

NBER WORKING PAPER SERIES

APPROXIMATING HIGH-DIMENSIONAL DYNAMIC MODELS:
SIEVE VALUE FUNCTION ITERATION

Peter Arcidiacono
Patrick Bayer
Federico A. Bugni
Jonathan James

Working Paper 17890
<http://www.nber.org/papers/w17890>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2012

Arcidiacono, Bayer, and Bugni thank the National Science Foundation for research support via grant SES-1124193. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2012 by Peter Arcidiacono, Patrick Bayer, Federico A. Bugni, and Jonathan James. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Approximating High-Dimensional Dynamic Models: Sieve Value Function Iteration
Peter Arcidiacono, Patrick Bayer, Federico A. Bugni, and Jonathan James
NBER Working Paper No. 17890
March 2012
JEL No. C13,C14,C54,C61,C63,C73

ABSTRACT

Many dynamic problems in economics are characterized by large state spaces which make both computing and estimating the model infeasible. We introduce a method for approximating the value function of high-dimensional dynamic models based on sieves and establish results for the: (a) consistency, (b) rates of convergence, and (c) bounds on the error of approximation. We embed this method for approximating the solution to the dynamic problem within an estimation routine and prove that it provides consistent estimates of the model's parameters. We provide Monte Carlo evidence that our method can successfully be used to approximate models that would otherwise be infeasible to compute, suggesting that these techniques may substantially broaden the class of models that can be solved and estimated.

Peter Arcidiacono
Department of Economics
201A Social Sciences Building
Duke University
Durham, NC 27708
and NBER
psarcidi@econ.duke.edu

Patrick Bayer
Department of Economics
Duke University
213 Social Sciences
Durham, NC 27708
and NBER
patrick.bayer@duke.edu

Federico A. Bugni
Department of Economics
Duke University
Box 90097
Durham, NC 27708
federico.bugni@duke.edu

Jonathan James
Department of Economics, Duke University
Box 90097
Durham, NC 27708
jwj8@duke.edu

1 Introduction

Dynamic problems in economics (and many other disciplines) are characterized by large state spaces that routinely make the model difficult or impossible to compute. Faced with these challenges, researchers proceed in a number of different ways including analyzing a static version of the problem, simplifying the dynamic problem to a setting where the state space is manageable, or employing a number of techniques for approximating the value and policy functions that characterize the solution to the problem. The problems analyzed in many broad literatures in economics ranging from models of market equilibrium (which are almost always treated in a static framework) to strategic or social interactions in networks (static or myopic dynamics) to dynamic games (typically characterized by a small number of players and states) to matching problems (almost always analyzed as static problems) remain limited by this large state space problem.

In recent years, techniques for approximating value functions in large state space problems have been developed using simulation and interpolation (Keane and Wolpin (1994)) or parametric policy iteration (PPI) (Rust (2000) and BÉNITEZ-SILVA, HALL, HITSCH, PAULETTO, and RUST (2000)). The main idea behind these methods is to approximate the value function using flexible functions of the relevant state variables. Simulation and interpolation has been applied by Keane and Wolpin (1997) and Crawford and Shum (2005) among many others. PPI has been applied to a single agent setting in Hendel and Nevo (2006) and to a dynamic game in Sweeting (2011). While the potential application of these methods is great, the literature contains very few formal results regarding the quality of the approximation and, therefore, provides little formal or practical foundation for researchers to use as they apply these methods (see Powell (2008) for a summary of the related literature and a discussion of the lack of formal results).

In this paper, we consider an alternative, but related, method for approximating value functions that we refer to as sieve value function iteration (SVFI).¹ We develop the method in the context of a generic single agent dynamic programming problem that can have either a finite or infinite horizon. The SVFI method involves approximating the integrated value function with a nonparametric sieve function. For any sieve function (i.e., for a particular choice of parameters), one can evaluate the Bellman equation and compute a notion of distance between the approximation and its contraction, and thus characterize how close the Bellman equation comes to holding. We approximate the value function by choosing the parameters of the sieve function that come as close as possible to making the Bellman equation hold exactly. Since the sieve space is a simple space, this minimization problem is relatively easy to solve. Moreover, as the sequence of sieve spaces becomes a better and better approximation of the original functional space \mathcal{F} , our approximation converges to the true value function.

In order to analyze the formal properties of the SVFI method, we assume that the complexity of the sieve space, n , increases to infinity. In this sense our SVFI approximation technique becomes non-parametric and we establish a number of results:

¹For an excellent review of the method of sieves see Chen (2007).

1. *Consistency*: we show that the sieve approximation converges to the true value function as the richness of the sieve space increases.
2. *Rates of convergence*: we provide the first results in the literature for the rate at which the approximating function converges to the true value function.
3. *Iteration of Bellman operator*: we characterize how rates of convergence are affected by iterating the Bellman operator, pointing out for the first time that this has the potential to improve performance in certain applications.
4. *Bound on the error of approximation*: following arguments in Rust (2000), we provide an upper bound to the error of approximation to the unknown true value function that is feasible to compute.

To the best of our knowledge, this is the first paper to formally establish the first three points for the SVFI approximation. While the consistency of the SVFI might be an expected result, establishing the rate of convergence of the SVFI approximation and understanding the effects of increasing the number of iterations on the quality of the approximation are two useful formal contributions to this literature. The fourth result, already present in Rust (2000), is included for the sake of completeness. This is an incredibly valuable feature of SVFI because it ensures that one can bound the extent of approximation error relative to the true value function even when the true value function cannot be computed, i.e., the case of relevance in very large state space problems. This also serves as a natural guide to researchers as they implement the SVFI, providing a clear sense of whether the approximation is reasonable for any given specification of the sieve function. Taken together, these results provide a formal foundation for the use of SVFI methods for approximating large state space problems.

The SVFI approach is quite straightforward to understand and implement and, thus, has the potential to be widely applied in economics and other disciplines. The framework can also be flexibly implemented. It is possible, for example, to estimate the parameters of the sieve approximation function by minimizing the distance in the Bellman operator for only a large subset of the states in the state space. This is attractive for problems with large finite or infinite state spaces.

The method can be applied equally well to infinite and finite horizon problems. For finite horizon problems, we develop two strategies for approximating value function. First, we show how SVFI can be used to approximate the value functions at each time period using a traditional backwards recursion solution method. More interestingly, by including the time to the horizon as another state variable in the sieve function, we show how it is possible to approximate a single sieve function that provides an approximation of the value function at each point in time without solving the problem backwards. The particular features of the application in question will generally determine which of these approaches is computationally lighter or easier to implement.

Having developed this general method for approximating value functions, we then formally show how SVFI can be used to estimate the structural parameters of large state space dynamic

problems in empirical applications. Conceptually it is easy to see that for any particular guess of the model's structural parameters, SVFI can be used to approximate the solution to the dynamic problem and thus compute the approximate value and policy functions as well as conditional choice probabilities. By comparing these objects to their empirical counterparts in the data it is possible to compute a wide variety of objective functions that would be appropriate for estimating the structural parameters. We show that it is possible to consistently estimate the model's structural parameters by embedding SVFI within the estimation algorithm.

We close the paper by demonstrating the performance of SVFI in a particular Monte Carlo application, inspired by the famous bus-engine replacement problem of Rust (1987), in which a firm must dynamically manage its entire fleet of buses. We begin by analyzing an infinite horizon problem with a state space that is finite and relatively large but manageable, i.e., for which it is still possible to compute the value function exactly at each state. We demonstrate that SVFI approximates value functions for this problem very closely in a tiny fraction of the time that it takes to compute the exact solution to the problem. We show that an accurate approximation can be obtained even when the Bellman operator is evaluated at a small, randomly, drawn subset of the full set of states and compare the speed and accuracy of alternative methods for minimizing this distance, including non-linear least squares, an iterative least squares method, and methods that iterate the Bellman operator. We extend this problem to an infinite state space by adding a continuous state variable, again demonstrating that it is possible to approximate the solution exceptionally well in a reasonably short amount of computation time. We then analyze an analogous finite horizon problem, comparing the speed and accuracy of SVFI approximations using (i) the traditional backwards recursion solution method in which we compute a separate sieve approximation function at each point in time and (ii) an approach that treats time to the horizon as a state variable, yielding a single time-interacted sieve approximation function.

We complete the Monte Carlo section of the paper, by evaluating the performance of SVFI approximation within the context of the estimation of the structural parameters of the model. Returning to the infinite horizon problem, we demonstrate that SVFI accurately estimates with only a minimal impact on the effective standard errors while being much faster to compute. In particular, we propose estimating the sieve parameters outside of the estimation of the structural parameters in a manner similar to Aguirregabiria and Mira (2002) swapping of the nested fixed point algorithm. Further, it is straightforward to examine how the estimated structural parameters are affected by the dimension of the sieve and, therefore, to ensure that the sieve is sufficiently rich such that its impact on the estimated structural parameters is minimal.

The rest of the paper proceeds as follows. In section 2 we describe the optimization problem. Sections 3 and 4 develop formal results for infinite horizon problems and finite horizon problems, respectively. In particular, we establish properties of the approximation including consistency and rates of convergence. In section 5, we show how to use SVFI in estimation as well as properties of the estimator. Section 6 investigates the small sample properties of our approximations, establishing both the speed and reliability of our methods. Section 7 concludes.

2 The dynamic decision problem

In this paper, we are interested in approximating the value function of a single agent solving a dynamic decision problem. We begin by introducing the general problem. In every period $t \in \mathbb{N}$, the agent observes the current value of a state (x_t, ε_t) and chooses an action a_t in a finite choice set $A(x_t)$. The first component x_t is observed by the researcher whereas the second component $\varepsilon_t \equiv \{\varepsilon_t(a) : a \in A(x_t)\}$ is not. We assume that $x_t \in X$ and $\varepsilon_t \in E$. Conditional on this information, the agent solves the following optimization problem:

$$V_t(x_t, \varepsilon_t) = \sup_{\Pi_t} E \left\{ \sum_{j=t}^T \beta^{j-t} [u(x_j, a_j) + \varepsilon_j(a_j)] \middle| x_t, \varepsilon_t \right\}, \quad (2.1)$$

where $\Pi_t = \{\{a_j\}_{j=t}^{\infty} : a_t = A(x_t)\}$ and $u(x_t, a_t) + \varepsilon_t(a_t)$ denotes the period utility of making decision a_t in state (x_t, ε_t) , $u(x_t, a_t)$ representing the structural part (possibly known up to a finite dimensional parameter) and $\varepsilon_t(a_t)$ representing the residual part, unobserved by the researcher.

The objective of this paper is to provide a computationally feasible approximation to the value functions (i.e. $\{V_t(\cdot)\}_{t=1}^T$) and study its properties. In general, we might be interested in approximating these function because we want to do welfare analysis (i.e. which can be conducted directly using these functions) or because we are interested in any other feature of the problem that can be computed from these functions (i.e. optimal decision rules, conditional choice probabilities (CCPs), etc.).

The formulation in Equation (2.1) encompasses both finite horizon problems (i.e. $T = \infty$) and infinite horizon problems (i.e. $T < \infty$). Following the dynamic programming literature, our strategy to approximate the value function in Equation (2.1) is to show that it is the unique solution to a (functional) Bellman equation, and construct an approximation based on this representation. Given that the Bellman equation representation of the finite and infinite horizon problems are fundamentally different, it is convenient to divide the rest of the discussion into these two cases.

3 Approximation in infinite horizon problems

This section describes the application of non-parametric sieve methods to approximate the value functions in Equation (2.1) when the dynamic problem has an infinite horizon (i.e. $T = \infty$). The distinctive feature of the infinite horizon problem is that the value function of the problem is stationary (i.e. $V_t(\cdot) = V(\cdot), \forall t \in \mathbb{N}$), provided that we impose mild additional assumptions to the problem. In particular, we follow Rust (1987) and Rust (1988) and assume that the joint stochastic process $\{x_t, \varepsilon_t, a_t\}$ is a controlled Markov process that satisfies the following conditional

independence (CI) assumption:

$$dP(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, a_t, x_{t-1}, \varepsilon_{t-1}, a_{t-1}, \dots) = dP(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, a_t), \text{ (Markov)} \quad (3.1)$$

$$= dP(\varepsilon_{t+1} | x_{t+1}, x_t, \varepsilon_t, a_t) dP(x_{t+1} | x_t, \varepsilon_t, a_t), \quad (3.2)$$

$$= dP(\varepsilon_{t+1} | x_{t+1}) dP(x_{t+1} | x_t, a_t). \text{ (CI)} \quad (3.3)$$

Under these assumptions, the literature provides several ways of formulating the value function in Equation (2.1) as the recursive solution of a Bellman equation. In particular, these include: (a) the *social surplus function* formulation discussed in Rust (1987) and Rust (1988), (b) the *conditional value function* formulation of Rust (1987), and (c) the *choice-specific value function* formulation of Rust (1988). Rather than describing each of these formulations in the main text, we now introduce a single *unified formulation* that encompasses all of these formulations.²

According to our unified formulation, the value functions in Equation (2.1) are stationary and solve the following Bellman equation:

$$V(s) = \max_{a \in A(s)} \{u(s, a) + \beta E(F(V, s') | s, a)\}, \forall s \in S, \quad (3.4)$$

where s is the current value of the state, s' is the future value of the state, S is the state space, a represents the action chosen by the agent, $A(s)$ is the set of actions available to the agent when the state is s , and F is a known functional of the value function $V(\cdot)$ and the future state s' that satisfies certain known properties.³

3.1 Approximating the value function

In order to discuss the approximation of V , we must first define the space to which this function belongs. In particular, assume that V belongs to a functional space, denoted \mathcal{F} . For example, we can take \mathcal{F} to be the space of measurable, bounded, real-valued functions from S to \mathbb{R} . We define a metric d in this space, making (\mathcal{F}, d) a normed vector space. If we do not indicate otherwise, the metric is the sup-norm metric, i.e., for any $f_1, f_2 \in \mathcal{F}$:

$$d(f_1, f_2) = \sup_{s \in S} |f_1(s) - f_2(s)|. \quad (3.5)$$

Furthermore, we assume that this metric space is complete, i.e., it is a Banach space. Consider the following (functional) operator $\Gamma : \mathcal{F} \rightarrow \mathcal{F}$:

$$(\Gamma\theta)(s) = \max_{a \in A(s)} \{u(s, a) + \beta E(F(\theta, s') | s, a)\}, \forall s \in S. \quad (3.6)$$

²Section B of the appendix describes each of these formulations and shows that all of them are special cases of our unified formulation.

³Among other properties, $F : \mathcal{F} \times S$ satisfies: (a) monotonicity: for functions f_1, f_2 with $f_1(s) \leq f_2(s)$ for all $s \in S$, then $F(f_1, s') \leq F(f_2, s')$ and (b) discounting: for any function f and any $\alpha \in \mathbb{R}$, $\beta F(f + \alpha, s') = \beta F(f, s') + \beta\alpha$.

According to the definition in Equation (3.4), the value function is the fixed point of this operator. Furthermore, under Blackwell (1965) sufficient conditions (see, e.g., Stokey and Lucas (1989), Theorem 3.3), the operator Γ can be shown to be a contraction mapping. Thus, as a consequence of the contraction mapping theorem (see, e.g., Stokey and Lucas (1989), Theorem 3.2), it follows that the value function is the *unique* fixed point of the contraction mapping. In particular, if we consider \mathcal{F} to be a space of functions that map $S \rightarrow \mathbb{R}$ and we consider any metric $d : \mathcal{F} \times \mathcal{F} \rightarrow \mathbb{R}$, then, the value function is the *unique* solution to the following minimization problem:

$$\inf_{\theta \in \mathcal{F}} d(\theta, \Gamma\theta) \tag{3.7}$$

As a consequence, if it is possible to solve the minimization problem in Equation (3.7), then the solution has to be unique and equal to the value function. Unfortunately, there are several situations of practical relevance in which this minimization problem is computationally infeasible. In this paper we focus on the difficulties that arise when the state space is too large, i.e., the cardinality of the set S is either infinity or finite but too large to permit the use of traditional methods.

