Comments on "Identification and Semiparametric Estimation of a Finite Horizon Dynamic Discrete Choice Model with a Terminating Action"

Øystein Daljord[*], Denis Nekipelov[†] & Minjung Park[‡]

February 18, 2019

1 Introduction

Many important marketing and economic problems, such as purchases and pricing of durable goods and consumer behavior in credit markets, crucially hinge on agents' time preferences. However, it has been shown that in the standard dynamic discrete choice (DDC) models, the discount factor is not identified jointly with the utility function without further assumptions (Rust (1994); Magnac and Thesmar (2002)). Magnac and Thesmar further showed that under certain exclusion restrictions, the discount factor is point identified. These exclusion restrictions are however hard to interpret.

[*]Booth School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL, USA. E-mail: Oeystein.Daljord@chicagobooth.edu

[†]Departments of Economics and Computer Science, 254 Monroe Hall University of Virginia, VA, USA. E-mail: denis@virginia.edu

[‡]Department of Economics, Ewha Womans University, 52 Ewhayeodaegil, Seoul, Korea. E-mail: minjung_park@ewha.ac.kr. Park gratefully acknowledges the support provided by the National Research Foundation of Korea (NRF) Grant 2018S1A5A2A01029529.

Bajari et al. (2016) showed that in a finite horizon optimal stopping model, a class of DDC models that includes marketing relevant applications such as adoption of durable goods, the discount factor is point identified under an assumption of stationary utilities, an easily interpretable assumption that can be economically motivated in many settings. In this note, we show that the Bajari et al. approach gives identification for a much larger class of problems.

We first reformulate the identification result in Bajari et al. in terms of exclusion restrictions on time as a state variable. We then extend the identification result to both finite horizon and infinite horizon optimal stopping models under more general exclusion restrictions, as in Abbring and Daljord (2018). Finally, we show how a similar approach gives identification of general discount functions in a finite horizon optimal stopping problem.

For all of the models we consider, the discount functions are shown to be closed form solutions to well-behaved moment conditions that are functions of identified features of the data distribution. The results directly suggest estimators that are robust to biases from finite sample approximations to the unknown utility function and can be implemented as simple linear regressions using pooled choice data.

2 Model

We first develop a single agent, finite horizon, optimal stopping version of Rust (1994) to reformulate the identification result in Bajari et al., and will later consider an extension of the results to infinite horizon optimal stopping models. Time is discrete and indexed by $t = 1, \ldots, T$. The choice set is discrete and indexed by $a \in \mathcal{A} = \{0, 1, \ldots, A\}$, where $a = 0$ denotes a terminating action. The panel may or may not include the final period T. The exogenous states $s \in \mathcal{S}$, which are observable to both the agent and the econometrician, follow a stationary, first-order Markov process with probability distribution $F(s_{t+1}|a_t, s_t)$, without absorbing states.

The states $\epsilon_t = \{\epsilon_{0,t}, \ldots, \epsilon_{A,t}\}$, which are observed by the agent, but not by the econometrician, are drawn independently of s from a stationary and absolutely continuous distribution G with infinite support. The agent's per-period utility function is $u_t(a, s, \epsilon) = u_t(a, s) + \epsilon_{a,t}$. The per-period utility of the reference choice is assumed state invariant and normalized to

$u_t(0,s) = 0$ for all $s \in \mathcal{S}$ and $t = 1, \ldots, T$. The agent's expectations are assumed to be rational and coincide with $F(s_{t+1}|a_t, s_t)$.

The agent's choices are assumed to maximize the expected, geometrically discounted sum of lifetime utilities

$$E_{s,\epsilon}\left[\sum_{\tau=t}^{T} \beta^{\tau-t}(1 - \max_{q<\tau} \mathbf{1}\{a_q = 0\})(u_\tau(a_\tau, s_\tau) + \epsilon_{a_\tau,\tau})|s_t\right],$$

where β is a finite and non-negative discount factor.[1] The primitives of the problem are the utility functions u_t, the discount factor β, the state transition process F, and the distribution of unobservable states G. We assume that ϵ_t is i.i.d. EV1 distributed throughout for ease of exposition and computation.

