so $C = 0$. But this is impossible since $C \neq 0$ by assumption.

Finally, for statement 4, suppose $C \neq 0$, $D \neq 0$. Recall that \tilde{p} as a function of π_J is the unique solution of

$$\tilde{b}_{-J}(\tilde{p}) = C \pi_J + H_J b(p)$$

⁴⁷The equality makes use of the identity $(I - BA)^{-1} = I + B(I - AB)^{-1}A$.

Table 6: Data Sources

Dataset	Description	Source
Cropland Data Layer	Land cover	http://nassgeodata.gmu.edu/CropScape/
DataQuick	Real estate transactions, assessments	DataQuick
US Counties	County boundaries	http://www.census.gov/cgi-bin/geo/shapefiles2010/layers.cgi
GAEZ Database	Protected land, soil type	http://www.gaez.iiasa.ac.at/
SRTM	Topographical – altitude and slope	http://dds.cr.usgs.gov/srtm/
NASS Quick Stats	Yields, prices, pasture rental rates	http://www.nass.usda.gov/Quick_Stats/
ERS	Operating costs	http://www.ers.usda.gov/data-products/commodity-costs-and-returns.aspx
NOAA Urban Centers	Urban center locations and populations	% http://www.nws.noaa.gov/geodata/catalog/national/html/urban.htm

The solution set is X -dimensional. Thus, (47) contains X^2 equations and X unknowns (π_J). By Sard’s Theorem, the system has generically no solution. ■

A.4.2 Proof of Corollary 4

Lack of identification of θ_0 is represented by the free parameter θ_0^J . So, applying the same argument as in Lemma 5, but differentiating ΔV with respect to θ_0^J , we prove the claim.

B Appendix: Data (for online publication)

Table 6 lists our data sources. All are publicly available for download save DataQuick’s land values. Our main sample is based on a sub-grid of the the Cropland Data Layer (CDL), a high-resolution (30-56m) annual land-use data that covers the entire contiguous United States since 2008. We took a 840m sub-grid of the CDL within those counties appearing in our DataQuick database.⁴⁸

DataQuick collects transaction data from about 85% of US counties and reports the associated price, acreage, parties involved, field address and other characteristics. The coordinates of the centroids of transacted parcels in the DataQuick database are known. To assign transacted parcels a land use, we associate a parcel with the nearest point in the CDL grid.

A total of 91,198 farms were transacted between 2008 to 2013 based on DataQuick. However, we dropped non-standard transactions and outliers from the data. First, because we are interested in the agricultural value of land (not residential value), we only consider transactions of parcels for which the municipal assessment assigned zero value to buildings and structures. Additionally, we drop transactions featuring multi-parcels, transactions between family members, properties held in trust, and properties owned by companies. Finally, we drop transactions with extreme prices:

⁴⁸The 840m grid scale was chosen for two reasons. First, it provides comprehensive coverage (i.e., most large agricultural fields are sampled) without providing too many repeated points within any given parcel. Second, the CDL data changed from a 56m to a 30m grid, and the 840 grid size allows us to match points across years where the grid size changed while matching centers of pixels within 1m of each other. The CDL features crop-level land cover information. See Scott (2013) for how “crops” and “non-crops” are defined.

Table 7: Summary Statistics

Statistics	Mean	Std Dev	Min	Max
In Cropland	0.147	0.354	0	1
Switch to Crops	0.0162	0.126	0	1
Keep Crops	0.849	0.358	0	1
Crop Returns (\$)	228	112	43	701
Slope (grade)	0.049	0.063	0	0.702
Altitude (m)	371	497	-6	3514
Distance to Urban Center (km)	79.8	63.7	1.22	362
Nearest commercial land value (\$/acre)	159000	792000	738	73369656
Land value (\$/acre)	7940	9720	5.23	50000

A slope of 1 refers to a perfect incline and a slope of 0 refers to perfectly horizontal land.

those with value per acre greater than \$50,000, total transaction price greater than \$10,000,000, or total transaction price less than \$60; these are considered measurement error. After applying the selection criteria, there remained 24,643 observations (transactions).