The approximation method we propose is inspired by the sieves non-parametric estimation method. Instead of solving the original minimization problem (Equation (3.7)), we replace the original (possibly infinite dimensional) parameter space \mathcal{F} with a sequence of simpler (often finite dimensional) parameter spaces, called sieves. Throughout this paper, the sequence of sieve spaces will be denoted by $\{\Theta_n\}_{n \geq 1}$, where $n \in \mathbb{N}$ is an index that represents the computation power of the problems that we can handle. In order for this replacement to produce an accurate approximation of the unknown value function, it will be required that the sieve space sequence $\{\Theta_n\}_{n \geq 1}$ to become increasingly more complex (i.e. for any $n \in \mathbb{N}$, $\Theta_n \subset \Theta_{n+1} \subseteq \mathcal{F}$) and dense in \mathcal{F} as $n \rightarrow \infty$.

For a given computational power (i.e. for a given sieve space Θ_n), our method produces an approximation, denoted $\hat{\theta}_n$. In essence, we replace the complex parameter space \mathcal{F} by the sieve parameter space Θ_n and, loosely speaking, our approximation will be given as:

$$\hat{\theta}_n \approx \underset{\theta \in \Theta_n}{\operatorname{argmin}} d(\theta, \Gamma\theta). \tag{3.8}$$

That is, we will seek to choose the parameters of the sieve to get as close to a fixed point of the Bellman operator as possible. Naturally, the quality of the approximation will be determined by the sieve space Θ_n used to approximate the original parameter space \mathcal{F} .

We introduce a definition of consistency of the approximation and rate of convergence of the approximation.

Definition 3.1 (Consistency). $\hat{\theta}_n$ is a consistent approximation to V if and only if:

$$d(\hat{\theta}_n, V) = o_p(1), \text{ as } n \rightarrow \infty. \tag{3.9}$$

To be precise, this is an approximation problem and not an estimation problem. In other words,

there are no data and, hence, no random sampling error⁴.

Definition 3.2 (Rate of Convergence). $\hat{\theta}_n$ converges to V at a rate of γ_n^{-1} if and only if:

$$d(\hat{\theta}_n, V) \leq O_p(\gamma_n), \quad (3.10)$$

where $\gamma_n = o(1)$ as $n \rightarrow \infty$.

3.2 Assumptions

We now provide a list of the assumptions used in this section. As we show in Example 3.1, all of these assumptions are satisfied in dynamic decision problems that have a very large but finite state space.

Assumption 3.1. (\mathcal{F}, d) is a complete metric space of functions that map S onto \mathbb{R} and $\Gamma : \mathcal{F} \rightarrow \mathcal{F}$ is a contraction mapping with modulus $\beta \in (0, 1)$.

Assumption 3.2. For any $n \in \mathbb{N}$, d_n is a pseudo-metric in (\mathcal{F}, d) such that $\exists K_1, K_2 > 0$,

$$K_1 d_n(f_1, f_2) - \eta_{1,n} \leq d(f_1, f_2) \leq K_2 d_n(f_1, f_2) + \eta_{1,n}, \quad (3.11)$$

where $\eta_{1,n} = O_p(\gamma_{1,n})$ and $\gamma_{1,n} = o(1)$, uniformly in $f_1, f_2 \in \mathcal{F}$.

Assumption 3.3. For any $n \in \mathbb{N}$ and some $k \in \mathbb{N}$ there is enough computational power such that:

- a. (**Computation of the CM**) For some $k \in \mathbb{N}$ and any $\theta \in \Theta_n$, we can compute $\Gamma^k \theta$, where $\Gamma \theta$ is given by:

$$(\Gamma \theta)(s) = \sup_{a \in A(s)} \{u(s, a) + \beta E(F(\theta, s') | s, a)\}, \quad \forall s \in S. \quad (3.12)$$

- b. (**Approximate minimization**) We can find $\theta_n \in \Theta_n$ that satisfies:

$$d_n(\theta_n, \Gamma^k \theta_n) \leq \inf_{\theta \in \Theta_n} d_n(\theta, \Gamma^k \theta) + \eta_{2,n}. \quad (3.13)$$

where $\eta_{2,n} = O_p(\gamma_{2,n})$ and $\gamma_{2,n} = o(1)$.

Assumption 3.4. For any $f \in \mathcal{F}$:

$$\inf_{\theta \in \Theta_n} d(\theta, f) = \eta_{3,n}(f), \quad (3.14)$$

where $\eta_{3,n}(f) = O_p(\gamma_{3,n}(f))$ and $\gamma_{3,n}(f) = o(1)$.

⁴The only reason why we use $O_p(1)$ notation instead of $o(1)$ notation is that, in general, we allow for randomness in the approximation. The randomness might occur, e.g., in: (a) the choice of the sieve space or (b) the solution to the approximation problem.

We now briefly comment on each of the assumptions. Even though Assumption 3.1 might look innocuous, it is not. The definition of a contraction mapping is associated to a metric space (\mathcal{F}, d) . In particular, note that the definition is associated to the particular metric d . In other words, if (\mathcal{F}, d_1) and (\mathcal{F}, d_2) are two metric spaces, it is possible that a mapping $\Gamma : \mathcal{F} \rightarrow \mathcal{F}$ is a contraction mapping with respect to d_1 but not with respect to d_2 . This is the relevant case in the context of single agent dynamic decision problems. According to Blackwell (1965) sufficient conditions, the mapping Γ in Equation (3.6) can be shown to be a contraction mapping with respect to the sup-norm metric. Unfortunately, it is not possible to show that the same mapping is a contraction mapping with respect to other metrics, such as the l_p -metric (for some $p > 0$)⁵.

In cases in which the state space S is too large, it might not be computationally possible to work with the metric d directly, but we might be able to compute an associated pseudo-metric d_n . In order for this replacement to produce interesting results, we need to assume that the difference between the metric d and the pseudo-metric d_n vanishes at a certain rate. This is what Assumption 3.2 accomplishes.

If we had the computational capability to solve the minimization problem in Equation (3.7), then the solution to that problem would be the object of interest. The motivation for this paper, however, is to consider cases in which the size of the state space S makes solving the minimization problem in Equation (3.7) impossible. Assumption 3.3 describes precisely which operations are within our computational possibilities. Let n denote the capability of our computing power, which by assumption is limited. For any $n \in \mathbb{N}$, we choose a space of functions (sieve), denoted Θ_n so that certain mathematical operations can be approximately solved. In particular, we assume that we can: (a) compute the k^{th} iteration of the Bellman equation mapping for any function in the sieve space Θ_n and (b) minimize an objective function within the sieve space Θ_n , possibly up to a small error denoted by $\eta_{2,n}$. In case the original objective function can be exactly minimized, then we can always take $\eta_{2,n} = 0$ in Assumption 3.3.

The strategy to approximate the fixed point is to replace the original parameter space \mathcal{F} with a sequence of approximating sieve spaces $\{\Theta_n\}_{n \geq 1}$. In order to guarantee that this replacement does not affect the asymptotic properties, Assumption 3.4 requires that any function v that belongs to the original parameter space \mathcal{F} can be approximated by an element in the sieve space, possibly up to a small error denoted by $\eta_{3,n}(f)$. In order to guarantee that this error is, in fact, small, it is convenient to understand the properties of the original parameter space \mathcal{F} . The better we understand this space, the better we can choose the sieve space Θ_n to approximate it. One way to understand the parameter space \mathcal{F} is to be able to restrict it as much as possible. For a very useful result to achieve this goal, see Corollary 1 in page 52 of Stokey and Lucas (1989).

In order to illustrate these assumptions, we consider the following example.

Example 3.1 (Large but finite state space). Suppose that the agent solves the value function in Equation (3.4) where the state-space S is finite but large, i.e., $\#S = N < \infty$. By means of Corollary 1, page 52 in Stokey and Lucas (1989) we can show that the value function belongs to

⁵This is true for all formulations of the problem.

$B(S)$, i.e., the space of bounded functions that map S onto $[-B, B] \in \mathbb{R}$, for some $B < \infty$. This implies that we can take the relevant metric space to be (\mathcal{F}, d) , with $\mathcal{F} = B(S)$ and d equal to the sup-norm metric, i.e.,

$$d(f_1, f_2) = \sup_{s \in S} |f_1(s) - f_2(s)| = \max_{i=1, \dots, N} |f_1(s_i) - f_2(s_i)|, \text{ for all } f_1, f_2 \in \mathcal{F}. \quad (3.15)$$

We now verify all of the assumptions. We begin with Assumption 3.1. By arguments in Stokey and Lucas (1989), page 47, (\mathcal{F}, d) is a complete metric space of functions that map S onto \mathbb{R} . By the Blackwell (1965) sufficient conditions, it is easy to see that $\Gamma : \mathcal{F} \rightarrow \mathcal{F}$ is a contraction mapping with modulus $\beta \in (0, 1)$.

Given that $\#S = N$ is a large number, we might not be able to compute d exactly. Instead, for any $n \in \mathbb{N}$ with $n < N$, we might be able to compute:

$$d_n(f_1, f_2) = \max_{i=1, \dots, n} |f_1(s_i) - f_2(s_i)|, \text{ for all } f_1, f_2 \in \mathcal{F}. \quad (3.16)$$

This is a pseudo-norm in (\mathcal{F}, d) . Notice that for any $f_1, f_2 \in \mathcal{F}$:

$$|d_n(f_1, f_2) - d(f_1, f_2)| = \eta_{1,n}, \quad (3.17)$$

with $\eta_{1,n} = \max_{i=n+1, \dots, N} |f_1(s_i) - f_2(s_i)| \leq \max\{N - n, 0\} B$, and so $\eta_{1,n} = O_p(\gamma_{1,n})$ and $\gamma_{1,n} = o(1)$. This verifies Assumption 3.2 with $K_1 = K_2 = 1$.

Another possibility is to make the sieve space Θ_n depend only on a finite number $n \in \mathbb{N}$ of points with $n < N$ and work directly with the original metric d . In this case, $d_n = d$ and consequently, $\eta_{1,n} = 0$. A third possibility is to use an alternative pseudo-metric given by:

$$d_n(f_1, f_2) = \left(\sum_{i=1, \dots, n} |f_1(s_i) - f_2(s_i)|^p \right)^{1/p}, \text{ for all } f_1, f_2 \in \mathcal{F}, \quad (3.18)$$

for some $p > 0$ and $n \in \mathbb{N}$. In order to verify Assumption 3.2 we notice that:

$$d_n(f_1, f_2) \leq \max_{i=1, \dots, n} |f_1(s_i) - f_2(s_i)| \leq N^{1/p} d_n(f_1, f_2), \quad (3.19)$$

and, therefore:

$$d_n(f_1, f_2) + \eta_{1,n} \leq d(f_1, f_2) \leq N^{1/p} d_n(f_1, f_2) + \eta_{1,n}, \quad (3.20)$$

with $\eta_{1,n} = \max_{i=n+1, \dots, N} |f_1(s_i) - f_2(s_i)| \leq \max\{N - n, 0\} B$, and so it follows that $\eta_{1,n} = O_p(\gamma_{1,n})$ and $\gamma_{1,n} = o(1)$. This verifies Assumption 3.2 with $K_1 = 1$ and $K_2 = N^{1/p}$.

Assumption 3.3 can be assumed to hold. Part *a* is a very mild assumption. Part *b* is a much stronger assumption as it assumes that we can (approximately) solve a minimization problem with respect to the sup-norm pseudo-metric. It becomes a milder assumption if we use the l_2 pseudo-metric.

Finally, Assumption 3.4 requires that the sieve space Θ_n can approximate the original space \mathcal{F} . This is a relatively mild assumption. Then, for any $n \in \mathbb{N}$, we take Θ_n to be the set of polynomial functions defined in S of degree n . Then, Weierstrass Theorem (see, e.g., Royden (1988), page 128, proposition 8) implies that:

$$\inf_{\theta \in \Theta_n} d(\theta, v) = o(1), \quad \forall v \in \mathcal{F}, \quad (3.21)$$

as required by Assumption 3.4. Results on the rate of convergence of the $o(1)$ term are available from Chen (2007), and references therein (e.g. Lorentz (1966) and Powell (1981), Chapters 15-16.)

3.3 An iterative procedure to implement the minimization

We close this section by discussing a final practical computational issue: solving the approximate minimization problem described in Equation (3.13). For some problems the solution to Equation (3.13) might be too computationally difficult to solve directly. In order to compute a solution in these circumstances, we now introduce an iterative procedure to implement this minimization problem.

Let $\{\eta_{1,n}\}_{n \in \mathbb{N}}$ be as in Assumption 3.2 and let $k \in \mathbb{N}$ and $\{\eta_{2,n}\}_{n \in \mathbb{N}}$ be as in Assumption 3.3. Furthermore, assume that $\gamma_{1,n} = O(\gamma_{2,n})$, where $\gamma_{1,n}$ and $\gamma_{2,n}$ are as in Assumptions 3.2 and 3.3.

For an arbitrary $f \in \Theta_n$, consider the following iterative procedure:

1. Choose a function $\theta_n \in \Theta_n$ such that:

$$d_n(\theta_n, \Gamma^k f) = \inf_{\theta \in \Theta_n} d_n(\theta, \Gamma^k f). \quad (3.22)$$

2. If $\max\{d_n(\theta_n, \Gamma^k f), d_n(\theta_n, f)\} \leq \eta_{2,n}$, then stop the algorithm. Otherwise, set $f = \theta_n$ and repeat the previous step.

If the algorithm stops, then it is easy to see that $\theta_n \in \Theta_n$ satisfies Equation (3.13). To see why, consider the following argument:

$$d_n(\theta_n, \Gamma^k \theta_n) \leq d_n(\theta_n, \Gamma^k f) + d_n(\Gamma^k f, \Gamma^k \theta_n) \quad (3.23)$$

$$\leq d_n(\theta_n, \Gamma^k f) + K_1^{-1} d(\Gamma^k f, \Gamma^k \theta_n) + K_1^{-1} \eta_{1,n} \quad (3.24)$$

$$\leq d_n(\theta_n, \Gamma^k f) + K_1^{-1} \beta^k d(f, \theta_n) + K_1^{-1} \eta_{1,n} \quad (3.25)$$

$$\leq d_n(\theta_n, \Gamma^k f) + K_1^{-1} \beta^k K_2 d_n(f, \theta_n) + K_1^{-1} \beta^k K_2 \eta_{1,n} + K_1^{-1} \eta_{1,n} \quad (3.26)$$

$$\leq \eta_{2,n} (1 + K_1^{-1} \beta^k K_2) + (K_1^{-1} \beta^k K_2 + K_1^{-1}) \eta_{1,n} \quad (3.27)$$

$$= \eta'_{2,n} \leq \inf_{\theta \in \Theta_n} d_n(\theta, \Gamma^k \theta) + \eta'_{2,n}, \quad (3.28)$$

where $\eta'_{2,n} \equiv \eta_{2,n} (1 + K_1^{-1} \beta^k K_2) + (K_1^{-1} \beta^k K_2 + K_1^{-1}) \eta_{1,n}$ and, thus, under our assumptions, $\eta'_{2,n} = O_p(\gamma_{2,n})$ and $\gamma_{2,n} = o(1)$. The first inequality follows from the triangular inequality (applied to the pseudo-metric d_n), the second and fourth inequalities follow from Assumption 3.2, the third

inequality follows from the fact that Γ is a contraction mapping with modulus β , the fifth inequality follows from the stopping rule in the algorithm, the following equality follows by definition of $\eta'_{2,n}$, and the final inequality follows from the fact that the pseudo-metric d_n is positive. The result is completed by noticing that, under our assumptions, $\eta'_{2,n} = O(\eta_{2,n})$.