The choice specific value function gives the expected discounted value of making choice a in state s in period t:

$$V_t(a,s) = u_t(a,s) + \beta E_s\left[V_{t+1}(s_{t+1})\middle| a_t = a, s_t = s\right], \tag{1}$$

where the ex ante value function

$$V_t(s) = E_\epsilon\left[\max_{a\in\mathcal{A}}\{V_t(a,s) + \epsilon_{a,t}\}\right] = \ln\left(\sum_{a\in\mathcal{A}} \exp(V_t(a,s))\right) \tag{2}$$

gives the expected value of making the optimal choice in state s in period t, prior to learning the unobserved states ϵ_t.[2] The choice probabilities $\sigma_t(a,s) = Pr[a_t = a|s_t = s]$ can be written in terms of the choice specific value functions as

$$\sigma_t(a,s) = \frac{\exp(V_t(a,s))}{\sum_{k\in\mathcal{A}} \exp(V_t(k,s))}. \tag{3}$$

The choice specific value contrasts can be uniquely recovered from the choice probabilities as

$$\ln\left(\frac{\sigma_t(a,s)}{\sigma_t(0,s)}\right) = V_t(a,s) - V_t(0,s). \tag{4}$$

The moment conditions in (4) are the reduced form of the model.[3] They map the endogenous objects, on the right hand side, to known functions of the data, on the left hand side. Using the

[1] The discount factor is not restricted to the unit in finite horizon models for the model to be well-defined.

[2] We omit the Euler's constant as it is just a constant and does not affect any of our analysis.

[3] Hotz and Miller (1993) and Norets and Takahashi (2013) show the value contrasts can be uniquely recovered from the choice data for any absolutely continuous distribution with infinite support.

choice specific value function in (1), the reduced form can be written in terms of the primitives

$$\ln\left(\frac{\sigma_t(a,s)}{\sigma_t(0,s)}\right) = u_t(a,s) - V_t(0,s) + \beta E_s\left[V_{t+1}(s_{t+1})\,\middle|\, a_t = a, s_t = s\right]. \tag{5}$$

The normalization $u_t(0,s) = 0$ for the terminating action implies $V_t(0,s) = u_t(0,s) = 0$. Then using (2) and (3), the moment conditions in (5) become known functions of the data and the primitives u_t and β

$$\ln\left(\frac{\sigma_t(a,s)}{\sigma_t(0,s)}\right) = u_t(a,s) - \beta E_s\left[\ln(\sigma_{t+1}(0,s_{t+1}))\,\middle|\, a_t = a, s_t = s\right]. \tag{6}$$

Magnac and Thesmar shows that even if F, σ, and G are known, the discount factor is not identified in DDC models without further restrictions on the primitives. The underidentification can be seen directly from (6): for each finite and non-negative β, and for each period $t \in 1,\ldots,T$, we can find a different function u_t that satisfies these moment conditions.

2.1 Identification under the assumption of stationary utilities

Bajari et al. shows that assuming stationary utilities point identifies the discount factor in the model developed in the previous section. Stationarity can be cast as exclusion restrictions of the form

$$u_t(a,s) = u_{t'}(a,s) \tag{7}$$

for some pair of periods t and t', some choice $a \in \mathcal{A}\backslash\{0\}$ and some subset of the state space $\mathcal{S}$. The following theorem is based on Bajari et al.[4]

Theorem 1 *Suppose that the exclusion restriction in (7) is satisfied for some pair of periods t and t', with $t, t' \in 1,\ldots,T$, and with $t \neq t'$, some choice $a \in \mathcal{A}\backslash\{0\}$ and some point $s \in \mathcal{S}$. Then β is point identified, subject to a rank condition.*

Proof:

