Identifying Present Bias from the Timing of Choices*

Paul Heidhues Philipp Strack

DICE Yale

February 22, 2021

Abstract

A (partially naive) quasi-hyperbolic discounter repeatedly chooses whether to complete a task. Her net benefits of task completion are drawn independently between periods from a time-invariant distribution. We show that the probability of completing the task conditional on not having done so earlier increases towards the deadline. Conversely, we establish non-identifiability by proving that for *any time-preference parameters* and *any data set* with such (weakly increasing) task-completion probabilities, there exists a stationary payoff distribution that rationalizes the agent's behavior if she is either sophisticated or fully naive. Additionally, we provide sharp partial identification for the case of observable continuation values.

*We thank seminar and conference audiences, and especially Nageeb Ali, Ned Augenblick, Stefano DellaVigna, Drew Fudenberg, Ori Heffetz, Botond Kőszegi, Muriel Niederle, Philipp Schmidt-Dengler, Charles Sprenger, Dmitry Taubinsky, and Florian Zimmermann for insightful and encouraging comments. Part of the work on this paper was carried out while the authors visited *briq*, whose hospitality is gratefully acknowledged. Philipp Strack was supported by a Sloan fellowship.

1 Introduction

Intuition and evidence suggests that many individuals are present-biased (e.g. Frederick et al., 2002; Augenblick et al., 2015; Augenblick and Rabin, 2019). Building on work by Laibson (1997) and others, O'Donoghue and Rabin (1999, 2001) illustrate within the quasi-hyperbolic discounting model that present bias, especially in combination with a lack of understanding thereof, leads individuals to procrastinate unpleasant tasks and to precrastinate pleasant experiences. Since excessive procrastination is a robust prediction of (naive) hyperbolic discounting models, it seems natural to use task-completion data to identify present bias from the pattern of completion times. In line with this idea, previous research uses completion near the deadline as an indication of present bias (Brown and Previtero, 2018; Frakes and Wasserman, 2016) or estimates the degree of time inconsistency from completion times under parametric assumptions (Martinez et al., 2017).[1]

This paper asks whether time preferences can be inferred by an outside observer—referred to as the analyst—when *only* task completion is observed *absent* parametric assumptions on the (unobservable) cost and benefit of task completion.[2] The key difficulty in doing so is to separate naivete or time-preference-based explanations of delay from those due to the option value of waiting (Wald, 1945; Weisbrod, 1964; Dixit and Pindyck, 1994): whenever the cost of doing a certain task is stochastic, an individual may wait in the hope of getting a lower cost draw tomorrow.[3] This challenge for the identification of time preferences is present even if the analyst imposes that the agent discounts payoffs exponentially.

Section 2 introduces our task-completion model. We consider an analyst who, from observing task completion times of a partially-naive quasi-hyperbolic discounter, tries to learn about some or all of the following parameters: the long-run discount factor δ, the present-bias parameter β, and the degree of sophistication $\hat{\beta}$. To facilitate learning by the

[1]Brown and Previtero (2018) classify individuals that select their health care plan close to the deadline as procrastinators and look for correlated behavior in other financial domains. Frakes and Wasserman (2016) investigate the behavior of patent officers that have to complete a given quota of applications supposing that the cost of working on a patent are deterministic and identical across days. In their model, for conventional discount rates the empirically observed bunching close to the deadline is inconsistent with exponential discounting. Slemrod et al. (1997) observe that random opportunity cost can explain a distribution over different stopping times, and Martinez et al. (2017) parametrically identify time preferences in such a model.

[2]De Oliveira and Lamba (2019) raise a closely related question in a different setting. See Section 8.

[3]To highlight this identification challenge, we abstract from forgetting. Conceptually, think of the agent as getting a non-intrusive reminder at the beginning of every period. This is not to say that limited memory and the strategic response to it are unimportant in determining task-completion behavior in the field. See, for example, Heffetz et al. (2016) for how reminders determine when parking fines are payed, Altmann et al. (2019) for how deadlines and reminders determine the probability of making a check-up appointment at the dentist, and Ericson (2017) for how present bias and limited memory interact.

analyst, we assume that the agent's task-completion payoffs are drawn each period from the same payoff distribution. Without this restriction, it is straightforward to rationalize any observed stopping behavior independently of the agent's taste for immediate gratification and degree of sophistication, leaving no hope for identification thereof.[4] Furthermore, to make identification easier, we suppose that the analyst observes the agent's exact stopping probability in each period. Again, this assumption strongly favors the analyst's ability to learn about underlying parameters. Finally, we impose that the agent can be described as a (partially) naive quasi-hyperbolic discounter. We are agnostic as to the nature of the task, so our analysis applies when task-completion leads to immediate benefits, costs, or both.

Section 3 contains motivating examples. The first highlights that, even when the parametric form of the underlying unobservable payoff distribution is known, stopping probabilities that increase close to the deadline are insufficient to distinguish a time-consistent from a present-biased agent. The second example illustrates how the estimated present bias can depend crucially on common parametric assumptions about the unobservable payoff distribution—such as that it is an extreme value distribution—, which are often imposed for analytical ease. This holds in our example even though the analyst knows (or guesses correctly) the long-run discount factor, as well as the mean and variance of the underlying stationary payoff distribution. Furthermore, the functional-form-driven time-preference estimates are not only incorrect but the squared error associated with some them is extremely low—suggesting that with finite noisy data it is difficult for the analyst to realize when she picks an incorrect functional form. Motivated by the importance of the parametric assumptions in the example as well as variants thereof, we turn to the main focus of the paper: what lessons about present bias and naivete thereof can be learned non-parametrically?

As a useful preliminary step, Section 4 establishes that the agent's perceived continuation value is characterized by a simple recursive equation. Section 5 establishes that for any quasi-hyperbolic discounter—independently of whether she is sophisticated or (partially) naive and of her degree of impatience—the perceived continuation value decreases over time. Because the agent in our model completes the task when the current benefit is greater than her perceived continuation value, Theorem 1 implies that a quasi-hyperbolic discounter becomes

[4]For example, suppose in every period the cost of doing the task is either one or zero, allowing for time-varying probability that the cost are zero. Simply setting the probability that the cost are zero in each period equal to that period's observed task completion probability rationalizes the data for any time-separable utility function. More generally, adopting well-known arguments from the dynamic discrete choice literature (e.g., Section 3.5 in Rust, 1994; Magnac and Thesmar, 2002), Section 8 shows that otherwise even parametric identification is infeasible.

more and more likely to complete the task the closer she is to the deadline.[5]

Section 6 establishes our main result: if the agent is either sophisticated or fully naive, for *any given* long-run discount factor δ and present-bias parameter β, any given penalty of not completing the task, and *any* weakly increasing profile of task completion, there exists a stationary payoff distribution that rationalizes the agent's behavior (Theorems 2 and 3, respectively). This implies that for *any* dataset the analyst may observe, absent parametric assumptions it is impossible for her to learn *anything* about the agent's long-run-discount factor, degree of present bias or level of sophistication. Importantly, this complete absence of partial identification continues to hold even if the analyst imposes a priori restrictions on permissible long-run discount factors, or imposes that the agent is time consistent. A very rough intuition for this fact is as follows: whether a self prefers to do a task today or tomorrow depends on her time preferences and on the perceived option value of waiting. The option value of waiting, in turn, depends on the payoff distribution. Through changing the unobservable payoff distribution, we can hence undo a change in the present bias or long-run discount factor of the agent.