In practice, the minimization problem in Equation (3.22) can be relatively easy to solve and, in particular, much easier than the approximate minimization problem in Equation (3.13). For example, when d_n is the l_2 -metric (i.e. sum of squared differences) and Θ_n is the space of polynomials then the minimization problem in Equation (3.22) is a least-squares problem.

Having defined the dynamic decision problem and laid out basic definitions and assumptions, we now introduce our method for approximating the value function and prove a series of results regarding the properties of this approximation. We begin by analyzing infinite horizon problems, taking up finite horizon problems in the following section of the paper.

3.4 Definition of the approximation

Under Assumption 3.1, the Contraction Mapping Theorem (see, e.g., Stokey and Lucas (1989), Theorem 3.2, page 50) indicates that the contraction mapping Γ defined by Equation (3.6) has a unique fixed point. As we have explained, Equation (3.4) implies this fixed point is the value function of interest, which we have denoted by $V \in \mathcal{F}$. As we have explained, this motivates us to consider the (unfeasible) optimization problem in Equation (3.7) or, more generally, for some $k \in \mathbb{N}$, the following (equally unfeasible) optimization problem:

$$\inf_{\theta \in \mathcal{F}} d(\theta, \Gamma^k \theta). \quad (3.29)$$

Even though the objective function can be computed, the sample space of the problem under consideration makes the domain of the optimization, \mathcal{F} , too complex to handle. In order to circumvent this issue, we consider replacing the space \mathcal{F} with a sequence of approximating spaces or sieves. According to Assumption 3.4, the sequence of spaces $\{\Theta_n\}_{n \geq 1}$ becomes a good approximation of \mathcal{F} as our computational possibilities, n , diverge to infinity. In particular, this motivates us to consider the following alternative optimization problem:

$$\inf_{\theta \in \Theta_n} d(\theta, \Gamma^k \theta). \quad (3.30)$$

In certain situations, minimizing with respect to the metric d might not be computationally easy or even possible. In those cases, it is convenient to consider replacing the metric d with a suitable pseudo-metric d_n according to Assumption 3.2. This leads us to consider the following optimization problem:

$$\inf_{\theta \in \Theta_n} d_n(\theta, \Gamma^k \theta). \quad (3.31)$$

In certain settings, the above minimization problem might only be feasible up to a certain residual term that will vanish as n diverges to infinity. This is precisely the case described in Assumption

3.3 (part *b*). This progression naturally leads to the definition of our SVFI approximation.

Definition 3.3 (Sieve Value Function Approximation). Assume Assumption 3.3. Then the SVFI approximation of V is any $\hat{\theta}_n \in \Theta_n$ that satisfies:

$$d_n(\hat{\theta}_n, \Gamma^k \hat{\theta}_n) \leq \inf_{\theta \in \Theta_n} d_n(\theta, \Gamma^k \theta) + \eta_{2,n}. \quad (3.32)$$

where $\eta_{2,n} = O_p(\gamma_{2,n})$ and $\gamma_{2,n} = o(1)$.

Unlike the optimization problems in Equations (3.7) or (3.29), Assumption 3.3 guarantees that the SVFI approximation is the (approximate) solution to the feasible minimization problem in Equation (3.32).

3.5 Properties of the approximation

All of the findings of this section are corollaries of the following lemma.

Lemma 3.1. *Assume Assumptions 3.1-3.4. Then, the SVFI approximation satisfies:*

$$d(\hat{\theta}_n, V) \leq \frac{(1 + K_2 K_1^{-1})\eta_{1,n} + K_2 \eta_{2,n} + K_1^{-1} K_2 (1 + \beta^k) \eta_{3,n}(V)}{1 - \beta^k}, \quad (3.33)$$

where $V \in \mathcal{F}$ is the unique fixed point of Γ in (\mathcal{F}, d) .

Lemma 3.1 is the key result to establish the consistency of the SVFI approximation, derive its rate of convergence, and investigate the finite sample properties of its approximation error. The following result establishes the consistency and the rate of convergence of the approximation.

Theorem 3.1. *Assume Assumptions 3.1-3.4. Then, the SVFI approximation satisfies:*

$$d(\hat{\theta}_n, V) = O_p(\max\{\gamma_{1,n}, \gamma_{2,n}, \gamma_{3,n}(V)\}), \quad (3.34)$$

where $\max\{\gamma_{1,n}, \gamma_{2,n}, \gamma_{3,n}(V)\} = o(1)$ as $n \rightarrow \infty$, and $V \in \mathcal{F}$ is the unique fixed point of Γ in (\mathcal{F}, d) . This implies that the SVFI approximation:

1. is consistent approximation of V , i.e., $d(\hat{\theta}_n, V) = o_p(1)$, as $n \rightarrow \infty$.
2. converges to V at a rate of $\min\{\gamma_{1,n}^{-1}, \gamma_{2,n}^{-1}, \gamma_{3,n}^{-1}(V)\}$.

Theorem 3.1 implies that the rate of convergence of the approximation depends on the rate at which three errors converge to zero. These errors are:

1. The error of approximating the metric d with the approximate pseudo-metric d_n according to Equation (3.11), denoted by $\eta_{1,n}$, which converges to zero at a rate of $\gamma_{1,n}^{-1}$.
2. The error of approximation when minimizing the objective function $d_n(\theta, \Gamma^k \theta)$ in Equation (3.32), denoted by $\eta_{2,n}$, which converges to zero at a rate of $\gamma_{2,n}^{-1}$.

Table 1: Value of constants associated to the upper bound on the error of approximation for different number of iterations.^a

	Number of iterations: k						
	1	2	3	4	5	6	7
$1/(1 - \beta^k)$	10	5.26	3.69	2.91	2.44	2.13	1.92
$(1 + \beta^k)/(1 - \beta^k)$	19	9.53	6.38	4.82	3.88	3.27	2.83

^a The discount factor β is set to 0.9.

3. The error of approximating the value function $V \in \mathcal{F}$ with an element in the sieve space Θ_n , denoted by $\eta_{3,n}(V)$, which converges to zero at a rate of $\gamma_{3,n}^{-1}(V)$.

The slowest of these three rates determines the rate of convergence of the approximation. The motivation for introducing sieve approximations was the fact that working with the original space \mathcal{F} was computationally infeasible, i.e., the third source of error $\eta_{3,n}(V)$ cannot be avoided. However, it might be possible to avoid the other sources of error. In other words, it might be possible to use $d_n = d$, leading to $\gamma_{1,n} = 0$, or it might be possible to solve the minimization of $d_n(\theta, \Gamma^k \theta)$ exactly, leading to $\gamma_{2,n} = 0$. If so, then, the convergence rate of this source of error is infinity, effectively disappearing from the expression for the rate of convergence $\min\{\gamma_{1,n}^{-1}, \gamma_{2,n}^{-1}, \gamma_{3,n}^{-1}(V)\}$.

It is interesting to notice that the findings in Theorem 3.1 do not depend on the number of contraction mapping iterations k in Assumption 3.3 (in particular, they hold even if $k = 1$). The choice of k affects several constants associated with the rates of convergence, which are “hidden” in the O_p notation. While these constants are not relevant for the asymptotic results, they can be very relevant for finite values of the computational power n . The right tool for this analysis is Equation (3.33) in Lemma 3.1, which provides a concrete upper bound on the error of approximation that can be used to study the effect of changes in k that is valid for any value of n . This result reveals that the three sources of error are each associated with a constant that depend on k . In particular, the error terms $\eta_{1,n}$, $\eta_{2,n}$, and $\eta_{3,n}$ are associated to the constants $(1 + K_2 K_1^{-1})/(1 - \beta^k)$, $K_2/(1 - \beta^k)$, and $K_1^{-1} K_2 (1 + \beta^k)/(1 - \beta^k)$, respectively. A corollary of this is that, *ceteris paribus*, an increase in the number of value function iterations k reduces the value of the upper bound of the error of approximation. In particular, Table 1 illustrates that there are significant gains in precision from raising the value of k when the discount factor is $\beta = 0.9$. For example, changing the number of iterations from $k = 1$ to $k = 2$ reduces the (upper bound on the) error of approximation by approximately 50%.

We illustrate the tradeoffs associated with increasing the number of contraction mapping iterations k in the Monte Carlo analysis in Section 6.

4 Approximation in finite horizon dynamic problems

This section describes the application of non-parametric sieve methods to approximate the value functions in Equation (2.1) when the dynamic problem has a finite horizon (i.e. $T < \infty$). The distinctive feature of the finite horizon problem is that the value function of the problem is no longer stationary (i.e. $V_t(\cdot)$ depends on the time index). As time passes and the terminal period T approaches, the agent's value function changes.

Using backward induction, the value function in any given period can be expressed the optimized choice between instantaneous utility and the discount value for the immediately proceeding period. By repeating the arguments used to develop the unified formulation used in the infinite horizon case, we use s to denote the current value of the state, s' to denote the future value of the state, S to denote the state space, a to denote the action chosen by the agent, $A(s)$ to denote the space of actions available to the agent when the state is s , and F to be a known functional of $V_{t+1}(\cdot)$ and s' that satisfies certain known properties. Based on this notation, in every period $t = 1, \dots, T$, the agent solves an optimization problem characterized by the following value function:

$$V_t(s) = \max_{a \in A(s)} \{u(s, a) + \beta E(F(V_{t+1}, s')|s, a)\}, \quad \forall s \in S, \quad (4.1)$$

where the variables and functions are defined as in Equation (3.4) and:

$$V_{T+1}(s) = 0, \quad \forall s \in S. \quad (4.2)$$

Using the notation developed in Equation (3.6), the sequence of value functions $\{V_t\}_{t=1}^T$ can be defined as follows: $V_T = \Gamma V_{T+1}$ with a zero terminal value, i.e., $V_{T+1}(s) = 0, \forall s \in S$.⁶

The approximation procedure developed for the infinite horizon problem must be modified to accommodate several distinct features of the finite horizon setup. First, the finite horizon problem requires an approximation for the value function at each point in time; the infinite horizon problem is stationary, i.e., the agent solves the same problem every period, and thus only requires the approximation of a single value function V . Second, the approximation procedure developed for the infinite horizon problem required the value function V to be a fixed point in a contraction mapping. Clearly, this will not be true for the non-stationary finite horizon problem. As time progresses, the last period of the game approaches and this affects the value of participating in the game. Thus the value functions of the finite horizon problem are not fixed points to any mapping, but are instead a finite sequence of functions that are sequentially related.

With enough computational power, the set of functions $\{V_t\}_{t=1}^T$ could be computed exactly using backward induction. Nevertheless, for economic models with large state spaces, the exact implementation of backward induction might be too computationally challenging or even impossible.

⁶It should be noted that each of the objects in Equation (4.1), i.e., the set of possible actions $A(\cdot)$, the period utility function $u(\cdot)$, the expectation operator $E(\cdot)$, and the functional $F(\cdot)$, could be allowed to be time-specific without affecting any of the theoretical results to follow. We opted to keep these elements time invariant to simplify the exposition and to relate them easily to elements in Equation (3.4).

The objective of this section is to propose a sieve value function approximation for such settings. For a given computational power (i.e. for a given sieve space Θ_n), our approximation method produces a sequence of approximations: $\{\hat{\theta}_{n,t}\}_{t=1}^T$ where, for all $t = 1, \dots, T$, $\hat{\theta}_{n,t}$ is the sieve value function approximation for V_t .

We consider two approximation procedures. The first is, essentially, the sieve approximation version of a traditional backward induction. The value function is first approximated for the last period and this approximate function is used to solve for the (approximate) value function for the previous period. Continuing to work backwards yields an approximate value function for each period. Implementing this procedure requires using computational routines that are specifically tailored for the finite horizon setup.

The second procedure entails expanding the state space to include time as a state variable. To the best of our knowledge, this procedure is novel to our paper. While this procedure is less intuitive than backward induction, it has the advantage of being implemented with the same computational routines developed above for the infinite horizon problem.

4.1 Approximation using backward induction

The computation of the sequence of value functions $\{V_t\}_{t=1}^T$ by backward induction is well understood and requires no further discussion. This section proposes an approximation to these value functions using sieves. The approximation requires the following assumptions. By repeating previous arguments, it is easy to see that these assumptions are satisfied in dynamic decision problems that have a very large but finite state space.

Assumption 4.1. (\mathcal{F}, d) is a complete metric space of functions that map S onto \mathbb{R} , where d is the sup-norm metric, i.e.,

$$d(f_1, f_2) = \sup_{s \in S} \|f_1(s) - f_2(s)\|, \quad \forall f_1, f_2 \in \mathcal{F}. \quad (4.3)$$

Assumption 4.2. For any $n \in \mathbb{N}$, d_n is a pseudo-metric in (\mathcal{F}, d) such that $\exists K_1, K_2 > 0$,

$$K_1 d_n(f_1, f_2) - \lambda_{1,n} \leq d(f_1, f_2) \leq K_2 d_n(f_1, f_2) + \lambda_{1,n}, \quad (4.4)$$

where $\lambda_{1,n} = O_p(v_{1,n})$ and $v_{1,n} = o(1)$, uniformly in $f_1, f_2 \in \mathcal{F}$.

Assumption 4.3. For any $n \in \mathbb{N}$ there is enough computational power such that:

a. (**Computation of the CM**) For any $\theta \in \Theta_n \cup V_{T+1}$, we can compute $\Gamma\theta$ given by:

$$(\Gamma\theta)(s) = \sup_{a \in A(s)} \{u(s, a) + \beta E(F(\theta, s')|s, a)\}, \quad \forall s \in S. \quad (4.5)$$

b. (**Approximate minimization**) For any $f \in \mathcal{F}$, we can find $\theta_n \in \Theta_n$ that satisfies:

$$d_n(\theta_n, f) \leq \inf_{\theta \in \Theta_n} d_n(\theta, f) + \lambda_{2,n}. \quad (4.6)$$

where $\lambda_{2,n} = O_p(v_{2,n})$ and $v_{2,n} = o(1)$, uniformly in $f \in \mathcal{F}$.

Assumption 4.4. For any $f \in \mathcal{F}$:

$$\inf_{\theta \in \Theta_n} d(\theta, f) = \lambda_{3,n}, \quad (4.7)$$

where $\lambda_{3,n} = O_p(v_{3,n})$ and $v_{3,n} = o(1)$, uniformly in $f \in \mathcal{F}$.

Assumption 4.5. For all $f_1, f_2 \in \mathcal{F}$, $a \in A(s)$, and $s \in S$:

$$E[F(f_1, s') - F(f_2, s') | s, a] \leq d(f_1, f_2). \quad (4.8)$$

We now briefly comment on each of the assumptions with focus on the differences with Assumptions 3.1-3.4. With respect to Assumption 3.1, Assumption 4.1 eliminates the contraction mapping requirement with the requirement that d is the sup-norm metric. As we have already explained, these two assumptions are not that different, as the mapping Γ in Equation (3.6) can be shown to be a contraction mapping with respect to the sup-norm metric but not with respect to other metrics. Assumption 4.2 is identical to Assumption 3.2. Assumption 4.3 is very similar to Assumption 3.3. The differences between the two are the following. First, Assumption 3.3(a) assumed that one could compute the k^{th} iteration of the Bellman equation mapping for any function in the sieve space Θ_n , whereas Assumption 4.3(a) explicitly requires this to be true for $k = 1$. Second, Assumption 3.3(b) assumed that one could (approximately) minimize a specific objective function within the sieve space Θ_n , whereas Assumption 4.3(b) assumes that one can (approximately) find the best approximation within the sieve space Θ_n for any function in \mathcal{F} . The difference between Assumption 3.3(b) and Assumption 4.3(b) is that the error of minimization has to converge to zero (in probability) uniformly in $f \in \mathcal{F}$. Assumption 4.4 strengthens Assumption 3.4 as it requires that $\inf_{\theta \in \Theta_n} d(\theta, f)$ converges to zero (in probability) uniformly in $f \in \mathcal{F}$. In other words, instead of requiring a vanishing error of approximation of any particular function, we require that the a vanishing error of approximation for the *worst* function in the class of functions \mathcal{F} . For references on these stronger results see, e.g., chapter 8 in Lorentz (1966). Finally, Assumption 4.5 is a mild assumption about the properties of the mapping F . In particular, Lemma A.1 verifies that this assumption holds for all possible the formulations of the problem.

