Pseudo-Value Functions and Closed-Form CCP Estimation of Dynamic Discrete Choice Models*

Adam Dearing†

January 15, 2019

[Work in Progress]

Abstract

Conditional choice probability (CCP) estimation of dynamic discrete choice models involves construction of a pseudo-value function using first-stage estimates of the conditional choice probabilities. Focusing on single-agent models, this paper provides two primary contributions to the literature. First, we show that the different pseudo-value functions used in the nested pseudo-likelihood estimator of Aguirregabiria and Mira (2002) and the finite-dependence-based estimators of Arcidiacono and Miller (2011) are both special cases of a general form, providing a unifying framework for these estimation methods. This result broadens the class of models to which nested pseudo-likelihood estimation can be easily applied and helps provide insight into the (in)efficiency properties of finite-dependence-based estimators. Second, we propose a minimum distance-style estimator that admits a closed-form solution, and we show that an appropriate choice of weighting matrix makes this estimator asymptotically equivalent to pseudo-maximum likelihood. Furthermore, we show that iterations on the estimator can be used to compute the (partial) MLE in finite samples.

1 Introduction

Many fields in economics implement stationary dynamic discrete choice models in applied work and, although these models are very useful for capturing relevant economic forces,

*We thank Christopher Sullivan, Sarah Johnston, Jason Blevins, and mini-seminar participants at NYU Stern for many helpful comments. All errors are our own.

†Department of Economics, The Ohio State University, 401 Arps Hall, 1945 N. High St., Columbus OH 43210. Email: dearing.30@osu.edu.

estimation can be computationally cumbersome.[1] Maximum likelihood estimation via the nested fixed-point algorithm of Rust (1987; 1988) is perhaps the most well-known and widely-used method, although it can be slow, especially as the state space grows in size. Several methods have since been proposed to reduce the computational burden of estimation while preserving asymptotic consistency - and in some cases, efficiency - of the estimator. Many these estimators fall under the broad class of conditional choice probability (CCP) estimators first introduced by Hotz and Miller (1993). Examples include their original estimators, the simulation-based estimator of Hotz et al. (1994), the nested pseudo-likelihood estimator of Aguirregabiria and Mira (2002), and estimators exploiting finite dependence (Altug and Miller, 1998; Arcidiacono and Miller, 2011).[2]

In general, CCP-based algorithms proceed by maximizing either a pseudo-likelihood or a minimum-distance-type objective function where the (integrated) value function is replaced by a pseudo-value function constructed from a consistent estimate of the CCPs in each state. Thus, they can be differentiated along two dimensions: (i) the choice of objective function (i.e., likelihood vs. minimum distance); and (ii) the choice of pseudo-value function. These choices have a substantial impact on the estimators, both in theory and in practice. For example, the specific choice of objective function in the simulation-based estimator of Hotz et al. (1994) allows for consistency of the estimator even when the number of simulation draws is fixed, whereas a likelihood formulation does not achieve consistency with a fixed number of draws. On the other hand, the NPL estimator achieves both consistency and efficiency after any number of iterations when initialized with a consistent estimate of the choice probabilities. A third class, the finite-dependence-based estimators, achieve consistency but not efficiency when used with standard objective functions. Note that the latter two examples often differ only in their construction of the pseudo-value function.[3] Pesendorfer and Schmidt-Dengler (2008) focus on the choice of objective function and show that likelihood and method-of-moments objective functions can be represented as special cases of minimum-distance estimators, establishing the minimum-distance objective function as a unifying framework. While the analysis in Pesendorfer and Schmidt-Dengler (2008) also applies to dynamic games, we limit our analysis to single-agent models. We further restrict attention to infinite-horizon (stationary) models without permanent unobserved heterogeneity or serially correlated unobserved shocks.[4]

[1] For reviews of the literature and applications, see Ackerberg et al. (2007); Aguirregabiria and Mira (2010); Keane et al. (2011), among others.

[2] These papers all deal with single-agent decision models. Some examples of applications to dynamic games include Aguirregabiria and Mira (2007); Pakes et al. (2007); Pesendorfer and Schmidt-Dengler (2008).

[3] Finite-dependece-based estimators are often implemented in settings where a minimum distance objective function becomes intractable without restrictions on the weighting matrix.

[4] Some examples of models that include these features can be found in Pakes (1986); Keane and Wolpin

The first contribution of this paper is to provide a unifying framework for the various known types of pseudo-value functions. We accomplish this by using a result analogous to those in Arcidiacono and Miller (2011; 2016) to construct a general-form of pseudo-value function, which we call the Arcidiacono-Hotz-Miller Pseudo-Value Function (AHM-PVF). We then show that the pseudo-value functions most commonly used in CCP-based estimators, including the one used in nested pseudo-likelihood (NPL-PVF) are special cases of the AHM-PVF. In stationary models, the AHM-PVF is not computationally tractable in its general form, so focusing on special cases ensures tractability. We provide several examples of its applicability.

A notable consequence of establishing this general framework is that it broadens the applicability of the NPL-PVF, which to the best of our knowledge has until now only been applied to multinomial logit and binary probit models due to the need for closed-form expressions. The AHM-PVF does not require such closed-form expressions and can be applied to arbitrary distributions.[5] Additionally, we use our framework to provide a new representation of the NPL-PVF as a first-order approximation to the true value function. We then use this representation to give alternative proof of the celebrated "zero Jacobian" property of the NPL-PVF that strengthens the intuition behind the result. We then present conditions for a similar zero jacobian property for the AHM-PVF and use these conditions to discuss the (in)efficiency properties of other applications, including finite-dependence-based estimators.

The second contribution of our paper is to provide a useful closed-form estimator for parameters that enter the flow utility function linearly. This closed-form estimator is a minimum-distance estimator that is asymptotically equivalent to the minimum-distance estimator of Pesendorfer and Schmidt-Dengler (2008) by construction. Indeed, our paper is complementary to theirs: they provide a unifying framework for objective functions, whereas we provide a new objective function and a unifying framework for pseudo-value functions. Asymptotic equivalence of our estimator to theirs is in turn used to establish asymptotic equivalence to pseudo-likelihood estimators by appropriate choice of the weighting matrix. Furthermore, we show that the closed-form estimator can be iterated upon in order to compute pseudo-likelihood estimates and even the (partial) MLE in finite samples and describe the local convergence properties of such iterations.

While ours is not the first closed-form estimator in the dynamic discrete choice literature (see Kalouptsidi et al. (2018) among others), it is to the best of our knowledge the first closed-form estimator that is asymptotically equivalent to MLE for models with an arbitrary

(1997); Norets (2009); Arcidiacono and Miller (2011).

[5] In particular, NPL-PVF requires the researcher to compute $e(P)$, which to date can only be computed via a closed-form expression. Whereas the AHM-PVF requires $\psi(P)$, which can be computed numerically. We discuss this further in the body of the paper.

number of actions and distribution of the utility shocks. The closest estimator to ours is the closed form estimator of Bray (2018), which is asymptotically equivalent to MLE but requires a multinomial logit framework, whereas our estimator allows for arbitrary distributions of the utility shocks.[6]

The remainder of the paper proceeds as follows. Section 2 describes the stationary dynamic discrete choice model and introduces some useful notation. Section 3 establishes the unifying framework for pseudo-value functions and its implications. Section 4 describes the closed-form estimators and provides analysis of their asymptotic and iterative properties. Section 5 provides Monte-Carlo simulations to illustrate finite sample performance of estimators. Section 6 concludes. Any proofs not appearing in the body of the text appear in the appendix. Throughout, let $\nabla_x f = \partial f/\partial x'$ and $\nabla_{xy} f = \partial^2 f/(\partial x \partial y')$ where appropriate.

2 Dynamic Discrete Choice Model

In this section, we describe a standard stationary (infinite-horizon) dynamic discrete choice (DDC) model and the re-characterization of the integrated value function. Time is indexed by t and there are three types of variables: observed state variables, x_t, unobserved state variables, ε_t, and action variable a_t. The decision maker observes both x_t and ε_t when choosing a_t, whereas the econometrician only observes the resulting pair (x_t, a_t). The choice comes from a finite set of mutually exclusive alternatives: $a_t \in \mathcal{A} = \{0, 1, 2, ..., J\}$. There is a vector of structural parameters, θ, and the state variables evolve according to a Markov transition process, $p(x_{t+1}, \varepsilon_{t+1}|x_t, \varepsilon_t, a_t, \theta)$, which is known to the decision maker. The per-period flow utility function is $u(x_t, \varepsilon_t, a_t, \theta)$ and there is a discount factor, $\beta \in (0, 1)$. We partition $\theta = (\theta_u, \theta_f, \theta_g)$ and impose some further assumptions, which are ubiquitous in the dynamic discrete choice literature.

Assumption 1. *(Additivity) The flow utility function is additively separable in the observed and unobserved state variables:*

$$u(x_t, \varepsilon_t, a_t, \theta) = u(x_t, a_t, \theta_u) + \varepsilon(a_t),$$

where $\varepsilon(a_t)$ is the corresponding element of ε_t, a $(J+1) \times 1$ vector.

Assumption 2. *(Conditional Independence) The Markov transition probability is given by*

[6] Bray's closed-form estimator arises as a special case of an estimation algorithm that "unnests" the fixed point algorithm even when utility is not linear in parameters. Applying the algorithm when utility is linear in parameters then results in his closed-form estimator.

$p(x_{t+1}, \varepsilon_{t+1}|x_t, \varepsilon_t, a_t, \theta) = g(\varepsilon_{t+1}|x_{t+1}, \theta_g) f(x_{t+1}|x_t, a_t, \theta_f)$. *Furthermore,* $g(\cdot|\theta_g)$ *has full support on* $\mathbb{R}^{J+1}$*, finite first moments, and is twice continuously differentiable.*

Assumption 3. *(Finite State Space)* $x_t \in \mathcal{X} = \{1, 2, ..., M\}$.

The first assumption is standard. Assumption 2 is the standard conditional independence assumption introduced by Rust (1987). Usually, the elements of ε_{t+1} are i.i.d. Type 1 Extreme Value, and their distribution therefore does not depend on x_{t+1}. Assuming that θ_u only affects the flow utilities while (θ_g, θ_f) only affects the transition probabilities is also standard. Assumption 3 is essential to the analysis presented here, although it may not be too restrictive because continuous state spaces are often discretized in practice.

The decision maker's choice maximizes the expected sum of current and (discounted) future utility. The problem is stationary, so we now replace x_t and x_{t+1} with x and x', respectively. A similar adjustment is made for the other variables of interest. This yields the value function, $V(x, \varepsilon, \theta)$, which is the unique solution to the standard Bellman equation:

$$V(x, \varepsilon, \theta) = \max_a \left\{ u(x, a, \theta_u) + \varepsilon(a) + \beta E_{x', \varepsilon'} \left[V(x', \varepsilon', \theta) \right] \right\}.$$

Now consider, instead, a Bellman equation based on the "integrated value function," $V(x, \theta)$. For this function, we integrate over ε:

$$V(x, \theta) = \int \max_a \left\{ u(x, a, \theta_u) + \varepsilon(a) + \beta \sum_{x'} V(x', \theta) f(x'|x, a, \theta_f) \right\} g(\varepsilon|x, \theta_g) d\varepsilon.$$

Alternatively, the Bellman equation can be defined on the expected continuation value, $EV(x, a, \theta)$, as in Rust (1987):[7]

$$EV(x, a, \theta) = \sum_{x'} \left(\int \max_a \left\{ u(x, a, \theta) + \varepsilon(a) + \beta EV(x, a, \theta) \right\} g(\varepsilon|x', \theta_g) d\varepsilon \right) f(x'|x, a, \theta_f).$$

Both of these functions are known to contain the same information, but the expected value function has $M \times (J + 1)$ terms, whereas the integrated value function has only M terms. Some of the results in this section of the paper apply to either function, but we focus on the integrated value function, which will allow us to explicitly link the various CCP estimators in the next section. From here on out, we will use the terms "integrated value function" and "value function" interchangeably. We will also assume that (θ_f, θ_g) are known and will often omit them for notational simplicity. The remainder of the analysis will solely focus

[7] Rust (1987) calls this the "expected value function." The integrated value function is also sometimes referred as the ex-ante value function.

on estimation of θ_u. In practice, values for (θ_f, θ_g) are usually (i) imposed from the outset, (ii) estimated in the first stage, or (iii) estimated jointly with θ_u.[8] That is, we are only interested in estimating the parameter vector that influences the flow utility values, and so our estimation procedures are amenable to (i) or (ii) but not (iii). However, this requires some careful use of the word "efficiency." When we say an estimator of θ_u is "efficient," we mean that it is efficient when (θ_f, θ_g) are known in advance. But frequently, θ_g is estimated in a first stage. In this case, "efficiency" of $\hat{\theta}_u$ means that it is equivalent to the partial maximum likelihood estimate. We remind the reader of this distinction sporadically throughout the paper.

