A Simple and Robust Estimator for Discount Factors in Optimal Stopping Dynamic Discrete Choice Models

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

August 2, 2019

Abstract

We propose a simple and robust two-step estimator for discount factors in a class of dynamic discrete choice models. The estimator follows from constructive identification results, including a new identification result for a general, time separable discount function. The estimator is derived as the solution to well-behaved sample moment conditions which are linear in the discount factors and are independent of the utility function. The estimator is therefore easy to implement, computationally light, and in contrast to existing estimators, robust to biases from finite sample approximations to the unknown utility function. We apply the estimator to data on mortgage defaults under an identifying assumption of time homogeneity of the utility function. We compare the performance of the proposed estimator to alternative two-step estimators that jointly estimate the discount factor and the utility function. The results show that our proposed estimator's robustness to finite sample approximation bias and its computational ease do not necessarily come at a material expense of precision.

[*]Booth School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL 60637, USA. E-mail: Oeystein.Daljord@chicagobooth.edu

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

[‡]Department of Economics, Ewha Womans University, 52 Ewhayeodaegil, Seoul, Korea. E-mail: minjung_park@ewha.ac.kr

Keywords: time preferences; dynamic discrete choice models; optimal stopping problems; general discount function

JEL Classifications: C14; C50; M30

1 Introduction

Optimal stopping problems form a class of dynamic discrete choice (DDC) models with many applications in marketing. For instance, they are workhorse models in the durable goods literature. DDC models have theoretically well-defined parameters and allow evaluation of marketing strategies that are not observed in the data. One example is Nair (2007) which evaluates counterfactual dynamic pricing strategies in the video game market using an optimal stopping model.

A common criticism of DDC models, e.g., in Heckman and Vytlacil (2007), is that these models are costly to compute and that the conditions for identification of DDC models are generally hard to establish and are less economically transparent than in the program evaluation literature. The criticism reflects a concern about the credibility of the management relevant inferences that can be drawn from these models.

It is well known that the utility function and the discount factor are not jointly identified from choice data in the standard DDC model of Rust (1994). Recent papers by Bajari et al. (2016) and Abbring and Daljord (2017) have however shown that the discount factor in DDC models is identified under intuitively appealing and transparent economic assumptions. The focus on identification of the discount factor in these papers is motivated by Magnac and Thesmar (2002) which shows that conditional on the discount factor being known, the utility function is identified up to an arbitrary normalization.

Bajari et al. shows that in a finite horizon optimal stopping model, time homogeneity of the per-period utility function point identifies the discount factor. The discount factor is shown to be a closed form solution to a well-behaved moment condition with a clear economic interpretation as the current choice response to variation in future values. Abbring and Daljord shows that the discount factor is point identified for a larger class of models, including infinite horizon optimal stopping models, under similar identifying assumptions and interpretations. The results formalize conjectures and common intuition about assumptions that identify the discount factor in dynamic models used in marketing, e.g., Yao et al. (2012) and Chung et al. (2013).[1]

We develop an estimator of the discount factor based on the identification result in

[1] It is for instance commonly believed that the discount factor is identified in finite horizon models since the terminal period is an identified, static problem. Bajari et al. formalizes conditions under which this intuition is correct.

Bajari et al.. The estimator has many desirable properties. It can be estimated independently of the utility function and is therefore fully robust to finite sample approximations to the utility function. It uses all the information the data carry on the discount factor and nothing more. It has an intuitive and transparent interpretation: it recovers the discount factor as the sensitivity of current period choice probabilities to shifts in the continuation values. It can be implemented as a simple linear regression and it is fast and easy to compute.

We then show how the estimator applies to infinite horizon optimal stopping problems as well under Abbring and Daljord's identification results. We establish a new identification result for a fully general, time separable discount function, which includes hyperbolic discounting, for a class of finite horizon optimal stopping problems, and show that our estimator extends to this case as well.

In our application of the estimator to the mortgage data used in Bajari et al., we compare the performance of our proposed estimator to alternative two-step estimators that recover the discount factor jointly with the utility function under the same identifying assumptions. We find that the discount factor estimate from our proposed estimator is similar to those from the alternative estimators. The estimate is slightly higher in value and slightly less precise, but statistically indistinguishable from those produced by joint estimation. The results show that our proposed estimator's robustness to finite sample approximation bias and its computational ease do not necessarily come at a material expense of precision.