The approximation considered in this section is defined as follows:

Definition 4.1 (Sieve Value Function Approximation). Assume Assumption 4.3. Then the approximation of $\{V_t\}_{t=1}^T$ is $\{\hat{\theta}_{n,t}\}_{t=1}^T$ constructed in the following iterative manner. For $z = 1, \dots, T$, let $t = T + 1 - z$ and complete the following steps:

1. Define $\hat{\theta}_{t,n}^* : S \rightarrow \mathbb{R}$ as follows:

$$\hat{\theta}_{t,n}^*(s) = (\Gamma \hat{\theta}_{t+1,n})(s) = \max_{a \in A(s)} \{u(s, a) + \beta E(F(\hat{\theta}_{t+1,n}, s') | s, a)\}, \quad \forall s \in S, \quad (4.9)$$

where either: $t = T$ and $\hat{\theta}_{t+1,n}(s) = V_{T+1}(s) = 0 \quad \forall s \in S$, or $t < T$ and $\hat{\theta}_{t+1,n}$ has been defined in a previous iteration of the algorithm.

2. Define $\hat{\theta}_{t,n} : S \rightarrow \mathbb{R}$ to be any $\hat{\theta}_{t,n} \in \Theta_n$ that satisfies:

$$d_n(\hat{\theta}_{t,n}, \hat{\theta}_{t,n}^*) \leq \inf_{\theta \in \Theta_n} d_n(\theta, \hat{\theta}_{t,n}^*) + \lambda_{2,n}, \quad (4.10)$$

where $\lambda_{2,n} = O_p(v_{2,n})$ and $v_{2,n} = o(1)$.

It is evident from the description of the procedure that this method implements the traditional backward induction procedure using sieve approximations, i.e., it performs an approximation of the value function of the terminal period and uses the approximation of the value function in a given period to conduct an approximation for the value function in the immediately preceding period. Under Assumption 4.3, this approximation is feasible and is completed by solving T sequential approximation problems. The following theorem establishes the asymptotic properties of the approximation.

Theorem 4.1. *Assume Assumptions 4.1-4.5. Then, the approximation satisfies:*

$$\max_{t=1, \dots, T} d(\hat{\theta}_{t,n}, V_t) = O_p(\max\{v_{1,n}, v_{2,n}, v_{3,n}\}), \quad (4.11)$$

where $\max\{v_{1,n}, v_{2,n}, v_{3,n}\} = o(1)$ as $n \rightarrow \infty$. This implies that for all $t = 1, \dots, T$, the sieve value function approximation $\hat{\theta}_{t,n}$:

1. is a consistent approximation of V_t , i.e., $d(\hat{\theta}_{t,n}, V_t) = o_p(1)$, as $n \rightarrow \infty$.
2. converges to V_t at a rate of $\min\{v_{1,n}^{-1}, v_{2,n}^{-1}, v_{3,n}^{-1}\}$.

As in Theorem 3.1, Theorem 4.1 indicates that the rate of convergence of the approximation depends on the rate at which three errors converge to zero. The slowest of these three rates determines the rate of convergence of the approximation.

4.2 Approximation using time as a state variable

The approximation considered in this section entails considering the time dimension as part of the state of the problem. In some sense, it may seem counterintuitive to “increase” the state space for a problem in which the size of space was already deemed to large to compute directly. However, as we demonstrate in this section, the approximation is computationally feasible and can be implemented using the exact same computational tools as in the infinite horizon case.

Consider the state space that results from the cartesian product of the (time invariant) state space S with the time dimension $\{1, \dots, T + 1\}$, i.e., $\tilde{S} = S \times \{1, \dots, T + 1\}$. Throughout this section, we superscript with the symbol \sim to denote objects in the new state space that includes the time dimension. For example, the new metric space is denoted by $(\tilde{\mathcal{F}}, \tilde{d})$, where $\tilde{\mathcal{F}}$ denotes a set functions from \tilde{S} onto \mathbb{R} and \tilde{d} denotes the corresponding norm in this space. In this enlarged state space, the sequence of value functions $\{V_t\}_{t=1}^{T+1}$ defined by Equations (4.1) and (4.2) can be equivalently re-written as follows:

$$V(s, t) \equiv V_t(s). \quad (4.12)$$

In order to conduct the approximation in the state space \tilde{S} , we use a sequence of sieve spaces, denoted by $\{\tilde{\Theta}_n\}_{n \geq 1}$, where each $\tilde{\Theta}_n$ is a space of (simple) functions that map \tilde{S} onto \mathbb{R} . We consider the following assumptions.

Assumption 4.6. $(\tilde{\mathcal{F}}, \tilde{d})$ is a complete metric space of functions that map \tilde{S} onto \mathbb{R} , where \tilde{d} is the sup-norm metric, i.e.,

$$\tilde{d}(f_1, f_2) = \sup_{(s,t) \in \tilde{S}} \|f_1(s, t) - f_2(s, t)\|, \quad \forall f_1, f_2 \in \tilde{\mathcal{F}}. \quad (4.13)$$

Assumption 4.7. For any $n \in \mathbb{N}$, \tilde{d}_n is a pseudo-metric in $(\tilde{\mathcal{F}}, \tilde{d})$ such that $\exists K_1, K_2 > 0$,

$$K_1 \tilde{d}_n(f_1, f_2) - \lambda_{1,n} \leq \tilde{d}(f_1, f_2) \leq K_2 \tilde{d}_n(f_1, f_2) + \lambda_{1,n}, \quad (4.14)$$

where $\lambda_{1,n} = O_p(v_{1,n})$ and $v_{1,n} = o(1)$, uniformly in $f_1, f_2 \in \tilde{\mathcal{F}}$.

Assumption 4.8. For any $n \in \mathbb{N}$ and some $k \in \mathbb{N}$ there is enough computational power such that:

a. **(Computation of the CM)** For some $k \in \mathbb{N}$ and any $\theta \in \Theta_n$, we can compute $\Lambda^k \theta$, where $\Lambda \theta$ is given by:

$$(\Lambda \theta)(s, t) = \sup_{a \in A(s)} \{u(s, a) + \beta E(F(\theta, (s', t + 1)|(s, t), a))\} \times 1[t < T + 1]. \quad (4.15)$$

b. **(Approximate minimization)** For any $f \in \tilde{\mathcal{F}}$, we can find $\theta_n \in \tilde{\Theta}_n$ that satisfies:

$$\tilde{d}_n(\theta_n, \Lambda \theta_n) \leq \inf_{\theta \in \tilde{\Theta}_n} \tilde{d}_n(\theta, \Lambda \theta) + \lambda_{2,n}. \quad (4.16)$$

where $\lambda_{2,n} = O_p(v_{2,n})$ and $v_{2,n} = o(1)$.

Assumption 4.9. For any $f \in \tilde{\mathcal{F}}$:

$$\inf_{\theta \in \tilde{\Theta}_n} \tilde{d}(\theta, f) = \lambda_{3,n}(f), \quad (4.17)$$

where $\lambda_{3,n}(f) = O_p(v_{3,n}(f))$ and $v_{3,n}(f) = o(1)$.

Assumption 4.10. For all $f_1, f_2 \in \tilde{\mathcal{F}}$, $a \in A(s)$, and $(s, t) \in \tilde{S}$:

$$E[F(f_1, (s', t + 1)) - F(f_2, (s', t + 1)) | (s, t), a] \leq d(f_1, f_2). \quad (4.18)$$

With the exception of the fact that the state space has been enriched with the time index, Assumptions 4.6-4.10 are analogous to assumptions that have already been discussed in the paper. On the one hand, Assumptions 4.6, 4.7, and 4.10 are analogous to Assumptions 4.1, 4.2, and 4.5 used to consider the sieve approximation to the backward induction solution in finite horizon problems. On the other hand, Assumptions 4.8 and 4.9 are analogous to Assumptions 3.3 and 3.4 used to consider the SVFI approximation in infinite horizon problems.

In the context of the SVFI approximation in infinite horizon problems, we described an iterative procedure for (approximately) minimizing the objective function. A similar iterative procedure can be developed for finite horizon problems. In particular, let $\{\lambda_{1,n}\}_{n \in \mathbb{N}}$ be as in Assumption 4.7 and let $k \in \mathbb{N}$ and $\{\lambda_{2,n}\}_{n \in \mathbb{N}}$ be as in Assumption 4.8. Furthermore, assume that $v_{1,n} = O(v_{2,n})$, where $v_{1,n}$ and $v_{2,n}$ are as in Assumptions 4.7 and 4.8.

For an arbitrary $f \in \tilde{\Theta}_n$, consider the following iterative procedure:

1. Choose a function $\theta_n \in \tilde{\Theta}_n$ such that:

$$\tilde{d}_n(\theta_n, \Lambda^k f) = \inf_{\theta \in \tilde{\Theta}_n} \tilde{d}_n(\theta, \Lambda^k f). \quad (4.19)$$

2. If $\max\{\tilde{d}_n(\theta_n, \Lambda^k f), \tilde{d}_n(\theta_n, f)\} \leq \lambda_{2,n}$, then stop the algorithm. Otherwise, set $f = \theta_n$ and repeat the previous step.

If the algorithm stops, then it implies that $\theta_n \in \tilde{\Theta}_n$ satisfies Equation (4.16). To show this, we can repeat the argument used in Section 3.3. In practice, the minimization problem in Equation (4.19) can be relatively easy to solve and, in particular, much easier than the approximate minimization problem in Equation (4.16).

The following result is the key to the asymptotic findings of this section.

Lemma 4.1. Assume Assumptions 4.6 and 4.10. Let Λ be the mapping in Assumption 4.8 and let $V \in \tilde{\mathcal{F}}$ be the function defined by: $V(t, s) \equiv V_t(s)$ for all $S \times \{1, \dots, T + 1\}$. Then:

1. Λ is a contraction mapping with modulus β on $(\tilde{\mathcal{F}}, \tilde{d})$.
2. V is the unique fixed point of the contraction mapping Λ , i.e., $\Lambda V = V$.

In the context of infinite horizon problems, Assumption 3.1 indicated that the value function was the unique fixed point of a certain contraction mapping. This result was the key to proposing the SVFI approximation method. Lemma 4.1 indicates that an analogous result holds for in the context of finite horizon problem. In fact, if we combine this result with the remaining assumptions, the current setup satisfies all of the conditions required for the SVFI approximation. As a consequence, an analogous approximation to the the SVFI approximation will have the same asymptotic properties, i.e, consistency and rates of convergence. This analogous approximation is defined next.

Definition 4.2 (Sieve Value Function Approximation). Assume Assumption 4.8. Then the sieve approximation of $\{V_t\}_{t=1}^T$ is $\{\hat{\theta}_{t,n}\}_{t=1}^T$ where, for every $(s, t) \in S \times \{1, \dots, T\}$, $\hat{\theta}_{t,n}(s) \equiv \hat{\theta}_n(s, t)$ and $\hat{\theta}_n \in \tilde{\Theta}_n$ is any function that satisfies:

$$\tilde{d}_n(\hat{\theta}_n, \Lambda^k \hat{\theta}_n) \leq \inf_{\theta \in \tilde{\Theta}_n} \tilde{d}_n(\theta, \Lambda^k \theta) + \lambda_{2,n}. \quad (4.20)$$

where $\lambda_{2,n} = O_p(v_{2,n})$ and $v_{2,n} = o(1)$.

Based on the previous discussion, the following result is a simple corollary of Theorem 3.1 and Lemma 4.1.

Theorem 4.2. *Assume Assumptions 4.6-4.10. Then, the function $\hat{\theta}_n \in \tilde{\Theta}_n$ in Definition 4.2 satisfies:*

$$\tilde{d}(\hat{\theta}_n, V) = O_p(\max\{v_{1,n}, v_{2,n}, v_{3,n}(V)\}), \quad (4.21)$$

where $\max\{v_{1,n}, v_{2,n}, v_{3,n}(V)\} = o(1)$ as $n \rightarrow \infty$ and $V \in \tilde{\mathcal{F}}$ is the function defined by: $V(t, s) \equiv V_t(s)$ for all $S \times \{1, \dots, T+1\}$. This implies that for all $t = 1, \dots, T$, the sieve approximation $\hat{\theta}_{t,n}$:

1. is a consistent approximation of V_t , i.e., $\sup_{s \in S} |\hat{\theta}_{t,n}(s) - V_t(s)| = o_p(1)$, as $n \rightarrow \infty$.
2. converges to V_t at a rate of $\min\{v_{1,n}^{-1}, v_{2,n}^{-1}, v_{3,n}^{-1}(V)\}$.

As in Theorems 3.1 and 4.1, Theorem 4.2 indicates that the rate of convergence of the approximation depends on the rate at which three errors converge to zero. Again, the slowest of these three rates determines the rate of convergence of the approximation.

5 Estimation

The results to this point in the paper characterize the (approximate) computation of the value function (and associated object of interests) for a known vector parameters π that characterize the agent's dynamic decision problem. In this section, we now consider the problem of estimating π in a parameter space Π when the researcher has data on dynamic decisions and, again, the state space is too large to permit the direct computation of V for a given value of π .⁷ As before, the associated value function V incorporates all the information that is relevant to the decision problem and depends on the parameter π , i.e., $V(\cdot|\pi)$. In this setup, the approximation problem of previous sections entails the approximation of the value function V for a particular value of the parameter $\pi \in \Pi$. Let π^* denote the true parameter value, i.e., $V(\cdot) \equiv V(\cdot|\pi^*)$.

For concreteness, consider an agent solving the value function in Equation (3.4) with:

$$AS \equiv \{(a, s) : a \in A(s) \text{ and } \forall s \in S\}, \quad (5.1)$$

⁷Throughout this section, we pretend that the dynamic problem we refer to is the infinite horizon single agent decision problem. Nevertheless, by making slight notational changes, the results of this section can also be applied to a finite horizon problem.

instantaneous utility function given by:

$$u(s, a) = u(s, a|\pi_1), \quad \forall (a, s) \in AS, \quad (5.2)$$

transition probabilities given by:

$$dP(\varepsilon', x'|x, a) = dP(\varepsilon', x'|x, a, \pi_2), \quad \forall (a, s) \in AS, \quad (5.3)$$

and a discount factor $\beta = \pi_3$. Set $\pi \equiv (\pi_1, \pi_2, \pi_3) \in \Pi$. For each parameter value $\pi \in \Pi$, the corresponding value function is denoted $V(\cdot|\pi)$.

The SVFI approximation procedure described in Definition 3.3 provides a method for approximating the value function for any given set of primitives (i.e. choice set, utility function, discount factor, and transition probabilities). In this setting, these primitives are functions of the unknown parameter value π . To be consistent with this interpretation, we denote: $\theta_n \equiv \theta_n(\cdot|\pi)$.

