Robust Estimation of Finite Horizon Dynamic Economic Models*

Thomas H. Jørgensen† Maxime Tô‡

June 30, 2017

Abstract

We study an estimation approach that is robust to misspecifications of the dynamic economic model being estimated. Specifically, the approach allows researchers to focus on a particular sub-period of the optimizing agent's finite horizon and thus alleviates the need for assumptions regarding expectation formation about the (distant) future. We propose to approximate a *pseudo* terminal period policy- or value function non-parametrically rather than fully specifying the remaining economic environment anticipated by agents until the terminal period. We illustrate through two Monte Carlo experiments the superior robustness of our approximate estimator compared to a standard full solution estimator (JEL.: C51, C61, C63).

Keywords: Robust, Structural Estimation, Finite Horizon Dynamic Programming.

*We thank Jeppe Druedahl, Anders Munk-Nielsen, Itay Saporta-Eskten and Mette Ejrnæs for helpful comments and discussions. Financial support from the Danish Council for Independent Research in Social Sciences is gratefully acknowledged (FSE, grant no. 4091-00040). The current research is part of and funded by an European Research Council Starting Grant (No. 312474).

†Department of Economics, University College London, Gower Street, London, United Kingdom. E-mail: thomas.h.jorgensen@econ.ku.dk. Webpage: www.tjeconomics.com.

‡Department of Economics, University College London. E-mail: maxime.to@ucl.ac.uk. Website: https://sites.google.com/site/maximeto9.

1 Introduction

In dynamic models with forward looking agents the current and future (expected) environment has to be specified. To do so, researchers often have to make strong assumptions on the evolution of the environment in which decisions are made. Manski (1993), for example, points out that without expectations data, all structural work hinges on assumptions about the expectation formation of agents. In this paper, we focus on cases where the researcher is interested in learning about the structure of individual decisions (like instantaneous payoff and transition functions) related to a given period, but is reluctant to make too many assumptions about distant future periods.

We investigate an estimation approach that is *independent of the distant future*, and thus, is robust to misspecification of the economic model in these periods. In particular, our approach alleviates the need for researchers to take a stand on whether agents foresee changes in their future economic environment, such as policy reforms. We distinguish between two sets of periods: The first set is the periods of interest to the researcher, and is typically the set of periods observed in the data. The other set is the distant future, and is not of direct interest to the researcher, but is taken into account by the agent when performing optimal choices in earlier periods. These later periods potentially contain complex dynamic mechanisms or involve expected policy changes, and thus prone to misspecifications.[1] An example could be a researcher aiming at estimating a model of educational choice. In doing so all later periods of life, involving for example house purchases, marriage and fertility decisions, retirement etc. potentially influence educational choices without being of particular interest to the researcher. The approach discussed herein would alleviate the need to explicitly model these later life periods while still allowing educational choices to be affected by these subsequent decisions.[2]

We define a pseudo-terminal period, $\check{T}$, as the threshold period between the two subsets of periods previously defined. We propose to non-parametrically approximate the value- or policy function at this pseudo-terminal period removing the need to fully specify the economic environment in *all* periods faced by the agent until the terminal period $\overline{T} \geq \check{T}$, as it is usually done in more traditional "full-solution" approaches.[3] The solution of the approx-

[1]By misspecification we refer to any situation in which either the researcher or the optimizing agent have a misperception about the distant future.

[2]A parametric approach similar to the one discussed here has been successfully employed in Keane and Wolpin (2001) who study the postsecondary educational choices of US youth and the role of parental transfers.

[3]While we focus on a full-solution or nested fixed point (NFXP) type estimator throughout, our approach would also apply to alternative estimators such as the mathematical programming with equilibrium constraints (MPEC) proposed by Su and Judd (2012).

imate model is then obtained by backward induction starting from this pseudo-terminal period and the associated approximate estimator is thus independent of the expected future environment faced by agents in periods after $\check{T}$. In the educational choice example, the pseudo-terminal period could for example be around age 35.

Monte Carlo experiments in Section 3 confirm the robustness of our approach. We investigate the performance of the approximate estimator using a tax reform in a discrete choice model of labor market supply and human capital accumulation and time-varying preferences in a continuous choice model of optimal intertemporal consumption allocation. In both experiments the approximate estimator is robust to misspecifications while the full solution estimator is not robust to even mild misspecifications.

The proposed approach has implications for many fields in economics. As we conclude in Section 4, the approximate estimator has the potential to allow researchers to estimate rich models of a particular decision problem while remaining relatively agnostic about the following environment facing the economic agents in subsequent periods. Particularly, the approach facilitates out-of-sample policy evaluation using structural models, similarly to the microfinance intervention studied in Kaboski and Townsend (2011) without having to rely on strong assumptions about agents' beliefs regarding the intervention.[4] Using quasi-experiments to perform out-of-sample validation based on pre-experiement data requires researchers to take a stand on what agents in the data expected about the reform. Our approach lets researchers estimate the pre-reform model on pre-reform data while remaining relatively agnostic about agent's beliefs regarding the future reform. The approach also provides researchers with an approach to perform robustness checks of their full-solution estimation results: significant differences across the two approaches could indicate misspecification of the full-solution estimator.

Related to our approach is a strategy sometimes applied in empirical applications: To iterate backwards from the terminal period $\overline{T}$ to period $\check{T}$ using a simple model and then use that solution as an approximate pseudo-terminal value function. Such a strategy is employed in for example Gourinchas and Parker (2002) where a simple linear consumption rule is assumed post retirement and two parameters adjusting this rule is estimated along side other parameters of the model. While such and similar approaches are motivated

[4] An alternative strategy is to use reforms as exogeneous variation to identify the structural parameters in the model. Blundell, Dias, Meghir and Shaw (2016), for example, use tax and welfare reforms to estimate a model of female labor force participation. In doing so, the authors must take a stand on household's expectation formation regarding institutional settings and assume that all reforms are completely *unanticipated*. While this assumption is computationally attractive, the amount of reforms to the tax and welfare system suggests that households would be better of by putting some positive probability to alternative schemes in the future.

by reduced computation time, we motivate the use of our approach from a robustness perspective.[5]

Although approaches similar in spirit to the one discussed herein might be familiar to some researchers, to the best of our knowledge, the robustness of the approach has not been studied and is not generally appreciated. Furthermore, the use of this approach to mitigate effects on estimation results from agents' expectation formation about policy reforms has not been acknowledged at all. Therefore the purpose of this paper is to illustrate the robustness properties of this estimator and highlight the potential use in future research.