We derive the model and the identifying moment condition in Section 2 and the estimator in Section 3. In Section 4, we compare the proposed estimator to joint estimators in an application to the mortgage data. In Section 5, we show that our estimator applies to a class of infinite horizon optimal stopping models, while maintaining the assumption of standard exponential discounting. We then give our new identification result for a general, time separable discount function in finite horizon optimal stopping problems and show that our proposed estimator extends to this case as well.

2 Model

The choice model is a single agent, finite horizon optimal stopping version of Rust (1994). Time is discrete and indexed by $t = 1, \ldots, T$. The choice set is discrete and indexed by

$a \in \mathcal{A} = \{0, 1, \ldots, A\}$, where $a = 0$ denotes a terminating action. The panel may or may not include the final period T. The exogenous states $s \in \mathcal{S}$, which are observable to both the agent and the econometrician, follow a stationary, first-order Markov process with probability distribution $F(s_{t+1}|a_t, s_t)$ without absorbing states.

The states $\epsilon_t = \{\epsilon_{0,t}, \ldots, \epsilon_{A,t}\}$, which are observed by the agent, but not by the econometrician, are drawn independently of s from a stationary and absolutely continuous distribution G with infinite support. The agent's per-period utility function is $u_t(a, s, \epsilon) = u_t(a, s) + \epsilon_{a,t}$. The per-period utility of the reference choice is assumed state invariant and normalized to $u_t(0, s) = 0$ for all $s \in \mathcal{S}$ and $t = 1, \ldots, T$. The agent's expectations are assumed to be rational and coincide with $F(s_{t+1}|a_t, s_t)$.

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

$$E\left[\sum_{\tau=0}^{T-t} \beta^\tau \max_{a \in \mathcal{A}} u_{t+\tau}(a_{t+\tau}, s_{t+\tau}, \epsilon_{a,t+\tau}) | a_t, s_t, \epsilon_t\right],$$

where β is a finite and non-negative discount factor.[2] The primitives of the problem are the utility functions u_t, the discount factor β, the state transition process F, and the distribution of unobservable states G. We assume that ϵ_t is i.i.d. EV1 distributed throughout for ease of exposition and computation.

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

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

where the ex ante value function

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

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

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

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

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

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

The moment conditions in (4) are the reduced form of the model.[3] They map the endogenous objects, on the right hand side, to known functions of the data, on the left hand side. Using the choice specific value function in (1), the reduced form can be written in terms of the primitives

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

The normalization $u_t(0,s) = 0$ for the terminating action implies $V_t(0,s) = u_t(0,s) = 0$. Then using (2) and (3), the moments in (5) are functions of the data and the unknown primitives u and β

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

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

2.1 Identification

Bajari et al. shows that assuming time homogeneity of the utility function point identifies the discount factor in finite horizon optimal stopping models developed in the previous section. Time homogeneity means that the utility function is time invariant for some subset of the state space and choice set. Time homogeneity can be cast as exclusion restrictions of the form

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

The following theorem, based on Bajari et al., shows that this exclusion restriction point identifies the discount factor.[4]

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

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

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

Proof:

Difference (6) between t and t' to get

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

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

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

Q.E.D.

The moment condition in (8) gives the discount factor as a closed form function of choice probabilities and expectations, which are both identified features of the data distribution. As a result, the moment condition in (8) is independent of the utility function. It also contains all the information in the data about the discount factor, in the sense that the discount factor is underidentified from the moment conditions in (6) alone, but identified from (8) independently of (6). This fact motivates a robust estimator for the discount factor based on the moment condition in (8) only.

Since the continuation values in the denominator on the right hand side of (8) are known mappings from the data distribution, point identification holds even for panels which do not include the final period, i.e., short panels, as long as the data distribution contains information on at least three consecutive periods (i.e., $t' = t + 1$).

The discount factor is the ratio of the difference in log choice probabilities between periods t and t', in the numerator, to the difference in continuation values between the same periods, in the denominator. The stronger the current period choice response (the numerator) is to shifts in future rewards (the denominator), the more forward looking the agent is. In the extreme case, if the behaviour of the economic agent does not respond to shifts in the continuation values, we can infer that the agent is myopic, i.e., $\beta = 0$.

We can think of the time periods t and t' as an 'instrument' that shifts the continuation values without shifting the current period utilities by the assumption of time homogeneous utilities. The rank condition in (9) requires that time shift the continuation values. If not, there is no information about time preferences in the model, and identification is lost.

3 Estimator

We develop a simple and robust two-step estimator based on the moment conditions in (8). In the empirical application in Section 4, we compare our proposed estimator to alternative two-step estimators that jointly recover the utility function and the discount factor based on the moment conditions in (6).