In our proofs of Theorems 2 and 3, we freely construct an iid payoff distribution F. One may hope to identify present bias through economically meaningful restrictions on this distribution. Arguably, the most natural assumptions are those regarding the moments of the payoff distribution; for example, an analyst may have an idea regarding the possible expected payoff of doing the task—that is regarding the mean of F—or may be willing to impose that payoffs do not vary to much between periods—that is to restrict the variance of F. Example 2 in Section 3, however, already highlights that even fixing these moments, common parametric assumptions can lead to widely varying estimates of the agent's time preferences. To expand on this point, in Section 6.3 we establish that as long as the penalty is unobservable or the task is mandatory, we can find a payoff distribution with *any* given mean and non-zero variance that rationalizes the observed stopping behavior for a time-consistent agent with $\delta = 1$. Any identification of the present-bias parameter β in this case, thus, follows from restrictions on higher-order moments.

Section 7 asks whether non-parametric identification is feasible with richer data in which the analyst in each period observes the agent's willingness to pay for continuing with the stopping problem in addition to the stopping probability. In the case of tax-filing, for example, this amount to eliciting the willingness to pay for having someone else file one's taxes immediately with zero hassle.[6] For the case of a sophisticated agent whose contemporaneous

[5] This result relies on payoffs each period being drawn independently from the same distribution.

[6] As we explain carefully in Section 7, our procedure does not explicitly or implicitly rely on the agent

utility function is quasi-linear in money, Theorem 4 provides an analytical answer in closed form. Indeed, to check whether or not the data is consistent with a given pair of parameters β, δ, the analyst only needs to check a simple set of inequalities. Intuitively, observing the continuation values allows the analyst to distinguish between a taste for immediate gratification and option-value-of-waiting-based delays because a high option value requires the unobservable payoffs to differ significantly. As a consequence, as the deadline approaches and the agent foresees less future draws, the option value must decrease quickly. In contrast, an impatient agent's continuation value decreases at a slower rate. At the cost of relying on numerical techniques commonly used in applied work, our sharp set-identification result can be extended straightforwardly to cover partial naivete and non-linear utility in money.

Applying our Theorem 4 to Example 2, however, illustrates that the analyst may need to observe a large number of continuation values to be able to tightly identify the present-bias parameter β. In the example, there is no meaningful identification with 5 periods of data. With 20 periods, however, β can be tightly identified if the analyst—based on other data or theory—is willing to impose a particular value for δ. Given that we made a number of assumptions facilitating identification—such as that the exact stopping probabilities and continuation values are observable to the analyst—, our results suggests that absent parametric assumptions a substantial amount of additional data is needed to identify a taste for immediate gratification or the degree of sophistication.

Section 8 relates our finding to work on identification in dynamic discrete choice models (e.g., see Rust, 1994).[7] Most of this literature focuses on allowing for flexible functional forms that map actions and observable states into mean utilities, while imposing that the unobservable state affects utility through an additively separable shock drawn from a distribution with known functional form.[8] When shocks are drawn iid from a known functional form, the literature provides conditions under which time preferences can be identified. In contrast, we prove for our stopping model that the chosen parametric form of the unobservable shock fully determines the time-preference estimates.

Section 9 concludes by discussing broader implications of our analysis.

---

comparing monetary rewards at different points in time, so it is robust to standard critiques of eliciting time-preference via monetary rewards (Augenblick et al., 2015; Ericson and Laibson, 2019; Ramsey, 1928).

[7]Technically, our setup is not a special case of the usual dynamic discrete choice framework. Furthermore, the dynamic discrete choice literature overwhelmingly imposes time consistency ($\hat{\beta} = \beta = 1$), thereby precluding issues of present bias we are interested in. See Section 8 for details.

[8]Exceptions are Norets and Tang (2014) who consider partial identification of δ for a time-consistent agent in a dynamic binary discrete choice model that rules out stopping problems, and Christensen and Connault (2019) who consider robust estimation in a general framework. See Section 8.

2 Setup

Time-Preferences Let time $t = 1, 2, \cdots, T+1$ be discrete. We consider an agent with quasi-hyperbolic preferences who can choose when and whether to complete a task before some deadline T. The agents' utility is time-separable, and we denote a level of instantaneous utility the agent receives in period t by u_t; let

$$U^t = \mathbb{E}\left[u_t + \beta \sum_{s=t+1}^{T+1} \delta^{s-t} u_s\right] \tag{1}$$

denote the expected utility Self t derives from the random sequence $(u_t, \cdots, u_{T+1})$. Following O'Donoghue and Rabin (1999), we allow the agent to have incorrect beliefs regarding future selves' behavior. The agent believes that all future selves $r > t$ maximize

$$\hat{U}^r = \mathbb{E}\left[u_r + \hat{\beta} \sum_{s=r+1}^{T+1} \delta^{s-r} u_s\right]. \tag{2}$$

We allow for any vector of preference and belief parameters $(\delta, \beta, \hat{\beta}) \in (0, 1]^3$. In case $\hat{\beta} = \beta = 1$, the agent has time-consistent preferences with an exponential discount factor δ. In case $\beta < 1$, she has a taste for immediate gratification. We say she is sophisticated—i.e. perfectly predicts her future behavior—when $\hat{\beta} = \beta$, she is fully naive—i.e. believes that her future selves behave according to her current preference—if $\hat{\beta} = 1$, and otherwise say that she is partially naive. Our setup allows the agent to overestimate her own future taste for immediate gratification $\hat{\beta} < \beta$ as well as to underestimate it $\hat{\beta} > \beta$.[9]

Task-Completion Environment The agent can complete the task once during the periods $t = 1, \cdots, T$, so that T is the deadline before which the task needs to be completed. If the agent does not undertake the task in a given period $t \leq T$, we assume her instantaneous utility u_t equals zero.[10] If she completes the task, she gets an instantaneous utility of zero in period $T+1$. If she did not complete the task by the end of period T, the agents gets

[9]The special case of the model where the agent is fully sophisticated $\hat{\beta} = \beta$ and the time-horizon is infinite $T = \infty$ is studied in Section 4 of Fudenberg and Levine (2006).

[10]We thus focus on task-completion problems—such as paying a parking ticket or writing a referee report—in which current payoffs in periods prior to the deadline do not depend on whether the task has already been undertaken. This rules out problems in which completing the task subsequently generates a positive flow payoff, which we discuss in Section 9.

a known utility of $y_{T+1} \in \mathbb{R}_- \cup \{-\infty\}$ in period $T+1$.[11] Setting $y_{T+1} = -\infty$ captures a *mandatory* task that the agent is forced to complete by the end of period T; $y_{T+1} = 0$ captures an *optional* task; $y_{T+1} \in (-\infty, 0)$ captures a task with a finite penalty for not completing it. We denote by $\underline{y} = \beta\delta y_{T+1}$ the agent's continuation value in period T. Finally, we suppose that in every period $t < T+1$ the instantaneous utility of completing the task is drawn independently from a given payoff distribution F, which is known to the agent. Prior to taking a decision at time t, the agent observes y_t, but does not know the payoffs $y_{t+1}, \ldots, y_T$ from completing the task in later periods.

Strategies and Equilibrium We look for *perception-perfect equilibria* in which each self t chooses an optimal strategy given its prediction of future selves' behavior, and a self t's prediction of future selves' behavior is consistent with how future selves with preference parameter $\hat{\beta}$ would optimally behave (O'Donoghue and Rabin, 1999, 2001). Formally, a pure strategy σ_t for Self t maps a history of payoff realizations $y_1, \ldots, y_t$ to $\{0, 1\}$, with the interpretation that 1 means Self t completes the task. A perception-perfect equilibrium is a pair of strategies $(\sigma_1, \cdots, \sigma_T)$ and $(\hat{\sigma}_2, \cdots, \hat{\sigma}_T)$ such that for all $t \in \{1, \cdots, T\}$, σ_t maximizes U^t under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$, and for all $t \in \{2, \cdots, T\}$, the strategy $\hat{\sigma}_t$ maximizes $\hat{U}^t$ under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$. We restrict attention to perception-perfect equilibria in which all selves that are indifferent between completing the task and waiting choose to wait.