Because the observed state, x, is discrete-valued, standard notions of differentiability, continuity, and convexity will not apply. It is not very intuitive to think of the integrated value function as a function of x. Standard notions of differentiability, continuity, and convexity will not apply. Instead, it is more natural to think of both x and a as indices, and to interpret the integrated value function as a function of the $(M \times (J+1))$ vector of choice-specific flow utilities, u, with elements $u_{(x,a)}$. This is also useful for characterizing transition probabilities. Define the vector $F_{x,a}$ as the vector of one-period-ahead transition probabilities, whose elements are $F_{x,a}(j) = Pr(x' = j|x, a)$. Define the matrix F_a as the $M \times M$ matrix of one-period-ahead transition probabilities conditional on choice d. That is, $F_a\{i, j\} = Pr(x' = j|x = i, a)$, so that

$$F_a = \begin{bmatrix} F'_{x=1,a} \\ F'_{x=2,a} \\ \vdots \\ F'_{x=M,a} \end{bmatrix}.$$

Because θ_f is known (or estimated), the elements of F_a are known (or estimated) in advance. We can now define the choice-specific value functions in terms of the flow utilities and the integrated value function:

$$v_{x,a} = u_{x,a} + \beta F'_{x,a} V.$$

Allowing V to violate the Bellman equation for the moment, these choice-specific value functions can be interpreted as a function of u and V. That is, $v = \nu(u, V)$ where $\nu : \mathbb{R}^{(M\times(J+1))\times M} \to \mathbb{R}^{M\times(J+1)}$ is linear in (u, V).

The social surplus function in observed state x is defined over the choice-specific value

[8] We do not discuss identification in this paper. Instead, we will assume that appropriate normalizations are imposed prior to estimation where necessary. See Magnac and Thesmar (2002) for a discussion of model identification. Identification of counterfactuals is discussed in Norets and Tang (2014); Kalouptsidi et al. (2017).

functions:

$$S_x(v) = \int \max_a \{v_{x,a} + \varepsilon_{x,a}\}\, g(\varepsilon|x, \theta_g) d\varepsilon.$$

Two important results of the Williams-Daly-Zachary Theorem are (i) $S_x(\cdot)$ is convex, and (ii) its gradient is the vector of conditional choice probabilities: $Pr(a|x, v) = \partial S_x(v)/\partial v_{x,a}$. Furthermore, it is clear by inspection that $S_x(\cdot)$ is *isotone*: if $v^1 \geq v^2$, then $S_x(v^1) \geq S_x(v^2)$. Stacking the values of $S_x(v)$ into an $M \times 1$ vector gives $S(v)$, which is also convex and isotone. Its derivative is given by $\nabla_v S(v) = \mathbf{P}(v)$ where $\mathbf{P}(v)$ has rows $P_x'(v)$ with elements

$$P_x(v)\{j\} = \begin{cases} Pr(a|x, v), & j = (y, a)\ s.t.\ y = x \\ 0, & j = (y, a)\ s.t.\ y \neq x. \end{cases}$$

A simple example helps illustrate the form of this matrix. Suppose that $(x, a) \in (\{1, 2\}, \{0, 1\})$ and let $v = (v_{1,0}, v_{1,1}, v_{2,0}, v_{2,1})$. Then we have

$$\boldsymbol{P}(v) = \begin{bmatrix} Pr(a = 0|x = 1) & Pr(a = 1|x = 1) & 0 & 0 \\ 0 & 0 & Pr(a = 0|x = 2) & Pr(a = 1|x = 2) \end{bmatrix}.$$

It is also useful to organize the non-zero elements of $\mathbf{P}(v)$ for all but some excluded choice into a $(M \times J)$ vector, $P(v)$. Without loss of generality, the excluded choice is $a = 0$. All choice probabilities are nonzero for finite v, so $P_x(v) \in \text{int}(\mathcal{S}^J)$ where $\mathcal{S}^J$ is the J-dimensional unit simplex. And we remark on some well-known properties of the social surplus and choice probability functions.

Remark 1. Let w be some vector such that $\kappa_{x,a} = \kappa_{x,a'} = \kappa_x$ for all x, a, a', where κ_x is a scalar. That is, the elements of w are constant across choices within a state. Then,

1. $S_x(v + \kappa) = S_x(v) + \kappa_x$.
2. $P(v + \kappa) = P(v)$.

We are now almost ready to introduce the representation of the integrated value function used in this paper. Define $T(u, V) \equiv S(\nu(u, V))$ so that $T : \mathbb{R}^{(M\times(J+1))\times M} \to \mathbb{R}^M$. Furthermore, define $G(u, V) \equiv V - T(u, V)$. Then the integrated value function, $V(u)$, can be defined implicitly as the solution to

$$G(u, V(u)) = 0.$$

The following proposition summarizes some important results for $T(\cdot)$, $G(\cdot)$, and $V(\cdot)$.

Proposition 1. *Under Assumptions 1-3, the following properties hold.*

- *T(u,V) is isotone and convex, with partial derivatives*

$$\nabla_u T = \mathbf{P}(\nu(u,V)), \quad \nabla_V T = \beta F^U(P(\nu(u,V))$$

where $F^U(P(\nu(u,V))$ is a matrix of unconditional transition probabilities with elements $F^U(P)\{i,j\} = P_{x,a}F_{x,a}(j) = Pr(a|x=i)Pr(x'=j|x=i,a)$.

- *G(u,V) is concave in (u,V), antitone in u, and inverse-isotone in V. Its partial derivatives are*

$$\nabla_u G = -\mathbf{P}(\nu(u,V)), \quad \nabla_V G = I - \beta F^U(P(\nu(u,V))).$$

- *V(u) is well-defined, surjective, isotone, and convex. Its derivative is*

$$\nabla_u V = \left(I - \beta F^U(P(\nu(u,V))\right)^{-1}\mathbf{P}(\nu(u,V)).$$

Furthermore, for any flow utility vector, $\overline{u}$, we can compute $\overline{V} = V(\overline{u})$ to arbitrary tolerance via iterative methods. Global convergence is guaranteed from any starting guess, V^0, using either of the following two methods:

1. *Bellman iterations: $V^{k+1} = T(\overline{u}, V^k)$.*
2. *Newton iterations: $V^{k+1} = V^k - \nabla_V G(\overline{u}, V^k)^{-1} G(\overline{u}, V^k)$.*

Most of the results of Proposition 1 are well-known. Global convergence of the Newton iterations was shown indirectly by Aguirregabiria and Mira (2002), although it can be proven directly using the properties of $G(\cdot)$. To the best of our knowledge, this is the first paper to establish that $V(u)$ convex. This property has important implications, which will be discussed later.

In addition to the representation considered so far, an alternative representation of the social surplus function can be used to obtain an alternative representation of the value function, as in Aguirregabiria and Mira (2002). The alternative social surplus function is:

$$S_x(v) = \sum_a P_{x,a}(v)(v_{x,a} + e_{x,a}(v))$$

where $e_{x,a}(v) = E[\varepsilon_{x,a}|x,a,v]$, the expectation of the unobservable $\varepsilon_{x,a}$ conditional on a being the optimal action in state x with choice-specific values function vector v. Stacking these again, we get

$$S(v) = \sum_a P_a(v) * (v_a + e_a(v)) = \mathbf{P}(v)\,[v + e]$$

where P_a, v_a, and e_a are appropriately defined vectors and $*$ is the Hadamard product. The second equality here is particularly important, as it translates the representation into notation used in Proposition 1. We can then define the (integrated) value function as

$$V(u) = \left[I - \beta F^U(P(u))\right]^{-1} \sum_d P_d(u) * [u_d + e_d(u)]$$

$$= \left[I - \beta F^U(P(u))\right]^{-1} \mathbf{P}(u) \left[u + e(u)\right],$$

with $P(u) = P(\nu(u, V(u)))$ and other terms are defined similarly. We note that for a square matrix, A, we have $(I - A)^{-1} = \sum_{t=0}^{\infty} A^t$ when the series converges. This Neumann series representation can be applied in our setting to obtain

$$V(u) = \sum_{t=0}^{\infty} \beta^t \left[F^U(P(u))\right]^t \mathbf{P}(u) \left[u + e(u)\right].$$

The formulas and definitions presented so far are valid for any u (and V, where appropriate). Of particular interest is the "true" flow utility vector; i.e., the one corresponding to the actual data-generating process. Let θ_u^* be the true structural parameter vector for the underlying DGP. The flow utility vector corresponding to this parameter is u^*, with value function and conditional choice probabilities $V^* = V(u^*)$ and $P^* = P(u^*)$, respectively.

3 Pseudo-Value Functions

In this section, we describe techniques for using CCPs to construct pseudo-value functions (PVFs). Some of these PVFs have been either explicitly or implicitly used in previous literature, while this is paper marks the first introduction of some others. In general, a PVF is a continuous, usually differentiable function that depends on both the flow utilities and conditional choice probabilities: $PVF(u, P)$. They are constructed so that $PVF(u, P(u)) = V(u)$, but in general $PVF(u, P) \neq V(u)$ when $P \neq P(u)$. This implies that $PVF(u^*, P^*) = V^*$, which is the key to consistency of (most) estimators that use PVFs instead of $V(u)$. The idea is, essentially, to utilize Slutsky's theorem to find that $V(u^*, \hat{P})$ is a consistent estimator of V^* when $\hat{P}$ is a consistent estimator of P^*. A well-known example of this is the PVF first introduced by Hotz and Miller (1993) and later used in the NPL estimator of Aguirregabiria and Mira (2002):

$$V_{NPL}(u, P) = \left(I - \beta F^U(P)\right)^{-1} \sum_d P_d * [u_d + e_d(P)]$$

$$= \left(I - \beta F^U(P)\right)^{-1} \mathbf{P} \left[u + e(P)\right],$$

where $\mathbf{P}$ is the appropriate matrix corresponding to choice probabilities, P.[9] We will eventually show that this is a special case of a more general PVF.

3.1 The Arcidiacono-Hotz-Miller Pseudo-Value Function (AHM-PVF)

This PVF is based on a representation of the value function similar reflecting theorems from Arcidiacono and Miller (2011; 2016).[10] The main difference is that while they consider a representation of the expected continuation value, we are interested in constructing the pseudo-value function. We begin by defining $\psi(P)$ such that $S(-\psi(P)) = 0$ and $\nabla_v S(-\psi(P)) = \mathbf{P}$, where $\mathbf{P}$ is the appropriate matrix corresponding to conditional choice probabilities, P. Chiong et al. (2016) show that ψ exists and is unique for any P, and they show how to compute it using a mass-transit approach. In some cases, a closed-form expression exists. In others, an alternative algorithmic approach can be used.

