On the Non-Identification of Counterfactuals in Dynamic Discrete Games

Myrto Kalouptsidi,[*] Paul T. Scott,[†] Eduardo Souza-Rodrigues[‡§]
Princeton University and NBER, Toulouse School of Economics, University of Toronto

January 2016

Abstract

In single-agent dynamic discrete choice models, counterfactual behavior is identified for some (but not all) counterfactuals despite the fact that the models themselves are generally under-identified. We review recent results on the identification of counterfactuals in dynamic discrete choice settings. When it comes to dynamic discrete games, we conjecture that counterfactuals are generally not identified, even when analogous counterfactuals of single-agent models are identified. Using the example of a duopoly entry game, we explain why strategic considerations undermine the identification of counterfactual equilibria in dynamic games.

KEYWORDS: Identification, Dynamic Discrete Choice, Dynamic Games, Counterfactual

1 Introduction

The main reason applied researchers estimate structural econometric models is to perform counterfactuals, such as hypothetical, or factual, policy interventions to the environment under study. Despite a vast literature on the identification of structural models, we know very little about the identification of counterfactuals associated with them. This issue is particularly important when the employed model is not identified, as in the case of dynamic discrete choice models (DDC). DDC models have proven useful to analyze public policies in a variety of contexts (e.g. labor markets, firm dynamics, health choices). In this paper, we review recent results exploring when counterfactual behaviour of dynamic discrete choice models are identified. While such results have been developed only for the case of single-agent models, they have strong implications for the identification of counterfactuals of dynamic games. We then present one such implication in the context of a simple duopoly dynamic entry game.

? and **?** showed that DDC models, under very general assumptions, are nonparametrically not identified; one source of under-identification is that there are always many different utility functions which can rationalize observed choice behavior. Given this fact, **?** claim that "the entire dynamic discrete choice project thus appears to be without empirical content, and the evidence from it at the whim of investigator choices about functional forms of estimating equations and application of ad hoc exclusion restrictions."

Whimsically imposed or not, the restrictions that allow researchers to estimate a DDC model might not matter for the counterfactuals the researchers are ultimately interested in. Counterfactuals typically involve a change in utility functions, in the process governing state transitions, and/or in the set of actions and states available to agents. If every utility function consistent with the observed data generates the

*Department of Economics, Princeton University, 315 Fisher Hall, Princeton, NJ 08544-1021, myrto@princeton.edu
†Toulouse School of Economics, 21 allée de Brienne, 31015 Toulouse, France, paul.scott@tse-fr.eu
‡Department of Economics, University of Toronto, Max Gluskin House, 150 St. George St., Toronto, Ontario M5S 3G7, Canada, e.souzarodrigues@utoronto.ca
§This paper is based on existing work in progress, part of which was presented at EARIE 2015 Munich and we thank the organizer, Michelle Sovinsky, for inviting Myrto Kalouptsidi. We are also grateful to Charis Katsiardis for excellent research assistance.

same behavioral response to a given counterfactual, then that counterfactual can be said to be identified. In other words, identified counterfactuals are not sensitive to restrictions that are necessary to identify the model. Recent work has shown that some (but not all) counterfactuals of single-agent dynamic models are indeed identified. **?**, **?**, **?**, and **?** give examples of counterfactuals which are identified and examples which are not. **?** offer a full characterization of a broad class of counterfactuals.

After reviewing some of these recent results, we then investigate how previous results extend to a dynamic entry game. We consider a duopoly game and the firms' counterfactual response to a change in entry costs. While the change in question produces an identified response in a single-agent context (or if the behavior of one of the firms was held fixed), strategic considerations prevent the identification of the firms' responses. In other words, the counterfactual equilibrium is sensitive to necessary identifying restrictions imposed on the game's payoff function. Lack of identification is not the result of multiplicity of equilibria. Non-identification occurs even under a simple and correctly specified selection rule. Given this result, we conjecture that strategic considerations imply that all (practically relevant) counterfactuals of dynamic discrete games will fall into a class of counterfactuals that are not identified.

The rest of the paper is organized as follows. Section 2 provides an overview of the literature on the identification of payoffs and counterfactuals in dynamic discrete choice models. Section 3 considers the example of a duopoly entry game.

2 Identification in dynamic discrete single-agent models

In this seciton, we review the standard dynamic discrete choice modeling framework and summarize some recent results on the identification of counterfactuals in single-agent models.