Both our robust estimator and the joint estimators require estimating the conditional choice probabilities $\sigma_t(a, s)$ and the continuation values $E\left[-\ln(\sigma_{t+1}(0, s_{t+1}))\middle| a_t = a, s_t = s\right]$ in the first step. In the second step, the first step estimates are plugged into the respective sample moment conditions to recover the discount factor.

Non-parametric estimation of the conditional choice probabilities has become standard in the literature and can be implemented in standard statistical software such as R, STATA or MATLAB, for example, by running a multinomial logit regression using orthogonal polynomials of the state variables s as basis functions, see e.g., Judd (1998). The same standard procedures can be used here, but one noteworthy feature is that estimation needs to account for non-stationarity of the problem due to finite horizon, for example by estimating conditional choice probabilities separately for each period t or by including orthogonal polynomials of both t and s in the basis function.

Once we construct the fitted values $\hat{\sigma}_t(a, s)$ for each observation in the sample, we can recover the continuation values as a function of current states and actions. Typically, the continuation value is recovered by estimating the state transitions $F\left(s_{t+1}|a_t, s_t\right)$ from the data and then computing $\int -\ln(\sigma_{t+1}(0, s_{t+1}))dF(s_{t+1}|a_t = a, s_t = s)$ using the estimated state transitions and conditional choice probabilities.

Following Bajari et al., we can alternatively recover the continuation values by running a linear regression of $-\ln\left(\hat{\sigma}_{t+1}(0, s_{t+1})\right)$, which we compute for each observation in the data, on orthogonal polynomials of state variables s_t, separately for each action a. As before, the non-stationarity of the problem can be accounted for either by running the estimation separately for each period t as well as each action a or by including orthogonal polynomials of both t and s_t in the basis function. Compared to the conventional approach which separately estimates the state transitions, the key advantage of this approach is computational.

The exact procedures of the first step estimation (for conditional choice probabilities and continuation values) are not key to our estimator, but for concreteness we briefly discuss the approaches used in the empirical application of Section 4. In the application, we estimate conditional choice probabilities using a multinomial logit where the basis function includes restricted cubic splines for the continuous state variables and time periods t. The basis function also includes the binary state variables as well as interactions among all the state variables, both binary and continuous, and t.[5]

For construction of the continuation values, we regress $-\ln\left(\hat{\sigma}_{t+1}(0, s_{t+1})\right)$ on the same regressors that we use in the estimation of the conditional choice probabilities, namely restricted cubic splines for continuous state variables and t, binary state variables, and interactions among them. The fitted values from this regression give us estimates of the continuation values $\hat{E}\left[-\ln(\sigma_{t+1}(0, s_{t+1}))\middle| a_t = a, s_t = s\right]$ for each observation in the sample.

For the second step, we need to decide which set of time periods we impose the time homogeneity assumption on. For identification, it is sufficient that the exclusion restriction in (7) holds for only one pair of periods $t, t' \in 1, \ldots, T$, one choice $a \in \mathcal{A}\backslash\{0\}$, and some subset of the state space $\mathcal{S}$. The economics of the problem may however justify that the exclusion restriction holds more generally. Bajari et al. assumes time homogeneity for all observed periods, states, and choices, and our empirical application in the next section will use that same identifying assumption for all three estimators.

Our proposed estimator relies on differencing the moment conditions in (6) between two periods t and t' to obtain (8). There are many possible ways to do the differencing under the assumption of time homogeneity for all observed t. In the empirical application, we use first differencing, i.e., we set $t' = t + 1$ in (8).[6]

We construct our proposed estimator and joint estimators as the sample equivalents to the population moment conditions in (8) and (6), respectively, by plugging in the first step estimates of the conditional choice probabilities and the continuation values. The second

[5] STATA code with detailed comments is provided as an online appendix.

[6] Different pairs t and t' generally contain different and unique information about time preferences.

step of our robust estimator is hence

$$\ln\left(\frac{\hat{\sigma}_t(a,s)}{\hat{\sigma}_t(0,s)}\right) - \ln\left(\frac{\hat{\sigma}_{t+1}(a,s)}{\hat{\sigma}_{t+1}(0,s)}\right) = \tag{10}$$
$$\beta\left(\hat{E}\left[\ln(\sigma_{t+2}(0,s_{t+2}))\middle| a_{t+1}=a, s_{t+1}=s\right] - \hat{E}\left[\ln(\sigma_{t+1}(0,s_{t+1}))\middle| a_t=a, s_t=s\right]\right)$$

which can be seen a system of equations, where we have one equation for each non-terminating action, each equation contains one regressor and no constant, and cross-equation restriction of common β is imposed. The discount factor can therefore be estimated using seemingly unrelated regressions, which will improve efficiency by assigning a greater weight to equations with lower variance and accounting for cross-equation correlation in residuals.