Algorithm 1. *We can compute ψ using the following steps:*

1. Set $v^0_{x,a=1} = 0$.
2. Use a (static) discrete choice inversion algorithm to compute $v^0_{x,a\neq 1}$ from P.
3. Compute $\psi_{x,a} = S_x(v^0) - v^0_{x,a}$.

Algorithm 1 can be used in settings where a closed-form is not readily available. For example, if ε has a random coefficients logit specification, then the contraction-mapping of Berry et al. (1995) can be used in Step 2. Furthermore, if the distribution of ε is log-concave, then Newton's method can be used as shown in [Authors' note: I want to reference a paper I have not yet finished here]. Neither these iterative approaches nor the mass-transit approach will compute ψ exactly, although they can become arbitrarily close. The iterative approaches require that the researcher choose a tight tolerance so that the iterations terminate close to the true value of ψ. The mass-transit approach uses a discretized grid, so ψ is set-identified, and the grid must be fine so that the points in this set are sufficiently close to the true value.

[9] This expression, $e(P)$, is valid because $e(\cdot)$ depends only on the differences in choice-specific value functions which, by the inversion theorem of Hotz and Miller (1993), are in turn a function of the choice probabilities.

[10] Theorem 1 in both papers.

Turning attention to the alternative representation of the social surplus function, we obtain a useful equality. First, notice that $P = P(-\psi(P))$, by construction. Furthermore, from the alternative representation of the social surplus function we obtain

$$0 = \boldsymbol{P}\left[-\psi(P) + e(P)\right].$$

Defining $e(P) = e(-\psi(P))$, this leads immediately to one of the results in the next proposition, which will be useful in the coming sections.

Proposition 2. *For a given vector of choice probabilities,* $\boldsymbol{P}\psi(P) = \boldsymbol{P}e(P)$ *and* $\boldsymbol{P}\nabla_P\psi(P) = 0$.

While it is not generally true that $e_{x,a}(P) = \psi_{x,a}(P)$, this does not matter for our purposes. In our implementations, only $\sum_a P_{x,a}e_{x,a}(P)$ is ever needed. Proposition 2 gives us a convenient way to compute this term. Proposition 2 also suggests that $e(P)$ can be replaced by $\psi(P)$ in procedures that appear to require individual values of $e_{x,a}(P)$ to be known. A particularly salient example is the simulation-based estimator of Hotz et al. (1994). Our analysis suggests that the asymptotic theory for their estimator will hold when $e(P)$ is replaced with $\psi(P)$, even if they are not equal. There is, however, one notable case in which the stronger condition that $e_{x,a}(P) = \psi_{x,a}(P)$ holds. It is described in the next example.[11]

Example 1. Suppose that ε is distributed i.i.d. Type 1 extreme value. Then, $e_{x,a}(P) = \psi_{x,a}(P) = \gamma - \ln(P_{x,a})$.

Now, recall that the value function satisfies $V = S(\nu(u, V))$. Along with Remark 1, this implies that

$$\psi_{x,a} = V_x - v_{x,a},$$

so that $\psi_{x,a}$ is the expected amount of present discounted utility given up by pre-committing to action a in state x in the current period (and only the current period) instead of following the optimal policy, as described in Arcidiacono and Miller (2011).[12] With $\psi(P)$ in hand, we can construct the PVF based on a sequence of pre-committed decision weights, rather than on optimal choice probabilities. The pre-commitment is for all periods, and the pre-committed choice weights need not be stationary; i.e., they can depend on both the state and time period. We index the current period as $t = 0$. Define $\omega_{t,x} = \{\omega_{t,x,a=0}, \ldots, \omega_{t,x,a=J}\}$ for $0 \leq t \leq \infty$ as the pre-committed decision weights, with $|\omega_{t,x,a}| < \infty$ and $\sum_a \omega_{t,x,a} = 1$ for all t, x, a. We also

[11] We provide an example where $e_{x,a}(P) \neq \psi_{x,a}(P)$ in the appendix.

[12] See Lemma 1 and the ensuing discussion in their paper.

impose that there is some $\bar{t} \geq 0$ such that $0 \leq \omega_{t,x,a} \leq 1$ for all $t \geq \bar{t}$. When the weights are non-negative, this has an intuitive interpretation: if the decision maker is in state x in period t, then they are pre-committed to choosing a with probability $\omega_{t,x,d}$, regardless of the realized value of $\varepsilon_{t,x}$. We can then define the time-dependent (weighted) transition matrix, $F^U(\omega_t)$ with elements $F^U(\omega_t)\{i,j\} = Pr(x_{t+1} = j | x_t = i, \omega_t) = \sum_a \omega_{t,x,a} f(x_{t+1} = j | x_t = i, a_t)$. We are now ready to define the Arcidiacono-Hotz-Miller PVF:

$$V_{AHM}(u, P, \{\omega\}_{t=0}^{\infty}) = \boldsymbol{W}_0 \left[u + \psi(P)\right] + \sum_{t=1}^{\infty} \left\{ \beta^t \left(\prod_{\tau=0}^{t-1} F^U(\omega_\tau) \right) \boldsymbol{W}_t \left[u + \psi(P)\right] \right\},$$

where $\boldsymbol{W}_t$ has the same dimensions as $\mathbf{P}$ but corresponds to choice probabilities ω_t. The restriction to $0 \leq \omega_{t,x,a} \leq 1$ after a finite number of time periods ensures that the series converges. The next theorem is central to the usefulness of the AHM-PVF.

Theorem 1. $V_{AHM}(u, P(u), \{\omega\}_{t=0}^{\infty}) = V(u)$.

The proof of Theorem 1 is nearly identical to the proof of Theorem 1 from Arcidiacono and Miller (2011) and is omitted. However, it is helpful to provide an intuitive discussion for non-negative weights, which can be interpreted as pre-committed choice probabilities. That is, the decision maker commits to choosing a_t in state x_t with probability $w_{t,a,x}$ regardless of the value of ε_t. Consider the term $\boldsymbol{W}_t \left[u + \psi(P(u))\right]$ and assume for momentary simplicity that $E[\varepsilon_t] = 0$. Then there are two parts to this expression. The first is $\boldsymbol{W}_t u$, which is the expected flow utility in period t from the pre-committed choice probabilities, ω_t. The second part, $\boldsymbol{W}_t \psi(P(u))$, is a correction term that adds back in the expected utility loss from the sub-optimal pre-commitment strategy. The presence of the correction term enforces the equality in the theorem.

The presence of the infinite sum renders the AHM-PVF computationally intractable in its general form. One solution to this problem is to consider a stationary sequence of choice probabilities, $\omega_t = \omega \in \Delta$, so that $F^U(\omega_t) = F^U(\omega)$ and $\boldsymbol{W}_t = \boldsymbol{W}$. We then have

$$\begin{aligned} V_{AHM}(u, P, \omega) &= \sum_{t=0}^{\infty} \left\{ \beta^t \left[F^U(\omega)\right]^t \boldsymbol{W} \left[u + \psi(P)\right] \right\} \\ &= \left(I - \beta F^U(\omega)\right)^{-1} \boldsymbol{W} \left[u + \psi(P)\right], \end{aligned}$$

which can be computed in finite time. In order to do so, we need to solve the linear system $\left[I - \beta F^U(\omega)\right] y = \boldsymbol{W} \left[u + \psi(P)\right]$ for y. When the state space is moderately-sized, so that the computational complexity of this operation is not an issue, we are then free to choose ω. Perhaps an obvious choice is $\omega = P$, and this turns out to be a convenient one.

Proposition 3. $V_{AHM}(u, P, \omega = P) = V_{NPL}(u, P)$.

Proof. With $\boldsymbol{W} = \boldsymbol{P}$, this follows from Proposition 2. □

Proposition 3 shows that the NPL-PVF is a special case of the AHM-PVF. Aguirregabiria and Mira (2002) show that the NPL-PVF can be used to construct a pseudo-likelihood estimator that is asymptotically efficient in single-agent models, with large computational savings over the (also efficient) NFXP maximum likelihood estimator. In this sense, the choice of $\omega = P$ is theoretically optimal.

When the state space is large, solving the linear system becomes cumbersome, requiring essentially $O(M^3)$ computation time without any additional restrictions on the form of $I - \beta F^U(\omega)$. Intriguingly, Bray (2018) devises a method of implementing the NPL-PVF that requires solving only one fixed point equation per iteration via contraction mappings, reducing the computational burden to $O(M^2)$. However, when M is large, this can still be very time consuming. Either way, choosing $\omega = P$ may not be feasible. Other choices of ω, however, can reduce the computational burden. Computation time is reduced from $O(M^3)$ to $O(M^2)$ if $F^U(\omega)$ is upper/lower triangular and can be reduced further to $O(M)$ if it is diagonal. The triangular case is particularly interesting. In some models, the state space can be ordered to yield triangularity under some decision path. One such example appears in Rust (1987). The state variable is mileage. $F^U(\omega)$ is then upper-triangular when the state space is ordered by increasing mileage and non-replacement is assigned a weight equal to 1 in every period. Including engine age in the observed state does not change this result. In this case, $x_t = (age_t, mileage_t)$, and age is non-decreasing when non-replacement is selected.[13] Ordering states lexicographically then preserves triangularity. Significant computational savings may be possible in models with both endogenous and exogenous variables determining the state, as long as the states can be ordered so that $F(\omega)$ has a convenient form for some ω.

A fully analytic expression for $V_{AHM}(u, P, \{\omega\}_{t=0}^{\infty})$ is also available in some cases. Two examples are those of a renewal action, as in Rust (1987), and terminal actions, as in Hotz and Miller (1993).

Example 2. (Renewal Action) Suppose that $a = 1$ is a renewal action: $F'_{x,a=1} = F'_{x',a=1}$ (i.e., all the rows of $F_{a=1}$ are identical). This implies that $GF_{a=1} = F_{a=1}$ for any $M \times M$ row-stochastic matrix, G. In particular, $F^k_{a=1} = F_{a=1}$. Using Neumann series representation,

[13] Finite state space implies that age must have a maximum value. I.e., "15 years and over" could be the highest age state.

it is straightforward to derive that

$$(I - \beta F_{a=1})^{-1} = I + \frac{\beta}{1-\beta} F_{a=1}.$$

So, defining $\omega_{a=1}$ to be the pre-committed path where $a = 1$ is chosen in every period, we obtain $V_{AHM}(u, P, \omega_{a=1}) = \left(I + \frac{\beta}{1-\beta} F_{a=1}\right)(u_{a=1} + \psi_{a=1})$. Furthermore, we can ignore $\frac{\beta}{1-\beta} F_{a=1}$ when calculating differences in expected continuation values:

$$\beta F_a V_{AHM}(u, P, \omega_{a=1}) - \beta F_{a'} V_{AHM}(u, P, \omega_{a=1}) = \beta(F_a - F_{a'})(u_{a=1} + \psi_{a=1}).$$

Only differences in the expected continuation values will matter in estimation, so it is not a problem to ignore $\frac{\beta}{1-\beta} F_{a=1}$ entirely.