To obtain a rich set of field characteristics, we use soil categories from the Global Agro-Ecological Zones database and information on protected land from the World Database on Protected Areas. Protected land was dropped from all analyses. The NASA’s Shuttle Radar Topography Mission (SRTM) database provides detailed topographical information. The raw data consist of high-resolution (approx. 30m) altitudes, from which we calculated slope and aspect, all important determinants of how land is used. Characteristics such as slopes and soil categories are assigned to fields/parcels using nearest neighbor interpolation.

To derive a measure of nearby developed property values, we find the five restaurants nearest to a field, and we average their appraised property values. For each field, we also compute the distance to the nearest urban center with a population of at least 100,000. Location of urban centers were obtained from the National Oceanic and Atmospheric Administration (NOAA).

Finally, we use various public databases on agricultural production and costs from the USDA. The final dataset goes from 2010 to 2013 for 515 counties and from 2008 to 2013 for 132 counties. Crop returns are based on information on yields, prices received, and operating expenditures; non-crop returns are based on more sparse information on pasture land rental rates (see Scott (2013)).

Table 7 presents some summary statistics. Table 8 compares the transacted fields (in DataQuick) to all US fields (in the CDL). Overall, the two sets of fields look similar. In particular, the probability of keeping (switching to) crops is very similar across the two datasets.

Table 8: Dataquick vs CDL Data – Time Invariant Characteristics

Mean by dataset	DataQuick	CDL
In Cropland	0.147	0.136
Switch to Crops	0.0162	0.0123
Keep Crops	0.849	0.824
Crop Returns (\$)	228	241
Slope (grade)	0.049	0.078
Altitude (m)	371	688
Distance to Urban Center (km)	79.8	103
Nearest commercial land value (\$/acre)	159000	168000

C Appendix: Estimation (for online publication)

C.1 Model with Unobserved States

Allowing for unobserved component to returns, the per period payoff becomes:

$$\pi(a, k_{imt}, w_{mt}, s_{im}, \varepsilon_{imt}) = \theta_0(a, k_{imt}, s_{im}) + \theta_1 Z(a, w_{mt}) + \xi(a, k_{imt}, w_{mt}, s_{im}) + \varepsilon_{imt} \quad (50)$$

where $\xi(a, k, w, s)$ captures unobservable variation in returns, and the idiosyncratic shock ε_{it} has a logistic distribution. We construct returns $Z_{mt}^a \equiv Z(a, w_{mt})$ using county-year information (expected prices and realized yields for major US crops, as well as USDA cost estimates) as in Scott (2013).⁴⁹ As described below, identification requires exclusion restrictions on $\xi(a, k, w, s)$ (see also Kalouptsi, Scott, and Souza-Rodrigues (2017b)).

C.2 Payoff Parameter Estimation

Throughout this section, we use t -subscripts in place of explicitly writing the aggregate state variable w_{mt} . We also omit the subscripts i (fields) and m (counties) to simplify. The derivation relies on two crucial assumptions: (a) agents are small; i.e., changing the action of any agent at time t does not affect the distribution of w_{t+1} , and (b) agents have rational expectations.

Here, we consider two estimators for the payoff function. Let $p_t^c(k, s)$ denote the probability of choosing action “crops” at time period t given state k for a field of type s , and let σ be a scale parameter (which we discuss below). We begin with Scott’s (2013) derivation of a linear estimating equation for a dynamic model with logit errors:

$$Y_t(k, s) = \tilde{\theta}_0(k, s) + \theta_1 (Z_t(c, s) - Z_t(nc, s)) + \tilde{\xi}_{k,s,t} + \tilde{\varepsilon}_{k,s,t} \quad (51)$$

⁴⁹We refer the interested reader to Scott (2013) for details of constructing the measure of observed returns Z . Due to data limitations, we restrict Z to depend only on (a, w_{mt}) . One important difference from Scott (2013) is that we have field level observable characteristics s_{im} and they affect land use switching costs.