Our approach is also related to two strands of literature on approximate dynamic programming. Particularly, parametric and non-parametric approaches to approximate dynamic programming is related to the approach discussed here. See, e.g. Keane and Wolpin (1994); Judd (1998); Rust (2000); and Rust, Hall, Benítez-Silva, Hitsch and Pauletto (2005) for parametric approaches and the recent studies by Kroemer and Peters (2011); Bhat, Farias and Moallemi (2012); and Arcidiacono, Bayer, Bugni and James (2013) for non-parametric approaches.[6] Common across these papers is that all time periods of the model is approximated through parametric or non-parametric approximations. In contrast, we focus on estimation of a model where only the pseudo terminal solution has been approximated. Supplemental material and code used to generate the results are available from the authors web-pages.

2 Framework

We focus on finite horizon dynamic programming models of the form[7]

$$\begin{aligned} V_t(\mathbf{s}_t) &= \max_{\mathbf{c}_t \in \mathcal{C}(\mathbf{s}_t)} U_t(\mathbf{c}_t, \mathbf{s}_t) + \mathbb{E}[\beta_{t+1} V_{t+1}(\mathbf{s}_{t+1}) | \mathbf{c}_t, \mathbf{s}_t], \; \forall t = 1, \ldots, \overline{T} - 1 \\ & V_{\overline{T}+1}(\mathbf{s}_{\overline{T}+1}) \text{ and } \mathbf{s}_0 \text{ given} \end{aligned} \tag{1}$$

where the time-heterogeneous value function depend on a state vector, $\mathbf{s}_t$, that could be any mix of discrete and continuous, and constant and time-varying state variables. Every period, agents choose a vector of controls $\mathbf{c}_t$ to maximize the sum of the instantaneous

[5] In fact, the computation time required by the approximate estimator is not necessarily lower than that of the full-solution estimator: Computational time is reduces in the approximate estimator because the model needs to be solved for $\overline{T} - \check{T}$ fewer periods but computational time might also increase because the number of parameters to estimate increases due to the introduced parameters in the approximation.

[6] Other approaches related to ours are those suggested in Geweke and Keane (2000) and Judd, Maliar, Maliar and Tsener (forthcoming).

[7] While this does not include the recursive Epstein–Zin-Weil preferences, we see no reason why not to believe our approach could be implemented with such preferences.

utility and the expected discounted continuation value. The distribution of future states is given by

$$F_t(\mathbf{s}_{t+1}) = \Gamma_t(\mathbf{c}_t, \mathbf{s}_t),\ \forall t = 1, \ldots, \overline{T}. \tag{2}$$

We collect the model formulation in compact notation as

$$\mathcal{M}_{\overline{T}}(\mathbf{s}, \mathbf{c}; \theta),$$

where θ contains all the parameters of the model, implicitly related to equations (1)–(2).

The aim of the researcher is to estimate the parameter vector θ. A subset of controls and endogenous state variables is observed for $t = 1, \ldots, T$ time periods. We denote the available data as $\mathbf{w}_{1:T}$ and define the objective function as $\mathcal{Q}(\theta, \mathbf{w}_{1:T}, \mathcal{M}_{\overline{T}}(\mathbf{s}, \mathbf{c}; \theta))$. The resulting "full solution" estimator is

$$\hat{\theta} = \arg\min_{\theta \in \Theta} \mathcal{Q}(\theta, \mathbf{w}_{1:T}, \mathcal{M}_{\overline{T}}(\mathbf{s}, \mathbf{c}; \theta)) \tag{3}$$

where the model is solved by backwards induction from period $t = \overline{T}$ for all trial guesses of θ. The general model formulation outlined here would typically require some numerical procedure to solve the model $\mathcal{M}_{\overline{T}}(\mathbf{s}, \mathbf{c}; \theta)$. We abstract throughout from any such numerical approximations and instead focus on an *approximate* model which we will describe below.

2.1 Robust Estimation

When implementing the estimator above, the researcher has to specify the utility function U_t and transition distributions F_t for all periods, even the ones that are not directly observed in the data. Our robust estimation approach is based on an alternative model in which no assumptions are made about utility and transitions after a chosen *pseudo* terminal period denoted $\check{T}$. Since our method is independent of model transitions between periods $\check{T}$ and $\overline{T}$, it is robust to any misspecification at these periods unlike the full solution estimation approach outlined above. By misspecification we refer to any situation in which either the researcher or the optimizing agent have a misperception about the subsequent periods, i.e., $t > \check{T}$. To be precise, we assume that $T \leq \check{T} \leq \overline{T}$.

Let the value functions based on the approximate pseudo-terminal value function be

$$\check{V}_t(\mathbf{s}_t) = \max_{\mathbf{c}_t \in \mathcal{C}(\mathbf{s}_t)} U_t(\mathbf{c}_t, \mathbf{s}_t) + \mathbb{E}[\beta_{t+1} \check{V}_{t+1}(\mathbf{s}_{t+1}) | \mathbf{c}_t, \mathbf{s}_t],\ \forall t = 1, \ldots, \check{T} \tag{4}$$

$$\check{V}_{\check{T}}(\mathbf{s}_t) = \gamma(\mathbf{s}_t), \tag{5}$$

where the state transition is unchanged as in (2) and $\gamma(\mathbf{s}_t)$ denotes a non-parametric approximation of the value function in the pseudo-terminal period $\check{T}$. A similar approach can be employed if the derivative of the value function is used (such as Euler equations), and we illustrate such an implementation in the first example in the Monte Carlo experiments below.