Difference (6) between t and t' to get

$$\beta = \frac{\ln\left(\frac{\sigma_t(a,s)}{\sigma_t(0,s)}\right) - \ln\left(\frac{\sigma_{t'}(a,s)}{\sigma_{t'}(0,s)}\right)}{E_s\left[\ln(\sigma_{t'+1}(0,s_{t'+1}))\,\middle|\, a_{t'} = a, s_{t'} = s\right] - E_s\left[\ln(\sigma_{t+1}(0,s_{t+1}))\,\middle|\, a_t = a, s_t = s\right]} \tag{8}$$

[4]We have adopted the term 'exclusion restriction' from Magnac and Thesmar for restrictions like (7).

which is a linear function of β, the only primitive in the equation. The moment condition therefore has a unique solution subject to the rank condition

$$E_s\left[\ln(\sigma_{t'+1}(0, s_{t'+1}))\middle| a_{t'} = a, s_{t'} = s\right] \neq E_s\left[\ln(\sigma_{t+1}(0, s_{t+1}))\middle| a_t = a, s_t = s\right]. \tag{9}$$

Q.E.D.

The moment condition in (8) gives the discount factor as a closed form function of choice probabilities and expectations, which are both identified features of the data distribution, and holds for an arbitrary utility function that satisfies the exclusion restrictions. The moment condition therefore suggests a simple linear estimator that recovers the discount factor only without jointly estimating the utility function, and the resulting estimator will be robust to biases from finite sample approximations to the unknown utility function. The moment condition contains all the information in the data about the discount factor in the sense that the discount factor is underidentified from the moment conditions in (6) alone, but identified from (8) independently of (6).

Since the continuation values in the denominator on the right hand side of (8) are known mappings from the data distribution, point identification holds even for panels which do not include the final period, i.e., short panels, as long as the data contain information on at least three consecutive periods.

The discount factor is the ratio of the difference in log choice probabilities between periods t and t' to the difference in continuation values between the same two periods. The stronger the current period choice response, in the numerator, is to shifts in future rewards, in the denominator, the more forward looking the agent is revealed to be. If the agent's choices do not respond to shifts in the continuation values, we infer that the agent is myopic, i.e., $\beta = 0$.

Under the assumption of stationary utilities, the continuation values differ between t and t', but the unobservable current period utilities are invariant between the two periods, conditional on s and a. We can therefore interpret the difference in choice probabilities between t and t' as entirely due to the shifts in continuation values, therefore providing information on time preferences. The rank condition in (9) requires that time non-trivially shift the continuation values. If not, there is no information about time preferences in the data, and identification is lost.

2.2 General exclusion restrictions with identifying power

There is nothing special about stationary utilities as an exclusion restriction. We next show that exclusion restrictions on states other than those involving time have identifying power as well.

Suppose that for some t there exist either a pair of choices $a_1 \in \mathcal{A}\backslash\{0\}$ and $a_2 \in \mathcal{A}$, or a pair of states $s_1, s_2 \in \mathcal{S}$, where either $a_1 \neq a_2$, or $s_1 \neq s_2$, or both, such that the following exclusion restriction holds.

$$u_t(a_1, s_1) = u_t(a_2, s_2) \tag{10}$$

In words, the exclusion restriction states that we can find a pair of states or a pair of choices, or both, for which the current period utility is invariant but the continuation values vary. It is easy to see that we can then express the discount factor as a closed form function of choice probabilities and expectations similar to (8), by taking appropriate differencing to eliminate current period utility. Therefore, the discount factor is point identified under the exclusion restrictions (10), subject to a rank condition similar to the one in (9).

Abbring and Daljord shows that this identification strategy reflects common intuition reported in the literature. One example is Lee (2013) which studies demand for video game consoles in an infinite horizon optimal stopping model where it is assumed that the expected quality and availability of future game releases shift the continuation values of owning a console, but not its current period utility. The assumption leads Lee to interpret current period demand responses to variation in the expected quality and availability of future releases as informative about time preferences.