Example 3. (Terminal Action) Suppose that $a = 1$ is a terminal action (i.e., exit or sterilization). We'll need to adjust the framework a bit here. Designate $x = M$ as the terminal state. If $a = 1$ is chosen at any state, then the terminal state is reached in the next period with probability equal to one. Once in the terminal state, the decision maker remains in the terminal state with probability equal to one next period no matter what choice is made. The distribution of $\varepsilon_{x=M}$ is degenerate at zero (i.e., there is no random component). Flow utilities are $u_{x=M,a=1} = 0$ and $u_{x=M,a\neq 1} = -\infty$. These conditions ensure that $a = 1$ is always chosen in the terminal state, implying that $V_{x=M} = 0$. The social surplus function for the terminal state is defined as $S_{x=M}(v) = v_{x=M,a=1}$.[14] Then, $\psi_{x=M,a=1} = 0$ and we can define $\psi_{x=M,a\neq 1} = -\infty$ for internal consistency. The setup here implies that $F_{a=1}$ has all ones in the last column and zeros everywhere else. So, $F_{a=1}(u_{a=1} + \psi_{a=1}) = 0$ and Neumann series representation can be used to obtain

$$\begin{aligned} V_{AHM}(u, P, \omega_{a=1}) &= (I - \beta F_{a=1})^{-1}(u_{a=1} + \psi_{a=1}) \\ &= u_{a=1} + \psi_{a=1}. \end{aligned}$$

If we want to allow for $u_{x=M,a=1}$ to depend on the last realized state before termination, as in Hotz and Miller (1993), then we can simply add $\frac{\beta}{1-\beta}$ times that flow utility to $u_{x,a=1}$ instead.

We see in Example 2 that there may be no need to fully compute $V_{AHM}(\cdot)$ for estimation purposes. This occurs when the ρ-period finite-dependence property from Arcidiacono and Miller (2011) holds, with Examples 2 and 3 exhibiting single-period finite-dependence ($\rho =$

[14] Notice that we have $\partial S_{x=M}(v)/\partial v_{x=M,a=1} = 1 = Pr(a = 1|x = M)$ and $\partial S_{x=M}(v)/\partial v_{x,a} = 0$ for any other (x, a).

1). Their implementation uses (potentially) multiple different decision weight sequences to construct differences in expected continuation values. Formally, construct differences in expected continuation values as $F_{x,a}V_{AHM}(u, P, \{\omega^1\}_{t=0}^{\infty}) - F_{x,a'}V_{AHM}(u, P, \{\omega^2\}_{t=0}^{\infty})$, where the two decision weight sequences are chosen so that for all $t \geq \rho$,

$$F'_{x,a}\left(\prod_{\tau=0}^{t-1} F^U(\omega_\tau^1)\right)\boldsymbol{W}_t^1 - F'_{x,a=0}\left(\prod_{\tau=0}^{t-1} F^U(\omega_\tau^2)\right)\boldsymbol{W}_t^2 = 0.$$

So, only the terms corresponding to $0 \leq t \leq \rho - 1$ in each series are needed. Finding the finite-dependence paths is a non-trivial exercise, and Arcidiacono and Miller (2016) provide an algorithm for finding finite-dependence paths when they are not obvious from the structure of the model.

3.2 Zero Jacobian Property and Efficiency

We now turn attention to estimation efficiency results. Aguirregabiria and Mira (2002, Proposition 2) show that $\nabla_P V_{NPL}(u, P(u)) = 0$, dubbing this the "zero jacobian" property. and use this to establish asymptotic efficiency of the pseudo-maximum likelihood estimate when using $V_{NPL}(\cdot)$ and starting with a $\sqrt{N}$-consistent $\hat{P}_0$. In this section, we first provide an alternative characterization of $V_{NPL}(\cdot)$ which yields a more intuitive proof (in our opinion) of the zero jacobian property. Second, we derive an alternative zero jacobian condition for $V_{AHM}(\cdot)$ with a fixed weighting matrix, which facilitates a discussion of efficiency properties of different choices of ω.

We begin by showing that the NPL-PVF can be interpreted as a first-order approximation of $V(u)$. That is, we will take a first-order Taylor approximation of $V(u)$ around a particular $u^0(u, P)$ and show that it is equivalent to NPL-PVF. The first step is to construct $u^0(u, P)$. We do so by first obtaining a AHM-PVF and then constructing a flow utility u^0 such that $V(u^0) = V_{AHM}(u, P, \{\omega\}_{t=0}^{\infty})$ and $P(u^0) = P$. If the corresponding choice specific value function vector is v^0, then we can use the relationship

$$v_a^0 = u_a^0 + \beta F_a V_{AHM}.$$

The associated choice probabilities should be P, implying that $V_{AHM} - v^0 = \psi(P)$. Some simple operations then give

$$u_a^0(u, P, \{\omega\}_{t=0}^{\infty}) = -\psi_a(P) + (I - \beta F_a)\, V_{AHM}(u, P, \{\omega\}_{t=0}^{\infty}).$$

From here, we will drop some of the function arguments for parsimony. This yields the

first-order PVF:

$$\begin{aligned} V_{FOA}(u,P) &= \nabla_u V(u^0(u,P)) \left[u - u^0(u,P)\right] + V(u^0(u,P)) \\ &= (I - \beta F^U(P))^{-1}\mathbf{P}\left[u - u^0(u,P)\right] + V(u^0(u,P)). \end{aligned}$$

The second equality follows from Proposition 1 and the construction of u^0. Notice that we have left $\{\omega\}_{i=0}^{\infty}$ out of these expressions. This is because FOA-PVF is actually equivalent to the NPL-PVF, so that the choice of initial weights does not matter. The next proposition formally states this equivalence and shows that they (weakly) underestimate the true value function.

Proposition 4. *$V_{NPL}(u,P) = V_{FOA}(u,P) \leq V(u)$ for all (u,P), with equality holding at $(u,P(u))$.*

Proof. The inequality follows immediately from convexity of $V(\cdot)$ (Proposition 1) and the form of $V_{FOA}(\cdot)$. This becomes an equality at $P(u)$ because $u^0(u,P(u)) = u$. The left equality follows by implementing the alternative representation of $V(u^0(u,P))$ and noticing some terms conveniently cancelling out:

$$\begin{aligned} V_{FOA}(u,P) &= (I - \beta F^U(P))^{-1}\mathbf{P}\left[u - u^0(P)\right] + V(u^0(u,P)). \\ &= (I - \beta F^U(P))^{-1}\mathbf{P}\left[u - u^0(u,P)\right] + (I - \beta F^U(P))^{-1}\mathbf{P}\left[u^0(u,P) + e(P)\right] \\ &= (I - \beta F^U(P))^{-1}\mathbf{P}u + (I - \beta F^U(P))^{-1}\mathbf{P}e(P) \\ &= V_{NPL}(u,P). \end{aligned}$$

□

Proposition 4 has several important implications. First, it gives us an intuitive interpretation of $V_{NPL}(u,P)$: it is a first-order approximation of $V(u)$ around some vectors of flow utilities. Those flow utility vector are jointly determined by u and P. To our knowledge, we are the first to characterize $V_{NPL}(u,P)$ in this way. This is also, to our knowledge, the first explicit proof of the result that $V_{NPL}(u,P) \leq V(u)$.

Proposition 4 also admits a very straightforward and intuitive proof of the zero jacobian property for the NPL-PVF. From the proposition, we have $P(u) = \text{argmax}_{P \in int(\mathcal{S}^J)}\{V_{NPL}(u,P)\}$. The set $int(\mathcal{S}^J)$ is open, so the first-order necessary condition requires $\nabla_P V_{NPL}(u,P(u)) = 0$. The intuition given by Aguirregabiria and Mira (2002, pp. 1526) is that "the optimal choice probabilities maximize the valuation operator locally." Our analysis strengthens this intuition: the optimal choice probabilities maximize this (pseudo-)value function operator *globally*.

While this analysis is illuminating, it does not directly shed light on the efficiency properties of pseudo-maximum likelihood estimation with the general AHM-PVF. One important distinction of the NPL-PVF as a special case of the AHM-PVF is that it requires ω to change with P. Then, efficiency in pseudo-maximum likelihood estimation seems to require ω to change with $\hat{P}_0$. However, other implementations of the AHM-PVF – such as one exploiting a terminal action – use a fixed ω that does not vary with P. Can we characterize conditions for efficiency of pseudo-maximum likelihood in this case? Perhaps surprisingly, the answer turns out to be yes.

The efficiency result of Aguirregabiria and Mira (2002) with $\sqrt{N}$-consistent $\hat{P}_0$ depends on two key features: (i) $\nabla_u V_{NPL}(u^*, P^*) = \nabla_u V(u^*)$, and (ii) $\nabla_P V_{NPL}(u^*, P^*) = 0$. These results are then used in standard first-order expansions to prove the asymptotic equivalence to the (partial) maximum likelihood estimator. We would need similar results for fixed ω. That is, we wish to find some ω such that $\nabla_u V_{AHM}(u^*, P^*, \omega) = \nabla_u V(u^*)$ and $\nabla_P V_{AHM}(u^*, P^*, \omega) = 0$. Such an ω does indeed exist, which we describe in the next proposition.

Proposition 5. *If $\omega = P^*$, then $\nabla_u V_{AHM}(u^*, P^*, \omega) = \nabla_u V(u^*)$ and $\nabla_P V_{AHM}(u^*, P^*, \omega) = 0$. Furthermore, the pseudo-maximum likelihood estimate using $V_{AHM}(u(\theta_u), \hat{P}_0, P^*)$ is asymptotically equivalent to the (partial) maximum likelihood estimator when $\hat{P}_0$ is $\sqrt{N}$-consistent.*

Proposition 5 establishes conditions for a zero jacobian property with fixed ω. Perhaps unsurprisingly, setting $\omega = P^*$ does the trick. This result is helpful, as it allows us to assess the potential efficiency loss from $\omega \neq P^*$. For values of ω that are near P^*, the efficiency loss should be minimal, whereas efficiency loss will be greater when ω is not close to P^*. However, the researcher will not know P^* from the outset. Instead, the researcher can compare the induced $\hat{P}_1$ with ω. If they are close, then the efficiency loss should indeed be minimal. If not, then it may still be possible to quantify the efficiency loss, but that is beyond the scope of this paper.

4 Closed-Form Estimation

So far, we have discussed estimation in the context of pseudo-maximum likelihood. In this section, we develop a closed-form estimator that is asymptotically equivalent to pseudo-maximum likelihood. Before constructing the closed-form estimator, we first introduce some additional notation. For notational simplicity, let $\theta_u = \theta$. Although we defined choice probabilities in Section 2 as a function of the choice-specific value functions, we note that choice probabilities actually only depend on *differenced* value functions. That is, $P(v) =$

$\Lambda(\tilde{v})$, where $\tilde{v}_a = v_a - v_0$ for $a \in \mathcal{A}\backslash\{0\}$. An important property of $\Lambda(\cdot)$ that it is invertible (see Hotz and Miller (1993) and Berry et al. (2013)). We can then define a function mapping θ and P to differenced value functions. Define this as

$$\varphi_a(\theta, P) = \tilde{u}_a(\theta) + \beta \left(F_a V_{AHM}(u, P, \{\omega^a\}_{t=0}^{\infty}) - F_0 V_{AHM}(u, P, \{\omega^0\}_{t=0}^{\infty}) \right)$$

for $a \in \mathcal{A}\backslash\{0\}$, where $\tilde{u}_a(\theta) = u_a(\theta) - u_0(\theta)$. The function is constructed in order to accommodate finite-dependence. And we can also define $\Psi(\theta, P) = \Lambda(\varphi(\theta, P))$. It is also useful to define $P(\theta) = P(u(\theta))$, etc. Furthermore, we note a few useful results.

Proposition 6. *The following properties hold for all θ:*

1. If $u(\theta)$ is linear in θ, then $\varphi(\theta, P)$ is linear in θ.
2. $\tilde{v}(\theta) = \varphi(\theta, P(\theta))$ and $P(\theta) = \Psi(\theta, P(\theta))$.
3. If $\varphi(u, P)$ is constructed with $\omega_t^0 = \omega_t^1 = \cdots = \omega_t^J = P$ for all $t \geq 0$, so that all pseudo-value functions equal to $V_{NPL}(u, P)$, then $\nabla_P \varphi(\theta, P(\theta)) = 0$ and $\nabla_P \Psi(\theta, P(\theta)) = 0$.