If the value function (or some function derived from it) were observed for an arbitrary set of values, then the SVFI approximation procedure could be used to approximate it. Furthermore, the consistency result in Theorem 3.1 suggests that the parameter π could be estimated as follows:

$$\hat{\pi}_n = \arg \min_{\pi \in \Pi} Q(\theta_n(\cdot|\pi)), \quad (5.4)$$

where $Q(\theta_n(\cdot|\pi))$ is an appropriately chosen function that measures the distance between $\theta_n(\cdot|\pi)$ and the value function V . For instance, Q could be the following function:

$$Q(\theta_n(\cdot|\pi)) = \int (V(s) - \theta_n(s|\pi))^2 d\mu(s), \quad (5.5)$$

where μ is any arbitrary positive measure over S . In practice, however, the estimator in Equation (5.4) is unfeasible because the value function is not observed by the researcher and, consequently, Q is unknown.

In practice, the value function, or some feature derived from it (e.g., CCPs), can be estimated from the data. With some abuse of notation, let V_I denote the estimated value function using a sample of size I (observations are indexed $i = 1, \dots, I$) and let Q_I denote the function that measures the distance between $\theta_n(\cdot|\pi)$ and V_I . Our previous discussion suggests that the parameter π could be estimated according to the following definition.

Definition 5.1 (Estimator based on the Sieve Value Function Approximation). Let $Q_I : \Pi \rightarrow \mathbb{R}_+$ be the function of the data that measures the distance between V_I and $\theta_n(\cdot|\pi)$ for any $\pi \in \Pi$ and let $n = n(I)$. The estimator of the true parameter value π^* , denoted $\hat{\pi}_I$, satisfies:

$$Q_I(\hat{\theta}_n(\cdot|\hat{\pi}_I)) \leq \inf_{\pi \in \Pi} Q_I(\hat{\theta}_n(\cdot|\pi)) + o_p(1), \text{ as } I \rightarrow \infty. \quad (5.6)$$

In order to clarify the structure of the problem, we consider an illustrative example.

Example 5.1. In this case, we estimate the CCPs from a sample of I observed choices. In particular, if the set S is finite, then the following estimator:

$$\hat{P}_I = \{\hat{P}_I(a|s), \forall (a, s) \in AS\}, \quad (5.7)$$

where:

$$\hat{P}_I(a|s) = \sum_{i=1}^I 1[a_i = a, s_i = s] / \sum_{i=1}^I 1[s_i = s], \quad (5.8)$$

is a \sqrt{I} -consistent estimator of the CCPs.

By definition, the CCPs can be derived from the value function. Let $J(\cdot|\pi) : \mathcal{F} \rightarrow [0, 1]^{\#AS}$ be the mapping between the value function and the CCPs, i.e.,

$$J(V|\pi) = \{P(a|s), \forall (a, s) \in AS\}, \quad (5.9)$$

where the conditioning on π indicates that the mapping itself could depend on the parameter value.⁸ This discussion suggests the estimation of π with $\hat{\pi}_I$ as in Equation (5.6), where the measure μ is set to be the empirical measure, i.e.,

$$Q_I(\theta_n(\cdot|\pi)) = \int (\hat{P}_I - J(\theta_n(\pi)|\pi))^2 d\hat{P}_I, \quad (5.10)$$

$$= I^{-1} \sum_{i=1}^I 1[(a_i, s_i) = (a, s)] (\hat{P}_I(a|s) - J_{(a,s)}(V|\pi))^2. \quad (5.11)$$

In other words, we choose the parameter value that minimizes the integrated squared distance between the observed CCPs and the approximated CCPs, where the empirical measure is used as the measure of integration.

The objective of the rest of the section is to provide conditions under which $\hat{\pi}_I$ is consistent. To this end, we now provide a list of the assumptions that are exclusively used in this section.

Assumption 5.1. (\mathcal{F}_1, d_1) and (\mathcal{F}_2, d_2) are metric spaces. The true value function $V = V(\cdot|\pi^*) \in \mathcal{F}_1$ and the true parameter value $\pi^* \in \Pi \subseteq \mathcal{F}_2$.

Assumption 5.2. The SVFI approximation satisfies:

$$\sup_{\pi \in \Pi} d_1(\hat{\theta}_n(\cdot|\pi), V(\cdot|\pi)) = o_p(1), \text{ as } n \rightarrow \infty. \quad (5.12)$$

Assumption 5.3. There is a function $Q(\cdot) : \mathcal{F}_1 \rightarrow \mathbb{R}$ such that:

$$a. \hat{c}(n, I) = o_p(\delta(n, I)) \text{ where } \hat{c}(n, I) \equiv \sup_{f \in \Theta_n} |Q_I(f) - Q(f)|.$$

⁸The function J will depend on the specific formulation of the dynamic decision problem. See Appendix B for a description of each of the formulations and the definition of the function J in each case.

- b. Q is uniformly continuous in \mathcal{F}_1 under d_1 , i.e., for any $\delta > 0$, there exists $\varepsilon > 0$ such that $f_1, f_2 \in \mathcal{F}_1$ with $d_1(f_1, f_2) < \varepsilon$ implies $Q(f_1) - Q(f_2) \leq \delta$.
- c. The function $Q(V(\cdot|\pi)) : \Pi \rightarrow \mathbb{R}$ is uniquely minimized at $\pi = \pi^*$.

Assumption 5.4. For every $I \in \mathbb{N}$, $n = n(I)$ with $n(I) \rightarrow \infty$ and $\delta(n(I), I) \rightarrow 0$ as $I \rightarrow \infty$.

We now briefly comment on each of the assumptions. Assumption 5.1 provides a label to relevant spaces, functions, and parameters. Assumption 5.2 is a high level assumption which indicates that $\hat{\theta}_n(\cdot|\pi)$ is a consistent approximation of $V(\cdot|\pi)$, uniform over $\pi \in \Pi$. By repeating the arguments in Theorem 3.1, it is not hard to see that Assumption 5.2 holds as a result of using the SVFI approximation procedure described in Definition 3.3 under Assumptions 3.1-3.4, with the exception that Assumption 3.4 is strengthened to hold uniformly in $f \in \mathcal{F}$ (i.e. as in Assumption 4.4). Assumption 5.3 are similar to assumptions used in the literature on extremum estimators (see, e.g., Theorem 3.1 in Chen (2007) for conditions pertaining to sieve estimators, and Theorem 4.1.1 in Amemiya (1985) or Theorem 2.1 in McFadden and Newey (1994) for finite dimensional estimators). Finally, Assumption 5.4 describes conditions that restrict the relationship between the data sample size and the complexity of the sieve.

Theorem 5.1. Assume Assumptions 5.1-5.4. Then, the estimator in Equation (5.4) is consistent, i.e.,

$$d_2(\hat{\pi}_I, \pi^*) = o_p(1), \text{ as } I \rightarrow \infty. \tag{5.13}$$

In this way, under Assumptions 5.1-5.4, it is possible to obtain consistent estimates of the parameters of large state space dynamic decision problems by embedding SVFI methods in the estimation algorithm.

6 Monte Carlo simulations

In this section we illustrate the small sample properties of the sieve value function approximation in a dynamic single agent setting for both finite and infinite horizon problems. We explore multiple aspects of the approximation, which will hopefully provide researchers guidance in custom tailoring the sieve value function approximation to their individual environment. After discussing the results of the Monte Carlo approximation exercises, we turn to the estimation of the structural parameters of the model.

The Monte Carlo experiments are conducted in an adaptation of the Rust bus engine problem. The important modification of the problem is that, rather than making the engine replacement decision of a single bus, the manager is now endogenizing the purchasing and replacement decision over an entire fleet of buses. This alteration substantially increases the state space of the problem, making it an ideal application to study the finite sample properties of the approximation method.

6.1 The model: bus fleet replacement

A single agent makes purchasing and replacement decisions over a fleet of buses to meet rider demand. Buses are assumed to last L periods, after which they are retired. The manager's problem is to decide when to make purchasing decisions of new buses to optimally replace the upcoming set of retiring buses. The manager's discrete decision in each period t is indicated by $a_t \in \{0, 1\}$, where $a_t = 0$ represents not buying a bus in period t , and $a_t = 1$ represents buying an additional bus in period t .

Since buses survive L periods, the entire age composition of the bus fleet, B , can be described by the purchasing decisions of the previous L periods. Formally, we can define the state of the bus fleet in period t as:

$$B_t = [a_{t-1}, a_{t-2}, \dots, a_{t-L}].$$

This representation reflects the fact that only one bus can be purchased in a single period, and thus there can be at most one bus of any particular age in the fleet.

Buses in the fleet are perfect substitutes across ages and are usable immediately when they are purchased. The total number of usable buses in period t is:

$$b(B_t, a_t) = a_t + \sum_{i=1}^L B_t(i),$$

where $B_t(i)$ represents the i^{th} element of vector B_t .

Demand for buses (ridership) is stochastic across time, meaning that when the agent is making purchasing decisions today, they take into account current ridership as well as expected future ridership. We assume ridership, r_t , evolves according to the following normal AR(1) process:

$$r_{t+1} = \psi_0 + \psi_1 r_t + \nu_t,$$

where $\nu_t \sim \mathcal{N}(0, \sigma_\nu^2)$ and $\psi_1 \in (0, 1)$.

The agent's objective is to possess enough buses in the fleet to accommodate demand in each period. Per-period profits, $u_t \equiv u(B_t, r_t, a_t)$, are a function of the current period demand, r_t , and the number of available buses, $b(B_t, a_t)$, defined as:⁹

$$u(B_t, r_t, a_t) = \alpha_1 (b(B_t, a_t) r_t) - \alpha_2 (b(B_t, a_t)^2 + r_t^2), \quad (6.1)$$

The per-period pay-off function, conditional on choice a_t , is characterized by: $u(x_t, a_t) + \varepsilon_t(a_t)$, where the state variable $x_t = \{B_t, r_t\}$ and $\varepsilon_t(a_t)$ is an additive pay-off shock, i.i.d. across choices and time.

The agent's problem is to make an optimal sequence of purchasing decisions to maximize the expected discounted flow of future per-period pay-offs. With discount factor, β , we can formulate

⁹In the Monte Carlo's we set $\alpha_1 = 0.2$ and $\alpha_2 = 0.1$, implying that profits are maximized when the number of buses equals the number of riders, and the agent is penalized if they have too many buses or too few buses.

the value function using using the recursive Bellman representation:

$$V(x_t, \varepsilon_t) = \max_{a_t \in \{0,1\}} \{u(x_t, a_t) + \varepsilon_t(a_t) + \beta E[V(x_{t+1}, \varepsilon_{t+1})|x_t, a_t]\} \quad (6.2)$$

In this formulation of the dynamic problem, it is convenient to write the choice specific value functions (excluding the shocks) for choice $a \in \{0, 1\}$ as follows:

$$v(x_t, a) = u(x_t, a) + \beta E[V(x_{t+1}, \varepsilon_{t+1})|x_t, a]. \quad (6.3)$$

The agent's optimal decision in period t is then to choose a_t^* such that:

$$a_t^* = \operatorname{argmax}_{a \in \{0,1\}} \{v(x_t, a) + \varepsilon_t(a)\}.$$

Given choice a_t , the bus fleet B_t evolves deterministically to B_{t+1} as follows:

$$B_{t+1} = [a_t, B_t(1 : L - 1)],$$

where $B_t(1 : L - 1)$ represents the 1st to $(L - 1)$ th elements of the vector B_t . Since the endogenous state variable B_t evolves deterministically, the expectation in Equation (6.2) is taken over next periods ridership r_{t+1} and next periods pay-off shock ε_{t+1} .

This represents a large state space problem for two reasons. First, ridership, r_t , is a continuous state variable. Second, although the contemporaneous profit function is invariant to the exact age composition of the buses (only the total number of buses matter), the dynamics of the problem imply that the fleet composition matters in the decision process. Therefore, the state space contains the total number of potential age combinations of the bus fleet. This results in a very large state space problem driven by L , the life span of a bus. Since there can be at most one bus of any given age, the number of possible bus fleet combinations is given by $\#\{B\} = 2^L$. As a consequence, the size of the state space is growing exponentially in L .

6.2 Sieve value function approximation

Letting θ_n denote the sieve approximation, our approximation is of the ex-ante value function or social surplus function $\theta_n(x_t) \approx E[\max_{a \in \{0,1\}} (v(x_t, a) + \varepsilon_t(a))|x_t]$ where the expectation is over the random pay-off shock ε . Assuming these shocks are distributed type-I extreme value, we have a known closed form representation of the expectation as:

$$E[V(x_t, \varepsilon_t)|x_t] = \ln \left(\sum_{a \in \{0,1\}} \exp(v(x_t, a)) \right) + \gamma, \quad (6.4)$$

where γ is Euler's constant. This setup correspond to the social surplus formulation of the dynamic decision problem.

Therefore, for a chosen sieve space Θ_n , we seek an approximation of the expression in Equation

(6.4), $\theta_n \in \Theta_n$, such that:

$$\begin{aligned}
\theta_n(x_t) &\approx \ln \left(\sum_{a \in \{0,1\}} \exp(v(x_t, a)) \right) + \gamma \\
&= \ln \left(\sum_{a \in \{0,1\}} \exp(u(x_t, a) + \beta E[V(x_{t+1}, \varepsilon_{t+1})|x_t, a]) \right) + \gamma \\
&\approx \ln \left(\sum_{a \in \{0,1\}} \exp(u(x_t, a) + \beta E[\theta_n(x_{t+1})|x_t, a]) \right) + \gamma.
\end{aligned} \tag{6.5}$$

The right hand side of Equation (6.5) represents one contraction on the approximation, θ_n , defined as $\Gamma^1 \theta_n(x_t)$. Section 3.5 demonstrates that we can actually reduce the upper-bound of the approximation error by increasing the number of iterations of the contraction operator on the approximation. Rather than immediately plugging the approximation on the right hand side into the expected future value term, we could alternatively write this term using Equation (6.4) as a function of next periods contemporaneous profit functions and a two-period away expected future value term. If we substitute the approximation into the two-period away expected future value function this corresponds to two iterations of the contraction ($k = 2$). We could continue in this way for $k > 2$, though the computational burden increases exponentially with k .

We construct the sieve approximation of size n using polynomials of the states and their interactions. The relevant state variables are contained in x and include the complete age composition of the bus fleet as well as current period demand:

$$x_t = [r_t, a_{t-1}, a_{t-2}, \dots, a_{t-L}] = [x(1), x(2), x(3), \dots, x(s)].$$

To form our n terms, we first take the state variables themselves. We then add in all pair-wise interactions, then all three-way interactions, etc. until n is reached. Should n exceed the number of, say, three-way interactions, the three-way interactions are randomly drawn such that the number of terms is n . Denote $w_j(x)$ as the j^{th} term in the polynomial. The sieve approximation is then:

$$\theta_n(x) = \rho(1)w_1(x) + \dots + \rho(n)w_n(x) = \rho_n W_n(x).$$

For $k = 1$ the objective function is now given by:

$$\hat{\rho}_n = \operatorname{argmax}_{\rho_n} \sum_{x \in X} \left[\rho_n W_n(x) - \ln \left(\sum_{a \in \{0,1\}} \exp(u(x, a) + \beta E[\rho_n W_n(x')|x, a]) \right) + \gamma \right]^2. \tag{6.6}$$

A key convenience of using polynomials to construct the sieve is that the approximation is linear in parameters. Consequently, although Equation (6.6) contains an expectation over next periods ridership (a continuous state variable), the expectation passes through to the elements of the polynomial, i.e., $E[\rho_n W_n(x')|x, a] = \rho_n E[W_n(x')|x, a]$. Since we know the distribution of the state transitions, evaluating this expectation is simply a matter of calculating the relevant moment conditions. For discrete variables this is easily done by hand and continuous variables often have

closed forms. We can then replace the expectation by its associated expected value transforming the objective function to:

$$\hat{\rho}_n = \operatorname{argmax}_{\rho_n} \sum_{x \in X} \left[\rho_n W_n(x) - \ln \left(\sum_{a \in \{0,1\}} \exp(u(x, a) + \beta \rho_n E[W_n(x'|x, a)]) \right) + \gamma \right]^2, \quad (6.7)$$

The important insight is that both functions $W_n(x)$ and $E[W_n(x'|x, a)]$, which may be very complex, are calculated outside of the maximization routine. For $k = 1$, this implies that the computational burden of solving the objective function is no greater for models with no stochastic variables than they are for models with many stochastic variables. To emphasize this point, with full solution methods, researchers are often forced to make undesirable simplifications to the problem (like discretizing state variables) for tractability reasons, but sieve value functions can accommodate continuous state variables with no additional computational complexity when $k = 1$.¹⁰

6.3 Results: infinite horizon problem

In practice, these approximation methods will be used in situations where a full solution method is computationally infeasible. However for the purposes of benchmarking the performance of the approximation, we first consider a model that is feasible (although slow) to solve using value function iteration over the entire state space. To do so, we discretize the continuous state variable, ridership, r , to nine points. We set the maximum bus age of $L = 25$, which implies a state space of $9 \times 2^{25} = 301,989,888$. The approximations are based off of a random draw of 20,000 points and is held constant for each approximation.