In the standard dynamic discrete choice framework, an agent agent i chooses one action a_{it} from the finite set $\mathcal{A} = \{1, ..., A\}$ in each period $t \in \{1, 2, ...\}$. The current 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 $x_{it} \in \mathcal{X} = \{x_1, ..., X\}$, $X < \infty$; while $\varepsilon_{it} = (\varepsilon_{it}(1), ..., \varepsilon_{it}(A))$ is i.i.d. across agents and time and has joint distribution G.

A standard assumption on how state variables evolve is conditional independence, meaning the idiosyncratic error term ε does not affect the state variable x except through actions. Formaly, the conditional independence means that the transition distribution function for $(x_{it}, \varepsilon_{it})$ factors as

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

and the current payoff function is given by

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

Agent i chooses a sequence of actions to maximize the expected discounted payoff:

$$V(x_{it}, \varepsilon_{it}) = E\left(\sum_{\tau=0}^{\infty} \beta^{\tau} u(a_{it+\tau}, x_{it+\tau}, \varepsilon_{it+\tau}) | x_{it}, \varepsilon_{it}\right)$$

where $\beta \in (0, 1)$ is the discount factor.

Following the literature, we define the *ex ante value function*, which represents the expectation of the value function before idiosyncratic shocks are realized:

$$V(x_{it}) \equiv \int V(x_{it}, \varepsilon_{it}) \, dG(\varepsilon_{it}),$$

and the *conditional value function*, which represents the expected discounted payoffs conditional on particular action before the realization of idiosycratic shocks:

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

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

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

where $1\{\cdot\}$ is the indicator function.

Researchers typically have data on actions and states. From such data, it is straightforward (in principle) to construct estimates of CCPs – with sufficiently rich data, simple frequency estimates will do. Studies on identification typically take for granted that CCPs are known and then consider what can be inferred about the model primitives $(u, \beta, F, G, V,)$ or endogenous functions of the model primitives, such as counterfactual CCPs.

Identification results clarify the role assumptions play in recovering primitives of interest. As previously noted, **?** and **?** showed that payoffs in single-agent dynamic discrete choice models are nonparametrically not identified. In particular, **?** characterized the set of restrictions one needs to impose: suppose we have J actions and X states and we are solving for the $J \times X$ payoff elements. This system has infinitely many solutions so that the model is not identified; indeed, since the researcher only has $(J-1) \times X$ linearly independent CCP estimates, the dimension of the set of solutions is given by the cardinality of the state space X. As shown in **?**, one possible way to represent the set of solutions is to take the payoff vector of one action as the vector of free parameters (u_J), and write the payoffs for other actions (u_J) as linear functions of it:

$$u_a = A_a u_J + b_a(p) \tag{1}$$

where $a, J \in A$ are available actions , $u_a \in \mathbb{R}^X$ is the (nonparametric) payoff of action a at all states $x = 1, ..., X$, and

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

$$b_a(p) = A_a \psi_J(p) - \psi_a(p),$$

where $F_a \in \mathbb{R}^{X \times X}$ is the transition matrix conditional on action a, and $\psi_a(p)$ is a known function of observed choice probabilities.[1] Notice that the right-hand side of equation (1) consist of the transition matrix F, which can be estimated in a first stage, the discount factor β, which is typically prespecified, and a known function of CCPs, which can also be estimate in a first-stage. Thus, equation (1) shows how to estimate the full payoff function after fixing the payoffs for a reference action J.

To reiterate, in order to obtain point identification of payoff functions, one needs to add extra restrictions. This has been the typical solution adopted in practice. Common extra assumptions include (combinations of) parametric functional forms, exclusion restrictions, and "normalizations." Parametric assumptions reduce the number of parameters to be identified and impose some shape restrictions; e.g., payoffs that are linear in observable states. Exclusion restrictions assume that some payoffs do not depend on all state variables; e.g., firm entry and exit costs are often assumed state invariant. "Normalizations" fix the payoff of some action in some states at some known value, as described in equation (1). Typically, the payoff of the action we have least information about is set to zero; usually referred to as the "outside option." Note that fixing u_J equal to some pre-specified level is different from imposing a normalization in the traditional sense. For instance, if we set u_J equal to a zero vector, we implicitly assume that the payoffs to action J do not depend on the state – that is a substantive assumption, and not something that can be arrived at through a positive transformation of the payoff function.

Nonidentification of payoffs seemingly poses serious challenges for counterfactual analysis. If different restrictions imposed on the model can lead to different payoff functions which are both equally consistent with the observed data, then it might seem that the output of dynamic discrete choice models is necessarily sensitive to restrictions imposed by researchers. However, both of these models might agree in how the agents respond to a given policy intervention. If all models consistent with the observed data agree on the response to a given counterfactual change, then we can say that the counterfactual in question is identified.