Proof. The first result arises because $V_{AHM}(u, P, \omega)$ is linear in u. The second arises because $V_{AHM}(u(\theta), P(u(\theta))) = V(u(\theta))$ by Proposition 4. The third follows immediately from Proposition 2 in Aguirregabiria and Mira (2002). □

4.1 The Estimator

We motivate our estimator with a minimum-distance estimator à la Pesendorfer and Schmidt-Dengler (2008):

$$\hat{\theta}_{MD} = \arg\min_{\theta \in \Theta} \quad \left(\hat{P} - \Psi(\theta, \hat{P}_0) \right)' \hat{W} \left(\hat{P} - \Psi(\theta, \hat{P}_0) \right),$$

The first order condition for this problem is

$$\nabla_\theta \Psi' \hat{W} \left(\hat{P} - \Psi(\theta, \hat{P}_0) \right) = 0.$$

Now, suppose that $\hat{P}_0$ is near $\hat{P}$ and that $\tilde{v}(\theta)$ is near $\Lambda^{-1}(\hat{P})$ in a neighborhood of $\hat{\theta}$. Then we can use a first-order Taylor approximation to get

$$\Psi(\theta, \hat{P}_0) \approx \hat{P}_0 + \widehat{\nabla_{\tilde{v}} \Lambda_0} \left[\varphi(\theta, \hat{P}) - \Lambda^{-1}(\hat{P}_0) \right],$$

where $\widehat{\nabla_{\tilde{v}}\Lambda_0} = \nabla_{\tilde{v}}\Lambda(\Lambda^{-1}(\hat{P}_0))$. We note that $\widehat{\nabla_{\tilde{v}}\Lambda_0}$ is symmetric and positive definite because $\Lambda(\cdot)$ is the derivative of social surplus function evaluated at utility differences, which is strictly convex. So, $\nabla_{\tilde{v}}\Lambda$ is the Hessian of a strictly convex function. With this, we can propose an alternative estimator:

$$\begin{aligned}\hat{\theta}_{CF} = \arg\max_{\theta\in\Theta} \quad & \left(\hat{P} - \hat{P}_0 - \widehat{\nabla_{\tilde{v}}\Lambda_0}\left[\varphi(\theta,\hat{P}_0) - \Lambda^{-1}(\hat{P}_0)\right]\right)' \hat{W} \\ & \times \left(\hat{P} - \hat{P}_0 - \widehat{\nabla_{\tilde{v}}\Lambda_0}\left[\varphi(\theta,\hat{P}_0) - \Lambda^{-1}(\hat{P}_0)\right]\right).\end{aligned}$$

If $u(\theta)$ is linear in θ, then $\varphi(\theta,\hat{P})$ is linear in θ by Proposition 1, and therefore $\hat{\theta}_{CF}$ has a closed-form expression. That is, suppose that $\tilde{u}(\theta) = X\theta$ and $\varphi(\theta,P) = H(P)\theta + Z(P)$ and let $\hat{H}_0 = H(\hat{P}_0)$ and $\hat{Z}_0 = Z(\hat{P}_0)$. Then, it is straightforward to show that

$$\hat{\theta}_{CF} = \left(\hat{H}_0'\widehat{\nabla_{\tilde{v}}\Lambda_0}W\widehat{\nabla_{\tilde{v}}\Lambda_0}\hat{H}_0\right)^{-1}\hat{H}_0'\widehat{\nabla_{\tilde{v}}\Lambda_0}\hat{W}\left(\hat{P} - \hat{P}_0 - \widehat{\nabla_{\tilde{v}}\Lambda_0}\left[\hat{Z}_0 - \Lambda^{-1}(\hat{P}_0)\right]\right).$$

The key here is that none of the terms directly depend on θ. In particular, $\widehat{\nabla_{\tilde{v}}\Lambda_0}$ can be computed with only $\hat{P}_0$.

Our estimator is also a minimum-distance estimator. Although by rearranging terms in the objective function we can see that, instead of minimizing a distance in choice probability space, our estimator minimizes distance in differenced choice-specific value function space by utilizing the weighting matrix $\tilde{W} = \widehat{\nabla_{\tilde{v}}\Lambda}\hat{W}\widehat{\nabla_{\tilde{v}}\Lambda}$. That is, consider minimizing the objective function $(\Lambda^{-1}(\hat{P}) - \varphi(\theta,\hat{P}_0))'\tilde{W}(\Lambda^{-1}(\hat{P}) - \varphi(\theta,\hat{P}_0))$. A problem arises here if $\hat{P}$ is the frequency estimator because it may be on the boundary of $\mathcal{S}^J$, where $\Lambda^{-1}(\cdot)$ is undefined. We deal with this issue by using the identity $\Lambda^{-1}(\hat{P}) - \varphi(\theta,\hat{P}_0) = \Lambda^{-1}(\hat{P}) - \Lambda^{-1}(\hat{P}_0) - (\Lambda^{-1}(\hat{P}_0) - \varphi(\theta,\hat{P}_0))$ and the approximation $\Lambda^{-1}(\hat{P}) - \Lambda^{-1}(\hat{P}_0) \approx \widehat{\nabla_{\tilde{v}}\Lambda_0}^{-1}(\hat{P} - \hat{P}_0)$. Plugging in these approximations and performing some matrix algebra then gives the objective function used to compute $\hat{\theta}_{CF}$. And importantly, our closed-form estimator is asymptotically equivalent to the minimum-distance estimator under some regularity conditions.

Theorem 2. *Under standard regularity assumptions, $\hat{\theta}_{CF}$ and $\hat{\theta}_{MD}$ are asymptotically equivalent.*

What remains now is to use the appropriate choice of weighting matrix to make this estimator asymptotically equivalent to the pseudo-maximum likelihood estimator. Pesendorfer and Schmidt-Dengler (2008) show equivalence of $\hat{\theta}_{MD}$ and $\hat{\theta}_{PML}$ if $\hat{P}$ is the sample decision frequency and an appropriately chosen weighting matrix. The weighting matrix is block diagonal, with a separate block for each state. Each block is the covariance matrix of a multinomial random variable with choice probability vector $\Psi_x(\hat{\theta}_{PML},\hat{P}_0)$, multiplied by the

fraction of observations in which the state is visited in the sample.[15] Note that states not observed in the data receive zero weight, so we will need at least as many observed states as there are variables to be estimated in order to satisfy rank conditions. Furthermore, $\hat{P}$ can take values on the boundary of the unit simplex, whereas $\hat{P}_0$ must reside in the interior. Similarly, $\hat{\theta}_{CF}$ is asymptotically equivalent to $\hat{\theta}_{MD}$, we can make it asymptotically equivalent to $\hat{\theta}_{PML}$ by appropriate choice of block diagonal weighting matrix, with a separate diagonal for each state. However, because we do not know $\hat{\theta}_{PML}$ in advance, each block of the weighting matrix is the covariance matrix of $\hat{P}_{0x}$, denoted $\hat{\Sigma}_{px}^{-1}$, multiplied by the fraction of observations with with the state is visited in the sample. That is, each block is $\hat{Q}_x\hat{\Sigma}_{px}^{-1}$ where $\hat{Q}_x = (n_x/N)I$ where I is the identity matrix. The full weighting matrix is consequently $\hat{W}_{PL} = \hat{Q}\hat{\Sigma}_p^{-1} = \hat{\Sigma}_p^{-1}\hat{Q}$.[16] This gives what we call the "closed-form pseudo-likelihood" estimator:

$$\hat{\theta}_{CFPL} = \left(\hat{H}_0'\widehat{\nabla_{\tilde{v}}\Lambda_0}'\hat{W}\widehat{\nabla_{\tilde{v}}\Lambda_0}\hat{H}_0\right)^{-1}\hat{H}_0'\widehat{\nabla_{\tilde{v}}\Lambda_0}\hat{W}_{PL}\left(\hat{P} - \hat{P}_0 - \widehat{\nabla_{\tilde{v}}\Lambda_0}\left[\hat{Z}_0 - \Lambda^{-1}(\hat{P}_0)\right]\right).$$

Intuitively, our estimator is very similar to taking a single Newton-Raphson step on the minimum-distance objective function in probability space. With the choice of weighting matrix, this is then essentially equivalent to taking a Newton-Raphson step on the pseudo-likelihood function. However, instead of taking the first-order approximation around an initial $\hat{\theta}_0$, as would be done in the Newton-Raphson step, we instead take a first-order approximation around $\Lambda^{-1}(\hat{P}_0)$. With a strongly $\sqrt{N}$-consistent $\hat{P}_0$, this turns out to be sufficient for asymptotic equivalence to MLE, similar to how a Newton-Raphson or scoring step on the likelihood function from a strongly $\sqrt{N}$-consistent $\hat{\theta}_0$ is asymptotically equivalent to MLE.

Proposition 7. *Under standard regularity assumptions, when $\varphi(u, P)$ is constructed with $\omega_t^0 = \omega_t^1 = \cdots = \omega_t^J = P$ for all $t \geq 0$, so that all pseudo-value functions equal to $V_{NPL}(u, P)$, the estimators $\hat{\theta}_{CFPL}$ and $\hat{\theta}_{MLE}$ are asymptotically equivalent.*

While we have derived the estimator for a general number of actions and distribution of the utility shocks, we would like to take this opportunity to point out that the expression for the estimator can be simplified in certain special cases. The most notable of these is when there are binary actions and i.i.d. Type 1 extreme value utility shocks (binomial logit). In this case, the estimator can be simplified a bit.

Remark 2. In the case of binary actions and i.i.d. Type 1 Extreme Value utility shocks,

[15] Pesendorfer and Schmidt-Dengler (2008) show how to construct the covariance matrix in their appendix.

[16] Both $\hat{Q}$ and $\hat{\Sigma}_p^{-1}$ are symmetric.

$\widehat{\nabla_{\tilde{v}}\Lambda}\hat{\Sigma}_0^{-1} = I$ and therefore

$$\hat{\theta}_{CFPL} = \left(\hat{H}_0' Q \widehat{\nabla_{\tilde{v}}\Lambda}\hat{H}_0\right)^{-1} \hat{H}_0' Q \left(\hat{P} - \hat{P}_0 - \widehat{\nabla_{\tilde{v}}\Lambda}\left[\hat{Z}_0 - \Lambda^{-1}(\hat{P}_0)\right]\right)$$

is asymptotically equivalent to pseudo-maximum likelihood.

Proof. We have $\widehat{\nabla_{\tilde{v}}\Lambda} = \text{diag}(\hat{P}_0(1-\hat{P}_0))$ and $\hat{\Sigma}_0^{-1} = \text{diag}(1/\hat{P}_0 + 1/(1-\hat{P}_0)) = \text{diag}(1/\hat{P}_0(1-\hat{P}_0))$. So, $\widehat{\nabla_{\tilde{v}}\Lambda}\hat{\Sigma}_0^{-1} = I$ and $\widehat{\nabla_{\tilde{v}}\Lambda}W = \widehat{\nabla_{\tilde{v}}\Lambda}\hat{\Sigma}_0^{-1}Q = Q$. □

Remark 2 establishes that the term $\widehat{\nabla_{\tilde{v}}\Lambda}\hat{\Sigma}_0^{-1}$ can be omitted from computation in a binary logit setting. Furthermore, the form of the estimator in this case bears a striking similarity to the closed-form estimator of Bray (2018). [Authors' note: Indeed, we suspect that our estimator is identical to Bray's in the logit case, although we have not yet formally established such a relationship. It would be remarkable for such a relationship to hold, as Bray uses substantially different methods to motivate and derive his closed-form estimator.]