3 Examples on the Influence of Parametric Assumptions

To illustrate the difficulty of identifying present bias from an agent's stopping behavior, consider the following stylized example.

Example 1. A sophisticated agent receives a parking fine, which has to be paid within thirty days. In case she does not pay the fine, she incurs a known payoff $y_{T+1} = -5$. Furthermore, the agent's long-run (daily) discount rate is (well approximated by) $\delta = 1$. Figure 1 compares the stopping probability conditionally on not having stopped earlier of a time consistent agent who draws the cost of completing the task from a log-normal distribution whose underlying normal distribution has mean $\mu = 1$ and variance $\eta = 1$ (dark gray bars) to that of a sophisticated one with a present-bias parameter $\beta = 0.8$ who draws the cost from a log-normal distribution with parameters $\mu = 0.48, \eta = 1.52$ (light gray bars).

[11] For some results, we also discuss the case in which the agent gets a reward for not completing the task $y_{T+1} > 0$. When doing so, we explicity highlight that we allow for $y_{T+1} > 0$.

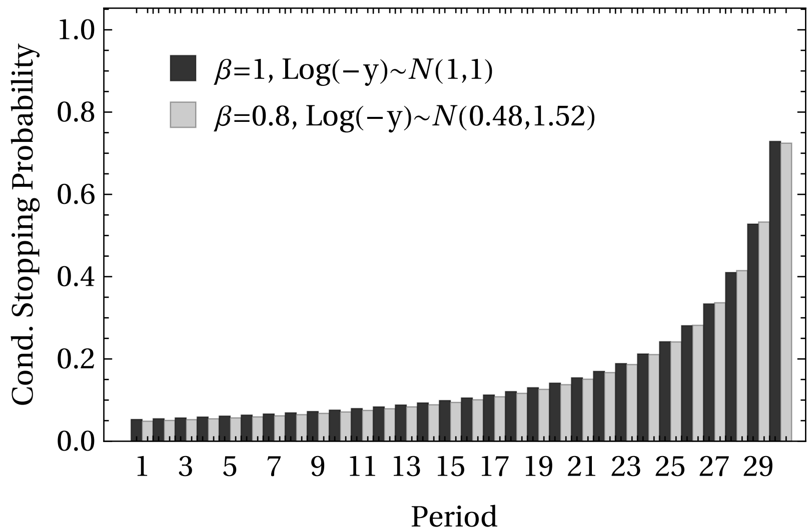

Figure 1: Conditional stopping probabilities in Example 1 for a time-consistent agent $\beta = 1$ (dark gray) and a sophisticated time-inconsistent agent with $\beta = 0.8$ (light gray).

An obvious first lesson from the example is that the likelihood of stopping close to the deadline is no reliable guide to identifying present bias: both agents probability of completing the task in the final period is just above 73%. Indeed, both agents observed stopping probabilities are remarkably similar and differ by less than 0.5% in any period, suggesting that in practise even an analyst who wants to test only between these two possible types may face a difficult problem.[12]

In the above illustrative example, the analyst knows or correctly guesses the parametric class of distributions (log-normal) from which the payoffs are drawn. Without knowing its exact parameters, however, it is hard to correctly identify the time-preference parameters in that example. In reality, however, payoffs are drawn from an unobservable payoff distribution and for typical field data—such as parking tickets—an analyst does *not* know the parametric form of the payoff distribution. The following example highlights how crucial functional form assumptions routinely imposed in applied papers can be in determining the analyst's findings.

[12]Independently of our work, Heffetz, O'Donoghue and Schneider observe that substantially different values of β can predict essentially equivalent day-to-day behavior in how people pay their parking tickets. In a simple model motivated by observed parking-ticket response behavior in New York City, which they analyze in Heffetz et al. (2016), they illustrate this supposing that the cost for paying the parking ticket is drawn from the small parametric family of distributions that has a mass point at zero and admit a constant density on an interval above zero. Their real-world application nicely demonstrates the practical importance of the identification challenge we illustrate in Example 1. We are very grateful to these authors for sharing their example with us during private communication.

Parametric Family	*Sq. Distance Minimzation*		*Likelihood Maximization*	
	β	Distance	β	Log-Likelihood
Uniform Naive	1.	0	1.	-1.59186
Uniform Sophisticate	1.	0	1.	-1.59186
Normal Naive	0.817	0.00231668	0.816	-1.59187
Normal Sophisticate	0.819	0.00267663	0.818	-1.59188
Extreme Value Naive	0.564	0.03968760	0.562	-1.59627
Extreme Value Sophisticate	0.570	0.04028880	0.570	-1.59638
Logistic Naive	0.757	0.00267137	0.756	-1.59188
Logistic Sophisticate	0.761	0.00331131	0.760	-1.59189
Laplace Naive	0.630	0.00806500	0.627	-1.59202
Laplace Sophisticate	0.638	0.00933172	0.634	-1.59207

Table 1: Parameter estimates of β and squared distance and log-likelihood.

Example 2. We suppose that the agent has 5 periods to complete the task and the payoffs from completing the task are drawn from a uniform distribution with mean 0 and standard deviation 1; in reality the agent is time-consistent with $\beta = \delta = 1$. The analyst can directly observe the exact stopping probabilities and, in addition, knows the true mean and standard deviation of the payoff distribution but not its exact functional form. Furthermore, suppose the analyst correctly imposes that $\delta = 1$ when analyzing the data. Let the analyst consider four standard parametric families of distributions: normal, extreme value, logistic, and Laplace. For each of these families, the analyst selects the parameter β that best fits—in the sense of squared distance or log-likelihood—the observed stopping probabilities allowing the agent to be either naive or sophisticated. Table 1 reports the parameter estimates for β and the squared distance/log-likelihood for the different parameterizations of the error distribution. The analyst's estimates of β range between $0.56 - 0.82$ even in this idealized situation in which she has infinite data, knows the mean and standard deviation of F, and knows the long-run discount factor δ. And if the analyst engaged in model testing selecting the model on the basis of minimizing squared distance or maximizing log-likelihood, she would conclude that the agent is naive time-inconsistent with $\beta \approx 0.816$ while in truth the agent is time-consistent and $\beta = 1$. Furthermore, for the normal distribution the squared difference in stopping probabilities in the sophisticated and naive case are remarkably small (less than 0.232%), so (in a finite dataset analogue) nothing would indicate to the analyst that these are bad distributional choices to model the unobservable shocks.[13]

The observed importance of functional form assumptions is robust to various changes of the example (see the Supplementary Appendix for details). As in the case with known mean and variance, the analyst also misestimates β to be substantially below 1 when she does not know (and thus estimates) the mean and standard deviation. Similarly, when the analyst uses 30 or 60 periods of data, she incorrectly estimates β to be substantially below 1. And in a variant of Example 2 in which the true distribution is logistic, the functional form assumption determines whether β is over- or underestimated. The analyst significantly underestimates β in a further variant of Example 2 in which the agent is truly present-biased and naive with $\beta = 0.9$. Furthermore, in this variant, we show that eventually as the number of periods the analyst observes increases, her estimates move further and further away from truth; contradicting the intuition that more data alleviates the identification problem.