where

$$\begin{aligned}
Y_t(k, s) &\equiv \ln\left(\frac{p_t^c(k, s)}{1-p_t^c(k, s)}\right) + \beta \ln\left(\frac{p_{t+1}^c(0, s)}{p_{t+1}^c(k'(nc, k), s)}\right) \\
\tilde{\theta}_0(k, s) &\equiv (\theta_0(c, k, s) - \theta_0(nc, k, s)) / \sigma \\
&\quad + \beta (\theta_0(c, 0, s) - \theta_0(c, k'(nc, k), s)) / \sigma \\
\theta_1 &\equiv 1 / \sigma \\
\tilde{\xi}_{k, s, t} &\equiv \xi_t(c, k, s) - \xi_t(nc, k, s) \\
&\quad + \beta (\xi_{t+1}(c, 0, s) - \xi_{t+1}(c, k'(nc, k), s)) \\
\tilde{e}_{k, s, t} &\equiv \beta (E_t[V_t(0, s)] - V_t(0, s)) \\
&\quad - \beta (E_t[V_{t+1}(k'(nc, k), s)] - V_{t+1}(k'(nc, k), s)).
\end{aligned}$$

Ultimately, this is a linear equation that can be used to estimate the parameters of the payoff function with no need to solve the agent's dynamic optimization problem.

On the left hand side of equation (51), we have a dependent variable which is a function of conditional choice probabilities (which are estimated in a first stage, described below in section C.3) and the discount factor (which is imputed; we assume it equals 0.95).

On the right hand side of (51), the intercept term $\tilde{\theta}_0$ is a combination of intercepts of the payoff function θ_0 . We discuss the identification of θ_0 in more detail below, for this is essentially where the two estimators differ.

The error term has two components, $\tilde{\xi}$ and \tilde{e} . The term $\tilde{\xi}$ is a function of ξ , representing unobservable variation in returns, while \tilde{e} is a function of expectational error terms. Because Z and ξ may be correlated, we follow Scott (2013) and implement an instrumental variable estimator. To do so, we need exclusion restrictions of the form

$$E\left[\nu_{k, s, t} \left(\tilde{\xi}_{k, s, t} + \tilde{e}_{k, s, t}\right)\right] = 0, \quad (52)$$

where $\nu_{k, s, t}$ is a vector of instrumental variables. Given that agents have rational expectations, $\tilde{e}_{k, s, t}$ is uncorrelated with any function of variables in the time- t information set by construction. For this reason, $E[\nu_{k, s, t} \tilde{e}_{k, s, t}] = 0$ holds for any $\nu_{k, s, t}$ in the time- t information set and the question of whether equation (52) is valid becomes a question of whether $E[\nu_{k, s, t} \tilde{\xi}_{k, s, t}] = 0$. Such a restriction is a substantive assumption as exclusion restrictions for instrumental variables typically are.⁵⁰

We take first-differences for each field and field state, implicitly allowing for $\tilde{\theta}_0(k, s)$ to have fixed effects for s and k (interacted).⁵¹ After taking first differences, the instruments we use are:

⁵⁰If we were willing to assume that $E[(Z_{s, t}^c - Z_{s, t}^{nc}) \tilde{\xi}_{k, s, t}] = 0$, then we could estimate equation (51) using ordinary least squares.

⁵¹Note that we predict CCPs for each field state k , not just for the field state actually observed on the field, so we can take these first differences for each k regardless of the actual path of k for the field.

a constant term, caloric yields, and the lagged value of $Z_{s,t}^c - Z_{s,t}^{nc}$.⁵² The moment restrictions are used to estimate θ_1 . We then form estimates of $\tilde{\theta}_0(k, s)$ by averaging over the residuals for each (k, s) pair.

Up to this point, our two estimators coincide; i.e., our two estimators agree on the estimates of θ_1 and $\tilde{\theta}_0(k, s)$. The estimators differ when it comes from mapping the estimates of $\tilde{\theta}_0(k, s)$ to estimates of $\theta_0(\cdot, k, s)$. Notice that for each type s , equation (51) includes one intercept parameter $\tilde{\theta}_0(k, s)$ for each field state k . However, the original payoff function involves two intercept parameters ($\theta_0(c, k, s)$ and $\theta_0(nc, k, s)$) for each (s, k) combination. Hence, the need for restrictions for the identification of the model (and our claim in Section 4.1 that θ_0 is not identified without restrictions).