For the second step of the joint estimators, we must choose a utility specification. Since the discount factor is point identified jointly with a non-parametric utility function under our identifying assumptions, we can think of the chosen utility specification as a sieve estimator of the unknown utility function, see e.g., Chen (2007). A sieve estimator estimates the infinite dimensional unknown utility function by fitting a lower-dimensional approximation, a sieve, where the dimensionality of the sieve space increases with the sample size. We use linear sieves, which are linear spans of finitely many known basis functions. For instance, we can interpret a linear utility specification as a first-order linear sieve to the unknown utility function.

The difference between the true utility function and the utility specification can be interpreted as finite sample approximation bias. In contrast, the difference is misspecification bias under identification by functional form.[7] Unlike misspecification bias, the approximation bias vanishes asymptotically as the dimensionality of the sieve increases with the sample size.

For the joint estimators, the sample analogs of (6) are

$$\ln\left(\frac{\hat{\sigma}_t(a,s)}{\hat{\sigma}_t(0,s)}\right) = g(s,\gamma_a) - \beta\hat{E}\left[\ln(\sigma_{t+1}(0,s_{t+1}))\middle| a_t=a, s_t=s\right],$$

where γ_a is a vector of action-specific parameters to be estimated, and $g(s,\gamma_a)$ is a linear sieve that approximates the unknown utility function $u(a,s)$.

We estimate two versions of the joint estimators which differ in the dimensionality of the sieve. One is a first order approximation, i.e., $g(s,\gamma_a) = s\gamma_a$ and the other is a higher

[7]The discount factor is typically identified by functional form under a linear utility assumption, see e.g., Rust (1987).

order approximation which uses restricted cubic splines of the state variables as in the first step procedure. To impose time homogeneity, we exclude t from the states s in $g(s, \gamma_a)$ and restrict γ_a to be time invariant. Both estimators are linear in the parameters γ_a and β, and they can be estimated using seeming unrelated regressions with one equation for each non-terminating action.

None of the three estimators is semi-parametrically efficient. The two joint estimators differ in the dimensionality of the approximation to the utility function and hence how they trade off bias to variance. The first order approximation has higher bias but lower variance compared to the specification with a higher order approximation. The key difference between the robust estimator and the joint estimators is that u is removed via differencing under the proposed estimator. As a result, while approximation bias from the utility function will in general affect the discount factor estimate under the joint estimators, our proposed estimator is robust to these biases. On the flip side, our proposed estimator will suffer from an efficiency loss compared to the joint estimators if the utility specification imposed under the joint estimators is correct.

For all three estimators, we use bootstraps for inference. We re-estimate both the first step and the second step in each bootstrap iteration to account for sampling variation in both the first step and second step estimates. In Section 2.1, we noted a similarity of our proposed estimator to an instrumental variable estimator. If there is little variation in the population continuation values over time, so that the denominator in (8) is close to zero, the estimator becomes disproportionately sensitive to small changes from sampling variation in the continuation value differences. If so, the estimator may not be asymptotically normal.

This case is similar to the problem of weak instruments: if time only weakly shifts the continuation values, holding current utility constant, then the estimates become numerically unstable, standard inference is invalid, and hence we cannot bootstrap the standard errors. In separate work, we show that the continuation value differences need to have sufficiently high density at values away from zero for the estimator to be asymptotically normal. In the next section, we verify the existence of sufficient variation in the continuation values in the data before proceeding with inference.

4 Empirical application

We use the subprime mortgage data in Bajari et al.. We give only a brief description of the data here, see Bajari et al. for a more extensive description. The data are from CoreLogic and cover mortgage loans that are securitized in the private-label market. The data report loan terms and borrower characteristics for each loan, such as the type of the mortgage (fixed rate, adjustable rate, etc.), the loan amount, the location of the property, the mortgage interest rate and the purchase price of the property.

There are three possible actions in each period for the borrower: default, prepay, or make the regular monthly payment only. The data track each loan over the course of its life, and stop tracking a loan if either the borrower defaults on the loan or the borrower makes a full prepayment. If the bank takes possession of the property or the loan has been delinquent for more than 90 days, the loan is considered to be in default. If the loan balance is observed going to zero before maturity, the loan is considered to be prepaid. If the borrower makes the regular monthly payment only, the loan survives into the next period and the data continue to track the loan.