As another example, in a study of household adoption of solar panel technology, De Groote and Verboven (2019) uses variation in subsidy programs that shifts future energy cost streams, but not the upfront installation cost, for identification of a discount factor. In its study, s_1 and s_2 may represent two different subsidy programs, for equal prices of installation, while $a_1 = a_2$ is choosing installation in the current period.

The exclusion restrictions in (10) are sufficient to point identify the discount factor, i.e., stationary utilities are not required (see Abbring and Daljord). Allowing for exclusion restrictions on state variables other than time broadens the scope for identification to infinite horizon problems as well, where the assumption of stationary utilities does not have identifying power. Abbring and Daljord shows that for infinite horizon problems where $V(0,s)$, the choice specific value of the reference choice, is a constant, the discount factor is point identified under the exclusion restrictions (10). These problems include some optimal stopping problems and renewal problems like in Rust (1987). For an infinite horizon optimal stopping problem, the moment condition corresponding to (8) is

$$\beta = \frac{\ln\left(\frac{\sigma(a_1,s_1)}{\sigma(0,s_1)}\right) - \ln\left(\frac{\sigma(a_2,s_2)}{\sigma(0,s_2)}\right)}{E_s\left[\ln(\sigma(0,s'))\middle| a=a_2, s=s_2\right] - E_s\left[\ln(\sigma(0,s'))\middle| a=a_1, s=s_1\right]}. \tag{11}$$

Again, the moment condition in (11) directly suggests a simple linear estimator for the discount factor that is robust to biases from finite sample approximations to the unknown utility function.

Both optimal stopping problems and renewal problems are examples of models with the finite dependence property of Arcidiacono and Miller (2011). A model satisfies single action one-period finite dependence if

$$F(s_{t+1}|s_t,a_1)F(s_{t+2}|s_{t+1},a_2) = F(s_{t+1}|s_t,a_2)F(s_{t+2}|s_{t+1},a_2) \tag{12}$$

holds for two choices a_1 and a_2, with $a_1 \neq a_2$, in the choice set. Single action one period finite dependence implies that in expectation, the choice sequence of a_1 and a_2 leads to the same states as the choice sequence of a_2 and a_2, if both sequences start in the same states. Models with finite dependence simplify estimation and are often used in applications, see Arcidiacono and Miller (2018) and Kalouptsidi et al. (2018) for examples and references. Finite dependence is also useful for identification. Theorem 2 in Abbring and Daljord shows that both finite horizon and infinite horizon models with single action one-period finite dependence are point identified under exclusion restrictions similar to (10).

2.3 General discount functions

In this extension, we consider identification of general discount functions in a finite horizon optimal stopping model under the assumption of stationary utilities, where the time preferences are expressed as multiplicative weights for the additively separable utilities in the future periods.

Time preferences are time-consistent if the choice between any two streams of utility does not change when the streams are forwarded by an equal period of time. The exponential discount function is the only time-consistent one, but is frequently rejected in experiments (Frederick et al. (2002)). The study of time-inconsistent preferences in the economics literature goes back to Strotz (1955) and Phelps and Pollak (1968). As is common in the literature on time-inconsistent preferences, we adopt the convention of considering the agent as a collection of selves, one in every period, whose preferences may be in conflict. O'Donoghue and Rabin (1999) distinguished between 'sophisticated' agents, who are fully aware of the time inconsistency of their preferences, versus those who are 'naïve,' as well as partial naïveté, an entire range in between.

In the analysis, we focus on the case of sophisticated agents who are fully aware of the time inconsistency of their preferences. We do so not because we believe naïveté is irrelevant in practice, but because identifying the degree of naïveté would require additional information that is unavailable in most datasets. For example, identifying the degree of sophistication may require observing whether individuals make appropriate use of self-commitment devices that constrain their future ability to take suboptimal actions, which only sophisticated agents would do. Assuming sophisticated agents enables two-step estimation, by ensuring that the observed choices of future selves are consistent with the current self's expectation about future selves' actions. The results in this extension crucially rely on the assumption of sophisticated agents. For an empirical model of partial naïveté, see Fang and Wang (2015).