Our first estimator (the CCP estimator) imposes the following restrictions on θ_0 :

$$\forall k, s : \quad \theta_0(nc, k, s) = 0. \quad (53)$$

After imposing (53), we can solve for $\theta_0(c, k, s)$ from our $\tilde{\theta}_0(k, s)$ estimates, recalling that

$$\tilde{\theta}_0(k, s) \equiv (\theta_0(c, k, s) - \theta_0(nc, k, s)) + \beta(\theta_0(c, 0, s) - \theta_0(c, k'(nc, k), s)), \quad (54)$$

noting that equations (53) and (54) represent six linearly independent equations in six unknowns for each (k, s) pair.

Our second estimator (the V-CCP estimator) does not impose equation (53), and instead using additional information in resale prices in place of a restriction like (53). In order to relate observed resale prices to farmer's payoff and value functions, we need a model of transaction prices. We assume that resale prices measure farmer's ex-ante value functions; i.e.,

$$\ln p_t^{RS} = \ln(\tilde{V}_t(k, s)) + \eta_t, \quad (55)$$

where p_t^{RS} is the resale price of a field, η_t is measurement error, and we will explain the reason for the tilde on the value function below. Using resale prices as signals of the value function can be justified by assuming that there is a competitive market for buying farms – see Kalouptside (2014) for further discussion of this assumption in the context of bulk shipping.

We estimate a flexible model of how resale prices depend on (k, s, t) , much like Kalouptside (2014) (see section C.4 for details about the implementation). Fitted values from this regression can be used as estimates of the value function, but an important caveat is that we must consider the scale of the utility function when interpreting the estimates. In econometric discrete choice models, we typically impose a scale normalization on the model that sets the variance of the idiosyncratic shocks equal to a convenient number (e.g., unity for a probit model of $\pi^2/6$ for a logit model). In our parametric land use model, the coefficient on returns, θ_1 , reflects this normalization: the parameter

⁵²See Scott (2013) for the measurement of caloric yields.

θ_1 can be understood as the scalar we need to multiply by to convert the units from dollars to utils. When we estimate a hedonic model of the value function, the value function is measured in dollars. Therefore, to convert from the estimated value function to the scale-normalized value function we should multiply by θ_1 :

$$V_t(k, s) = \theta_1 \tilde{V}_t(k_{it}, s_{it}).$$

A relationship between value functions and the payoff function can be derived as follows:

$$\begin{aligned} V_t(k, s) &= v_t(c, k, s) + \psi_c(p_t^c(k, s)) \\ &= \pi_t(c, k, s) + \beta E_t[V_{t+1}(k'(c, k), s)] + \psi_c(p_t^c(k, s)) \\ &= \pi_t(c, k, s) + \beta V_{t+1}(k'(c, k), s) + e_{k,s,t} + \psi_c(p_t^c(k, s)) \end{aligned}$$

where

$$e_{k,s,t} \equiv \beta (E_t[V_t(k'(c, k), s)] - V(k'(c, k), s)).$$

Ultimately, we can write the payoff function as a function of conditional choice probabilities (estimated in a first stage), value functions (estimated using retail prices in a first stage), and an expectational error term (mean zero):

$$\pi_t(c, k, s) = V_t(k, s) - \beta V_{t+1}(k'(c, k), s) - \psi_c(p_t^c(k, s)) - e_{k,s,t}. \quad (56)$$

Recalling that the measured version of the value function needs to be converted from dollars to utils to be on the same scale as the normalized payoff function, we have

$$\pi_t(c, k, s) = \theta_1 \left(\tilde{V}_t(k, s) - \beta \tilde{V}_{t+1}(k'(c, k), s) \right) - \psi_c(p_t^c(k, s)) - e_{k,s,t}. \quad (57)$$