The objective functions in Equation (6.7) represents a standard non-linear least squares problem that could be solved using direct minimization. However, solving the nonlinear minimization may not be practical when the number of terms in the sieve is large. We therefore consider an alternative iterative approach to solving for θ_n . Letting θ_n^m be the sieve approximation at the m^{th} iteration, we update the approximation at $m + 1$ by:

$$\theta_n^{m+1} = \operatorname{argmax}_{\theta_n \in \Theta} \sum_{x \in X} [\theta_n(x) - \Gamma^k \theta_n^m(x)]^2, \quad (6.8)$$

and then continually update the parameters until convergence.

The Monte Carlo results for the infinite horizon bus fleet replacement problem are summarized in Table 2. Each column corresponds to a different approach to the sieve value function approximation, considering different methods for minimizing the objective function or for different values of k . For each approach, we conduct the approximation for different sieve sizes, ranging from 500 points to 3,000 points. Finally the table is broken into three panels: Panel A reports the mean squared error of the approximation against the true values for the points in the sample. In practice, the truth

¹⁰This is not true for $k > 1$ because we are required to evaluate all possible evolutions of the state variables in the next period.

is unknown. Hence, Panel B reports the mean squared error of the approximation. This measure is easily computable in practice and will be the main insight researchers use to inform them as to how well the approximation is performing. The key theoretical result is that as we know the values of panel B are going to 0, then this implies the values in panel A are going to 0. Finally, Panel C reports the computational time.

Column (i) considers an infeasible estimator for the approximation, which is to regress the state variables on the true expected future value terms. Panel A of the table shows that our constructed sieve performs very well at approximating the true social surplus function and does exceedingly well as we increase the richness of the sieve space n . Column (ii) displays the results for the sieve value function iteration when $k = 1$ and we use direct minimization for the approximation. The table also shows that as the sieve space becomes richer, we achieve a substantial reduction in the objective of the minimization (panel B) and, more importantly, in the mean squared distance to the true value function (panel A).

Column (iii) reports the results for the sieve approximation when $k = 1$ using the iterative method. Although this method shares a similar objective function to the direct minimization approach, the results in Panel B show it is unable to achieve as small of a within model fit for a given n . However, as the sieve size increases to $n = 3,000$ the fit is nearly identical. The benefit to this method is that the computational times are about four times faster than direct minimization.

Column (iv) shows the results with $k = 2$, i.e., applying the contraction operator twice in order to construct the approximation. Because we need to evaluate the sieve at a significantly larger number of points, the approximation becomes more computationally intensive as k increases. However, as outlined by the theory, using $k = 2$ results in about a 50% better fit of the value function compared to using $k = 1$ for the same level of richness of the sieve space (panel A).

Finally, column (v) shows the results for the sieve approximation for the continuous state version of the problem. Comparing the results for the continuous state variable model with the discrete version in column (iii), highlights the approximation method's ability to accommodate an infinite dimensional state variable with an identical model fit and computational time.

6.4 Results: finite horizon problem

We now adapt the approximation method to a finite horizon setting. The bus fleet replacement problem now ends at date T , where the conditional value functions are defined by:

$$v_t(x_t, a) = \begin{cases} u(x_t, a) + \beta E[V(x_{t+1}, \varepsilon_{t+1}, t+1)|x_t, a_t], & \text{for } t = 1, \dots, T-1, \\ u(x_T, a), & \text{for } t = T. \end{cases}$$

The non-stationarity of the value functions introduced by presence of a finite horizon implies that there are T value functions that need to be approximated. The theoretical results of the sieve value

Table 2: Ininite horizon approximation^a

	Discrete state				Continuous state ^b
	Baseline ^c	Direct minimization ($k = 1$)	Iterative minimization ($k = 1$)	Iterative minimization ($k = 2$)	Iterative minimization ($k = 1$)
	(i)	(ii)	(iii)	(iv)	(v)
Panel A: Mean squared difference between EV and $\hat{\theta}_n$					
$n = 500$	0.3149	9.9359	3.7226	1.2402	–
$n = 1,000$	0.2450	5.5071	2.2300	0.8497	–
$n = 2,000$	0.1168	1.5910	0.9944	0.4225	–
$n = 3,000$	0.0073	0.0714	0.0711	0.0414	–
Panel B: Mean squared difference between $\Gamma^k \hat{\theta}_n$ and $\hat{\theta}_n$					
$n = 500$	–	0.0518	0.0641	0.1674	0.0641
$n = 1,000$	–	0.0379	0.0556	0.1399	0.0548
$n = 2,000$	–	0.0173	0.0583	0.0992	0.0576
$n = 3,000$	–	0.0017	0.0018	0.0043	0.0016
Panel C: Computation time (in minutes)					
$n = 500$	–	16	4	24	3
$n = 1,000$	–	50	10	41	9
$n = 2,000$	–	125	27	83	28
$n = 3,000$	–	194	50	124	54

^a The full model consists of 301,989,888 state points, with $L = 25$ and 9 ridership points. The full solution takes 316 minutes. The approximation is done on a single random draw of 20,000 states.

^b Since we cannot solve for the true value functions with continuous state variables, we cannot compare the approximation to the truth.

^c The baseline is the estimate when we regress the sieve on the truth.

Table 3: Finite horizon approximation^a

	$T = 15$	$T = 20$	$T = 25$
Panel A: Backwards Recursion ^b			
Mean sq. diff. between EV_t and $\hat{\theta}_{n,t}$	0.2119	0.3714	0.5025
Mean sq. diff. between EV_1 and $\hat{\theta}_{n,1}$	0.6417	0.9596	1.0485
Panel B: Time Interacted (2^{nd} order) ^c			
Mean sq. diff. between EV_t and $\hat{\theta}_{n,t}$	1.6201	3.4197	5.4476
Mean sq. diff. between EV_1 and $\hat{\theta}_{n,1}$	0.6883	1.1380	1.6291
Panel C: Time Interacted (3^{rd} order) ^d			
Mean sq. diff. between EV_t and $\hat{\theta}_{n,t}$	0.3063	0.6266	0.9886
Mean sq. diff. between EV_1 and $\hat{\theta}_{n,1}$	0.6613	0.8653	0.9762

^a In these Monte Carlos the size of the sieve over the bus and demand state is fixed to 1,000 parameters.

^b Computation times for backwards recursion are small and amount to conducting T least squares minimization routines on the $n = 1,000$ parameters.

^c Second order polynomial of time interacted with original 1,000 parameters, so $n = 2,000$. Computation times are comparable to those reported in Table 2 with similar total size n .

^d Third order polynomial of time interacted with original 1,000 parameters, so $n = 3,000$.

function approximation require us to minimize the following objective function:

$$\{\hat{\theta}_{n,t}\}_{t=1}^T = \operatorname{argmax}_{\{\theta_{n,t}\}_{t=1}^T} \sum_{j=1}^{T-1} \sum_{x \in X} [\theta_{n,j}(x) - \Gamma^k \theta_{n,j+1}(x)]^2. \quad (6.9)$$

Direct minimization of Equation (6.9) is computationally difficult because there are $n \times T$ unknown parameters, where the infinite horizon problem only had n unknown parameters. In previous sections, we considered two methods to conduct the approximation: (1) backwards recursion and (2) including time as part of the state space.

The first method is backwards recursion, where setting $\hat{\theta}_{n,T+1} = \theta_{n,T+1} = 0$, we work our way to $t = 1$, by recursively approximating the function $\theta_{n,t}$ using:

$$\hat{\theta}_{n,t} = \operatorname{argmax}_{\theta_{n,t}} \sum_{x \in X} [\theta_{n,t}(x) - \Gamma^k \hat{\theta}_{n,t+1}(x)]^2. \quad (6.10)$$

The benefit of solving it recursively is that it is extremely fast, only requiring the solution to T OLS regressions, and it still converges to the true value functions as $n \rightarrow \infty$.

The alternative approach exploits the fact that we can re-write the T individual approximating functions, $\theta_{n,t}$ as one large non-parametric function $\theta_{n \times T}$. A third order polynomial of time interacted with the ridership and bus states, n , reduces the number of parameters to a potentially feasible $n \times 3$. By bringing time into the approximating function, we obtain an objective function

which is similar to Equation (6.9),

$$\{\hat{\theta}_{n,t}\}_{t=1}^T = \operatorname{argmax}_{\{\theta_{n,t}\}_{t=1}^T} \sum_{j=1}^{T-1} \sum_{x \in X} [\theta_{n,j}(x) - \Gamma^k \theta_{n,j+1}(x)]^2. \quad (6.11)$$

Table 3 shows the Monte Carlo results for approximating the finite horizon problem. Throughout the table, we fix the size of the sieve over the bus and demand state to $n = 1,000$ and compare the results to the different approaches to approximating the value functions. Panel A shows the results of backwards recursion for three different models, $T = 15$, $T = 20$, and $T = 25$. The two rows correspond to the mean squared error of the true values and the approximation across all time periods and the mean squared error at $t = 1$. As seen in Panel A, for backwards recursion the error at $t = 1$ is two to three times larger than the average approximation error, due to backwards recursion not taking into account the effects of period j 's sieve on periods previous to j when solving for $\theta_n(j)$.

Panels B and C show the time interacted approach by interacting the sieve with a second and third order polynomial of time, respectively. Since we keep the sieve for bus composition and ridership fixed at 1,000, this approximation corresponds to a search over 2,000 or 3,000 parameters. This method is more computationally intensive than backwards recursion but results in similar approximation errors in earlier time periods as in later time periods.

6.5 Sieve value function approximation and estimation

We now apply sieve value function approximation to estimate dynamic models, where we simultaneously solve for the structural parameters and the associated sieve value function approximation that maximizes the likelihood of some observed choice data. Data is generated by simulating choices from the true choice probabilities from 20 markets for 15 periods each. The initial states (fleet composition and ridership) are drawn at random. The fact that the data is generated from the optimal policy function will be relevant for our approximation. Rather than randomly drawing from the state space for the approximation, we can improve our approximation (and thus our structural parameter estimates) by focusing on the state space near the generated data. Although our data only has 300 observations, we are free to draw as many state variables as we would like for the approximation. For these exercises we will draw 20,000 random points which will be perturbations around the data points.¹¹

Our model contains two unknown structural parameters, α_1 and α_2 , which are described in Equation (6.1). Rather than engage in a massive search over all structural parameters and sieve approximation, estimation is easily implemented by modifying the iterative algorithm used for approximation described in Section 3.3. The parameter of interest is denoted by $\pi \equiv (\alpha_1, \alpha_2)$.

¹¹To do this, note the state variables consists of 26 variables, one ridership state and 25 potential bus ages. Our method is to randomly draw states from the 300 data points and for each draw, randomly change 3 of the 25 bus variables by changing them from 1 to 0 or 0 to 1. We then do this 20,000 times.

Table 4: Estimation results^a

	Computation time ^b (in minutes)	Estimate of α_1 ^c ($\alpha_1 = 0.2$)	Estimate of α_2 ^c ($\alpha_2 = -0.1$)
$n = 500$	3	0.1745 (0.0397)	-0.0848 (0.0198)
$n = 1,000$	8	0.1894 (0.0473)	-0.0924 (0.0237)
$n = 2,000$	22	0.1922 (0.0443)	-0.0945 (0.0223)
$n = 3,000$	41	0.1992 (0.0416)	-0.0993 (0.0214)

^a Results from 1,800 replications, each with 300 observations. Initial values of structural parameters and sieve approximation were set to zero.

^b Numbers reported correspond to median computation time.

^c Numbers reported correspond to average over simulations (standard errors in parenthesis).

Let the observations be indexed by $i = 1, \dots, I$, where I denotes the sample size, a_i denotes the observed choice for observation i , and x_i denotes the observed state for observation i .

The estimator of π is denoted by $\hat{\pi}_I$ and is computed according to the following procedure. First, we choose arbitrary initial values of $(\hat{\pi}_I^0, \hat{\theta}_n^0) \in \Pi \times \tilde{\Theta}_n$. Then, we set $m = 1$ and use the following iterative algorithm:

1. Choose $\hat{\pi}_I^m$ to maximize the (approximate) likelihood:

$$\hat{\pi}_I^m = \arg \max_{\pi \in \Pi} \sum_{i=1}^I \ln(\Pr(a_i | \pi, x_i, \hat{\theta}_n^{m-1})).$$

2. Choose $\hat{\theta}_n^m$ according to a one-step sieve value function approximation, implemented by the following least squares minimization:

$$\hat{\theta}_n^m = \operatorname{argmax}_{\theta_n \in \Theta} \sum_{x \in X} \left[\theta_n(x) - \ln \left(\sum_{a \in \{0,1\}} \exp \left(u(x, a, \hat{\pi}_I^m) + \beta E[\hat{\theta}_n^{m-1}(x' | x, a)] \right) \right) + \gamma \right]^2.$$

3. If $\|\hat{\pi}_I^m - \hat{\pi}_I^{m-1}\| < 10^{-6}$, then stop the algorithm and set $\hat{\pi}_I = \hat{\pi}_I^m$. Otherwise, set $m = m + 1$ and return to the first step.

This algorithm parallels Aguirregabiria and Mira (2002) in the sense that it swaps the calculation of the value function to outside of the maximization of the likelihood over the structural parameters.

Results of the estimation exercise for different values of size of the sieve space are presented in Table 4. The results show that for even a very small sieve space ($n = 500$) the estimates are within one standard deviation of the true parameter values, although they appear to be significantly biased downwards. As we increase the size of the sieve to $n = 3,000$ this bias virtually disappears. In addition, as the sieve size increases we are also able to achieve mild improvements in the standard errors of the estimated value, meaning that the efficiency of the approximation brings some stability to the parameter estimates.

7 Conclusion

This paper proposes a methodology to approximate the value function in single agent dynamic problems where a large state space makes value function iteration unfeasible. Our method is based on non-parametric sieve estimation, and we refer to it as Sieve Value Function Iteration (SVFI). We provide a formal framework to analyze the approximation error. In particular, we show that the SVFI approximation converges to the value function as the complexity of the sieve increases, and we characterize the rate of this convergence. Furthermore, we provide a concrete upper bound on the error of approximation which can be used to analyze its contributing factors.

A Monte Carlo analysis reveals that the SVFI approximation is very successful in estimating the value function. These results suggest that our approximation can successfully be used to solve models that would otherwise be computationally infeasible, implying that these techniques may substantially broaden the class of models that can be solved and estimated.

Given the standard challenges with large state space problems, we expect SVFI to open up a wide variety of avenues of theoretical and empirical exploration of complex dynamic single agent and equilibrium problems. For example, in Arcidiacono, Bayer, Bugni, and James (2011), we consider sequential move dynamic games. Estimation of these games can be done via standard two-step procedures. Through the use of sieves, it is possible to calculate the conditional choice probabilities of the finite horizon game which has a unique equilibrium, the limit of which is also an equilibrium in the infinite horizon game. It is then possible to compare the conditional choice probabilities from the finite horizon equilibrium to those observed in the data, testing to see whether this equilibrium was played.