We denote this approximate model as

$$\check{\mathcal{M}}_{\check{T}}(\mathbf{s}, \mathbf{c}; \theta, \gamma)$$

and the associated robust estimator as

$$(\check{\theta}, \check{\gamma}) = \arg \min_{\theta \in \Theta, \gamma \in \Gamma} \mathcal{Q}(\theta, \mathbf{w}_{1:T}, \check{\mathcal{M}}_{\check{T}}(\mathbf{s}, \mathbf{c}; \theta, \gamma)) \tag{6}$$

Consistency of estimators based on a Sieve approximated value function is derived in Arcidiacono, Bayer, Bugni and James (2013) as the complexity of the Sieve increases with the sample size at a suitable rate. Throughout, we assume that all parameters (θ and γ) are identified from the available data. Identification will depend on the exact nature of the economic model but dynamic models are often highly over-identified through functional forms and the identification of the additional function γ should in many cases be possible using the last couple of periods of data.[8]

3 Monte Carlo Experiments

We illustrate the proposed robust estimation approach using two different structural models. The first example is a modern work horse model of incomplete markets pioneered by Deaton (1991) and Carroll (1992) and the second example is a discrete-choice labor supply model of human capital accumulation in the spirit of Keane and Wolpin (1997). To investigate the performance of the two estimators we simulate $N = 20,000$ individuals who live $\bar{T} = 40$ periods, and that we observe for $T = 10$ time periods. We simulate $S = 200$ of such independent data-sets and estimate the parameters using each simulated data set. For our robust approximation approach, we approximate the $\check{T} = \tau - 5$ pseudo terminal consumption or value function using Sieves. We let $\tau = 25$. The supplemental

[8]Some caution may be advised when performing counter-factual policy simulations based on the approximate model. Although the approximate model provides an acceptable approximation for *observed* data, the approximate model might provide a poorer approximation for simulated data under alternative contingencies. In many cases this is unlikely to be important but if the counter factual simulation changes the distribution of data radically, this might be a real concern.

material contains a detailed description of how we solve the models and formulate the two estimators.

In the two examples the misspecifications consists of simple changes in the model at period $\tau = \check{T} + 5$. Admittedly, if the researcher had information that the correct data generating process (DGP) was time-varying, it would be optimal to change the model to the correct one. However, what we want to emulate here is a situation in which the researcher does not have such information and/or there might be an array of things changing in the DGP at periods which the researcher does not have data on and/or is not interested in explicitly modeling.

3.1 Example 1: Intertemporal Consumption Allocation

In this example, households solve the dynamic problem of intertemporal consumption allocation summarized by the Bellman equation

$$\begin{aligned} V_t(M_t, P_t) &= \max_{0 \leq C_t \leq M_t} \frac{C_t^{1-\rho}}{1-\rho} + \beta_t \mathbb{E}_t[V_t(M_{t+1}, P_{t+1})] \\ &\text{s.t.} \\ M_{t+1} &= (1+r)A_t + Y_{t+1} \\ Y_{t+1} &= P_t \varepsilon_t \\ P_t &= G_t P_{t-1} \eta_t \\ \log \varepsilon_t &\sim \mathcal{N}(-.5\sigma_\varepsilon^2, \sigma_\varepsilon^2) \\ \log \eta_t &\sim \mathcal{N}(-.5\sigma_\eta^2, \sigma_\eta^2) \end{aligned}$$

for all $t = 1, \ldots \overline{T}$ and a terminal condition implying consumption of all remaining resources, $V_{\overline{T}+1} = \frac{M_{\overline{T}+1}^{1-\rho}}{1-\rho}$. M_t is beginning-of-period resources, A_{t-1} is beginning-of-period wealth and Y_t is beginning-of-period income. P_t denotes permanent income, G_t is permanent income growth and ε and η are transitory and permanent income shocks, respectively.

For the purpose of our analysis, we let the discount factor be time-varying

$$\beta_t = (1 - \delta \mathbf{1}\{t \geq \tau\})\tilde{\beta} \tag{7}$$

such that we allow the discount factor to differ by age.

We will assume throughout that the researcher insists on imposing a constant discount factor and thus imposes $\delta = 0$ while in reality $\delta \geq 0$. The key point we want to illustrate is that using data for the first $T < \tau$ periods, the "full solution" employing backwards induction from period $\overline{T}$ imposing a constant discount factor may bias estimation results. On the other hand, using our approach, approximating the pseudo terminal period $\check{T} < \tau$, should be robust to such misspecification.

The current model can be normalized by permanent income and we denote lower-case letters as such normalized quantities. We let the approximated consumption function (normalized by permanent income) in the pseudo-terminal period be

$$\check{c}_{\check{T}}(m) = \exp(\gamma_0) m^{\exp(\gamma_1)}$$

ensuring that consumption is an increasing and concave function of normalized resources.

Both estimators are based on the assumption that observed consumption is contaminated with additive mean zero measurement error. For simplicity we focus on estimation of the constant relative risk aversion and keep the remaining parameters fixed at their true values, reported in the footnote of Table 1. We observe consumption, market resources and permanent income and the estimators thus are

$$\begin{aligned} \hat{\rho} &= \arg\min_{\rho} \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} (C_{i,t} - C_t^{\star}(M_{i,t}, P_{i,t} | \rho, \delta = 0))^2 \\ (\check{\rho}, \check{\gamma}) &= \arg\min_{\rho, \gamma} \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} (C_{i,t} - \check{C}_t^{\star}(M_{i,t}, P_{i,t} | \rho, \gamma))^2 \end{aligned}$$

where $C_t^{\star}(M_{i,t}, P_{i,t} | \rho, \delta = 0)$ and $\check{C}_t^{\star}(M_{i,t}, P_{i,t} | \rho, \gamma)$ denotes respectively the "full" backwards induction solution assuming $\delta = 0$ and the robust approximated solution depending also on the parameters in the Sieve, $\gamma = (\gamma_0, \gamma_1)$, but not on any assumption about δ.