These examples illustrate that without knowledge of the functional form of the unobservable payoff distribution, the analyst can end up significantly misestimating β even in a best case scenario where she has access to arbitrarily much data and knows the other parameters of the model (mean and variance). Our general results below, which establish that non-parametrically the degree of time inconsistency is never identified from task completion data, prove that the above examples are not artefacts of the numbers we have chosen. For every set of model parameters $\delta, \beta, \hat{\beta}$ and any given dataset, there exists some unobserved stationary payoff distribution that perfectly fits the data. As a consequence, the analyst's conclusions are—in line with Example 2—solely determined by her parametric choice for the unobservable payoff distribution.

4 Preliminary Analysis: Recursive Structure

This section establishes that the agent's problem has a simple recursive structure. We say a strategy σ is a *cutoff strategy* if there exists cutoffs $z = (z_1, \ldots, z_T)$ such that the agent stops in period t if and only if $y_t > z_t$. It follows from backward induction that the agent uses a cutoff strategy: Self T completes the task if and only if her realized payoff is strictly greater than $\underline{y}$. Furthermore, selves $t < T$ believe that Self T will complete the task if and only if her realized payoff is strictly greater than $({\hat{\beta}}/{\beta})\underline{y}$. Hence both the perceived and actual

[13]If the analyst needs to estimate the mean and standard deviation of the shock distribution as well, she is able to fit the data even better, making it even harder to detect her misspecification.

strategy in the final period are cutoff strategies. Similarly, if all future selves are perceived to use cutoff strategies, Self t can calculate the perceived continuation value of waiting v_t, and will complete the task if and only if her current payoff is greater than this perceived continuation value. Hence, by induction, all selves use a cutoff strategy and perceive their future selves to use cutoff strategies.

Using the fact that any two selves t and t' believe that future selves $s > \max\{t, t'\}$ use the same perceived cutoff $\hat{v}_s$, in the Appendix we derive a simple recursive equation that characterizes the perceived continuation value (or actual cutoff) Self t uses as well as the perceived cutoff that earlier selves believe Self t uses.

Lemma 1 (Recursive Characterization)**.** *A pair of strategies $(\sigma, \hat{\sigma})$ constitute a perception-perfect equilibrium if and only if both are cutoff strategies with cutoffs $(v, \hat{v}) \in \mathbb{R}^T \times \mathbb{R}^T$ that satisfy the equations $\hat{v}_t = \left({}^{\hat{\beta}}/_{\beta}\right)\ v_t$ and*

$$v_t = \begin{cases} \beta\,\delta \int_{\hat{\beta}/\beta\ v_{t+1}}^{\infty} z\, dF(z) + F(\hat{\beta}/\beta\ v_{t+1})\,\delta\, v_{t+1} & \text{for } t < T \\ \underline{y} & \text{for } t = T \end{cases}. \tag{3}$$

To see the intuition behind Equation 3, suppose first that the agent is sophisticated in which case ${}^{\hat{\beta}}/_{\beta} = 1$. Then the first term is the expected discounted payoff from stopping tomorrow times the probability of doing so, which the agent does whenever the payoff from stopping falls above the continuation value that tomorrow's self uses. This payoff is discounted according to Self t's short-term discount factor $\beta\delta$. The second term captures the fact that with probability $F(v_{t+1})$ tomorrow's self continues because it prefers its perceived continuation value v_{t+1}. As today's self discounts payoffs that realize after period $t+1$ by a factor of δ more than tomorrow's self, this term is discounted with δ. When predicting future behavior, a partially naive agent uses the perceived cutoffs $\hat{v}_t = ({}^{\hat{\beta}}/_{\beta})\, v_t$ determined by the continuation value a former time $s < t$ self believes Self t has. If $\hat{\beta} > \beta$ (respectively $\hat{\beta} < \beta$) current selves overestimate (respectively underestimate) future selves' patience and, hence, the cutoff they use.

For a given distribution F, Lemma 1 allows for a straightforward calculation of stopping probabilities, which we used to solve the examples in Section 3. Indeed, when imposing a given parametric form, Lemma 1 allows to analytically solve for the log-likelihood of different parameters given observed stopping probabilities, which is how we computed the log-likelihood (and similarly minimum-distance) estimators in Section 3. This is noteworthy as prior empirical approaches based on dynamic discrete choice models—which simulate a naive agent's beliefs—are computationally costly and restricted to unobservable payoffs that are drawn from a logistic distribution (derived from two unobservable extreme-value-type-1 shocks).[14] Furthermore, for a given vector of parameters, Lemma 1 can be used to numerically check parametric identification, which for general distributions is an open question. An important exception is the case of logistic shocks with known variance for which Martinez et al. (2017, Section A.4.1.) and Daljord et al. (2019) provide parametric identification results.

5 Rate of Task Completion Increases Over Time

This section establishes that a partially-naive quasi-hyperbolic agent is (weakly) more likely to stop and complete the task, the closer she is to the deadline T. Hence, any dataset consistent with quasi-hyperbolic discounting and independent draws from a given payoff distribution has a weakly increasing conditional stopping probability.

Theorem 1 (Monotonicity of the Continuation Value)**.** *The perceived continuation values are non-increasing over time*

$$v_1 \geq v_2 \geq \ldots \geq v_T \,.$$

To understand intuitively why the perceived continuation value decreases, consider first the case in which doing the task is always costly—i.e., where the support of F is a subset of $\mathbb{R}_-$, and compare the penultimate Self $T-1$ to the ultimate Self T. The ultimate self's perceived continuation value is the discounted penalty. If the penultimate self decides to wait, there are two possible cases; either the ultimate self waits or she completes the task. In case the ultimate self waits, the penultimate self discounts the penalty by an extra δ making her (weakly) better off. In case the ultimate self completes the task, she strictly prefers it to waiting, implying that the realized costs are strictly lower than the penalty. But in this case the penultimate self is strictly better of as she, in addition, discounts these lower costs. Replacing the penalty in the above argument with the perceived continuation value of a future self, extends the argument to earlier selves.

Suppose now instead that the agent is sophisticated and that the payoff of completing the task is always positive—i.e., the support of F is a subset of $\mathbb{R}_+$. Because there is a

[14]Indeed, upon circulation of a first version of this paper, Martinez et al. (2017) developed an estimation technique based on our procedure for their tax-filing data in order to provide robustness tests under different parametric assumptions that complement their core development, which hews closely to the type-1 extreme value tradition that gives rise to logistic shocks in our setting.

penalty if the agent does not complete the task in period T, Self T never waits. Self $T-1$ waits whenever her benefit of doing the task is so low that she prefers the positive discounted expected benefit from completing it tomorrow. Thus, she has a (weakly) higher perceived continuation value. If Self $T-1$ always completes the task immediately, Self $T-2$ and Self $T-1$'s benefit from waiting coincide. But whenever the drawn benefit from doing the task in period $T-1$ is low enough for Self $T-1$ to prefer to wait, Self $T-2$ who is (at least weakly) more patient regarding tradeoffs between period $T-1$ and T, must prefer Self $T-1$ to wait rather than to complete the task immediately. As a result, Self $T-2$ benefits from the fact that Self $T-1$ sometimes waits and hence has a higher perceived continuation value than Self $T-1$. Iterating this argument, the perceived continuation values are weakly increasing in the distance to the deadline. Partially naive agents think they are sophisticated, and since the perceived continuation value falls for a sophisticated agent they predict tomorrow's self to have a lower continuation value—and hence to wait less often—the closer they are to the deadline. At the same time, because a partially naive self predicts that tomorrow's self is (weakly) less patient than she is, she realizes that she benefits whenever tomorrow's self chooses to wait. Hence, she predicts a lower benefit from waiting, i.e. her perceived continuation value decreases over time. Formally, our proof studies properties of solutions to the recursive equation (3) to extend the above intuitions to cases in which the support of the payoff distribution may contain positive and negative elements.