The data cover loans originated between January 2000 and September 2007 and track each loan's status until March 2013 on a monthly basis. Due to the relatively short span of the sample compared with the mortgage term, we do not have data on the behaviour of borrowers whose mortgages are close to maturity, i.e., we have a short panel.

The data are selected on the same criteria as in Bajari et al., focusing on 30-year fixed-rate mortgages with non-missing information on borrower income in 20 major Metropolitan Statistical Areas, and the selected sample includes about 11,500 borrowers. The state variables that are used for estimation are listed in Table 1. The summary statistics for the estimation sample are given in Table 2.

It is in general hard to make meaningful statements about which moment conditions identify individual parameters in non-linear models where all parameters are simultaneously determined in all equations. However, for our robust estimator, the discount factor is determined by moment conditions that must hold independently of the other parameters, which here represent the utility function. As a result, we can meaningfully point to the variation in the data that identifies the discount factor.

In particular, the discount factor is identified by variation in the continuation values between periods, which is the independent variable on the right hand side of (10), and

Table 1: State variables

Low Doc	= 1 if the loan was done with no or low documentation on wealth and income, = 0 otherwise.
Multiple Liens	= 1 if the borrower has other, junior mortgages, = 0 otherwise.
FICO	FICO score at loan origination, a credit score by Fair Isaac & Co. Scores range between 300 and 850, higher for higher credit quality.
Payment	Monthly payment due.
Prepayment Penalty	= 1 if the loan has prepayment penalty, = 0 otherwise.
Income	Borrower's monthly income, imputed from front-end debt-to-income ratio and monthly payment due.
Contractual Interest Rate	Interest rate specified by the mortgage.
MSA	Metropolitan Statistical Area of the loan property.
Housing Value	Current housing value, imputed by adjusting appraised property value at loan origination by a zip-code level home price index from CoreLogic.
Net Equity	Current housing value minus Outstanding loan balance.
Market Rate	Current market interest rate available to the borrower for refinancing.
Unemployment Rate	Monthly unemployment rate at the county level.

Table 2: Summary statistics

Time-invariant Variables	No. of Loans	Mean	Standard Deviation
Low Doc	11,685	0.396	0.489
Multiple Liens	11,685	0.140	0.347
FICO	11,685	672.40	744.87
Income ($)	11,685	6238.9	6560.9
Contractual Interest Rate	11,685	7.670	1.724
Time-varying Variables	**No. of Loan-Month Pairs**	**Mean**	**Standard Deviation**
Payment ($)	478,950	1209.6	900.93
Prepayment Penalty	478,950	0.406	0.491
Housing Value ($1000)	478,950	283.28	263.35
Net Equity ($1000)	478,950	108.14	165.15
Market Rate	478,950	6.812	2.076
Unemployment Rate	478,950	6.470	2.561

which matches with common intuition that time preferences reflect how an agent's current choices respond to changes in expected future values. The differences in continuation values between two adjacent months are plotted against loan age in Figure 1, conditional on state variables and action and where the periodization is in months. Each dot represents an observation in the sample. The graph shows that the differenced continuation values have sufficiently high positive density at values away from zero, mitigating concerns over insufficient variation in the continuation values.

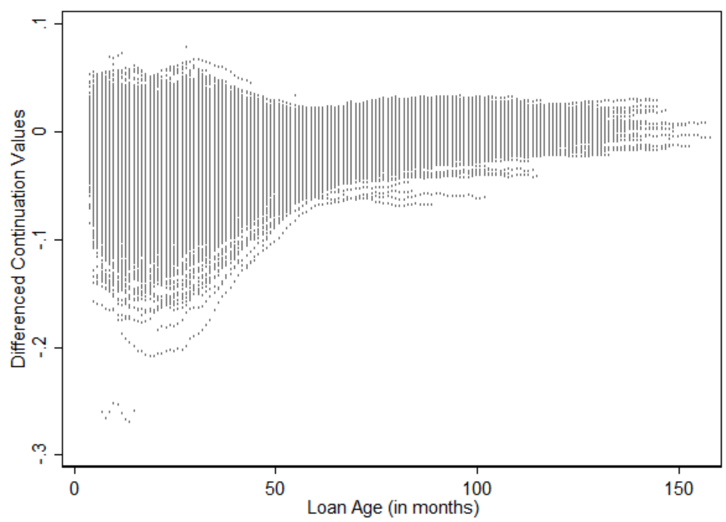

Figure 1: Identifying variation

The identifying assumption of time homogeneity of the per period utilities is substantive and allows us to interpret changes in the distribution of the default choice over time, holding states fixed, as informative about time preferences. In other words, once we control for state variables such as remaining loan amount, current home value, interest rates, etc., the observed difference in default behaviour over time will be entirely attributed to the fact that the loan gets closer to the maturity (the terminal period).