Table 1 reports the bias, standard error and the root mean square error (RMSE) of the two estimators under different values of $\delta \in \{0, 0.001, 0.005, 0.01\}$. In the first case ($\delta = 0$) the full solution approach is correctly specified and in the other cases it is not. As expected, we find that in the first case the full solution estimator performs better than the approximation but only in terms of variance of the estimator. In the misspecified cases, however, the performance of the full solution methods deteriorates significantly (as expected) with an increased bias and RMSE. On the other hand, our robust estimator is unaffected by this and the bias and RMSE are unchanged across the different values of δ. When $\delta \geq 0.005$ our robust estimator outperform the full solution estimator. The estimated nuisance parameters, γ, are reported in the supplemental material.

Table 1: MC results. Buffer-Stock model.

	$\delta = 0.000$		$\delta = 0.001$		$\delta = 0.005$		$\delta = 0.010$	
	Full[†]	Robust	Full[†]	Robust	Full[†]	Robust	Full[†]	Robust
BIAS	0.001	0.001	-0.018	0.001	-0.088	0.001	-0.159	0.001
STD	0.020	0.067	0.021	0.068	0.023	0.068	0.026	0.068
RMSE	0.020	0.067	0.027	0.067	0.090	0.068	0.161	0.068

Notes: The table reports the bias (BIAS $= \rho - \frac{1}{S}\sum_{s=1}^{S}\hat{\rho}_s$), standard deviation (STD $= \sqrt{\frac{1}{S-1}\sum_{s=1}^{S}(\hat{\rho}_s - \frac{1}{S}\sum_{s=1}^{S}\hat{\rho}_s)^2}$) and the root mean square error (RMSE $= \sqrt{\frac{1}{S}\sum_{s=1}^{S}(\rho - \hat{\rho}_s)^2}$) under different values of the true model, $\delta \in \{0, 0.001, 0.005, 0.01\}$ across $S = 200$ replications. The remaining model parameters are $\rho = 2$, $\tilde{\beta} = 0.98$, $R = 1.03$, $\sigma_\varepsilon^2 = \sigma_\eta^2 = 0.01$, $G_t = 1.02$, $\sigma_\xi^2 = 0.1$.

[†] The estimate refers to the full solution approach where $\delta = 0.0$ has been imposed under estimation, irrespective of the true value of this parameter.

3.2 Example 2: Education Choice and Labor Supply Model

We study a simplified version of the dynamic discrete choice model over education and labor supply proposed by Keane and Wolpin (1997). In this model, individuals live $\overline{T}$ periods and in each period they decide to work ($k = 1$), to go to school ($k = 2$), or to stay at home ($k = 3$). We note d_{kt} a dummy variable that equals 1 if the individual chooses alternative k and decisions are mutually exclusive ($d_{1t} + d_{2t} + d_{3t} = 1$). To make their decision, they compare the life-time utility of each alternative, which depends on their current set of state variables, $\mathcal{S}_t$. The direct payoff of alternative k, $U_t(k, \mathcal{S}_t)$, is a function of time spent at school S_t, labor market experience X_t, whether the individual was at school in the previous period d_{2t}, and on unobserved uncorrelated shocks $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t}, \varepsilon_{3t})$ which are *iid* mean zero Gaussian with variances σ_k^2. Collecting all these variables, the state vector is then $\mathcal{S}_t = (S_t, X_t, d_{2t}, \varepsilon_t)$.We use the following specification for per-period payoffs

$$U_t(k, \mathcal{S}_t) = \begin{cases} (1 - \nu_t)\exp(\alpha_0 + \alpha_1 S_t + \alpha_2 X_t + \varepsilon_{1t}) & \text{if } k = 1 \\ \alpha_3 + \alpha_4 \mathbf{1}\{S_t \geq 13\} + \alpha_5 d_{2,t-1} + \varepsilon_{2t} & \text{if } k = 2 \\ \alpha_6 + \varepsilon_{3t} & \text{if } k = 3, \end{cases}$$

where agents receive after tax income $w_t = (1 - \nu_t)\exp(\alpha_0 + \alpha_1 S_t + \alpha_2 X_t + \varepsilon_{1t})$, where ν_t is the tax rate if working ($k = 1$), a stipend of $\alpha_3 + \alpha_4 \mathbf{1}\{S_t \geq 13\} + \alpha_5 d_{2,t-1} + \varepsilon_{2t}$ if going to school ($k = 2$), or unemployment benefits $\alpha_6 + \varepsilon_{3t}$ if they stay at home ($k = 3$). For the purpose of our analysis, we suppose that the baseline tax rate is 0 ($\nu_t = 0$), and we introduce a tax-reform at time τ,

$$\nu_t = \delta \mathbf{1}\{t \geq \tau\},$$

where δ is the new tax rate. We suppose that agents fully anticipate the reform, but that the econometrician does not, and sets his model as if the tax reform was *unanticipated* by agents. So, in the full-solution estimation approach, agents are assumed to *believe* that $\nu_t = 0$ forever (or $\delta = 0$), while in reality they perfectly foresee the future reform.[9]

We define the value associated with state $\mathcal{S}_t$ as

$$\begin{aligned} V_t(\mathcal{S}_t) &= \max_{k \in \{1,2,3\}} U_t(k, \mathcal{S}_t) + \beta \mathbb{E}\left[V_{t+1}(\mathcal{S}_{t+1}) | \mathcal{S}_t, d_{kt} = 1\right] \\ &\text{s.t.} \\ X_{t+1} &= X_t + d_{1t} \\ S_{t+1} &= S_t + d_{2t} \end{aligned}$$

where school and labor market experience evolves deterministically.