Note that the probability p_t that the agent stops in period t conditional on not having stopped before equals the probability that the value of completing the task y_t is above the perceived continuation value v_t; i.e.

$$p_t = \mathbb{P}[y_t \geq v_t] = 1 - F(v_t)\,.$$

As the perceived continuation value v_t is non-increasing, we have that the objective probability p_t that the agent stops in period t is non-decreasing.

Corollary 1. *For any payoff distribution F and in every perception-perfect equilibrium, the probability with which the agent completes the task conditional on not having completed it before is non-decreasing towards the deadline, i.e.*

$$p_1 \leq p_2 \leq \ldots \leq p_T\,.$$

Importantly, Corollary 1 establishes that the probabilities of stopping *conditional* on not having stopped previously increase over time. The unconditional stopping probability, however, may either increase or decrease.[15]

Remark 1. Because Theorem 1 implies that the perceived continuation value decreases as the deadline approaches, each self prefers the deadline to be as far away from today as possible. Thus, perhaps surprisingly, the agent never wants to commit to a shorter deadline in an iid environment even if she is sophisticated and present biased. This is in contrast to examples from non-iid environments in which a present-biased quasi-hyperbolic discounter wants to commit to an earlier deadline to manage her conflict with future selves.[16]

Remark 2. If an analyst observes data from a heterogenous population, one can generate non-monotone aggregate conditional stopping probabilities.[17] Our later non-identification results (Theorem 2 and 3) extend to the case of a heterogeneous population as they establish non-identification for a better informed analyst.

6 Time Preferences are Unidentifiable from Task Completion

We showed that for any arbitrary preference profile β, δ and any belief $\hat{\beta}$, the conditional stopping probabilities p are non-decreasing. This section establishes the converse: absent (parametric) restrictions on the payoff distribution F, any non-decreasing profile of stopping probabilities is consistent with any arbitrary preference profile β, δ in case the agent is either sophisticated or fully naive. Hence, it is impossible, for example, to distinguish a naive time-inconsistent agent from a time-consistent one based on their task-completion behavior. Importantly, this impossibility continues to hold even if the analyst exogenously imposes that the long-run discount factor $\delta = 1$, as is plausible if she observes task completion on a frequent (e.g. daily) basis. Similarly, even if the researcher is willing to impose a priori restrictions on plausible levels of β, absent exogenous restrictions on F, no information on δ or β can be inferred from the task-completion data.

Intuitively, whether a self prefers to do a task today or tomorrow depends on her time preferences (as well as beliefs about future selves' time preferences) and on the perceived

[15] This difference is of practical relevance: for example, Figure 1 and 2 in Martinez et al. (2017) imply that the conditional stopping probabilities increase over time in their tax-filing data while the unconditional stopping probabilities decrease.

[16] See our working paper version or Bisin and Hyndman (2018) for examples.

[17] Suppose, for example, there are two types of agents in the population that face a three-period task-completion problem. The first type stops in each period with probability 1, while the second type stops only in the final period. If $\alpha > 0$ is the fraction of the first type, then the aggregate conditional stopping probability is α in the first period, 0 in the second, and 1 in the final period. Heffetz et al. (2016) discuss heterogeneity and provide evidence for its importance in determining when individuals pay their parking fines.

option value of waiting. The option value of waiting, in turn, depends on the payoff distribution. Through changing the unobservable payoff distribution, we can hence potentially undo a change in the agent's present bias or long-run-discount factor. Technically, however, a local change in the payoff distribution affects continuation values in every period in a highly non-linear way, so to establish that we can construct an appropriate payoff distribution, we need a non-local argument. When the agent is either sophisticated or fully naive—for different technical reasons—the analysis simplifies to a tractable fixed-point problem, which allows us to establish that we can rationalize the stopping behavior for any arbitrarily chosen β, δ.

If the penalty is unobservable, we furthermore illustrate that the data is rationalizable as the optimal behavior of a fully patient time-consistent agent ($\hat{\beta} = \beta = \delta = 1$) facing an unobservable payoff distribution F with *any given* expected value and (non-zero) variance of the distribution; any parametric identification of present bias in such a task-completion setting, therefore, is based on higher-order moments of the payoff distribution.

6.1 Time Preferences are Unidentifiable: Sophisticated Case

In this subsection, we establish that absent (parametric) restrictions on the payoff distribution F, any non-decreasing profile of stopping probabilities is consistent with any arbitrary preference profile β, δ of a sophisticated quasi-hyperbolic discounter. In particular, we have:

Theorem 2 (Non-identifiability)**.** *Suppose the agent is sophisticated $\hat{\beta} = \beta$. For every non-decreasing sequence of stopping probabilities $0 < p_1 \leq p_2 \leq \ldots \leq p_T < 1$, every $(\delta, \beta) \in (0, 1] \times (0, 1]$, and every penalty $y_{T+1} \in \mathbb{R}$, there exists a distribution F that rationalizes the agent's stopping probabilities as the outcome of the unique perception perfect equilibrium.*

Technically, to prove the theorem, we construct a distribution with $t + 2$ mass points, where each of the non-extreme values equals the perceived continuation value of a self $t \in \{1, \ldots, T\}$; i.e. the second lowest mass point is set at the value $v_T = \underline{y}$, and so on. The probability on each mass point is chosen so that the agent—who waits if and only if $y_t \leq v_t$—selects the exogenously given stopping probability. The constructions is feasible since when $\hat{\beta} = \beta$, the recursive representation (Lemma 3) takes a particular simple form, and together with the chosen construction of the distribution gives rise to a system of linear equations, which can be solved.

Remark 3. Theorem 2 extends to mandatory tasks for stopping probabilities $0 < p_1 \leq p_2 \leq \ldots \leq p_{T-1} < p_T = 1$ by simply considering the auxiliary tasks with $T - 1$ periods for which completion in period $T - 1$ is not mandatory.

6.2 Time Preferences are Unidentifiable: Naive Case

We now consider a fully naive agent and establish that for every non-decreasing profile of stopping probabilities and every preference profile β, δ, there exists a payoff distribution F that admits a piecewise constant density and induces the agent to choose the stopping behavior given by the data.

Theorem 3 (Non-identifiability)**.** *Suppose the agent is fully naive $\hat{\beta} = 1$. For every non-decreasing sequence of stopping probabilities $0 < p_1 \leq p_2 \leq \ldots \leq p_T < 1$, every $(\delta, \beta) \in (0,1) \times (0,1]$, and every penalty $y_{T+1} < 0$, there exists a distribution F that rationalizes the agent's stopping probabilities as the outcome of the unique perception perfect equilibrium.*

Our formal proof in the Appendix proceeds roughly as follows. Step (i). Fix the agent's time preference as well as period T's perceived continuation value (which equals $\underline{y}$). Step (ii). Take an arbitrary $(T-1)$-element vector of non-increasing perceived continuation values $v_1 \geq v_2 \geq \ldots \geq v_{T-1}$. Step (iii). Here, we generate a payoff distribution for these perceived continuation values that induces the desired stopping probabilities. To do so, we put a probability mass equal to the difference in the exogenously given stopping probability between periods t and $t+1$ between the corresponding periods' perceived continuation values, for simplicity using a uniform density. This step, hence, amounts to mapping perceived continuation values into distributions that lead to the correct stopping probabilities. Step (iv). Calculate the actual perceived continuation values that the new payoff distribution from the third step gives rise to. This maps the set of distributions back into the vector of perceived continuation values. By Theorem 1, these perceived continuation values are again non-decreasing, and thus the combined function maps a non-increasing sequence of perceived continuation values into a non-increasing sequence of perceived continuation values. Step (v). We show that this function is bounded and maps sequences from an appropriately chosen interval into itself. Furthermore, the function is monotone as higher perceived continuation values lead to a better distribution (in the sense of first-order stochastic dominance) and a better distribution increases the perceived continuation values for an agent who believes to be time-consistent. Thus, the mapping from perceived continuation values into perceived continuation values is a monotone mapping from a complete lattice into a complete lattice, and by Tarski's Theorem admits at least one fixed point. Any fixed point gives the desired distribution, since by Step (iii) the stopping probabilities are correct and by Step (iv) the perceived continuation values are those consistent with the limit distribution.