Consider a general discount function $\beta : \{1, \ldots, T\} \mapsto [0, 1]$ with normalization $\beta(0) = 1$. Under standard exponential discounting, we get $\beta(t) = \beta^t$. In the DDC problem with a terminating action ($a = 0$), the lifetime utility to the self in period t can be expressed as

$$U_t(a_t, \ldots, a_T, s_t, \ldots, s_T) = \sum_{\tau=t}^{T} \beta(\tau - t)(1 - \max_{q<\tau} \mathbf{1}\{a_q = 0\})(u_\tau(a_\tau, s_\tau) + \epsilon_{a_\tau,\tau}).$$

If $d_t(s,\epsilon)$ is the decision function of the self in period t when the state variable is equal to s, we can express the choice specific value function for choice a to the self in period t as

$$V_t(a,s) = u_t(a,s) + E_{s,\epsilon}\Bigg[\sum_{\tau=t+1}^{T} \beta(\tau-t)\prod_{q<\tau}\mathbf{1}\{d_q(s_q,\epsilon_q)\neq 0\} \\ \left(u_\tau(d_\tau(s_\tau,\epsilon_\tau),s_\tau) + \epsilon_{d_\tau(s_\tau,\epsilon_\tau),\tau}\right)\Bigg|\, a_t = a, s_t = s\Bigg].$$

Taking into account the recursive optimization by the economic agent and the fact that the decision in the last period is static and the decision in the period before the last is equivalent to that in case of standard exponential discounting, we can express the choice specific value of the self in period $T-2$ as follows

$$V_{T-2}(a,s) = u_{T-2}(a,s) + \beta(1)E_s\Big[V_{T-1}(s_{T-1})\,|\,a_{T-2}=a, s_{T-2}=s\Big] \\ + \left(\beta(2)-\beta(1)^2\right)E_s\Big[V_T(s_T)\,|\,a_{T-2}=a, s_{T-2}=s\Big]. \quad (13)$$

Note that we have written the choice specific value function of the self in period $T-2$ in terms of the ex ante value functions $V_{T-1}(s)$ of the self in period $T-1$ and $V_T(s)$ of the period T self. The expectations on the right hand side of (13) are formed by the self in period $T-2$, as it is the current self's expectations about future selves' behavior that influence the current self's actions. Under the assumption of sophisticated agents, the current self's expectations about future selves' behavior become consistent with the actual behavior of future selves, which allows us to write the expectations on the right hand side in terms of known mappings from the observed choices of the future selves, e.g., $E\Big[V_{T-1}(s_{T-1})\,|\,a_{T-2}=a, s_{T-2}=s\Big] = E\Big[-\ln(\sigma_{T-1}(0,s_{T-1}))\,|\,a_{T-2} = a, s_{T-2}=s\Big]$.

This representation generalizes to any $1\leq t\leq T$. We introduce a recursive sequence of coefficients $\delta(t)$ such that $\delta(1)=\beta(1)$ and $\delta(t)=\beta(t)-\sum\limits_{r=1}^{t-1}\beta(r)\delta(t-r)$. We note that there is a one-to-one correspondence between the discount function $\beta(\cdot)$ and the sequence $\delta(\cdot)$.

We write the generalized Bellman equation for the finite horizon problem with a general discount function as

$$V_t(a,s) = u_t(a,s) + \sum_{\tau=1}^{T-t}\delta(\tau)E_s\Big[V_{t+\tau}(s_{t+\tau})\,\Big|\,a_t=a, s_t=s\Big]. \quad (14)$$

The choice specific value function for the self in period t is seen to depend on the ex ante value functions of all future selves.