When implementing the robust estimator, we approximate the value function at the pseudo terminal period using a flexible function of the observed state variables. Given that the model predicts that any value function will be positive, we constrain our approximate solution to be positive too,

$$\check{V}_{\check{T}}(\mathcal{S}_t) = \exp(\gamma_0 + \gamma_1 S_t + \gamma_2 X_t + \gamma_3 S_t X_t + \gamma_4 X_t^2 + \gamma_5 d_{2t})$$

where $\check{T} = \tau - 5$ as before. The full solution estimator and the robust approximate estimator of this model is, respectively,

$$\begin{aligned} \hat{\alpha}_1 &= \arg\min_{\alpha_1} \mathcal{L}(\theta) \\ (\check{\alpha}_1, \check{\gamma}) &= \arg\min_{\alpha_1, \gamma} \check{\mathcal{L}}(\theta, \gamma) \end{aligned}$$

where the log-likelihood functions are derived in the supplemental material. We observe discrete choices and the associated income, stipend or unemployment benefits, depending on the choice taken.

Table 2 reports the bias, standard deviation and the root mean square error (RMSE)

[9] Note, we here frame the analysis as a cohort-analysis since time and age are on-to-one in the current formulation. If data from several cohorts are used in estimation, the age should be an independent state variable in the model and the pseudy-terminal period approximation should then also be a function of age.

Table 2: MC results. Labor Supply Model.

	$\delta = 0$		$\delta = 0.01$		$\delta = 0.05$		$\delta = 0.1$	
	Full†	Robust	Full†	Robust	Full†	Robust	Full†	Robust
BIAS	-0.000	-0.001	-0.004	-0.001	-0.022	-0.001	-0.047	-0.002
STD	0.003	0.005	0.003	0.005	0.003	0.006	0.003	0.006
RMSE	0.003	0.006	0.005	0.006	0.023	0.006	0.047	0.006

Notes: See notes for Table 1. The remaining model parameters are $\alpha_0 = 8.9$, $\alpha_1 = 0.085$, $\alpha_2 = 0.031$, $\alpha_3 = 5000$, $\alpha_4 = -8000$, $\alpha_5 = -8000$, $\alpha_6 = 20000$, $\sigma_1^2 = 0.4$, $\sigma_2^2 = 1600$, and $\sigma_3^2 = 2000$.

† The estimate refers to the full solution approach where $\delta = 0.0$ has been imposed under estimation, irrespective of the true value of this parameter.

of the two estimators under different values of the tax-rate $\delta \in \{0, 0.01, 0.05, 0.1\}$. The numbers are multiplied by 100 for readability. We report results for the return to schooling parameter only and report in the supplemental material results when estimating all parameters of the wage equation. As in the first example, under relatively small tax reforms the robust estimator has lower bias and RMSE than the full solution estimator. The estimated nuisance parameters are reported in the supplemental material.

4 Concluding Discussion

We proposed an estimator of finite horizon dynamic choice models that is robust to misspecifications in the future economic environment faced or anticipated by agents. We have illustrated through two Monte Carlo experiments that estimating an approximate model in which we approximate a pseudo terminal period is independent of the economic environment after this pseudo terminal period. This approach has very promising applications. Particularly, using structural models to evaluate quasi-experiments will potentially benefit greatly from our approach because the researchers can remain agnostic about the anticipation of future reforms or changes in the contingencies. The simple approximate estimator also facilitates robustness checks of concrete model assumptions. For example, using our approach researchers gain insight into whether assuming some simple law of motion after the periods of interest is a sensible thing to do.

References

Arcidiacono, P., P. Bayer, F. A. Bugni and J. James (2013): "Approximating High-dimensional Dynamic Models: Sieve Value Function Iteration," in *Advances in Econometrics*, ed. by T. B. Fomby, R. C. Hill, I. Jeliazkov, J. C. Escanciano and E. Hillebrand, vol. 31, pp. 45–95. Emerald Group Publishing Limited.

Bhat, N., V. F. Farias and C. C. Moallemi (2012): "Non-parametric Approximate Dynamic Programming via the Kernel Method," Unpublished working paper, Graduate School of Business, Columbia University.

Blundell, R., M. C. Dias, C. Meghir and J. M. Shaw (2016): "Female Labour Supply, Human Capital and Welfare Reform," *Econometrica*, 84(5), 1705–1753.

Carroll, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 2, 61–135.

Deaton, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1221–1248.

Geweke, J. F. and M. P. Keane (2000): "Bayesian inference for dynamic discrete choice models without the need for dynamic programming," in *Simulation-based inference in econometrics: Methods and Applications*, ed. by R. Mariano, T. Schuermann and M. J. Weeks, chap. 4, pp. 100–131. Cambridge University Press, Cambridge UK.

Gourinchas, P.-O. and J. A. Parker (2002): "Consumption Over the Life Cycle," *Econometrica*, 70(1), 47–89.

Judd, K. (1998): *Numerical Methods in Economics*. MIT Press, Cambridge.

Judd, K. L., L. Maliar, S. Maliar and I. Tsener (forthcoming): "How to Solve Dynamic Stochastic Models Computing Expectations Just Once," *Quantitative Economics*.

Kaboski, J. P. and R. M. Townsend (2011): "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Incentive," *Econometrica*, 79(5), 1357–1406.

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," *The Journal of Political Economy*, 105(3), 473–522.

——— (2001): "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment," *International Economic Review*, 42(4), 1051–1103.

Kroemer, O. B. and J. R. Peters (2011): "A Non-Parametric Approach to Dynamic Programming," in *Advances in Neural Information Processing Systems*, ed. by J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira and K. Weinberger, vol. 24. Neural Information Processing Systems Foundation, Inc.

Manski, C. F. (1993): "Adolescent Econometricians: How Do Youth Infer the Returns to Schooling?," in *Studies of Supply and Demand in Higher Education*, ed. by C. T. Clotfelter and M. Rothschild, chap. 2, pp. 43–60. University of Chicago Press.

Rust, J. (2000): "Parametric Policy Iteration: An Efficient Algorithm for Solving Multidimensional DP Problems?," unpublished working paper, Yale University.

Rust, J., G. Hall, H. Benítez-Silva, G. J. Hitsch and G. Pauletto (2005): "A Comparison of Discrete and Parametric Approximation Methos for Solving Dynamic Programming Problems in Economics," Unpublished working paper, University of Maryland.