Step (v) in the proof relies on the fact that a first-order dominance shift in the payoff distribution increases the perceived continuation value, which is where we use the assumption of a fully naive agent.[18]

6.3 Known Expected Value and Variance

For our general results, we have not restricted the class of permissible distribution functions. One may hope to rule out time consistency and find evidence through restricting features of the distribution. We first note that Theorems 2 and 3 extend to the case in which we only consider well-behaved (smooth) distributions:

Remark 4. The distributions we construct to establish Theorems 2 and 3 have bounded support and thus admit all moments. Furthermore, it follows from the construction in the respective proofs that the non-identifiability result extends if we restrict attention to distributions that admit a density.

In light of this remark, the perhaps most natural way of even further narrowing down the class of possible distribution is to add restrictions regarding the moments of F; an analyst may have an idea regarding the possible expected payoff of doing the task—that is regarding the mean of F—or may be willing to impose that payoffs do not vary to much between periods—restricting the variance of F.

If the continuation payoff $y_{T+1} \in \mathbb{R}$ is unobservable (or the task mandatory), however, even with a priori knowledge of the mean and variance of F it is impossible to rule-out time-consistent behavior. To see this, consider an agent for whom $\beta = \delta = 1$. Theorem 2 implies that a payoff distribution F exists that rationalizes any increasing profile of stopping probabilities. Furthermore, in this case the recursive formulation in Lemma 1 simplifies to

$$v_t = \mathbb{E}\left[\max\{y_{t+1}, v_{t+1}\}\right] \qquad \text{for all } t < T.$$

Hence, if the distribution F together with the penalty y_{T+1} rationalize the data, so does the distribution $F + \kappa_1$ together with the penalty $y_{T+1} + \kappa_1$ for any $\kappa_1 \in \mathbb{R}$. In other words, we can always select a payoff distribution with a given expected value. Furthermore for any $\kappa_2 > 0$, the stopping behavior remains optimal if we scale the payoffs and y_{T+1} by κ_2. This implies that we can not only select a distribution with a given mean but that we can at the same time select any desired variance and explain the observed stopping behavior.[19]

[18] See the working paper version for a counterexample in which a first-order dominance shift in the payoff distribution decreases the perceived continuation value when the agent is not fully naive.

[19] Indeed, since the construction of F in the proof of Theorem 2 uses bounded support, we can rationalize

Corollary 2. *Suppose the agent is time-consistent and fully patient $\hat{\beta} = \beta = \delta = 1$. For every non-decreasing sequence of stopping probabilities $0 < p_1 \leq p_2 \leq \ldots \leq p_T < 1$, and every $\mu \in \mathbb{R}$ and $\sigma^2 > 0$, there exists a distribution F with mean μ and variance σ^2 and a penalty $y_{T+1} \in \mathbb{R}$ that rationalizes the agent's stopping probabilities as the outcome of the unique perception perfect equilibrium.*

Corollary 2 does not preclude that knowledge regarding higher moments could rule out $\hat{\beta} = \beta = \delta = 1$ for *some* datasets. Knowledge regarding the skewness of the payoff distribution, such as that it is symmetric, might for example facilitate set identification. Given the importance of such an assumption, however, it should be carefully justified.

7 Non-Parametric Identification with Richer Data

Since stopping data is insufficient to identify time preferences, we now asks whether richer data helps. Intuitively, the analyst needs to disentangle whether the stopping behavior is driven by a desire to delay incurring costs or the option value of waiting. In the latter case, a considerable option value requires payoffs to differ significantly. Hence, as the deadline approaches and a waiting agent faces fewer future draws, the continuation value should drop considerably. In contrast, even with a (relatively) constant option value, an agent who is present biased is willing to delay a costly activity to the last minute. Thus, observing, in addition to task-completion times, continuation values should facilitate the non-parametric identification of the time-preference parameters $\delta, \beta, \hat{\beta}$.

This section, hence, analyzes the case in which the agent's stopping behavior as well as her exact willingness to pay for continuing with the task are observable. Conceptually, the analyst could elicit this information by selecting some stopping problems in which she offers the agent a mechanism at the end of period t that truthfully elicits her willingness to pay for continuing with the task from $t + 1$ onwards.[20] Denote the amount she is willing to pay at the end of period t by m_t. If the agent's utility is quasi-linear in money, which is a good approximation for a quasi-hyperbolic discounter when the involved stakes are small, observing m_t is equivalent to observing the continuation value v_t; otherwise, $v_t = u(m_t)$ for some monotonically increasing utility function m_t. We provide an analytical result regarding

the observed stopping behavior as resulting from a patient agent ($\beta = \delta = 1$) whose net benefits vary arbitrarily little.

[20] If the analyst sees infinitely many identical agents, she can randomly select T agents. Label these agents $k = 1, \ldots, T$. At the end of period k, the analyst then elicits agent k's willingness to pay for facing the task-completion problem from period $k + 1$ to T.

sharp partial identification for the case of linear utility in money and a sophisticated agent. But—at the cost of using numerical methods common in empirical work—our results can be readily extended in multiple directions, including partial naivete and non-linear utility in money.

Theorem 1 and Corollary 1 imply that the elicited continuation values must be non-increasing and the observed stopping probabilities non-decreasing. We refer to data (v, p) that has these properties as *plausible*.[21] Any data that is not plausible cannot be justified by our quasi-hyperbolic setup. Imposing that the agent is sophisticated, we now non-parametrically identify the set of β, δ consistent with the observed data. Using Lemma 1 and the fact that an agent stops whenever his payoff is strictly above the perceived continuation value, for a sophisticate the perceived continuation values v and conditional stopping probabilities p must satisfy

$$\begin{aligned} v_t &= u(m_t) && \text{for all } t \in \{1, \ldots, T\}, \\ \int_{v_{t+1}}^{\infty} z \, dF(z) &= \frac{\delta^{-1} v_t - (1 - p_{t+1}) v_{t+1}}{\beta} && \text{for all } t \in \{1, \ldots, T-1\}, \\ 1 - F(v_t) &= p_t && \text{for all } t \in \{1, \ldots, T\}. \end{aligned} \tag{4}$$

Conversely, if (u, F) satisfies (4) for a given plausible dataset, Lemma 1 implies that (v, p) are the perceived continuation values and stopping probabilities of a sophisticated agent.

Note that for $u(m_t) = m_t$ the right-hand-side of (4) is given by the data and hypothesized values of β and δ. Thus, the data is consistent with a given pair β, δ if and only if there exists a distribution F that solves (4). Lemma 4 in the Appendix shows that whenever (4) admits a solution, it also admits a solution that is a distribution consisting of $T + 1$ mass points. Using this insight, (4) becomes a non-linear system with finitely many real-valued unknowns. Theorem 4 exploits that this system can be transformed into a simple set of inequalities that identify the values of δ and β that are consistent with the observed stopping behavior and elicited continuation values.