The specification in (14) nests commonly used discount functions. Under exponential discounting, $\beta(t) = \beta^t$, we get $\delta(1) = \beta$ and $\delta(r) \equiv 0$ for all $r > 1$, and the generalized Bellman equation reduces to the conventional Bellman equation for all $t \in 1, \ldots T$.

For the case of hyperbolic discounting, let β denote the standard discount factor which captures long-run, time-consistent discounting, and let α denote the present bias factor, which captures short-term discounting, such that $\beta(t) = \alpha\beta^{t-1}$. Then, we have $\delta(t) = \alpha(\beta - \alpha)^{t-1}$. When α is smaller than β, the agent exhibits present bias, i.e., is more impatient in the short run than in the long run.

We can express the generalized Bellman equation in (14) as

$$\ln\left(\frac{\sigma_t(a,s)}{\sigma_t(0,s)}\right) = u_t(a,s) - \sum_{\tau=1}^{T-t} \delta(\tau) E_s\left[\ln \sigma_{t+\tau}(0, s_{t+\tau}) \,\middle|\, a_t = a, s_t = s\right],$$

for all states s and periods t. Taking the difference between periods t and t' under the assumption of stationary utilities $u_t(a,s) = u_{t'}(a,s)$ for some $a \in \mathcal{A}\backslash\{0\}$ and $s \in \mathcal{S}$, we can write

$$\ln\left(\frac{\sigma_t(a,s)}{\sigma_t(0,s)}\right) - \ln\left(\frac{\sigma_{t'}(a,s)}{\sigma_{t'}(0,s)}\right) = \tag{15}$$
$$\sum_{\tau=1}^{T-t'} \delta(\tau) E_s\left[\ln \sigma_{t'+\tau}(0, s_{t'+\tau}) \,\middle|\, a_{t'} = a, s_{t'} = s\right] - \sum_{\tau=1}^{T-t} \delta(\tau) E_s\left[\ln \sigma_{t+\tau}(0, s_{t+\tau}) \,\middle|\, a_t = a, s_t = s\right].$$

This expression allows us to formulate the following theorem.

Theorem 2 *Suppose that the data are generated by sophisticated agents who are aware of their time inconsistency while making their discrete choice. Also suppose that $u_t(a,s) = u_{t'}(a,s)$ for some pair of periods t and t' ($t \neq t'$), some $a \in \mathcal{A}\backslash\{0\}$ and all $s \in \mathcal{S}$, and that data are observed for all periods $\min(t,\, t'), ..., T$. Then the general discount function $\beta(\cdot)$ is identified on the subset of its support $1, \ldots, T - \min(t,t')$ subject to the following rank condition: all leading principal minors of matrix M with entries $M_{rp} = Cov(Z_r, Z_p)$ are non-zero, where*

$$Z_r = E_s\left[\ln \sigma_{t'+r}(0, s_{t'+r}) \,\middle|\, a_{t'} = a, s_{t'} = s\right] - E_s\left[\ln \sigma_{t+r}(0, s_{t+r}) \,\middle|\, a_t = a, s_t = s\right]$$

for r = 1, 2, ..., $T - \max(t,t')$ and

$$Z_r = (-1)^{\mathbf{1}\{t<t'\}} E_s\left[\ln \sigma_{\min(t,t')+r}(0, s_{\min(t,t')+r}) \,\middle|\, a_{\min(t,t')} = a, s_{\min(t,t')} = s\right]$$

for $T - \max(t,t') < r \leq T - \min(t,t')$.