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

Supplemental Material: Robust Estimation of Finite Horizon Dynamic Economic Models

Thomas H. Jørgensen   Maxime Tô

June 30, 2017

A Example 1: Intertemporal Consumption Allocation

The buffer-stock model pioneered by Carroll (1992) and Deaton (1991) and estimated first by Gourinchas and Parker (2002) and Cagetti (2003). The model can be normalized with permanent income, P_t, and re-formulated as

$$\begin{aligned}
V_t(m_t) &= \max_{m_t \geq c_t \geq 0} \frac{c_t^{1-\rho}}{1-\rho} + \beta_t \mathbb{E}_t[(G_{t+1}\eta_{t+1})^{1-\rho} V_t(m_{t+1})] \\
&\text{s.t.} \\
m_{t+1} &= (G_{t+1}\eta_{t+1})^{-1}(1+r)a_t + \varepsilon_{t+1} \\
\log \varepsilon_t &\sim \mathcal{N}(-.5\sigma_\varepsilon^2, \sigma_\varepsilon^2) \\
\log \eta_t &\sim \mathcal{N}(-.5\sigma_\eta^2, \sigma_\eta^2)
\end{aligned}$$

where small letters denotes normalized values.

A.1 Solving the Model

To solve the model, we employ the endogenous grid method proposed by Carroll (2006) and use the normalized consumption Euler equation

$$c_t^{-\rho} = \max\left\{ m_t^{-\rho} \,,\ \beta_t(1+r)\mathbb{E}_t\left[(c_{t+1}G_{t+1}\eta_{t+1})^{-\rho}\right]\right\}$$

such that when agents are off the credit constraint, we can invert the Euler equation to get for the jth grid point in $\vec{a} = (a^1, \ldots, a^\#)$,

$$c_t^j = \beta_t(1+r)\sum_{k=1}^{Q}\sum_{l=1}^{Q}\omega^k\omega^l(\mathring{c}_{t+1}(m_{t+1}^{jkl})G_{t+1}\eta_{t+1}^k)^{-\rho}$$

where $\mathring{c}_{t+1}(m_{t+1}^{jkl})$ is a linear interpolation function evaluated at grid point j for the kth and lth Gauss-Hermite nodes associated with the permanent and transitory income shock, respectively, $m_{t+1}^{jkl} = (G_{t+1}\eta_{t+1}^k)^{-1}(1+r)a^j + \varepsilon_{t+1}^l$. The weights associated with these quadrature nodes are denoted ω^k and ω^l. We use 100 discrete points to approximate the consumption function ($\# = 100$) and 8 quadrature points to approximate each of the income shock integrals ($Q = 8$). The endogenous level of resources is then

$$m_t^j = c_t^j + a^j$$

At the credit constraint we simply have that $a = 0$ and including a point $(m_{t+1}^0, c_{t+1}^0) = (0, 0)$ when interpolating the future consumption function using linear interpolation automatically handles the credit constraint.

In the full solution estimator we solve the model from the terminal period T assuming that everything is consumed. This gives $c_T^j = m_T^j$, where we construct an exogenous grid over $\vec{m}$. However, we assume that the researcher insists on using a fixed discount factor, not taking into account that the true data generating process (DGP) contains time-varying preferences. Particularly, the full solution is

$$\begin{aligned} c_{t|\delta=0}^j &= \tilde{\beta}(1+r)\sum_{k=1}^{Q}\sum_{l=1}^{Q}\omega^k\omega^l(\mathring{c}_{t+1|\delta=0}(m_{t+1}^{jkl})G_{t+1}\eta_{t+1}^k)^{-\rho},\ \forall t = 1,\ldots,\overline{T}-1 \\ c_{\overline{T}|\delta=0}^j &= m^j \end{aligned}$$

In the robust estimator we solve the model backwards from the pseudy-terminal period $\check{T}$ where we have imposed the approximate solution $c_{\check{T}}^j = \exp(\phi_0)(m_{\check{T}})^{\exp(\phi_1)}$ where we use the same exogenous grid over $\vec{m}$. In turn, the solution is

$$\begin{aligned} \check{c}_t^j &= \beta_t(1+r)\sum_{k=1}^{Q}\sum_{l=1}^{Q}\omega^k\omega^l(\mathring{\check{c}}_{t+1}(m_{t+1}^{jkl})G_{t+1}\eta_{t+1}^k)^{-\rho},\ \forall t = 1,\ldots,\check{T}-1 \\ \check{c}_{\check{T}}^j &= \exp(\gamma_0)(m^j)^{\exp(\gamma_1)} \end{aligned}$$

This functional form was chosen to impose that $c(0) = 0$ and that consumption is an increasing and concave function of normalized resources.

A.2 Estimation Details

We assume that consumption is contaminated with additive mean zero measurement error

$$C_{i,t} = C_t^{\star}(M_{i,t}, P_{i,t}|\theta) + \xi_{i,t},\ \mathbb{E}[\xi|M, P] = 0$$

such that we can estimate model parameters using non-linear least squares

$$\begin{aligned}
\hat{\rho} &= \arg\min_{\rho} \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}(C_{i,t} - C_t^{\star}(M_{i,t}, P_{i,t}|\rho, \delta = 0))^2 \\
(\check{\rho}, \check{\gamma}) &= \arg\hat{o}\min_{\rho,\gamma} \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}(C_{i,t} - \check{C}_t^{\star}(M_{i,t}, P_{i,t}|\rho, \gamma))^2
\end{aligned}$$

where $C_t^{\star}(M, P)$ denotes the model-implied optimal consumption. To get this, we simply have to multiply the found solution with permanent income, $C_t^{\star}(M, P) = c_t^{\star}(m)P$. We assume for simplicity that we observe all state variables, including permanent income. When simulating data, we draw measurement error from a normal distribution with mean zero and variance $\sigma_{\xi}^2 = 0.1$.