While preferences are commonly assumed time invariant in finite horizon models, e.g,. Dubé et al. (2014), we may be open to the idea that utilities evolve in ways that are predictable to the agent. For instance, in a study of the career decisions of young men, Keane and Wolpin (1997) allows the utilities to deterministically and predictably shift between early and later adolescence. Within the early and the later phases, time homogeneity is however implicitly imposed. The discount factor is therefore identified from variation within each phase of adolescence, but not between phases. If we abandon the

Table 3: Estimation results

Proposed Estimator	Joint Estimator 1	Joint Estimator 2
0.982 (0.017)	0.974 (0.011)	0.972 (0.012)

time homogeneity assumption altogether, other restrictions on the utility function must be imposed to ensure identification. We give examples on alternative identification strategies in Section 5.

The time homogeneity assumption comes in addition to the maintained assumptions of the model, such as rational expectations and the i.i.d. EV1 distribution of the unobservable states. While we recognize that these are also strong and substantive assumptions, they are standard. It is beyond the scope of this paper to weaken the maintained assumptions further, but see e.g., Aguirregabiria and Magesan (2016) for identification of non-rational beliefs and Buchholz et al. (2016) for identification of the distribution of the unobservable states.

The estimates are reported in Table 3. Joint Estimator 1 uses a first order approximation for the utility, i.e., assumes a linear utility function. Joint Estimator 2 uses restricted cubic splines of the state variables for the specification of the utility function. Since our focus is on the discount factor, we only report the discount factor estimates although the two joint estimators produce the estimates of the utility function as well.[8] Standard errors are reported inside the parentheses and are based on 500 bootstraps.

The discount factor estimates are very similar for the three estimators. We cannot reject a null hypothesis that the estimates are equal by conventional significance levels.[9] The precision of the estimates is also similar under the three approaches. As expected, the standard errors are greatest under the proposed estimator and smallest under the joint estimator with a first order approximation for the utility function, but the difference is only slight. Between the two joint estimators, the fact that the standard error is greater when the utility function is more flexibly specified is not surprising as it has fewer degrees

[8] Under the proposed estimator, the utility function can be also recovered by plugging the estimated discount factor in the moment conditions in (6).

[9] The monthly interest rate implied by the estimated discount factor is 1.8%-2.8%, much higher than the monthly interest rates of the mortgages observed in the data. This suggests that the usual approach of fixing the discount factor at a level implied by the market interest rate and estimating the remaining structural parameters could lead to biased estimates in the considered application.

of freedom.

The fact that the estimates and standard errors are very similar under the two joint estimators suggests that linear utility is a good approximation in this application. In settings where a linear utility assumption is a poor approximation to the unknown utility function, the first order joint estimator would expectedly differ materially from both the higher order joint estimator and the robust estimator. In settings where the utility function is highly non-linear and cannot be accurately approximated by the restricted cubic splines of the state variables, we would also expect the higher order joint estimator and the proposed estimator to give materially different estimates, with only the latter being bias-free. The fact that the proposed estimator produces only slightly greater standard error compared to the joint estimators, while being agnostic about the form of the utility function, shows that its robustness does not necessarily come at a material loss of efficiency.

In summary, we find that the proposed estimator gives a sensible estimate to a reasonable level of precision. Given the conceptual simplicity of the proposed estimator, its computational feasibility and its robustness, we believe the proposed estimator can be useful for many empirical researchers interested in robust estimation of time preferences.

5 Extensions

5.1 Infinite horizon optimal stopping models

Our robust estimator applies to infinite horizon optimal stopping problems as well, but the time homogeneity assumption does not have identifying power in infinite horizon, stationary DDC models. Abbring and Daljord shows that the discount factor is point identified if we can find a pair of states, or a pair of choices, for which the current period utility is invariant, but the continuation values differ.

This identification strategy reflects common intuition reported in the literature. In a study of housing, Bayer et al. (2016) argues that only current amenities, such as neighbourhood crime rates, affect the current period utility of housing, but lagged amenities are predictive of future amenities and can therefore shift continuation values without affecting current period utility, conditional on current amenities. If so, the responses of current period choice to lagged amenities can be informative about time preferences. In a study of demand of video game consoles, Lee (2013) assumes that the expected quality and availability of future game releases shift the continuation values of owning a console, but not its current period utility. The assumption leads Lee to interpret current period demand responses to variation in the expected quality and availability of future releases as informative about time preferences.