4.2 Iterating the Estimator

Having established the asymptotic equivalence to maximum pseudo-likelihood, we now turn attention to showing that our estimator can be iterated upon to get the actual maximum likelihood estimates. We begin by discussing iteration to (partial) the maximum likelihood estimator, as we believe this to be both of greater interest and a more intuitive procedure overall. In order to do so, it is first useful to introduce the notation $\hat{\theta}_{CFPL} = \Phi(\hat{P}, \hat{P}_k)$, with $k = 0$ in the analysis so far. Furthermore, we can define $\hat{P}_k = \Psi(\Phi(\hat{P}, \hat{P}_{k-1}), \hat{P}_{k-1})$ for $k \geq 1$. The next proposition then establishes that iterations can be used to compute the MLE.

Proposition 8. *Consider a fixed sample. When $\varphi(u, P)$ is constructed with $\omega_t^0 = \omega_t^1 = \cdots = \omega_t^J = P$ for all $t \geq 0$, so that all pseudo-value functions equal to $V_{NPL}(u, P)$, the following properties hold for the iterated CFPL estimator:*

1. *$\hat{\theta}_{MLE} = \Phi(\hat{P}, \hat{P}_{MLE})$ and $\hat{P}_{MLE} = \Psi(\Phi(\hat{P}, \hat{P}_{MLE}), \hat{P}_{MLE})$.*
2. *There exists some neighborhood, $\mathcal{B}$, of $\hat{P}_{MLE}$ such that for any from any $\hat{P}_0 \in \mathcal{B}$, we have $||\hat{\theta}_k - \hat{\theta}_{MLE}|| \leq c_1||\hat{P}_{k-1} - \hat{P}_{MLE}||^2$ and $||\hat{P}_{k-1} - \hat{P}_{MLE}|| \leq c_2||\hat{P}_{k-1} - \hat{P}_{MLE}||^2$ for some constants, c_1 and c_2.*

The first result of Proposition 2 establishes that $(\hat{\theta}_{MLE}, \hat{P}_{MLE})$ is a fixed point of the iterates. The second result establishes local quadratic convergence to the fixed point. These results are not asymptotic but rather for a fixed sample. Together, these results show that iterating on the CFPL provides a computationally light algorithm for computing the MLE.

5 Monte-Carlo Simulations

We now describe some Monte-Carlo simulations to illustrate the performance of the various PVFs in estimation procedures. We focus on pseudo-likelihood estimation, where $V(u)$ is replaced by PVFs in the log-likelihood function. We limit our analysis to estimating θ_u^* and use the true values of all other parameters in the estimation. We follow Su and Judd (2012) in generating a model using mostly parameters from Table X in Rust (1987), beginning by discretizing the mileage state space into 175 grid points so that $x \in \{0, 1, \ldots, 174\}$. This model is relatively small and does not fully capture the disparities in computational times for estimators using different PVFs, although some significant differences do still arise. On the other hand, low computation time allows us to run extensive comparisons between a wide variety of PVFs, focusing mainly on bias and root mean squared error of the estimates. There are two decisions, $a \in \{0, 1\}$, with $a = 1$ corresponding to replacement. Flow utilities are parameterized as

$$u_{x,a} = \begin{cases} \theta_{u,0} + \theta_{u,1}(0.001x), & a = 0 \\ 0, & a = 1 \end{cases}$$

with $(\theta_{u,0}^*, \theta_{u,1}^*) = (11.7257, -2.4569)$, and the transition parameter vector is

$$\theta_f^* = (0.0937, 0.4475, 0.4459, 0.0127, 0.0002).$$

We set the discount factor to $\beta = 0.975$. The unobserved shocks are distributed i.i.d Type 1 Extreme Value. All estimation routines use a log-likelihood objective function.

To estimate the parameters, we simulate 10,000 choices and transitions for a single bus. The initial state is drawn randomly from the stationary distribution of the model. The bias and MSE are computed using 1,000 replications. We obtain first-stage estimates of the CCPs via a binary logit model with a second-degree polynomial in x. The choice probabilities are quite close to the true values, as illustrated in Figure 1.

Table 1 presents results for several estimation strategies. Algorithms are either pseudo-maximum likelihood (PML) or its closed-form counterpart (CFPL). The columns for ω =MLE use iterations on the corresponding estimators to compute the maximum likelihood estimator. We terminate when either $||\hat{\theta}_k - \hat{\theta}_{k-1}||_\infty < 10^{-6}$. No problems with convergence were encountered. [17] We see that both methods produce identical bias and mean squared errors but that iterations on CFPL are about an order of magnitude faster. In other cases, we do not iterate. All PML algorithms use the true parameters as starting values for $k = 1$, and MLE uses $\hat{\theta}_{k-1}$ as starting values for $k \geq 1$. We characterize the AHM-PVF used in

[17] In pseudo-likelihood estimation, we use MATLAB's "fminunc" solver with an analytic gradient.

Figure 1: Replacement Probabilities for 500th Replication: Polynomial Degree 2

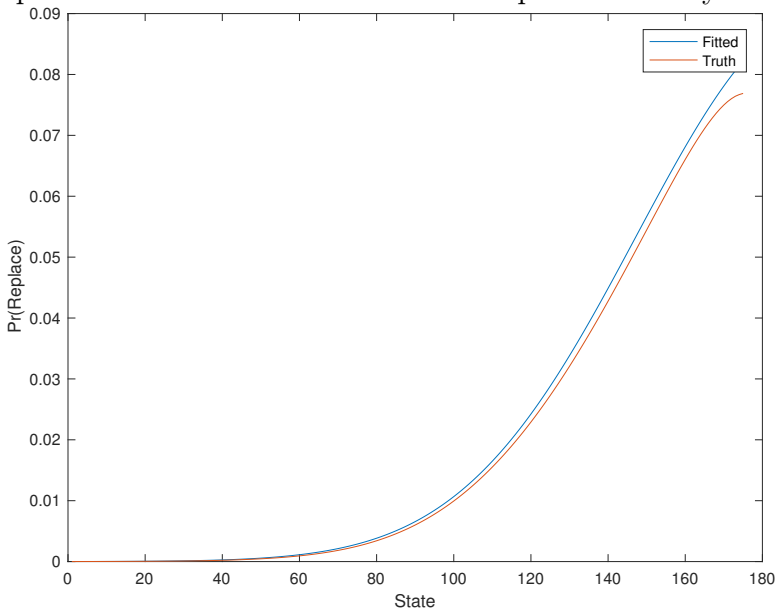

estimation using a small abuse of notation, with ω being the weight on replacement. So, $\omega = 1$ corresponds to always replacing as in Example 2, and we can obtain an analytic inverse of $I - \beta F^U(\omega = 1)$. Estimation times this case is fast, as expected. Choosing $\omega = \widehat{P}_0$ corresponds to the case where the PVF is $V_{NPL}(\cdot)$.

We see relatively little difference between MLE and using $\omega = \widehat{P}_0$ in both PML and CFPL. In this application, we obtain better performance from both PML and CFPL $\omega = \widehat{P}_0$ than MLE in terms of both bias and mean squared error. Although we do not report results for larger sample sizes, we found that difference in estimates vanishes as the sample size grows, which is expected because they are asymptotically equivalent. Intriguingly, there is little difference between the PML and CFPL estimates for both $\omega = \hat{P}_0$ and $\omega = 1$. This is likely an artifact of having $\hat{P}_0$ relatively close to P^*. We expect that there would be larger differences with poor starting values, and so we address this issue in further simulations.

We perform further simulations in order to investigate the impact of using poor starting values. To do so, we re-estimate $\hat{P}_0$ with a binary logit model with only a constant. This is equivalent to setting $\hat{P}_{0,x} = \bar{P}_0$ to the average replacement probability across *all* observations. We chose this starting value because it is in some sense a "worst case scenario" for the initial estimates. Figure 2 illustrates the new estimates. Having obtained starting values, we then set $\omega = \hat{P}_0$ and consider iteratively estimating $\hat{\theta}_k$ for $k \in \{1, \ldots, 5, \infty\}$ with both PML and CFPL.

Table 2 contains the resulting bias and mean squared error for $\theta_{u,0}$. Perhaps surprisingly, we find that both estimators still converge to the MLE ($k = \infty$) even from the poor starting

Table 1: Bias and MSE of $\hat{\theta}_u$ with Good Starting Values

		$\hat{\theta}_{u,0}$		$\hat{\theta}_{u,1}$		
Method	ω	Bias	MSE	Bias	MSE	Time (min.)
PML	MLE	0.2238	1.0484	-0.0641	0.0915	1.8159
CFPL		0.2238	1.0484	-0.0641	0.0915	0.1991
PML	$\hat{P}_0$	0.1649	1.1059	-0.0415	0.0885	0.6710
CFPL		0.1279	1.0122	-0.0338	0.0886	0.1346
PML	1	0.5110	2.7256	-0.0653	0.0930	0.4819
CFPL		0.5110	2.7257	-0.0656	0.0930	0.1273
Truth		11.7257		-2.4569		

Figure 2: Replacement Probabilities for 500th Replication: Polynomial Degree 0

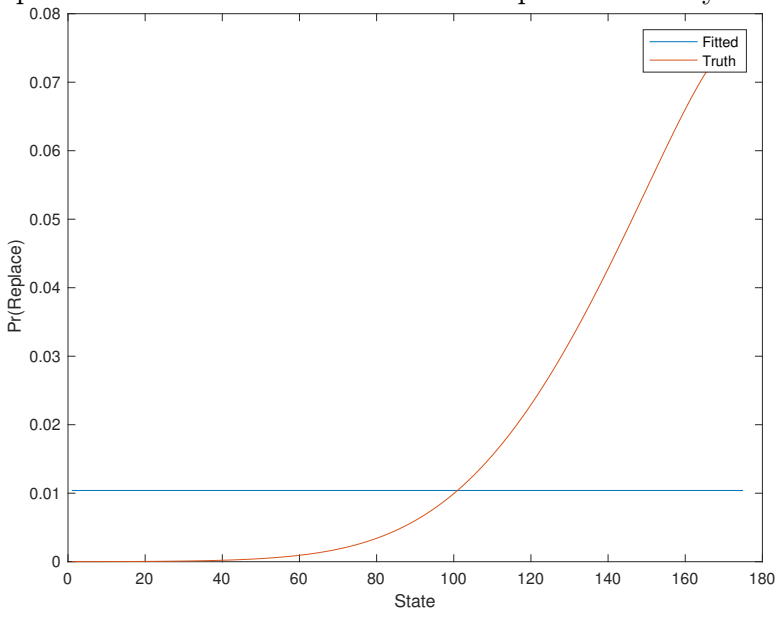

Table 2: Bias and MSE of $\hat{\theta}_{u,0}$ with $\omega = \hat{P}_0$ and Poor Starting Values

Method	Statistic	k					
		1	2	3	4	5	∞
PML	Bias	-1.8590	0.2158	0.2240	0.2238	0.2238	0.2238
	MSE	3.7800	1.0425	1.0486	1.0484	1.0484	1.0484
CFPL	Bias	-4.8444	-2.9241	-1.1636	-0.0638	0.2089	0.2238
	MSE	23.4747	8.5943	1.5911	0.6869	1.0113	1.0484

value of $\bar{P}_0$ for all states. However, we find substantial differences in their bias and MSE with a small number of iterations, with PML performing substantially better. Furthermore, while most of the bias and efficiency gains for PML come within the first two or three iterations, the convergence takes more iterations for CFPL.[18] However, the overall time to convergence reported in Table 1 shows that iterating to convergence with CFPL is ultimately faster, as each iteration takes far less time. In fact, Table 1 shows that iterating to convergence on CFPL is faster than even a single iteration of PML. Which of these methods the researcher wishes to use is then a matter of initial conditions. For small state spaces, we suggest iterating on CFPL until convergence, as it provides a very fast method to compute the MLE and is surprisingly robust to poor starting values. As the state space grows, a larger share of the computational burden arises from the matrix inversion needed to compute the pseudo-value function, whereas the actual PML estimation is still only as burdensome as estimating a static binary logit. In this case, fewer iterations of any sort will be preferred and we recommend iterating on PML.