Table 1 reports the mean and standard deviation of the estimated nuisance parameters $\gamma = (\gamma_0, \gamma_1)$ across the Monte Carlo runs.

Table 1: MC results, Nuisance Parameters. Buffer-Stock model.

	$\delta = 0.000$		$\delta = 0.001$		$\delta = 0.005$		$\delta = 0.010$	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
γ_0	-0.030	0.011	-0.029	0.011	-0.027	0.012	-0.026	0.014
γ_1	-2.019	0.163	-2.003	0.167	-1.939	0.182	-1.859	0.199

Notes: The table reports the mean and variance of γ for $\delta \in \{0, 0.001, 0.005, 0.01\}$ across $S = 200$ replications. The remaining model parameters are $\rho = 2$, $\tilde{\beta} = 0.98$, $R = 1.03$, $\sigma_{\varepsilon}^2 = \sigma_{\eta}^2 = 0.01$, $G_t = 1.02$, $\sigma_{\xi}^2 = 0.1$.

B Example 2: Discrete Labor Market Model

Here, we present details about the dynamic discrete choice model of education and labor supply. As mentioned in the paper, the Bellman equation of the model is:

$$\begin{aligned}
V_t(\mathcal{S}_t) &= \max_k U(k, \mathcal{S}_t) + \beta \mathbb{E}\left[V_{t+1}(\mathcal{S}_{t+1}) | \mathcal{S}_t, d_{kt} = 1\right] \\
&\text{s.t.} \\
X_{t+1} &= X_t + d_{1t} \\
S_{t+1} &= S_t + d_{2t} \\
\varepsilon_{kt} &\sim \mathcal{N}(0, \sigma_k^2),\ \forall k = 1, 2, 3
\end{aligned}$$

We first present how the solution of the dynamic programming was obtained. The model is estimated through maximum likelihood estimation, we thus provide elements for the derivation of the likelihood.

B.1 Dynamic programming solution

To describe the solution of the program, let $V_{kt}(\mathcal{S}_t)$, the alternative-specific value of decision k:

$$V_{kt}(\mathcal{S}_t) = U(k, \mathcal{S}_t) + \beta \mathbb{E}\left[V_{t+1}(\mathcal{S}_{t+1}) | \mathcal{S}_t, d_{kt} = 1\right]$$

Given a set of parameters, and knowing the next period value $V_{t+1}(\mathcal{S}_{t+1})$, we need to solve $\mathbb{E}\left[\max_k V_{k,t}(\mathcal{S}_t)\right]$ .To do so, several techniques can be used. Keane and Wolpin (1994) propose to get the expected value from simulations, but the computational cost of such estimate for a three dimensional integral may be substantial. Given the simple form of our model, we propose to take advantage of the knowledge of the cdf of each of the alternative specific value functions, and of the independence between random terms. Indeed, we can show that $M(\mathcal{S}_t) = \max_k V_{k,t}(\mathcal{S}_t)$ is a positive random variable which cdf is equal to the product of the cdfs of $V_{k,t}(\mathcal{S}_t)$, and for any positive random variable it is a well known result that:

$$\mathbb{E}(M(\mathcal{S}_t)) = \int_0^{+\infty} (1 - P(M(\mathcal{S}_t) < m))dm$$

So for any state S_t, we can obtain the expected value $V_{kt}(\mathcal{S}_t)$ at the cost of this unique integral, which we compute by adaptive quadrature.[1]

[1] To realize this, note first that since the labor income is always positive, the maximum over this and the remaining alternatives must also be positive. Further, note that since $(1 - P(M(\mathcal{S}_t) < m)) = P(M(\mathcal{S}_t) \geq m) =$

B.2 Likelihood function

Data consists in observed decisions (d_{1t}, d_{2t}, d_{3t}) and wages $w_t$if individuals are working. The derivation of the likelihood is obtained from the normality assumption of the ε_t random variables. At a given period given the individual state $\mathcal{S}_t$, we derive the likelihood contribution of each of the three decisions:

1. In the case of the labor market participation, the econometrician observes both the decision and the wage w_t. So the contribution is:

$$\ell_1(w_t|\mathcal{S}_t) = \frac{1}{\sigma_1}\phi\left(\frac{\log w_t - f_1(\mathcal{S}_t)}{\sigma_1}\right) P_1\left(d_{1t} = 1|w_t, \mathcal{S}_t\right)$$

   the probability of choosing to work conditional on the wage residual $\varepsilon_{1t} = \log w_t - f_1(\mathcal{S}_t)$ is

$$\begin{aligned} P_1(d_{1t} = 1|\varepsilon_{1t}, \mathcal{S}_t) &= P(V_{1t}(\mathcal{S}_t) > V_{2t}(\mathcal{S}_t), V_{1t}(\mathcal{S}_t) > V_{3t}(\mathcal{S}_t)|\varepsilon_1) \\ &= P(V_{1t}(\mathcal{S}_t) > V_{2t}(\mathcal{S}_t)|\varepsilon_1)P(V_{1t}(\mathcal{S}_t) > V_{3t}(\mathcal{S}_t)|\varepsilon_1) \end{aligned}$$

   where the second step obtained from the independence between unobserved heterogeneity terms. ε_2 and ε_3 appear separately in the two conditional probabilities, so these probabilities can be obtained from the functional form of the utility function and the normal distribution of ε_{2t} and ε_{3t}.