Counterfactuals consist of transformations of model primitives, notably payoffs and transitions. A counterfactual that changes payoffs u to $\widetilde{u}$, is described by a known function, $h : R^{|A||X|} \to R^{|A||X|}$ (so that $\widetilde{\pi} = h(\pi)$). A counterfactual can also change transitions F to $\widetilde{F}$. **?** show that counterfactual choice probabilities are identified *if and only if* specific relationships hold between the transition matrix and counterfactual function. In other words, identification hinges crucially on the counterfactual performed

[1] For a detailed derivation, see **?**.

and the primitive transitions; this is intuitive, given that all dynamics are captured in state transitions conditional on today's chosen action.

?, **?**, **?**, and **?** give examples of counterfactuals which are identified and examples which are not. The following two results summarize two important common findings from these papers.

Result 1 *In a single-agent setting, if* $\widetilde{F} = F$ *and* $\widetilde{u} = u + \Delta$*, where* Δ *is a* $J \times X$ *matrix, then counterfactual choice probabilities* $\widetilde{p}$ *are identified. I.e., the behavioral response to a lump-sum transfer is identified.*

In interpreting Result 1, it is important to note that the change in utility Δ is very flexible in one way, but quite limited in another. It is flexible in the sense that it allows for changes in utility to be of any finite size and conditioned arbitrarily on actions and states. This seemingly allows for arbitrary changes in the payoff function, but Δ is not allowed to depend on u; thus, the researcher must be able to specify Δ before estimating the model.

To see why Result 1 holds, it is easiest to use Arcidiacono and Miller's (**?**) Lemma 1, which states that we can write the difference between the ex ante value function $V(x)$ and conditional value function $v_a(x)$ for an arbitrary action a as a function of conditional choice probabilities:

$$V = v_a + \psi_a(p) \tag{2}$$

where V, v_a, and $\psi_a(p)$ are X-dimensional vectors.

The next step is to write out the definition of the conditional value function in vector notation,

$$v_a = u_a + \beta F_a V,$$

and then use Arcidiacono and Miller's Lemma to substitute for the ex ante value function V:

$$v_a = u_a + \beta F_a (v_a + \psi_a(p)).$$

This allows us to express the conditional value in terms of the payoff function and transition matrix:

$$v_a = (I - \beta F_a)^{-1} (u_a + \beta F_a \psi_a(p)). \tag{3}$$

Next, we consider equation (3) for the reference action J as well as another arbitrary action a, but for the latter, we use equation (1) to express u_a in terms of u_J:

$$\begin{aligned} v_a &= (I - \beta F_a)^{-1} (A_a u_J + b_a(p) + \beta F_a \psi_a(p)), \\ v_J &= (I - \beta F_J)^{-1} (u_J + \beta F_J \psi_J(p)). \end{aligned} \tag{4}$$

Finally, note that differencing equation (2) allows us to write $v_a - v_J = \psi_J(p) - \psi_a(p)$. Given this, we can difference the equations (4) to write

$$\begin{aligned} \psi_J(p) - \psi_a(p) &= (I - \beta F_a)^{-1} (A_a u_J + b_a(p) + \beta F_a \psi_a(p)) \\ &\quad - (I - \beta F_J)^{-1} (u_J + \beta F_J \psi_J(p)). \end{aligned} \tag{5}$$

Equation (5) expresses a difference in continuation values in terms of only model primitives and functions of conditional choice probabilities. It is satisfied by the original conditional choice probabilities p, and it must also be satisfied by counterfactual conditional choice probabilities $\widetilde{p}$:

$$\begin{aligned} \psi_J(\widetilde{p}) - \psi_a(\widetilde{p}) &= \left(I - \beta \widetilde{F}_a\right)^{-1} \left(\widetilde{A}_a h_J(u) + \widetilde{b}_a(\widetilde{p}) + \beta \widetilde{F}_a \psi_a(\widetilde{p})\right) \\ &\quad - \left(I - \beta \widetilde{F}_J\right)^{-1} \left(h_J(u) + \beta \widetilde{F}_J \psi_J(\widetilde{p})\right) \end{aligned} \tag{6}$$

where $h_J(u) = \widetilde{u}_J$ is the counterfactual vector of utilities for the reference action.

Equation (6) allows us to see why Result 1 holds. The question here is whether all payoff functions u consistent with the observed data imply the same counterfactual CCPs $(\widetilde{p})$. In other words, if we plug different payoff functions u and u' into equation (6), where both u and u' are consistent with the observed CCP data, will the equation be satisfied for the same counterfactual CCPs $(\widetilde{p})$?