Theorem 4 (Non-Parametric Identification). *Suppose $u(m_t) = m_t$ for all t and that $p_1 > 0$.*[22] *Plausible data (v, p) is consistent with β, δ and sophistication $\hat{\beta} = \beta$ if and only if (i)*

$$\beta < \frac{\delta^{-1} v_1 - (1 - p_2) v_2}{v_2(p_2 - p_1) + v_1 p_1}$$

[21] If y_{t+1} is observable, we require that $v_T = \beta\delta y_{t+1}$.

[22] We require $p_1 > 0$ only to simplify the statement.

and (ii) $v_{t+1}\beta < v_{t+1}a(\delta, t) \leq v_t\beta$ *for all* $t \in \{2, \ldots, T-1\}$*, where*

$$a(\delta, t) = 1 - \frac{\delta^{-1}(v_{t-1} - v_t) - (1 - p_t)(v_t - v_{t+1})}{v_{t+1}(p_{t+1} - p_t)}.$$

The theorem provides a sharp characterization of what time-preference parameters are consistent with the observed rich data. To illustrate its implications, reconsider Example 2 from Section 3 in which $T = 5$, the agent's payoff of completing the task are uniformly distributed with mean zero and variance 1, and $\beta = \delta = 1$. Figure 2 shows the set of parameters the analyst can identify non-parametrically for $T = 5$ and $T = 20$. When continuation values are observable, not all parameter combinations β, δ are consistent with the data. Figure 2, however, also illustrates that even if the analyst correctly imposes that $\delta = 1$, she cannot make precise inference when $T = 5$. Indeed, in the example any β between 0.82 and 1.28 is consistent with the data. This changes drastically for $T = 20$ in which case β is tightly identified once $\delta = 1$ is imposed. Without imposing $\delta = 1$, however, the inference about β remains imprecise even in the case of $T = 20$, as it is impossible to reject $\beta = 0.84$. Similarly, if the analyst imposes that the agent is time consistent, she can identify δ tightly for the case of $T = 20$ while absent this restriction a wide range of δ is consistent with the observed stopping behavior. Overall, the example suggests that rich data—including a significant number of continuation values—is needed for tight parameter estimates. It is also in line with conventional wisdom in behavioral economics that knowledge from theory or other datasets that pins down δ is useful in identifying present bias.

As the intuition at the beginning of this section highlights, it is the change in option value that allows identifying time-preference parameters. Since the willingness to pay for extending the deadline reflects the drop in continuation value, observing it would also give rise to a rich dataset that allows for non-parametric set identification.

Generalizations of the Methodology We think of Theorem 4 as a proof of concept that analysts can adopt to the data at hand and the assumption they are willing to make. For example, it is straightforward to incorporate partial naivete. One just needs to account not only for the probability mass and expectation of falling between two actual continuation values but also to differentiate whether a given probability mass falls above or below the anticipated continuation values $\hat{v}$. An analog to Lemma 4 implies that this can be done with $2T+1$ mass points. For intervals that are bounded by anticipated and not actual continuation values, however, the probability that y_t falls into this interval is unknown. Thus, the analyst

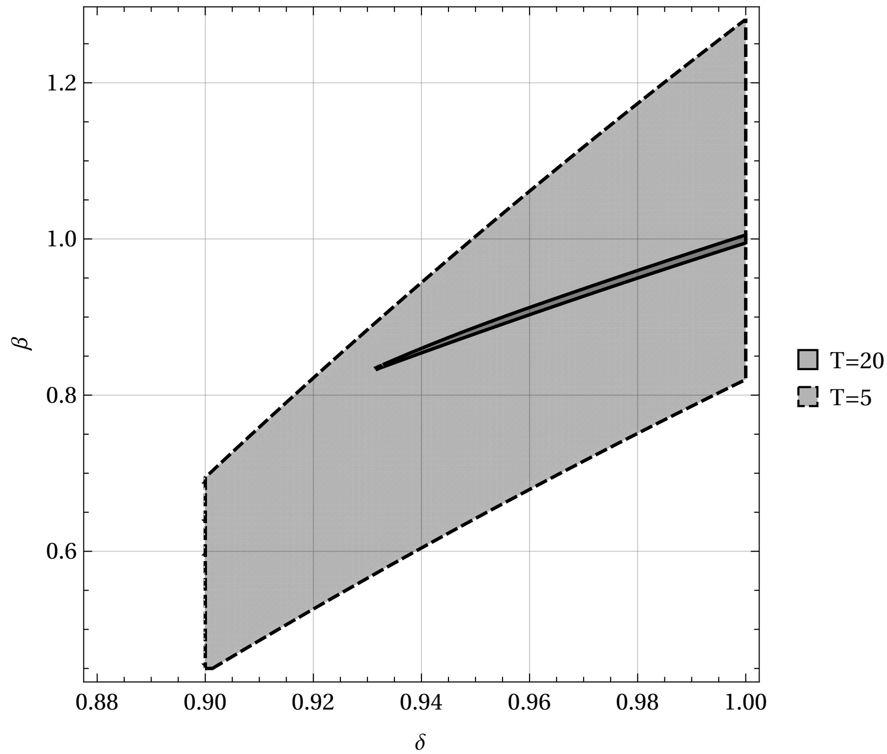

Figure 2: The above figures illustrates the set of parameters β, δ that the analyst can non-parametrically identify if she correctly imposes that the agent's instantaneous utility is linear in money. The agent's true values of completing the task are uniformly drawn with mean 0 and variance 1 and she is time-consistent with $\beta = \delta = 1$. Lightly shaded is the case of $T = 5$ periods of data and in dark shade $T = 20$ periods of data.

needs to choose both the mass point and the weight on it (with the appropriate constraints from the observed stopping behavior), giving rise to quadratic constraints. While this can be solved numerically using standard techniques, a closed-form solution as in the case of Theorem 4 is unavailable. Similarly, because we only need to consider a finite number of mass points, allowing for non-linear (strictly increasing) utility in money simply requires the analyst to choose increasing utility values $u(m_t)$ in addition to the mass points.[23]

Time-Preferences over Money Our procedure does not (explicitly or implicitly) impose constraints on how the agent handles monetary payments at different points in time. It is sufficient for contemporaneous utility to be separable in money, and the marginal utility

[23]If the analyst wants to impose risk-aversion in money, this adds simple (linear) constraints that ensure that the slope of u is non-increasing in m_t. Again, this can be solved using standard numerical techniques.

of receiving money to be the same across periods. This assumption is consistent with an intertemporal set-up in which the agent can borrow and lend at given interest rates—in which case the interest rate determines how she trades off monetary payments at different points in time (Augenblick et al., 2015; Ericson and Laibson, 2019; Ramsey, 1928). But it is also consistent with an agent narrow bracketing and consuming small monetary payments immediately—or reasoning as if she does so—so long as she trades of money and effort consistently over time.

8 Related Dynamic Choice Literature

There is a vast literature on identification in dynamic discrete choice models (for a survey see Rust, 1994). These papers typically focus on allowing flexible (non-parametric) functional forms for the mapping from actions and the observable state to the agent's mean utility while for analytical ease imposing that the unobservable state is an additively separable utility shock from some known distribution. Because, whenever the agent takes a decision, the action space is binary and time is excluded from affecting mean utilities, in our model there is no difference between a parametric or non-parametric mapping of the agent's observable state and action into instantaneous payoffs. But, perhaps surprisingly, in our setting the chosen parametric form on distributions of the unobservable shock fully determines time-preference estimates.