References

- AGUIRREGABIRIA, V., AND P. MIRA (2002): “Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models,” *Econometrica*, 70(2), 1519–1543.
- AMEMIYA, T. (1985): *Advanced Econometrics*. Harvard University Press.
- ARCIDIACONO, P., P. BAYER, F. BUGNI, AND J. JAMES (2011): “Sieve Value Function Iteration for Large State Space Dynamic Games’,” Mimeo: Duke University.
- BÉNITEZ-SILVA, H., G. HALL, G. HITSCH, G. PAULETTO, AND J. RUST (2000): “A comparison of discrete and parametric approximation methods for continuous-state dynamic programming problems,” Mimeo: Yale University, S.U.N.Y. Stony Brook and University Geneva.
- BLACKWELL, D. (1965): “Discounted Dynamic Programming,” *The Annals of Mathematical Statistics*, 36(1), 226–235.
- CHEN, X. (2007): “Large Sample Sieve Estimation of Semi-Nonparametric Models,” vol. 6 of *Handbook of Econometrics*.

- CRAWFORD, G., AND M. SHUM (2005): “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 73(4), 1137–1173.
- HENDEL, I., AND A. NEVO (2006): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74(6), 1637–1673.
- KEANE, M. P., AND K. I. WOLPIN (1994): “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence,” *The Review of Economics and Statistics*, 76(4), 648–672.
- (1997): “The Career Decisions of Young Men,” *Journal of Political Economy*, 105(3), 473–522.
- LORENTZ, G. (1966): *Approximation of functions*. Holt.
- MCFADDEN, D., AND W. NEWEY (1994): “Large sample estimation and hypothesis testing,” vol. 4 of *Handbook of Econometrics*, pp. 2111–2245.
- POWELL, M. (1981): *Approximation theory and methods*. Cambridge University Press.
- POWELL, W. (2008): *Approximate Dynamic Programming: Breaking the Curses of Dimensionality*. Wiley Series in Probability and Statistics.
- ROYDEN, H. L. (1988): *Real Analysis*. Prentice-Hall.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- (1988): “Maximum Likelihood Estimation of Discrete Control Processes,” *SIAM J. Control and Optimization*, 26, 1006–1024.
- RUST, J. (2000): “Parametric Policy Iteration: An Efficient Algorithm for Solving Multidimensional DP problems,” Mimeo: University of Maryland.
- STOKEY, N., AND R. LUCAS (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press.
- SWEETING, A. (2011): “Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry,” Mimeo: Duke University.

A Technical appendix

Proof of Lemma 3.1. We begin by showing that for any $\theta \in \mathcal{F}$ and $m \in \mathbb{N}$:

$$d(\theta, V) \leq d(\theta, \Gamma^m \theta) / (1 - \beta^m), \quad (\text{A.1})$$

To see this, consider the following derivation:

$$d(\theta, V) \leq d(\theta, \Gamma^m \theta) + d(\Gamma^m \theta, \Gamma^m V) + d(\Gamma^m V, V), \quad (\text{A.2})$$

$$= d(\theta, \Gamma^m \theta) + d(\Gamma^m \theta, \Gamma^m V), \quad (\text{A.3})$$

$$\leq d(\theta, \Gamma^m \theta) + \beta^m d(\theta, V), \quad (\text{A.4})$$

where the first inequality follows from the Triangle Inequality, the next equality follows from the fact that V is a fixed point, and the final inequality follows from the fact that Γ is a contraction mapping. Equation (A.1) is a straight-forward consequence of this result.

Let $\hat{\theta}_n \in \Theta_n \subseteq \mathcal{F}$ be the SVFI approximation in Definition 3.3. On the one hand, consider the following derivation:

$$d(\hat{\theta}_n, V)(1 - \beta^k) \leq d(\hat{\theta}_n, \Gamma^k \hat{\theta}_n), \quad (\text{A.5})$$

$$\leq K_2 d_n(\hat{\theta}_n, \Gamma^k \hat{\theta}_n) + \eta_{1,n}, \quad (\text{A.6})$$

$$\leq K_2 \inf_{\theta \in \Theta_n} d_n(\theta, \Gamma^k \theta) + \eta_{1,n} + K_2 \eta_{2,n}, \quad (\text{A.7})$$

$$\leq K_1^{-1} K_2 \inf_{\theta \in \Theta_n} d(\theta, \Gamma^k \theta) + (1 + K_2 K_1^{-1}) \eta_{1,n} + K_2 \eta_{2,n}, \quad (\text{A.8})$$

where the first inequality follows from Equation (A.1), the second and fourth inequalities follow from Equation (3.11), and the third inequality follows from Equation (3.32).

On the other hand, for any $\theta \in \mathcal{F}$, consider the following derivation:

$$d(\theta, \Gamma^k \theta) \leq d(\theta, V) + d(V, \Gamma^k V) + d(\Gamma^k V, \Gamma^k \theta), \quad (\text{A.9})$$

$$= d(\theta, V) + d(\Gamma^k V, \Gamma^k \theta), \quad (\text{A.10})$$

$$\leq (1 + \beta^k) d(\theta, V), \quad (\text{A.11})$$

where the first inequality follows from the triangle inequality, the next equality follows from the fact that V is a fixed point, and the final inequality follows from the fact that Γ is a contraction mapping in (d, Γ) . If we take infimum of $\theta \in \Theta_n \subseteq \mathcal{F}$ on both sides:

$$\inf_{\theta \in \Theta_n} d(\theta, \Gamma^k \theta) \leq (1 + \beta^k) \inf_{\theta \in \Theta_n} d(\theta, V) = (1 + \beta^k) \eta_{3,n}(V). \quad (\text{A.12})$$

The result follows directly from combining the previous results. \square

Proof of Theorem 3.1. By combining Lemma 3.1 with Assumption 3.4, it follows that:

$$d(\hat{\theta}_n, V) \leq \{K_1^{-1} K_2 (1 + \beta^k) \eta_{3,n}(V) + (1 + K_2 K_1^{-1}) \eta_{1,n} + K_2 \eta_{2,n}\} (1 - \beta^k)^{-1}, \quad (\text{A.13})$$

$$= O_p(\max\{\gamma_{1,n}, \gamma_{2,n}, \gamma_{3,n}(V)\}), \quad (\text{A.14})$$

where $\max\{\gamma_{1,n}, \gamma_{2,n}, \gamma_{3,n}(V)\} = o(1)$ as $n \rightarrow \infty$. Using elementary arguments, this result implies

that: (1) $d(\hat{\theta}_n, V) = o_p(1)$ as $n \rightarrow \infty$ and (2) $d(\hat{\theta}_n, V)$ converges in probability to zero at a rate of $\max\{\gamma_{1,n}, \gamma_{2,n}, \gamma_{3,n}(V)\}^{-1} = \min\{\gamma_{1,n}^{-1}, \gamma_{2,n}^{-1}, \gamma_{3,n}(V)^{-1}\}$. \square

Lemma A.1. *Under Assumption 4.1, Assumption 4.5 is satisfied for all the formulations of the dynamic decision problem described in Appendix B.*

Proof. We verify the result for each formulation described in Appendix B. We begin with the *conditional value function* formulation. In this case: $s = (x, \varepsilon)$ and $F(V, s') = V(s')$. Then:

$$E[F(f_1, s') - F(f_2, s')|s, a] = E(f_1(s') - f_2(s')|s, a) \leq d(f_1, f_2), \quad (\text{A.15})$$

where the equality holds by definition of F and the inequality holds because d is the sup-norm metric.

We now consider the *social surplus function* formulation. In this case: $s = (x, a)$, $F(V, s') = G(V(x')|x')$, and $A(s) = \{a\}$. Then:

$$E[F(f_1, s') - F(f_2, s')|s, a] = E[G(f_1(s')|s') - G(f_2(s')|s')|s, a], \quad (\text{A.16})$$

$$= E[E[\max_{a' \in A(x')} (f_1(x', a') + \varepsilon(a')) - \max_{a' \in A(x')} (f_2(x', a') + \varepsilon(a'))|x']|x, a], \quad (\text{A.17})$$

$$\leq E[|E[\max_{a' \in A(x')} (f_1(x', a') + \varepsilon(a')) - \max_{a' \in A(x')} (f_2(x', a') + \varepsilon(a'))|x']||x, a], \quad (\text{A.18})$$

$$\leq E[\max_{a' \in A(x')} |f_1(x', a') - f_2(x', a')||x, a] \leq d(f_1, f_2), \quad (\text{A.19})$$

where the first equality holds by definition of F , the second equality holds by definition of G , the first and second inequalities hold by elementary arguments, and the final inequality holds because d is the sup-norm metric.

We conclude with the *choice-specific value function* formulation. In this case: $s = (x, a)$, $F(V, s') = G(\beta^{-1}u(x') + V(x')|x')$, and $A(s) = \{a\}$. Using the same arguments as in the social surplus function formulation, it is not hard to verify the result. \square

Proof of Theorem 4.1. This proof proceeds by induction. Set $t = 1$. By definition, $\hat{\theta}_{T+2-t,n} = V_{T+1}$ and, thus, $d(V_{T+2-t}, \hat{\theta}_{T+2-t,n}) = 0$. We now prove the inductive step. Suppose that for some $t \geq 1$, $d(V_{T+2-t}, \hat{\theta}_{T+2-t,n}) = O_p(\max\{v_{1,n}, v_{2,n}, v_{3,n}\})$. We now show that: $d(V_{T+1-t}, \hat{\theta}_{T+1-t,n}) = o_p(1)$, as $n \rightarrow \infty$. By the triangular inequality:

$$d(V_{T+1-t}, \hat{\theta}_{T+1-t,n}) = d(V_{T+1-t}, \hat{\theta}_{T+1-t,n}^*) + d(\hat{\theta}_{T+1-t,n}^*, \hat{\theta}_{T+1-t,n}). \quad (\text{A.20})$$

We now analyze each of the terms on the right hand side satisfies the result in Equation (4.11).

We begin with $d(V_{T+1-t}, \hat{\theta}_{T+1-t,n}^*)$. Fix $s \in S$ arbitrarily. Let $a_{1,n}(s) \in A(s)$ be such that:

$$\hat{\theta}_{T+1-t,n}^*(s) = u(s, a_{1,n}(s)) + \beta E(F(\hat{\theta}_{T+2-t,n}, s')|s, a_{1,n}(s)), \quad (\text{A.21})$$

i.e., $a_{1,n}(s)$ is the maximizer, which exists due to the fact that $A(s)$ is a finite set. Then, consider the following derivation:

$$\hat{\theta}_{T+1-t,n}^*(s) = u(s, a_{1,n}(s)) + \beta E(F(\hat{\theta}_{T+2-t,n}, s')|s, a_{1,n}(s)), \quad (\text{A.22})$$

$$\leq u(s, a_{1,n}(s)) + \beta E(F(V_{T+2-t}, s')|s, a_{1,n}(s)) + \beta d(\hat{\theta}_{T+2-t,n}, V_{T+2-t}), \quad (\text{A.23})$$

$$\leq V_{T+1-t}(s) + \beta d(\hat{\theta}_{T+2-t,n}, V_{T+2-t}), \quad (\text{A.24})$$

where the first equality holds by Equation (A.21), the next inequality holds by Assumption 4.5, and the final inequality holds by Equation (4.1). As a consequence, it follows that: $\sup_{s \in S} (\hat{\theta}_{T+1-t,n}^*(s) - V_{T+1-t}(s)) \leq \beta d(\hat{\theta}_{T+2-t,n}, V_{T+2-t})$. By a similar argument, one can show that $\sup_{s \in S} (V_{T+1-t}(s) - \hat{\theta}_{T+1-t,n}^*(s)) \leq \beta d(\hat{\theta}_{T+2-t,n}, V_{T+2-t})$. By combining both inequalities with Assumption 4.1, it follows that:

$$d(\hat{\theta}_{T+1-t,n}^*, V_{T+1-t}) = \sup_{s \in S} |\hat{\theta}_{T+1-t,n}^*(s) - V_{T+1-t}(s)| \leq \beta d(\hat{\theta}_{T+2-t,n}, V_{T+2-t}). \quad (\text{A.25})$$

By the inductive assumption, it follows that: $d(\hat{\theta}_{T+1-t,n}^*, V_{T+1-t}) = O_p(\max\{v_{1,n}, v_{2,n}, v_{3,n}\})$.

We continue with $d(\hat{\theta}_{T+1-t,n}^*, \hat{\theta}_{T+1-t,n})$. Consider the following derivation:

$$d(\hat{\theta}_{T+1-t,n}^*, \hat{\theta}_{T+1-t,n}) \leq K_2 d_n(\hat{\theta}_{T+1-t,n}^*, \hat{\theta}_{T+1-t,n}) + \lambda_{1,n}, \quad (\text{A.26})$$

$$\leq K_2 \inf_{\theta \in \Theta_n} d_n(\theta, \hat{\theta}_{T+1-t,n}) + \lambda_{1,n} + \lambda_{2,n}(\hat{\theta}_{T+1-t,n}^*), \quad (\text{A.27})$$

$$\leq K_1^{-1} K_2 \inf_{\theta \in \Theta_n} d(\theta, \hat{\theta}_{T+1-t,n}) + \lambda_{1,n}(1 + K_1^{-1}) + \lambda_{2,n}, \quad (\text{A.28})$$

$$\leq K_1^{-1} K_2 \sup_{f \in \mathcal{F}} \lambda_{3,n}(f) + \lambda_{1,n}(1 + K_1^{-1}) + \lambda_{2,n} \quad (\text{A.29})$$

where the first and third inequalities hold by Assumption 4.2, the second inequality holds by Assumption 4.3, and the final inequality holds by Assumption 4.4. By the properties of $\lambda_{1,n}$, $\lambda_{2,n}$, and $\lambda_{3,n}$, it follows that $d(\hat{\theta}_{T+1-t,n}^*, \hat{\theta}_{T+1-t,n}) = O_p(\max\{v_{1,n}, v_{2,n}, v_{3,n}\})$. \square

Proof of Lemma 4.1. Part 1. Consider a pair of functions $f_1, f_2 \in \tilde{\mathcal{F}}$. First, consider $t = T+1$. By definition, $\Lambda f_1(s, t) = \Lambda f_2(s, t) = 0 \forall s \in S$, which implies that: $\sup_{s \in S} |\Lambda f_1(s, T+1) - \Lambda f_2(s, T+1)| = 0$. Next, consider $t = 1, \dots, T$. For any arbitrary $s \in S$:

$$|\Lambda f_1(s, t) - \Lambda f_2(s, t)| \quad (\text{A.30})$$

$$= \left\{ \begin{array}{l} \sup_{a \in A(s)} \{u(s, a) + \beta E(F(f_1, (s', t+1)) | (s, t), a)\} \\ - \sup_{a \in A(s)} \{u(s, a) + \beta E(F(f_2, (s', t+1)) | (s, t), a)\} \end{array} \right\} \quad (\text{A.31})$$

$$\leq \beta \sup_{a \in A(s)} E(F(f_1, (s', t+1)) - F(f_2, (s', t+1)) | (s, t), a) \quad (\text{A.32})$$

$$= \beta \sup_{a \in A(s)} E(F(f_{1,t+1}, s') - F(f_{2,t+1}, s') | s, a) \quad (\text{A.33})$$

$$\leq \beta \sup_{a \in A(s)} E(f_{1,t+1}(s') - f_{2,t+1}(s') | s, a) \quad (\text{A.34})$$

$$\leq \beta \sup_{s \in S} |f_{1,t+1}(s) - f_{2,t+1}(s)| = \beta \sup_{s \in S} |f_1(s, t+1) - f_2(s, t+1)|, \quad (\text{A.35})$$

where for any $f \in \tilde{\mathcal{F}}$ and $(s, t) \in S \times \{1, \dots, T+1\}$, we use $f_t(s) \equiv f(s, t)$. Notice that we are also using the fact that conditional expectations are time invariant, but that is assumed for simplicity of notation, i.e., the assumption can be eliminated by indexing expectations with a time index. By reversing the roles of f_1 and f_2 , we deduce that: $\forall t = 1, \dots, T$:

$$\sup_{s \in S} |\Lambda f_1(s, t) - \Lambda f_2(s, t)| \leq \beta \sup_{s \in S} |f_1(s, t+1) - f_2(s, t+1)|. \quad (\text{A.36})$$