Result 2 *In a single-agent setting, if $\widetilde{F} \neq F$ and $\widetilde{u} = u$, then counterfactual choice probabilities $\widetilde{p}$ are not identified.*

(preliminary description on how Result 2 suggests cfs of dynamic games are not identified)

? offer a full characterization of a broad class of counterfactuals.

(MORE on how KSS generalize above results)

3 An entry game with non-identified counterfactual equilibria

In this section, we consider a simple example to help explain why counterfactuals of dynamic games are generally non-identified even when analogous counterfactuals of single-agent models are identified.

Consider a duopoly entry game with two players indexed by $i = A, B$. In each period t, the players simultaneously choose whether to be active $(a_{i,t} = 1)$ or not $(a_{i,t} = 0)$. Associated with each player is a state variable which equals the player's action in the previous period: $s_{i,t} = a_{i,t-1}$. The state of the game is simply the pair of states, $(s_{A,t}, s_{B,t})$. [2]

The game is symmetric, and player i receives payoffs which may depend on her own actions and state variable as well as her opponent's action and state variable. As a baseline case, we consider the following payoff function:

$$u_i(a_i, a_{-i}, s_i, s_{-i}) = \begin{cases} 0 & \text{if } a_i = 0, s_i = 0 \\ x & \text{if } a_i = 0, s_i = 1 \\ \pi_1 - c & \text{if } a_i = 1, s_i = 0, a_{-i} = 0 \\ \pi_1 & \text{if } a_i = 1, s_i = 1, a_{-i} = 0 \\ \pi_2 - c & \text{if } a_i = 1, s_i = 0, a_{-i} = 1 \\ \pi_2 & \text{if } a_i = 1, s_i = 1, a_{-i} = 1 \end{cases} \tag{7}$$

We can interpret x as a scrap value, c as an entry cost, π_1 as monopoly profits, ad π_2 as duopoly profits. The baseline parameterization is $x = .1$, $c = .2$, $\pi_1 = 1.2$, and $\pi_2 = -1.2$.

Note that the above payoff function represents a restrictive parameterization. In principle, payoffs when a player exits ($a_{i,t} = 0$ and $s_{i,t} = 1$) might depend on the other player's behavior and/or state variable. Table 1 makes this ore clear – in principle, there are sixteen different combinations of $(a_i, a_{-i}, s_i, s_{-i})$ and therefore the payoff function's parameter space is potentially sixteen-dimensional.

On the other hand, an equilibrium of the game will only involve up to eight linearly independent choice probabilities (one for each player and state). For simplicity, we consider only symmetric equilibria, and symmetric equilibria only involve four linearly independent choice probabilities. Thus, in a symmetric equilibrium we will only be able to identify four parameters, and we need restrictions on the payoff function in order to identify the *model*. However, the question this note focuses on is whether we need restrictions in order to identify *counterfactual behavior*. Perhaps any restrictions we might make in order to identify the model will lead to the same results when we simulate counterfactual behavior; i.e., perhaps some counterfactuals are identified even though the model is clearly not.

We refer to the model with the payoff function described above as Model 1, and Table 1 specifies the payoffs for each combination of actions and states. Table 2 describes a symmetric equilibrium for this model. In this equilibrium, each player enters with probability .576 when both players were active in the previous period. Each player remains active with probability .595 when both players competed in the previous period. When only one player was active in the previous period, the incumbent remains active with probability .842 and the other firm enters with probability .305.

[2] This game was introduced by **?**.

Table 1: Entry Game: Payoff Functions

Given the under-identification of the model, there are many other payoff functions that could rationalize this baseline equilibrium. One such alternative is based on the following parameterization, where we restrict all payoffs to be zero except when $(a_i, a_{-i}) = (1,1)$:

$$u_i(a_i, a_{-i}, s_i, s_{-i}) = \begin{cases} 0 & \text{if } a_i = 0 \text{ or } a_{-i} = 0 \\ \pi_{00} & \text{if } a_i = 1, a_{-i} = 1, s_i = 0, s_{-i} = 0 \\ \pi_{01} & \text{if } a_i = 1, a_{-i} = 1, s_i = 0, s_{-i} = 1 \\ \pi_{10} & \text{if } a_i = 1, a_{-i} = 1, s_i = 1, s_{-i} = 0 \\ \pi_{11} & \text{if } a_i = 1, a_{-i} = 1, s_i = 1, s_{-i} = 1 \end{cases} \tag{8}$$

Table 1 also describes parameter values for (8) which rationalize the baseline equilibrium.