This intuition can be formalized as exclusion restrictions on the current period utility. If there exists either a pair of choices $a_1 \in \mathcal{A}\backslash\{0\}$ and $a_2 \in \mathcal{A}$, or a pair of states $s_1, s_2 \in \mathcal{S}$, where either $a_1 \neq a_2$, or $s_1 \neq s_2$, or both, such that the exclusion restriction

$$u(a_1, s_1) = u(a_2, s_2) \tag{11}$$

holds, then Theorem 1 in Abbring and Daljord shows that the discount factor is in general set identified, subject to a rank condition similar to the one in (9). Example 1 in Abbring and Daljord shows that for problems where $V(0, s)$, the choice specific value of the reference choice, is a constant, the discount factor is point identified. These problems include optimal stopping problems with a terminating action, like in Bajari et al., and regenerative optimal stopping problems like in Rust (1987). For these problems, the identifying moment condition corresponding to (8) is

$$\beta = \frac{\ln\left(\frac{\sigma(a_1,s_1)}{\sigma(0,s_1)}\right) - \ln\left(\frac{\sigma(a_2,s_2)}{\sigma(0,s_2)}\right)}{E\left[\ln(\sigma(0,s'))\middle| a = a_2, s = s_2\right] - E\left[\ln(\sigma(0,s'))\middle| a = a_1, s = s_1\right]}. \tag{12}$$

Adapting our proposed estimator to the moment condition in (12) is immediate. By extension, the joint estimators also apply to infinite horizon optimal stopping problems given the exclusion restriction in (11). Finally, the exclusion restrictions in (11) are sufficient to point identify the discount factor in finite horizon models (i.e., time homogeneity is not additionally required), see Abbring and Daljord. Adapting the finite horizon estimator to the exclusion restriction in (11) is again immediate.

5.2 General discount functions

While we developed our estimator for the case of standard exponential discounting, we next show that it extends to more general time preferences in a finite horizon optimal stopping model. As long as the time preferences are expressed as multiplicative weights for the utility in the future periods and the economic agent solves the dynamic program in each period by optimizing the flow of weighted future utilities, our approach can be used to estimate general discount function.

As is common in the literature on time-inconsistent preferences, we adopt the convention of considering the agent as a collection of selves, one in every period, whose preferences may be in conflict. The literature on time inconsistent preferences has distinguished between "sophisticated" agents, who are self-aware about the time inconsistency of their preferences, versus those who are "naïve," as well as an entire range in between, e.g., Malmendier and DellaVigna (2006). For our arguments, we focus on the case of sophisticated agents who are fully aware of the time inconsistency of their preferences.

We do so not because we believe naïveté is irrelevant in practice, but because identifying the degree of naïveté would require additional information that is unavailable in most datasets. For example, identifying the degree of sophistication may require observing whether individuals make appropriate use of self-commitment devices that constrain their future ability to take suboptimal actions, which only sophisticated agents would do. Assuming sophisticated agents enables two-step estimation, by ensuring that the observed choices of future selves are consistent with the current self's expectation about future selves' actions.

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

$$U_t(a_t, \ldots, a_T, s_t, \ldots, s_T) = \sum_{r=t}^{T} \beta(r-t)(1 - \max_{q<r} \mathbf{1}\{a_q = 0\})(u_r(a_r, s_r) + \epsilon_{a_r}).$$

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

$$V_t(a, s) = u_t(a, s) + E\Bigg[\sum_{r=t+1}^{T} \beta(r-t) \prod_{q<r} \mathbf{1}\{d_q(s_q, \epsilon_q) \neq 0\} \\ \Big(u_r(d_r(s_r, \epsilon_r), s_r) + \epsilon_{d_r(s_r,\epsilon_r)}\Big) \Bigg| a_t = a, s_t = s \Bigg].$$

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

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

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

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

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

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

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

Under exponential discounting, $\beta(t) = \beta^t$, we get $\delta(r) \equiv 0$ for all $r > 1$, and now the generalized Bellman equation reduces to the conventional Bellman equation for all $t \in 1, \dots T$.

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

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

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

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

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

This expression allows us to formulate the following theorem.