Noting that an estimate of θ_1 can be obtained from the CCP estimator, we can then obtain estimates of payoffs using equation (57), simply by plugging in the estimated values of θ_1 , \tilde{V} and p .⁵³ More to the point, we can obtain estimates of the intercept parameters:

$$\theta_0(c, k, s) = -\theta_1 Z_t(c, s) + \theta_1 \left(\tilde{V}_t(k, s) - \beta \tilde{V}_{t+1}(k'(c, k), s) \right) - \psi_c(p_t^c(k, s)) - e_{k,s,t}. \quad (58)$$

The V-CCP estimator uses equation (58) to estimate $\theta_0(c, k, s)$ by averaging the right-hand-side of (58) over time. Finally, the estimates of $\theta_0(nc, k, s)$ are then recovered from equation (54).

Note that we could alternatively estimate $\theta_0(nc, k, s)$ from an equation like (58), but using non-

⁵³Recall that estimating θ_1 with the CCP estimator does not require any identifying restrictions on θ_0 . Consider equation (51), a regression equation that allows us to estimate θ_1 and $\hat{\theta}_0$. The identifying restrictions are only needed if we want to map from $\hat{\theta}_0$ to θ_0 .

crops as the action instead of crops. Thus, we have over-identifying restrictions. As the primary reason we consider the V-CCP estimator is to replace the *a priori* identifying restrictions in the CCP estimator with a more data-driven approach, we only take as much information as we need from the resale prices to fully identify the payoff function. If we were to use more information from the resale prices, then the two estimators might not agree on the value of $\tilde{\theta}_0(k, s)$, an object that is identified from CCP data without imposing identifying restrictions. Our two estimators only differ when it comes to parameters that cannot be identified from CCP data without restrictions. Thus, by comparing these two estimators, we isolate the impact of identifying restrictions.

C.3 Conditional Choice Probabilities

We estimate conditional choice probabilities using a semiparametric logit model. The model is fully flexible over field states and year, but smooth across counties. In particular, we maximize the following log likelihood function:

$$\max_{\theta_{ckt}} \sum_{m' \in S_m} \sum_{i \in I_{m'}} w_{m,m'} I[k_{imt} = k] \left\{ \begin{array}{l} I[a_{imt} = c] \log(p_{mt}(c, k, s_{im}; \theta_{ckt})) \\ + I[a_{imt} = nc] \log(1 - p_{mt}(c, k, s_{im}; \theta_{ckt})) \end{array} \right\}$$

where S_m is the set of counties in the same US state as m , I_m is the set of fields in county m , $w_{m,m'}$ is the inverse squared distance between counties m and m' , and $I[\cdot]$ is the indicator function. The conditional choice probability is parameterized as follows:

$$p_{mt}(c, k, s_{im}; \theta_{ckt}) = \frac{\exp(s'_{im} \theta_{ckt})}{1 + \exp(s'_{im} \theta_{ckt})}.$$

Note that without fields' observable characteristics, this regression would amount to taking frequency estimates for each county, field state, and year, with some smoothing across counties. Including covariates allows for within-county field heterogeneity. The final specification for the conditional choice probabilities only uses $slope_{im}$ among regressors because it proved to be the most powerful predictor of agricultural land use decisions (after conditioning on county and field state).

The set of counties in S_m only includes counties which also appear in the DataQuick database. For some states, the database includes a small number of counties, so in these cases we group two or three states together. For example, only one county in North Dakota appears in our sample, and it is a county on the eastern border of North Dakota, so we combine North Dakota and Minnesota. Thus, for each county m in North Dakota or Minnesota, S_m represents all counties in both states in our sample.⁵⁴

⁵⁴In particular, we form a number of groups for such cases: Delaware and Maryland; North Dakota and Minnesota; Idaho and Montana; Arkansas, Louisiana, and Mississippi; Kentucky and Ohio; Illinois, Indiana, and Wisconsin; Nebraska and Iowa; Oregon and Washington; Colorado and Wyoming.