While the baseline equilibrium is an equilibrium of both Models 1 and 2, it is not obvious *a priori* whether the equilibria of the two models will coincide when the payoff functions are changed. We consider a counterfactual in which entry costs are increased by .25 (in levels, not in proportional terms). Formally, we consider the following transformation of the payoff function for each of the two models:

$$\widetilde{u}(a_i, a_{-i}, s_i, s_{-i}) = \begin{cases} u(a_i, a_{-i}, s_i, s_{-i}) - .25 & \text{if } a_i = 1, s_i = 0 \\ u(a_i, a_{-i}, s_i, s_{-i}) & \text{otherwise} \end{cases}$$

The counterfactual payoffs $\widetilde{u}$ increase the costs of entry relative to the original payoffs u. In the parameterization expressed in equation (7) we could simply say that c increases by .25.

Before considering equilibria of the counterfactual games, we first consider how an individual player's best response would change if her payoff function changed from u to $\widetilde{u}$ and her opponent's behavior remained fixed in the baseline equilibrium strategy. These best responses are described by the "CF – fixed opponent" columns of Table 2. They are identical for the two models – i.e., the different ways of rationalizing the baseline equilibrium are equivalent when we consider these interim best responses. This should not be surprising. As Result 1 tells us, lump-sum counterfactuals of single-agent models are identified, and this exercise could be described as a lump-sum counterfactual of a single agent model. We can always look at the problem of solving for an individual player's best response as a single-agent problem, and if we hold the opponent's strategy fixed, then the only change is the change in the payoff function from u to $\widetilde{u}$. The counterfactual $\widetilde{u}$ here involves only a lump-sum change, so the results from single-agent models apply, and we have identified interim best responses.

However, as the final two columns of Table 2 describe, the identification of interim best responses does not lead to identification of the counterfactual equilibrium. In the new equilibrium, the change from u to $\widetilde{u}$ is no longer the only change in player i's problem; player i must also consider the change in her opponent's strategy, and this is amounts to a change in the transition process for the dynamic problem the player is solving. As Result 2 tells us, counterfactuals involving changed transition functions are generally not identified. For this reason, we conjecture that counterfactuals of dynamic games are generally not identified.[3]

Looking at the "fixed opponent" counterfactual best responses, one can see why the two models lead to different counterfactual equilibria. Thus, before strategic considerations are taken into effect, increasing

[3] Formally, Kalouptsidi et al.'s (?) conditions could be applied to evaluate whether a counterfactual of a given dynamic discrete game is identified. However, those conditions depend on properties of the transition process in both the baseline and counterfactual setting. For a single-agent model, checking the conditions amounts to checking primitives of the model because the transition process can be taken as a primitive. But for a dynamic game, the transition process is typically an equilibrium object, and therefore checking whether the conditions are satisfied is not simply a property of the primitives of the model, and Kalouptsidi et al.'s (?) conditions cannot be checked for a dynamic game without first solving for counterfactual equilibria.

Table 2: Entry Game: Choice Probabilities

entry costs decreases the rate of entry regardless of whether the opponent is active or not. In Model 1, reduced entry makes incumbent monopolists even more likely to remain active than in the baseline, for the risk of ending up in a duopoly, which comes with much lower profits, is reduced. In equilibrium, the increased appeal of becoming a monopolist increases the rate of entry in the state where both players were inactive in the previous period. Ex ante, it is ambiguous whether increased entry costs will end up increasing or decreasing the rate of entry when both firms are inactive; in this case, the rate of entry when both firms have been inactive is increased.

In Model 2, the direct effect of increasing entry costs is the same before strategic considerations are considered. However, in this model, the firm earns the highest profits when both firms are active, but when the opponent was inactive before. This is perhaps a strange model, for incumbents like entrants. The decreased rate of entrants challenging incumbents makes being a monopolist less appealing. The decreased appeal of becoming an incumbent monopolist here lowers the rate at which incumbents stay active and also lowers further the rate at which firms enter when no firms are active.

In summary, while the direct effects of the increased entry costs are the same in both models, the strategic effects actually push entry rates in different directions. Taking a step back, counterfactuals of dynamic games typically involve strategic considerations, meaning that firms face changes in opponents' expected behavior. Changes in opponents' expected behavior means that the transition process a given agent faces changes. Thus, Result 2 suggests that counterfactuals of dynamic games are generally not identified.

Appendix

In this appendix, we sketch proofs of Results 1 and 2. First, we define some notation and review some basic results from the literature.