Theorem 2 *Suppose that the data is generated by sophisticated agents who are aware of their time inconsistency while making their discrete choice and that the following rank condition is satisfied: there exist periods t and t' such that the principal minors of matrix M with entries $M_{rp} = Cov(Z_r, Z_p)$ has full rank, where*

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

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

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

for $T-\max(t,t') < r \le T-\min(t,t')$. Also suppose that choice probabilities are observed for all periods $\min(t,t'),\ldots,T$ (i.e., the dataset is a long balanced panel). Then the general discount function $\beta(\cdot)$ is identified on the subset of its support $1,\ldots,T-\min(t,t')$.

Proof:

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

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

and

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

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

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

This means that

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

As a result

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

By assumption of the theorem M^{-1} exists and, thus, solution $\vec{\delta}$ is well-defined. Then we can find the discounting schedule by setting $\beta(1) = \delta(1)$ and then iterating to get

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

Q.E.D.

Note that if both t and t' are greater than 1, the discount function is only identified on the subset of support, and not in its entirety. The general discount function can be estimated by regressing the difference between the ratio of the probabilities on the expected differences of future choice probabilities for the reference choice, similar to the previous examples. Intuitively, the general discount function is recovered by checking how sensitive the current period choices are to variation in rewards one period later, two periods later, etc. By using the moment condition in (15), we can recover how quickly the weights on future values decline in explaining the current period choices and perform various hypothesis testing on time preferences, e.g., test whether the data are consistent with hyperbolic discounting.

6 Summary

We developed a simple two-step estimator for the discount factors in optimal stopping models. The proposed estimator builds on a transparent and intuitively appealing identification argument, is easy to compute, and is robust to approximation bias for the utility function. In an application to mortgage default data, the estimator delivers estimates of reasonable values and precision compared to alternative estimators. The estimator extends to infinite horizon problems as well as finite horizon problems with non-standard preferences.

References

Abbring, J. and O. Daljord (2017). Identifying the discount factor in dynamic discrete choice models. Mimeo, University of Chicago. 3, 4, 16, 17

Aguirregabiria, V. and A. Magesan (2016). Identification and estimation of dynamic games when players' beliefs are not in equilibrium. Technical report, University of Toronto. 15

Bajari, P., C. S. Chu, D. Nekipelov, and M. Park (2016). Identification and semiparametric estimation of a finite horizon dynamic discrete choice model with a terminating action. *Quantitative Marketing and Economics 14*(4), 271–323. 3, 4, 6, 8, 9, 12, 17

Bayer, P., R. McMillan, A. Murphy, and C. Timmins (2016, May). A dynamic model of demand for houses and neighborhoods. *Econometrica 84*(3), 893–942. 16

Buchholz, N., m. Shum, and H. Xu (2016). Semiparametric estimation of dynamic discrete choice models. Technical report, University of Texas. 15

Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Chapter 76, pp. 5550–5623. Elsevier. 10

Chung, D., T. Steenburgh, and K. Sudhir (2013). Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. *Marketing Science 33*(2), 165–187. 3

Dubé, J.-P., G. J. Hitsch, and P. Jindal (2014). The joint identification of utility and discount functions from stated choice data: An application to durable goods adoption. *Quantitative Marketing and Economics 12*(4), 331–377. 14

Heckman, J. and E. Vytlacil (2007). Econometric evaluation of social programs, part i: Causal models, structural models and econometric policy evaluation. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics* (1 ed.), Volume 6B, Chapter 70, pp. 4779–4874. Elsevier. 3

Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies 60*(3), 497–529. 6

Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press. 8

Keane, M. P. and K. I. Wolpin (1997, June). The career decisions of young men. *Journal of Political Economy 115*(3), 473–522. 14

Lee, R. (2013). Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review 103*(7), 2960–3000. 16, 17

Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete choice processes. *Econometrica 70*, 801–816. 3, 6

Malmendier, U. and S. DellaVigna (2006). Paying not go to the gym. *American Economic Review 96*(3), 674–719. 18

Nair, H. (2007). Intertemporal price discrimination with forward-looking consumers: Application to the us market for console video-games. *Quantitative Marketing and Economics 5*(3), 239–292. 3

Norets, A. and S. Takahashi (2013). On the surjectivity of the mapping between utilities and choice probabilities. *Quantitative Economics 4*, 149–155. 6

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica 55*, 999–1033. 10, 17

Rust, J. (1994). Structural estimation of Markov decision processes. In R. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, pp. 3081–3143. Amsterdam: North-Holland. 3, 4

Yao, S., C. F. Mela, J. Chiang, and Y. Chen (2012). Determining consumers' discount rates with field studies. *Journal of Marketing Research 49*(6), 822–841. 3