6 Conclusion and Future Extensions

This paper makes two contributions to the literature on estimation of dynamic discrete choice models. The first is a unification result for two frameworks for conditional choice probability (CCP) estimation, which have until now been treated as distinct. A byproduct of this unification is that we broaden the class of models to which the nested pseudo-likelihood estimator of Aguirregabiria and Mira (2002) can be applied. The unification result also enables us to discuss asymptotic (in)efficiency properties of, for example, estimators based on finite-dependence-type properties.

Our second contribution is proposing a method of closed-form estimation that is asymptotically equivalent the minimum distance estimator of Pesendorfer and Schmidt-Dengler

[18] Although it is not reported in the table, we found that convergence took fewer than 10 iterations for CFPL.

(2008). We also show that this closed-form estimator is asymptotically equivalent to (partial) MLE with the proper choice of pseudo-value function and weighting matrix. Furthermore, we show that iterating on the estimator produces the (partial) MLE, providing a very fast computational algorithm.

Our analysis is subject to some limitations. Perhaps most importantly, we do not allow for serial correlation in the unobservables in our framework, which seemingly rules out time-invariant unobserved heterogeneity. In some settings, incorporating such heterogeneity reduces to simply estimating more linear parameters and expanding the state space, which can be easily accommodated within our framework. In other cases, a researcher may wish to use EM-type algorithms (see Arcidiacono and Miller (2011)). While it may be possible to use our closed-form estimator to construct a similar sequence of iterations, we leave such an extension to future work. A second restriction of our analysis is that we have restricted attention to single-agent models. We are currently working on such an extension, but it is beyond the scope of this paper [Authors' note: might be able to include this in paper or appendix. Although, it won't be asymptotically equivalent to (partial) MLE.

References

Ackerberg, D., L. Benkard, S. Berry, and A. Pakes, "Econometric Tools for Analyzing Market Outcomes," in J.J. Heckman and E.E. Leamer, eds., *Handbook of Econometrics*, Vol. 6A, Amsterdam: North-Holland, 2007.

Aguirregabiria, V. and P. Mira, "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models," *Econometrica*, 2002, *70*, 1519–1543.

_ **and** _ , "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 2007, *75*, 1–53.

_ **and** _ , "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics*, 2010, *156*, 38–67.

Altug, S. and R.A. Miller, "The Effect of Work Experience on Female Wages and Labour Supply," *Review of Economic Studies*, 1998, *65*, 45–85.

Arcidiacono, P. and R.A. Miller, "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," *Econometrica*, 2011, *7*, 1823–2868.

_ **and** _ , "Nonstationary Dynamic Models with Finite Dependence," 2016. Working Paper.

Berry, S., A. Gandhi, and P. Haile, "Connected Substitutes and Invertibility of Demand," *Econometrica*, September 2013, *81* (5), 2087–2111.

_ , J. Levinsohn, and A. Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

Bray, R.L., "Unnesting the Fixed Point in Dynamic Program Estimation," December 2018. Working Paper.

Chiong, K.X., A. Galichon, and M. Shum, "Duality in Dynamic Discrete-Choice Models," *Quantitative Economics*, 2016, *7*, 83–115.

Hotz, V.J. and R.A. Miller, "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 1993, *60*, 497–529.

_ , _ , S. Sanders, and J. Smith, "A Simulation Estimator for Dynamic Models of Discrete Choice," *Review of Economic Studies*, 1994, *61*, 265–289.

Kalouptsidi, M., P.T. Scott, and E. Souza-Rodrigues, "Identification of Counterfactuals in Dynamic Discrete Choice Models," 2017. Working Paper.

_ , _ , and _ , "Linear IV Regression Estimators for Structural Dynamic Discrete Choice Models," October 2018. NBER Working Paper 25134.

Keane, M.P. and K.I. Wolpin, "The Career Decisions of Young Men," *Journal of Political Economy*, 1997, *105*, 473–522.

_ , P.E. Todd, and K.I. Wolpin, "The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications," in "Handbook of Labor Economics," Vol. 4, Elsevier, 2011, chapter 4, pp. 331–461.

Magnac, T. and D. Thesmar, "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 2002, *70*, 801–816.

McFadden, D.L. and W. Newey, "Large Sample Estimation and Hypothesis Testing," in R.E. Engle and D.L. McFadden, eds., *Handbook of Econometrics*, Vol. 4, Amsterdam: North-Holland, 1994.

Moré, J.J., "Global Convergence of Newton-Gauss-Seidel Methods," *SIAM Journal on Numerical Analysis*, June 1971, *8* (2), 325–336.

Norets, A., "Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables," *Econometrica*, 2009, *77*, 1665–1682.

_ **and X. Tang**, "Semiparametric Inference in Dynamic Binary Choice Models," *Review of Economic Studies*, 2014, *81*, 1229–1262.

Pakes, A., "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 1986, *54*, 1027–1057.

_ **, M. Ostrovsky, and S. Berry**, "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)," *RAND Journal of Economics*, Summer 2007, *38* (2), 373–399.

Pesendorfer, M. and P. Schmidt-Dengler, "Asymptotic Least Squares Estimators for Dynamic Games," *Review of Economic Studies*, 2008, *75*, 901–928.

Rust, J., "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55*, 999–1013.

_ , "Maxiumum Likelihood Estimation of Discrete Control Processes," *SIAM Journal on Control and Optimization*, 1988, *26*, 1006–1024.

Su, C-L. and K.L. Judd, "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 2012, *80* (5), 2213–2230.

A Proofs

A.1 Proof of Proposition 2

The first result follows from the analysis immediately preceding the proposition. For the second result, recall that $\psi(\cdot)$ is defined so that $0 = S(-\psi(P))$ and $\boldsymbol{P} = \nabla_v S(-\psi(P))$. Taking derivatives on both sides of the first equation, we get

$$\begin{aligned} 0 &= -\nabla_v S(-\psi(P))\nabla_P \psi(P) \\ &= -\boldsymbol{P}\nabla_P \psi(P). \end{aligned}$$

Multiplication by -1 then gives the result stated in the proposition.

A.2 Proof of Proposition 1

We begin with the properties of $T(u, V) = S(\nu(u, V))$. We first note that $\nu(u, V)$ is linear and isotone. The composite of isotone functions is also isotone, so isotonicity of $T(\cdot)$ follows from isotonicity of $S(\cdot)$ and $V(\cdot)$. Convexity of $T(\cdot)$ follows from the convexity of $S(\cdot)$ and linearity of $\nu(\cdot)$. The derivative formulas follow from the chain rule and echo those of Rust (1987; 1988).

For $G(u, V) = V - T(u, V)$, the derivative formulas follow immediately from those for $T(\cdot)$. Concavity of $G(\cdot)$ follows from convexity of $T(\cdot)$. Antitonicity of $G(\cdot)$ in u follows from isotonicity of $T(\cdot)$. For inverse-isotonicity of $G(\cdot)$ in V, it is useful to introduce $G_u(V)$ by fixing the value of u. We will show that $G_u(V)$ is inverse-isotone. Note that $dG_u/dV' = I - \beta F_u^U(V)$. This matrix is diagonally dominant, with positive diagonal elements and non-positive off-diagonal elements. So, it is an M-matrix and is therefore nonsingular with $G_u'(V)^{-1} \geq 0$. This last inequality is clear from the Neumann series representation: $(I - \beta F^U)^{-1} = \sum_{t=0}^{\infty}(\beta F^U)^t$. We use the next lemma to establish inverse-isotonicity.

Lemma 1. *(Moré (1971), Lemma 3.3) Let $f : D \subset \mathbb{R}^n \to \mathbb{R}^n$ be convex (concave) and continuously differentiable on the convex set $D_0 \subset D$. If $f'(x)^{-1} \geq 0$ for x in D_0, then f is inverse-isotone on D_0. Conversely, if f is inverse isotone on D_0, then $f'(x)^{-1} \geq 0$ for each x in the interior of D_0.*

Proof. Moré (1971) provides a proof for convex f. The proof for concave f is nearly identical. □

$G_u(V)$ is concave, with $G_u'(V)^{-1} \geq 0$. So, inverse-isotonicity of G_u follows from Lemma 1.

We now turn attention to the properties of $V(u)$, which is defined implicitly: $G(u, V(u)) = 0$. The derivative formula follows from implicit differentiation and some algebra. That $V(u)$ is well-defined is a well-established result. We present an alternative proof here that does not rely on showing $T_u(V)$ is a contraction mapping. Instead, we show that $G_u : \mathbb{R}^M \to \mathbb{R}^M$ is a homeomorphism, so that it has a unique root.

Lemma 2. *(Moré (1971), Lemma 3.6) Let $f : \mathbb{R}^n \to \mathbb{R}^n$ be continuously differentiable, and assume there is a constant, c, such that $||f'(x)^{-1}|| \leq c$ for all x in $\mathbb{R}^n$. Then f is a homeomorphism of $\mathbb{R}^n$ onto $\mathbb{R}^n$.*

We note that row-stochasticity of $F_u^U(V)$ implies $||F_u^U(V)^t||_\infty = 1$ for $t \geq 0$. Using the Neuman series representation of $G_u'(V)$, we obtain

$$\begin{aligned}
\left\|G_u'(V)^{-1}\right\|_\infty &= \left\|\sum_{t=0}^{\infty}(\beta F_u^U(V))^t\right\|_\infty \\
&\leq \sum_{t=0}^{\infty}\left\|\beta^t F_u^U(V)^t\right\|_\infty \\
&= \sum_{t=0}^{\infty}\beta^t\left\|F_u^U(V)^t\right\|_\infty \\
&= \frac{1}{1-\beta}.
\end{aligned}$$

So, $G_u(V)$ is a homeomorphism by Lemma 2, and $V(u)$ is therefore well-defined. In Section 3.2, we showed how to construct flow utilities that result in any value function and choice probability pair, (V, P). Thus, $V(u)$ is surjective. Isotonicity and convexity of $V(u)$ follow from the next lemma.

Lemma 3. *Suppose $f(x, y) : \mathbb{R}^{n\times m} \to \mathbb{R}^m$ is concave in (x, y), antitone in x, and inverse-isotone in y. Define the implicit function $y(x)$ such that $f(x, y(x)) = k$ and suppose that $y(x)$ is well-defined for all $x \in \mathbb{R}^n$. Then, $y(x)$ is isotone and convex.*

Proof. For isotonicity, consider $x_0 \leq x_1$. Then, $f(x_1, y(x_0)) \leq k = f(x_1, y(x_1))$ by antitonicity of f in x. Inverse-isotonicity of f in y then implies that $y(x_0) \leq y(x_1)$, which establishes the isotonicity of $y(x)$. For convexity, consider x_1, x_2 and $0 < \lambda < 1$. Define $\overline{x} = \lambda x_1 + (1-\lambda)x_2$ and $\overline{y} = \lambda y(x_1) + (1-\lambda)y(x_2)$. We show that $y(\overline{x}) \leq \overline{y}$. Then, using

the concavity of f and the definition of $y(x)$ gives the following:

$$\begin{aligned} f(\overline{x}, y(\overline{x})) &= k \\ &= \lambda k + (1-\lambda)k \\ &= \lambda f(x_1, y(x_1)) + (1-\lambda) f(x_2, y(x_2)) \\ &\leq f(\overline{x}, \overline{y}). \end{aligned}$$

and it follows from inverse-isotonicity of f in y that $y(\overline{x}) \leq \overline{y}$. □

Finally, we consider global convergence of the iterative algorithms. For Bellman iterations, global convergence is well-established in the literature. For Newton iterates, we appeal to the properties of $G_u(V)$, we another result from Moré (1971) give sufficient conditions for global convergence.