By combining information from all values of $t = 1, \dots, T+1$, it follows that:

$$\tilde{d}(\Lambda f_1, \Lambda f_2) = \max_{t=1, \dots, T+1} \sup_{s \in S} |\Lambda f_1(s, t) - \Lambda f_2(s, t)| \leq \beta \max_{t=1, \dots, T+1} \sup_{s \in S} |f_1(s, t) - f_2(s, t)| = \tilde{d}(f_1, f_2). \quad (\text{A.37})$$

Part 2. By the Contraction Mapping Theorem (see, e.g., page 50 in Stokey and Lucas (1989)) the mapping Λ has a unique fixed point in $\tilde{\mathcal{F}}$. It suffices to show that V is such fixed point, i.e., $\tilde{d}(V, \Lambda V) = 0$. For $t = T+1$, the definition of ΛV is: $(\Lambda V)(s, T+1) = 0 = V(T+1, s) \equiv V_{T+1}(s)$. For any other $t = 1, \dots, T$, it follows that:

$$(\Lambda V)(s, t) = \sup_{a \in A(s)} \{u(s, a) + \beta E(F(V, (s', t+1)) | (s, t), a)\} \quad (\text{A.38})$$

$$= \sup_{a \in A(s)} \{u(s, a) + \beta E(F(V_{t+1}, s') | s, a)\} = V_t(s) = V(s, t). \quad (\text{A.39})$$

This completes the proof. \square

Proof of Theorem 4.2. By Lemma 4.1, the analogue of Assumptions 3.2-3.4 hold for the state space \tilde{S} . Under these assumptions, the result is a corollary of Theorem 3.1. \square

Proof of Theorem 5.1. Fix $\varepsilon > 0$ arbitrarily. Then, $\exists \delta > 0$ such that $\forall \pi \in \Pi: d_2(\pi, \pi^*) > \varepsilon \Rightarrow Q(V(\cdot | \pi)) - Q(V(\cdot | \pi^*)) > \delta$. This implies that $\exists \delta > 0$ such that:

$$P(d_2(\hat{\pi}_I, \pi^*) \leq \varepsilon) \geq P(Q(V(\cdot | \hat{\pi}_I)) - Q(V(\cdot | \pi^*)) \leq \delta). \quad (\text{A.40})$$

The strategy of the proof is to show that the RHS converges to one as $I \rightarrow \infty$. To this end, for a fixed $\delta > 0$, consider the following argument:

$$P(Q(V(\cdot | \hat{\pi}_I)) - Q(V(\cdot | \pi^*)) \leq \delta) \quad (\text{A.41})$$

$$= P \left(\begin{array}{l} Q(V(\cdot | \hat{\pi}_I)) - Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) + Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q_I(\hat{\theta}_n(\cdot | \hat{\pi}_I)) \\ + Q_I(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q(V(\cdot | \pi^*)) \leq \delta \end{array} \right), \quad (\text{A.42})$$

$$\geq P \left(\begin{array}{l} \{Q(V(\cdot | \hat{\pi}_I)) - Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) \leq \delta/3\} \cap \{Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q_I(\hat{\theta}_n(\cdot | \hat{\pi}_I)) \leq \delta/3\} \\ \cap \{Q_I(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q(V(\cdot | \pi^*)) \leq \delta/3\} \end{array} \right), \quad (\text{A.43})$$

$$\geq \left\{ \begin{array}{l} P(Q(V(\cdot | \hat{\pi}_I)) - Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) \leq \delta/3) + P(Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q_I(\hat{\theta}_n(\cdot | \hat{\pi}_I)) \leq \delta/3) \\ + P(Q_I(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q(V(\cdot | \pi^*)) \leq \delta/3) - 2 \end{array} \right\}. \quad (\text{A.44})$$

The last expression on the RHS includes three probability expressions. The proof is completed by showing that these expressions converge to one as $I \rightarrow \infty$.

Consider the first expression. By the uniform continuity of Q , there is $\eta > 0$ such that:

$$\sup_{\pi \in \Pi} d_1(\hat{\theta}_n(\cdot | \pi), V(\cdot | \pi)) < \eta \Rightarrow d_1(\hat{\theta}_n(\cdot | \hat{\pi}_I), V(\cdot | \hat{\pi}_I)) < \eta \Rightarrow Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q(V(\cdot | \hat{\pi}_I)) \leq \delta/3. \quad (\text{A.45})$$

It follows then that:

$$P(Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q(V(\cdot | \hat{\pi}_I)) \leq \delta/3) \rightarrow 1, \text{ as } n \rightarrow \infty, \quad (\text{A.46})$$

or, equivalently, as $I \rightarrow \infty$ because $n(I) \rightarrow \infty$ as $I \rightarrow \infty$.

Consider the second expression. Notice that:

$$Q(\hat{\theta}_n(\cdot | \hat{\pi}_I)) - Q_I(\hat{\theta}_n(\cdot | \hat{\pi}_I)) \leq \sup_{f \in \Theta_n} |Q_I(f) - Q(f)| = \hat{c}(n, I), \quad (\text{A.47})$$

and since $\hat{c}(n, I) = o_p(\delta(n(I), I)) = o_p(1)$ as $I \rightarrow \infty$, it follows that:

$$P(Q(\hat{\theta}_n(\cdot|\hat{\pi}_I)) - Q_I(\hat{\theta}_n(\cdot|\hat{\pi}_I)) \leq \delta/3) \rightarrow 1, \text{ as } I \rightarrow \infty. \quad (\text{A.48})$$

Finally, consider the third expression. By definition of the estimation:

$$Q_I(\hat{\theta}_n(\cdot|\hat{\pi}_I)) \leq \inf_{\pi \in \Pi} Q_I(\theta_n(\cdot|\pi)) + o_p(1) \leq Q_I(\theta_n(\cdot|\pi^*)) + o_p(1), \quad (\text{A.49})$$

as $I \rightarrow \infty$. This implies that:

$$Q_I(\hat{\theta}_n(\cdot|\hat{\pi}_I)) - Q(V(\cdot|\pi^*)) \leq Q_I(\hat{\theta}_n(\cdot|\pi^*)) - Q(V(\cdot|\pi^*)) + o_p(1), \quad (\text{A.50})$$

$$= Q_I(\hat{\theta}_n(\cdot|\pi^*)) - Q(\hat{\theta}_n(\cdot|\pi^*)) + Q(\hat{\theta}_n(\cdot|\pi^*)) - Q(V(\cdot|\pi^*)) + o_p(1), \quad (\text{A.51})$$

$$\leq \sup_{f \in \Theta_n} |Q_I(f) - Q(f)| + Q(\hat{\theta}_n(\cdot|\pi^*)) - Q(V(\cdot|\pi^*)) + o_p(1). \quad (\text{A.52})$$

Repeating arguments used in previous expressions, it follows that:

$$P(Q_I(\hat{\theta}_n(\cdot|\hat{\pi}_I)) - Q(V(\cdot|\pi^*)) \leq \delta/3) \rightarrow 1, \text{ as } I \rightarrow \infty. \quad (\text{A.53})$$

This completes the proof. \square

B Alternative formulations of dynamic decision problem

As described in the main text of Section 2, there are three possible ways to formulate the dynamic decision problem of Equation (2.1) as a recursive problem. The objective of this section is to describe each of these formulations and show that each of them is a special case of the unified formulation described in the main text.

B.1 Value function formulation

This formulation of the dynamic decision problem is the most obvious one but, as we will argue, the most computationally demanding in terms of the size of the state space. It is easy to verify that Equation (2.1) can be represented as follows:

$$V(x, \varepsilon) = \max_{a \in A(x)} \{u(x, a) + \varepsilon(a) + \beta \int V(x', \varepsilon') dP(x', \varepsilon' | x, \varepsilon, a)\}, \quad (\text{B.1})$$

$$= \max_{a \in A(x)} \{u(x, a) + \varepsilon(a) + \beta E[V(x', \varepsilon') | x, \varepsilon, a]\}. \quad (\text{B.2})$$

If we let $s \equiv (x, \varepsilon)$ and $S \equiv X \times E$, $V : S \rightarrow \mathbb{R}$ is the fixed point of the following mapping:

$$(\Gamma_1 m)(x, \varepsilon) = \max_{a \in A(x)} \{u(x, a) + \varepsilon(a) + \beta E[m(x', \varepsilon') | x, \varepsilon, a]\}, \quad \forall (x, \varepsilon) \in S. \quad (\text{B.3})$$

This corresponds to the conditional value function formulation of the problem (Equation (4.4) in Rust (1987)). Notice that our derivations so far did not require the Conditional Independence (CI) Assumption. It is straight-forward to verify that it is a contraction mapping (in the sup-norm) by using the Blackwell (1965) sufficient conditions for a contraction (see Stokey and Lucas (1989), Theorem 3.3). This is an obvious

object of interest, as it allows us to deduce optimal decision rules in the following way:

$$f(x, \varepsilon) = \arg \max_{a \in A(x)} \{u(x, a) + \varepsilon(a) + \beta E[V(x', \varepsilon') | x, \varepsilon, a]\}, \quad (\text{B.4})$$

with a related conditional choice probability (CCP):

$$P(a = \bar{a} | x) = \int_{\varepsilon} \left(\arg \max_{a \in A(x)} \{u(x, a) + \varepsilon(a) + \beta E[V(x', \varepsilon') | x, \varepsilon, a]\} = \bar{a} \right) dP(\varepsilon | x). \quad (\text{B.5})$$

We now verify that this formulation is a particular case of the unified formulation. This is shown in Table 5, which explains how the objects in the unified notation are defined in terms of the primitive objects in Equation (2.1). According to the table, F is defined as follows: $F(V, s') = V(s')$. It is straightforward to verify that this function satisfies all of the properties required for F .

Primitive object	V	(x, ε)	$X \times E$	$A(x)$	a	$u(x, a) + \varepsilon(a)$	$V(s')$
Unified formulation	V	s	S	$A(s)$	a	$u(s, a)$	$F(V, s')$

Table 5: Representation of the primitive object of the conditional value function formulation in terms of the unified formulation.

Notice that the state space in this formulation is given by $s = (x, \varepsilon)$, which is relatively large as it includes $\varepsilon = \{\varepsilon(a) : a \in A(x)\}$ (i.e. has the cardinality of $A(x)$). By the nature of the type of problems under consideration it is reasonable to consider alternative formulations of the problem, where ε is not part of the state space.

B.2 Social surplus function formulation

This formulation of the problem is developed in Rust (1988). In order to explain the formulation, it is necessary to define certain objects. Given a vector $v(x) = \{v(x, a) : a \in A(x)\}$, define the social surplus function as follows:

$$G(v(x) | x) \equiv \int_{\varepsilon} \max_{a \in A(x)} \{v(x, a) + \varepsilon(a)\} dP(\varepsilon | x) = E \left(\max_{a \in A(x)} \{v(x, a) + \varepsilon(a)\} \middle| x \right). \quad (\text{B.6})$$

If we let $s \equiv (x, a)$ and $S \equiv \{(x, a) : a \in A(x), x \in X\}$, Theorem 3.2 in Rust (1988) indicates that:

$$V(x, \varepsilon) = \max_{a \in A(x)} [v(x, a) + \varepsilon(a)], \quad (\text{B.7})$$

where $v : S \rightarrow \mathbb{R}$ is the unique fixed point of the following contraction mapping:

$$(\Gamma_2 m)(x, a) = u(x, a) + \beta E[G(m(x') | x') | x, a], \forall (x, a) \in S. \quad (\text{B.8})$$

Notice that, in the formula, $m(x')$ is shorthand for the vector $\{m(x', i') : i' \in A(x')\}$. Our assumptions imply that the mapping is a contraction mapping (in the sup-norm) by using the Blackwell (1965) sufficient conditions for a contraction (in this case, it is important to establish that G satisfies the properties in Theorem 3.2 in Rust (1988)). Notice that the fixed point v is the object of interest because it allows us to recover the original value function V (using Equation (B.7)) and, with it, we can then construct the rest of the objects of interest such as optimal decision rules and CCPs.

The advantage of working with this formulation of the dynamic decision problem is that the state space is (x, a) , whose dimension is smaller than for the condition value function formulation. Another advantage is that the CCPs can be written simply as (see Theorem 3.4 in Rust (1988)):

$$P(a = \bar{a}|x) = G_{\bar{a}}(v(x)|x), \quad (\text{B.9})$$

where $G_{\bar{a}}$ denotes the derivative of the social surplus function (i.e. Equation (B.6)) with respect to $v(x, \bar{a})$.

We now verify that this formulation is a particular case of the unified formulation. This is shown in Table 6, which explains how the objects in the unified notation are defined in terms of the primitive objects in Equation (2.1). According to the table, F is defined as follows: $F(V, s') = G(V(x')|x')$, where G is the social surplus function (i.e. Equation (B.6)) with its own choice set $A(x')$. Theorem 3.1 in Rust (1988) shows that the function G satisfies all of the desired properties.

Unified formulation	v	s	S	$A(s)$	a	$u(s, a)$	$F(V, s')$
Primitive object	V	(x, a)	$\{(x, a) : a \in A(x), x \in X\}$	$\{a\}$	a	$u(x, a)$	$G(V(x') x')$

Table 6: Representation of the primitive object of the social surplus function formulation in terms of the unified formulation.

B.3 Choice-specific value function formulation

This formulation is related to the social surplus function formulation and is explained in Rust (1988), page 1014. As in the previous formulation, let $s \equiv (x, a)$ and $S \equiv \{(x, a) : a \in A(x), x \in X\}$. Under our assumptions, Rust (1988) indicates that:

$$v(x, a) = u(x, a) + \beta EV(x, a), \quad (\text{B.10})$$

where $v : S \rightarrow \mathbb{R}$ is the function in Equation (B.7) and $EV : S \rightarrow \mathbb{R}$ is the fixed point in the following contraction mapping:

$$(\Gamma_3 m)(x, a) = E[G(u(x') + \beta m(x')|(x'))|x, a], \quad \forall (x, a) \in S. \quad (\text{B.11})$$

This formulation is intimately related to the social surplus function formulation and shares all its advantages relative to the conditional value function formulation, i.e., the state space is (x, a) , which have a relatively smaller dimension, and the CCPs are relatively easy to compute using Equation (B.9).

We now verify that this formulation is a particular case of the unified formulation. This is shown in Table 7, which explains how the objects in the unified notation are defined in terms of the primitive objects in Equation (2.1). According to the table, F is defined as follows: $F(V, s') = G(\beta^{-1}u(x') + V(x')|x')$, where G is the social surplus function (i.e. Equation (B.6)) with its own action set $A(x')$. Theorem 3.1 in Rust (1988) shows that the function G satisfies all of the desired properties.

Unified formulation	V	s	S	$A(s)$	a	$u(s, i)$	$F(V, s')$
Primitive object	EV	(x, a)	$\{(x, a) : i \in A(x), x \in X\}$	$\{a\}$	a	0	$G(\beta^{-1}u(x') + V(x') x')$

Table 7: Representation of the primitive object of the choice-specific value function formulation in terms of the unified formulation.