For the sake of precision, rather than only estimating CCPs using the CDL sample that was merged with resale data, we used the full 840m sub-grid of fields from the CDL (848,384 fields) for the CCP estimation. We then predicted CCPs and estimated payoff functions using fields that were merged with the resale data.

C.4 Resale Price Regression

Next, we discuss how we estimate the value function from resale prices. We view that our resale market assumptions are not overly restrictive in the context of rural land which features a large number of small agents. The land resale market is arguably thick, with a large number of transactions taking place every year.⁵⁵ Moreover, we are able to control for a rich set of field characteristics. Finally, we did not find evidence of selection on land use changes upon resale, as discussed below.

As our transaction data is much more sparse than our choice data, we adopt a more restrictive (parametric) form for modeling land values. We estimate the following regression equation:

$$\ln p_{it}^{RS} = X_{it}'\theta_V + \eta_{it}, \quad (59)$$

where p_{it}^{RS} is a transaction price (in dollars per acre), and X_{it} is a vector of characteristics for the corresponding field. The covariates X_{it} include all variables in Table 7 (i.e. k , slope, altitude, distance to urban centers, nearby commercial values). They also include year dummies, returns interacted with year dummies, field state dummies interacted with year dummies, and county dummies.

Table 9 presents the estimated coefficients. Although not shown in the table, the estimated coefficients of k are significant and have the expected signs (the large number of interactions makes it difficult to add them all in the table). This is important for the second stage estimation, as k is the main state variable included in the switching cost parameters $\theta_0(a, k)$.

Note that, because field acreage is available only in the DataQuick dataset, when merging with the CDL and remaining datasets we lose this information. This implies, for example, that acreage cannot be a covariate in the choice probabilities. For this reason, we choose a specification for the value function that regresses price per acre on covariates. The value of our R^2 in our regression is a direct consequence of this fact. When we use total land prices as the dependent variable and include acres on the covariates we obtain R^2 as high as 0.8.

Finally, we briefly discuss the possibility of selection on transacted fields. As shown previously in Table 8 of Appendix B, the characteristics of the transacted fields (in DataQuick) look similar to all US fields (in the CDL). Furthermore, we investigate whether land use changes upon resale.

⁵⁵Comparing DataQuick with the CDL data we see that 1.4-2% of fields are resold every year. Moreover, the USDA reports that in Wisconsin there are approximately 100 thousand acres transacted every year (about 1000 transactions) out of 14.5 million acres of farmland (seemingly information on other states is not available).

Table 9: Land Resale Price Regression

VARIABLES	(1) log(land value)
log(distance to urban center)	-0.471*** (0.0297)
commercial land value	0.102*** (0.00930)
slope	-1.654*** (0.160)
alt	-0.000226** (9.00e-05)
Observations	24,643
R-squared	0.318

Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Omitted: soil, county, year, and field state dummies

as well as interactions with returns.

Using a linear probability model we find no such evidence (see Table 10). We regress the land use decision on dummy variables for whether the field was sold in the current, previous, or following year as well as various control variables. In regressions within each cross section, ten of the eleven coefficients on the land transaction dummy variables are statistically insignificant, and the estimated effect on the probability of crops is always less than 1% (see Table 10). We have tried alternative specifications such as modifying the definition of the year to span the planting year rather than calendar year, and yet we have found no evidence indicating that there is an important connection between land transactions and land use decisions.

Table 10: Land use and transactions

VARIABLES	(1) incrops2010	(2) incrops2011	(3) incrops2012	(4) incrops2013
soldin2009	0.000647 (0.00604)			
soldin2010	0.000116 (0.00326)	0.00364 (0.00334)		
soldin2011	-0.00117 (0.00316)	0.000629 (0.00324)	-0.00159 (0.00330)	
soldin2012		-0.000620 (0.00306)	-0.00472 (0.00313)	0.00411 (0.00265)
soldin2013			-0.00962*** (0.00306)	-0.000445 (0.00256)
Observations	23,492	23,492	23,492	23,492
R-squared	0.666	0.698	0.717	0.757

Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Linear probability model. Omitted covariates include current returns, field state, US state, slope, local commercial land value, distance to nearest urban center, and interactions.