Lemma 4. *(Moré (1971), Corollary 3.5) Let $f : \mathbb{R}^n \to \mathbb{R}^n$ be continuously differentiable and convex (concave) on $\mathbb{R}^n$, and suppose that $f'(x)^{-1} \geq 0$ for all x in $\mathbb{R}^n$. If $f(x) = 0$ has a solution x^*, then this solution is unique, and for any $x^0 \in \mathbb{R}^n$, the sequence*

$$x^{k+1} = x^k - f'(x^k)^{-1} f(x^k)$$

converges to x^*.

Proof. Moré (1971) proves the result for the convex case. The concave case is similar. □

We have already shown that $G_u(V)$ satisfies the sufficient conditions of Lemma 4, proving global convergence for Newton iterates.

A.3 Proof of Proposition 5

With $\omega = P^*$, we have $V_{AHM}(u, P, P^*) = \left(I - \beta F^U(P^*)\right)^{-1} \boldsymbol{P}^* \left[u + \psi(P)\right]$. We immediately see that $\nabla_u V_{AHM}(u^*, P^*, P^*) = \left(I - \beta F^U(P^*)\right)^{-1} \boldsymbol{P}^* = \nabla_u V(u^*)$, where the second inequality follows from Proposition 1. Proposition 2 gives $\boldsymbol{P}^* \nabla_P \psi(P^*) = 0$, so we have $\nabla_P V_{AHM}(u^*, P^*, P^*) = \left(I - \beta F^U(P^*)\right)^{-1} \boldsymbol{P}^* \nabla_P \psi(P^*) = 0$. The equivalence to (partial) MLE arises from analysis similar to Aguirregabiria and Mira's (2002) proof of the same result for $V_{NPL}(u(\theta_u), \hat{P}_0)$.

A.4 Proof of Theorem 2

[Authors' note: the regularity assumptions are standard consistency assumptions such as $\hat{W} \to_p W$ and so on.] We first derive the asymptotic distribution of $\hat{\theta}_{MD}$. In order to do so,

we consider the first order condition,

$$\begin{aligned}0 &= \nabla_\theta \Psi(\hat{\theta}_{MD}, \hat{P}_0)' \hat{W} \left(\hat{P} - \Psi(\hat{\theta}_{MD}, \hat{P}_0) \right) \\ &= H(\hat{P}_0)' \nabla_{\tilde{v}} \Lambda(\varphi(\hat{\theta}_{MD}, \hat{P}, \hat{P}_0)) \hat{W} \left(\hat{P} - \Psi(\hat{\theta}_{MD}, \hat{P}_0) \right) \\ &= S_{MD}(\hat{\theta}_{MD}, \hat{P}, \hat{P}_0).\end{aligned}$$

where the second equality follows from $\nabla_\theta \Psi(\theta, P)' = \nabla_\theta \varphi(\theta, P)' \nabla_{\tilde{v}} \Lambda(\varphi(\theta, P))$ and $\varphi(\theta, P) = H(P)\theta + Z(P)$ so that $\nabla_\theta \varphi(\theta, P) = H(P)$. Note that $S_{MD}(\theta, P(\theta), P(\theta)) = 0$ for all θ, so that $S(\hat{z}) \to_p 0$ by Slutsky's theorem. Expansion of $S(\cdot)$ in θ around $\hat{\gamma} = (\theta^*, \hat{P}, \hat{P}_0)$ an abuse of notation yields

$$0 = S_{MD}(\hat{\gamma}) - H(\hat{P}_0)' \nabla_{\tilde{v}} \Lambda(\varphi(\hat{\gamma})) \hat{W} \nabla_{\tilde{v}} \Lambda(\varphi(\bar{\theta}, \hat{P}, \hat{P}_0)) H(\hat{P}_0) \left(\hat{\theta}_{MD} - \theta^* \right)$$

Asymptotic normality of $(\hat{P}, \hat{P}_0)$ gives asymptotic normality of $S(\hat{\gamma})$, so that $\sqrt{N}(S(\hat{\gamma})) \to_d \mathcal{N}(0, \Omega)$. So, the asymptotic distribution of $\hat{\theta}_{MD}$ is then derived by multiplying both sides of the equation by $\sqrt{N}$, performing some algebraic manipulations, and applying Slutzky's Theorem (see the proof of Theorem 3.2 in McFadden and Newey (1994)) to get

$$\sqrt{N}(\hat{\theta}_{MD} - \theta^*) \to_d H^{*\prime} \nabla_{\tilde{v}} \Lambda^* W \nabla_{\tilde{v}} \Lambda^* H^* \mathcal{N}(0, \Omega)$$

Now consider deriving the asymptotic distribution of $\hat{\theta}_{CF}$. The first order condition for the minimization is

$$\begin{aligned}0 &= H(\hat{P}_0)' \widehat{\nabla_{\tilde{v}} \Lambda_0} \hat{W} \left(\hat{P} - \hat{P}_0 - \widehat{\nabla_{\tilde{v}} \Lambda_0} \left[\varphi(\theta^*, \hat{P}_0) - \Lambda^{-1}(\hat{P}_0) \right] \right) \\ &= S_{CF}(\hat{\theta}_{CF}, \hat{P}, \hat{P}_0).\end{aligned}$$

We now show that $S_{MD}(\hat{z})$ and $S_{CF}(\hat{z})$ have the same asymptotic distribution. To do so, it suffices to show that $\Psi(\theta^*, \hat{P}_0)$ and $\hat{P}_0 + \widehat{\nabla_{\tilde{v}} \Lambda_0}[\varphi(\theta, \hat{P}_0) - \Lambda^{-1}(\hat{P}_0)]$ have the same asymptotic distribution. First, note that $\Psi(\theta^*, P^*) = P^*$ and $\nabla_P \Psi(\theta^*, P^*) = \nabla_{\tilde{v}} \Lambda^* \nabla_P \varphi^*$. Now, define $R(\theta, P) = P - \nabla_{\tilde{v}} \Lambda(\Lambda^{-1}(P)) \left[\varphi(\theta, P) - \Lambda^{-1}(P) \right]$ and recall that $\widehat{\nabla_{\tilde{v}} \Lambda_0} = \nabla_{\tilde{v}} \Lambda(\Lambda^{-1}(\hat{P}_0))$, so that $R(\theta, \hat{P}_0) = \hat{P}_0 + \widehat{\nabla_{\tilde{v}} \Lambda_0}[\varphi(\theta, \hat{P}_0) - \Lambda^{-1}(\hat{P}_0)]$. It is straightforward to show that $R(\theta^*, P^*) = P^*$ and $\nabla_P R(\theta^*, P^*) = \nabla_{\tilde{v}} \Lambda^* \nabla_P \varphi^*$. So, applying the delta method will give the same asymptotic distribution for $\Psi(\theta^*, \hat{P}_0)$ and $R(\theta^*, \hat{P}_0)$.

Mean value theorem expansion of $S_{CF}(\cdot)$ in θ around $\hat{z} = (\theta^*, \hat{P}, \hat{P}_0)$ and now gives

$$0 = S_{CF}(\hat{z}) - H(\hat{P}_0)' \widehat{\nabla_{\tilde{v}} \Lambda_0} \hat{W} \widehat{\nabla_{\tilde{v}} \Lambda} H(\hat{P}_0) \left(\hat{\theta}_{CF} - \theta^* \right)$$

and we can again use algebraic manipulations and Slutsky's Theorem to establish that

$$\sqrt{N}(\hat{\theta}_{CF} - \theta^*) \to_d H^{*\prime}\nabla_{\tilde{v}}\Lambda^* W \nabla_{\tilde{v}}\Lambda^* H^* \mathcal{N}(0, \Omega),$$

which is the same distribution derived for $\hat{\theta}_{MD}$.

A.5 Proof of Proposition 7

We have $\hat{\theta}_{CFPL}$ asymptotically equivalent to $\hat{\theta}_{MD}$ with $\hat{W} = \hat{W}_{PL}$ by Theorem 2. Asymptotic equivalence of $\hat{\theta}_{MD}$ with $\hat{W} = \hat{W}_{PL}$ to $\hat{\theta}_{PML}$, where $\hat{\theta}_{PML}$ is the pseudo-maximum likelihood estimator, follows from the asymptotic equivalence result of Pesendorfer and Schmidt-Dengler (2008) and from Slutsky's theorem giving $\hat{\Sigma}_p - \Sigma_p = o_p(1)$. Finally, asymptotic equivalence of $\hat{\theta}_{PML}$ to $\hat{\theta}_{MLE}$ when the pseudo-value function is $V_{NPL}(\cdot)$ is shown in Aguirregabiria and Mira (2002).

A.6 Proof of Proposition 8

We are considering a fixed sample, so $\hat{P}$ does not vary with iteration. Define $\phi(P) = \Phi(\hat{P}, P)$ and $\Upsilon(P) = \Psi(\Phi(\hat{P}, P), P)$ so that $\hat{\theta}_k = \phi(\hat{P}_{k-1})$ and $\hat{P}_k = \Upsilon(\hat{P}_{k-1})$. By definition, $\hat{P}_{MLE} = P(\hat{\theta}_{MLE})$. so from Proposition 6 and the definitions of $\phi(\cdot)$ and $\Upsilon(\cdot)$, we can see that $\hat{P}_{MLE} = \Psi(\hat{\theta}_{MLE}, \hat{P}_{MLE})$. This is enough to establish that $\hat{\theta}_{MLE} = \phi(\hat{P}_{MLE})$ and $\hat{P}_{MLE} = \Upsilon(\hat{P}_{MLE})$. Additionally, we again invoke Proposition 6 to get $\nabla_P \phi(\hat{P}_{MLE}) = 0$ and $\nabla_P \Upsilon(\hat{P}_{MLE}) = 0$. Both $\phi(\cdot)$ and $\Upsilon(\cdot)$ are twice-continuously differentiable, so the second result follows immediately.

B Example of $\psi(P) \neq e(P)$

Consider binary choice, $a \in \{0, 1\}$. We focus on only one state here, so we only use a choice subscript. The error distribution is

$$\varepsilon_a \overset{i.i.d.}{\sim} \mathcal{N}\left(0, \frac{1}{2}\right).$$

Define $\tilde{v} = v_1 - v_0$ and $\tilde{\varepsilon} = \varepsilon_1 - \varepsilon_0$, where $\tilde{\varepsilon} \sim \mathcal{N}(0, 1)$. It is easy to show that $P_0(v) = Pr(a = 0|v) = \Phi(-\tilde{v})$. We use equation (13) from Aguirregabiria and Mira (2007) to derive

$$e_a(P) = \frac{\phi\left(\Phi^{-1}(P_a)\right)}{2P_a}.$$

Some algebraic manipulations based on Algorithm 1 then yield

$$\psi_0(P) = P_1\Phi^{-1}(P_1) + \phi(\Phi^{-1}(P_0))$$

$$\psi_1(P) = P_0\Phi^{-1}(P_0) + \phi(\Phi^{-1}(P_0)).$$

Plugging in $(P_0^0, P_1^0) = (0.9, 0.1)$ gives $e(P^0) = (0.0975, 0.8775)$ and $\psi(P^0) = (0.0473, 1.3289)$.