2. If the individual chooses to go to school, none of the three unobserved heterogeneity terms are observed. As in the previous case, one can use the independence conditional on ε_2 to

$\int_m^{+\infty} \frac{\partial}{\partial M(\mathcal{S}_t)} P(M(\mathcal{S}_t) \geq k)dk$

$$\begin{aligned} \int_0^{+\infty}(1 - P(M(\mathcal{S}_t) < m))dm &= \int_0^{+\infty}\int_m^{+\infty}\frac{\partial P(M(\mathcal{S}_t) \geq k)}{\partial M(\mathcal{S}_t)}dkdm \\ &= \int_0^{+\infty}\int_0^{k}\frac{\partial P(M(\mathcal{S}_t) \geq m)}{\partial M(\mathcal{S}_t)}dmdk \\ &= \int_0^{+\infty}\left[\frac{\partial P(M(\mathcal{S}_t) \geq m)}{\partial M(\mathcal{S}_t)}\right]_0^k dk \\ &= \int_0^{+\infty} k\frac{\partial P(M(\mathcal{S}_t) \geq k)}{\partial M(\mathcal{S}_t)}dk \\ &= \mathbb{E}(M(\mathcal{S}_t)) \end{aligned}$$

where the second equality sign follows from changing the order of integration. It is also evident from the fourth equality that this only holds if $M(\mathcal{S}_t)$ is non-negative.

derive a simple expression for the likelihood:

$$
\begin{aligned}
P_2(d_{2t}=1|\mathcal{S}_t) &= P(V_{2t}(\mathcal{S}_t) > V_{1t}(\mathcal{S}_t), V_{2t}(\mathcal{S}_t) > V_{3t}(\mathcal{S}_t)) \\
&= \int P(V_{2t}(\mathcal{S}_t) > V_{1t}(\mathcal{S}_t), V_{2t}(\mathcal{S}_t) > V_{3t}(\mathcal{S}_t)|\varepsilon_{2t}) dF(\varepsilon_{2t}) \\
&= \int P(V_{2t}(\mathcal{S}_t) > V_{1t}(\mathcal{S}_t)|\varepsilon_{2t}) P(V_{2t}(\mathcal{S}_t) > V_{3t}(\mathcal{S}_t)|\varepsilon_{2t}) dF(\varepsilon_{2t})
\end{aligned}
$$

Both terms in the integral can be obtained from functional forms. It is still necessary to compute the integral. In practice, we evaluate this integral using 20 Gauss-Hermite quadrature nodes.

3. The probability to choose to stay out of labor force $P_2(d_{2t}=1|\mathcal{S}_t)$ is derived analogously to the schooling case and does not need further details.

The complete log-likelihood is then

$$
\mathcal{L} = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T} d_{1,i,t}\log \ell_1(w_{i,t}\mathcal{S}_{it}) + d_{2,i,t}\log P_2(d_{2,i,t}=1|\mathcal{S}_{i,t}) + d_{3,i,t}\log P_3(d_{3,i,t}=1,\mathcal{S}_{i,t})
$$

B.3 Pseudo-Terminal Value Function

In the case of the robust estimate, the backward induction start from period $\check{T}$, instead of the terminal period $\bar{T}$. The pseudo-terminal value is set as a function of the additional set of parameters ϕ, which are estimated jointly with the rest of the parameters of the model. Table 2 reports the mean and standard deviation of the estimated nuisance parameters $\gamma = (\gamma_0, \gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5)$ across the Monte Carlo runs.

Table 2: MC results. Labor Supply Model.

	$\delta = 0$		$\delta = 0.01$		$\delta = 0.05$		$\delta = 0.1$	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
γ_0	11.4076	0.0020	11.4030	0.0021	11.3849	0.0021	11.3629	0.0024
γ_1	0.8424	0.0011	0.8421	0.0012	0.8405	0.0011	0.8379	0.0013
γ_2	0.2972	0.0007	0.2963	0.0007	0.2921	0.0007	0.2862	0.0008
γ_3	0.0010	0.0004	0.0013	0.0004	0.0025	0.0004	0.0042	0.0005
γ_4	0.0029	0.0004	0.0031	0.0004	0.0039	0.0004	0.0050	0.0005
γ_5	0.0106	0.0067	0.0106	0.0068	0.0095	0.0078	0.0074	0.0111

Notes: The table reports the average absolute bias (BIAS) and the root mean square error (RMSE) under different values of the true model, $\delta \in \{0, 0.1, 0.25, 0.5\}$ across $S = 200$ replications. The remaining model parameters are $\alpha_0 = 8.9$, $\alpha_2 = 0.031$, $\alpha_3 = 5000$, $\alpha_4 = -8000$, $\alpha_5 = -8000$, $\alpha_6 = 20000$, $\sigma_1^2 = 0.4$, $\sigma_2^2 = 1600$, and $\sigma_3^2 = 2000$.

† The estimate refers to the full solution approach where $\delta = 0.0$ has been imposed under estimation, irrespective of the true value of this parameter.

B.4 Estimating the whole wage equation

Table 3 shows results for the model if we were estimating the whole wage equation (i.e. estimating the wage intercept as well as estimating the returns to experience).

Table 3: MC results. Labor Supply Model.

	$\delta = 0$			$\delta = 0.01$			$\delta = 0.05$			$\delta = 0.1$		
	BIAS	STD	RMSE	BIAS	STD	RMSE	BIAS	STD	RMSE	BIAS	STD	RMSE
Full Solution												
β_1	-0.002	0.143	0.143	0.065	0.157	0.169	0.431	0.231	0.489	0.958	0.294	1.002
β_2	0.000	0.006	0.006	-0.006	0.007	0.009	-0.036	0.010	0.037	-0.077	0.013	0.078
β_3	0.001	0.005	0.005	-0.003	0.006	0.006	-0.018	0.007	0.019	-0.038	0.009	0.039
Robust Solution												
β_1	0.032	0.089	0.095	0.023	0.090	0.093	0.029	0.089	0.093	0.031	0.103	0.107
β_2	-0.002	0.004	0.005	-0.002	0.004	0.005	-0.002	0.005	0.005	-0.002	0.006	0.007
β_3	-0.002	0.012	0.012	-0.002	0.011	0.011	-0.002	0.011	0.011	-0.001	0.012	0.012

Notes: See notes for Table 2.

References

CAGETTI, M. (2003): "Wealth Accumulation Over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics*, 21(3), 339–353.

CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 2, 61–135.

——— (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91(3), 312–320.

DEATON, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1221–1248.

GOURINCHAS, P.-O. AND J. A. PARKER (2002): "Consumption Over the Life Cycle," *Econometrica*, 70(1), 47–89.

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.