Proof:

Let $\vec{\delta} = (\delta(1), \ldots, \delta(T - \min(t,t')))'$,

$$Y = \ln\left(\frac{\sigma_t(a,s)}{\sigma_t(0,s)}\right) - \ln\left(\frac{\sigma_{t'}(a,s)}{\sigma_{t'}(0,s)}\right)$$

and

$$\vec{Z} = (Z_1, \ldots, Z_{T-\min(t,t')})'.$$

Then the generalized Bellman equation can be written in the vector form as

$$Y = \vec{Z}'\vec{\delta}.$$

This means that

$$Y - E[Y] = (\vec{Z} - E[\vec{Z}])'\vec{\delta}.$$

As a result

$$\vec{\delta} = M^{-1}E[(\vec{Z} - E[\vec{Z}])(Y - E[Y])].$$

By assumption, M is invertible, and therefore $\vec{\delta}$ is well-defined. The discounting schedule can be found by setting $\beta(1) = \delta(1)$ and iterating to get

$$\beta(t) = \delta(t) + \sum_{r=1}^{t-1} \beta(r)\delta(t-r).$$

Q.E.D.

Note that if both t and t' are greater than 1, the discount function is only identified on the subset of support, and not in its entirety. The general discount function can be estimated by regressing the difference in the ratio of the choice probabilities between two periods on the expected differences in all future choice probabilities for the reference choice, similar to the previous cases. Intuitively, the general discount function is recovered as the sensitivity of the current period choices to variation in rewards one period later, two periods later, etc., and the rank condition ensures that the variations in future rewards are not perfectly collinear with each other. From the moment condition in (15), we can recover the weights on future values that rationalize the current period choices. The recovered weights can be subjected to hypothesis tests, e.g., whether the data are consistent with exponential or hyperbolic discounting. As in previous cases, the estimator recovers the discount function independently of the utility function, and therefore is robust to biases from finite sample approximations to the unknown utility function.

3 Summary

This note provides two extensions of the identification results in Bajari et al. First, we show that similar identification results hold under a broader set of exclusion restrictions. These general exclusion restrictions do not require stationary utilities, and therefore extend to infinite horizon optimal stopping models as well. Second, we show that identification of general discount functions is possible in a finite horizon optimal stopping problem under the assumption of sophisticated agents. Under both extensions, the identification proof directly suggests estimators of the discount factor (or discount function) that are robust to biases from finite sample approximations to the unknown utility function and can be implemented as simple linear regressions.

References

Abbring, J. and O. Daljord (2018). Identifying the discount factor in dynamic discrete choice models. Mimeo, University of Chicago. 2, 6, 7

Arcidiacono, P. and R. A. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica 79*(6), 1823–1867. 7

Arcidiacono, P. and R. A. Miller (2018). Identifying dynamic discrete choice models off short panels. Working paper, Cambridge, MA. 7

Bajari, P., C. S. Chu, D. Nekipelov, and M. Park (2016). Identification and semiparametric estimation of a finite horizon dynamic discrete choice model with a terminating action. *Quantitative Marketing and Economics 14*(4), 271–323. 1, 2, 4, 12

De Groote, O. and F. Verboven (2019). Subsidies and myopia in technology adoption: Evidence from solar photovoltaic systems. Forthcoming american economic review. 6

Fang, H. and Y. Wang (2015). Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. *International Economic Review 56*(2), 565–596. 8

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature 40*(2), 351–401. 8

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. 3

Kalouptsidi, M., P. T. Scott, and E. Souza-Rodrigues (2018). Linear iv regression estimators for structural dynamic discrete choice models. Working paper, Harvard University. 7

Lee, R. (2013). Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review 103*(7), 2960–3000. 6

Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete choice processes. *Econometrica 70*, 801–816. 1, 4

Norets, A. and S. Takahashi (2013). On the surjectivity of the mapping between utilities and choice probabilities. *Quantitative Economics 4*, 149–155. 3

O'Donoghue, T. and M. Rabin (1999). Doing it now or later. *American Economic Review 89*(1), 103–124. 8

Phelps, E. and R. A. Pollak (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies 35*(2), 185–199. 8

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica 55*, 999–1033. 7

Rust, J. (1994). Structural estimation of Markov decision processes. In R. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, pp. 3081–3143. Amsterdam: North-Holland. 1, 2

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies 23*(3), 165–180. 8