While theoretically stopping problems are a special case of a dynamic discrete choice problem, the assumptions imposed in the (applied) econometric literature on the topic often rule out our setting.[24] In particular, the usual paper assumes: i) infinite horizon; ii) additive separability between the observable state-dependent part of utility and the unobservable component; iii) the set of feasible actions does not depend on past actions; iv) shocks are drawn from some given distribution with unbounded support; and v) the agent is time-consistent. Assumptions ii), iii), and iv) jointly rule out that the modeled dynamic discrete choice problem is equivalent to an optimal stopping problem.[25] Of course, this does not preclude insights from the dynamic discrete choice literature to carry over to our setting.

Classic parametric non-identification results in the dynamic discrete choice literature (e.g., Section 3.5 in Rust, 1994; Magnac and Thesmar, 2002) consider the case in which

[24] See, for example, Rust (1994) or Magnac and Thesmar (2002) for a statement of these assumptions.

[25] To see why, suppose the agent stopped. Then it must be optimal for her not to do the task again even though—in contrast to an optimal stopping problem—the action is available by iii). Yet, by assumption ii) and iv), there is a realization of the shock for which it is optimal to do the task again.

there is state-dependent unobserved utility and show that it possible to rationalize any data with any time-preference parameter for any distribution of additive shocks using some state-dependent utility. As the state in our model equals calendar time, this is equivalent to the observation that if one allows for non-iid payoffs any time-preference $(\delta, \beta, \hat{\beta})$ can be rationalized — even when the analyst knows the distribution of shocks.[26]

To circumvent this identification problem, as Levy and Schiraldi (2020) point out,[27] "[a] typical approach to identification in the exponential discounting model adds exclusion restrictions on utility (or conditional value function) across states, a choice restriction such as an absorbing choice, (e.g. Magnac and Thesmar, 2002; Abbring and Daljord, 2019), or restricts attention to finite horizon model (e.g. Yao et al., 2012; Chung et al., 2014; Bajari et al., 2016; Chou, 2016), usually coupled with a strong normalization on the utility of a reference alternative."

The task-completion setting with iid payoffs imposes *all* of the above restrictions simultaneously. Indeed, focusing on present bias, Martinez et al. (2017) prove that β is identified when the agent is naive, the analyst knows δ, and shocks are logistic with know variance. For sophisticated agents and logistic shocks, Daljord et al. (2019) provide identification conditions. When we relax the assumption that the analyst knows the parametric form of the unobservable error distribution, in contrast, identification of time preference becomes infeasible—even when restricting attention to a time-consistent agent as (most of) the dynamic discrete choice literature does or focusing on the present-biased and fully naive case with known δ.

Closely related to Theorem 4 of our paper, Norets and Tang (2014) investigate identification of time preferences when the stationary error distribution is unknown. Under Assumptions i) to v) above, they provide a system of equations for such dynamic discrete binary choice models with time-consistent agents that allows one to check (numerically) whether a distribution F exists that rationalizes the data for a given parameter δ.[28] Recently, for a

[26]To see this, suppose $F_t = \mu_t + G$ for some given distribution G that is invertible and known to the analyst. To match the data, simply proceed recursively and choose μ_T for any given v_T so that the observed stopping probability in the final period T matches the data, i.e. so that $G\left(v_T \hat{\beta}/\beta - \mu_T\right) = 1 - p_T$. Once μ_T is determined, use the analog to the Bellman equation in Lemma 1 to calculate v_{T-1} and then choose μ_{T-1} to match the observed stopping probability in the penultimate period (i.e. so that $G\left(v_{T-1}\hat{\beta}/\beta - \mu_{T-1}\right) = 1 - p_{T-1}$); and so forth.

[27]For a sophistcated β, δ agent, Levy and Schiraldi investigate parametric identification of the time preferences when the additive shock to utility follows an extreme-value distribution. Their identification result, however, rules out our setup as it assumes at least four actions. Relatedly, Mahajan et al. (2020) exploit exclusion restrictions that do not apply in our setting to establish parametric identification of present bias in their dynamic discrete choice problem.

[28]Their setup differs from ours since they do not allow for present-bias, exclude optimal stopping problems,

general set of problems Christensen and Connault (2019) develop a novel, generally applicable, "robust-estimation" technique that also allows one to numerically check the importance of parametric assumptions for a given dataset.[29] Economically related, we highlight the importance of parametric assumption by proving that absent these, in task-completion environments identification of time preferences is impossible *for any* dataset.

Imposing time consistency, De Oliveira and Lamba (2019) ask what an analyst can infer about δ when she observes an agent who chooses actions over time.[30] The analyst does not know the agent's information structure but knows payoffs conditional on the state and the chosen action profile. They characterize rationalizable δ's for a *single* sequence of actions in a dynamic decision problem. In contrast, we allow the analyst to observe the *distribution* over action profiles and consider iid payoffs absent which identification of time-preference is impossible in task-completion environments (see Footnote 4).

9 Discussion

Our formal non-identification results apply to the specific task-completion setting analyzed, and should not be misconstrued as implying complete non-identifiability of the quasi-hyperbolic discounting model in other settings. In richer and different datasets, it is possible to identify $\beta, \hat{\beta}$ more directly. For example, lotteries (or contracts) that pay off differently depending on the agent's own future behavior can be used to reveal whether the agent missperceives her own future behavior and, hence, whether she is (partially) naive (see, for example, DellaVigna and Malmendier, 2006; Spiegler, 2011). Similarly, if the agent is willing to pay for reducing her choice set or for imposing a fine for certain future actions, she values commitment and—within the quasi-hyperbolic discounting framework—must be time inconsistent (see, for example, Strotz, 1956). Such identification strategies, however, rely on data that is fundamentally different from the task-completion data for which we establish the impossibility of non-parametric identification.

In the closely related, but different, problem of task timing (Carroll et al., 2009; Laibson,

and do not establish an analog of our non-identification result.

[29] They also provide two examples—one for a dynamic discrete choice dataset—that highlight the importance of parametric assumption on the unobservable shocks.

[30] More distantly related, Echenique et al. (forthcoming) provide a non-parametric revealed-preference test to identify time preferences when consumption choices are observed for different intertemporal budget sets and Dziewulski (2018) provides a revealed preference characterization of general discounted utility for datasets that consists of finitely many choices over delayed rewards. As our setting corresponds to a game between the different selves, it is conceptually related to Bergemann and Morris (2013, 2016) who derive predictions that are robust to changes of the private information structure of players.

2015) in which the benefit from doing the task start accumulating as soon as the agent finishes it, it is possible to construct examples in which an agent wants to commit to an earlier deadline, implying that at least partial identification of perceived present bias ($\hat{\beta} \neq 1$) is theoretically feasible. While agents may theoretically benefit from imposing a deadline in such task-timing problems, however, the calibration of the example in Laibson (1997) suggests that their willingness to do so is small, implying that identifying time inconsistency can nevertheless be challenging in real-world data.

Let us emphasize the obvious: even though present-bias is non-identifiable in our task-completion settings absent data on continuation values, it may still be a major driver for why some agents complete tasks last minute. Our results simply caution that the observed task-completion behavior on its own is not enough to infer that present bias is widespread. More broadly, in our setting the chosen parametric form of the unobservable shock fully determines time-preferences' estimates, raising the question in which other economic environments parametric assumptions on the unobservable shocks determine parameter estimates and economic conclusions.

Appendix

Proof of Lemma 1: We first show that the conditions are necessary for a perception-perfect equilibrium. We already argued that any equilibrium must be in cutoff strategies and that the cutoffs used by each self must equal their perceived continuation value v. To characterize these cutoffs, note that for a partially-naive quasi-hyperbolic discounter, the time t and t' selves have the same beliefs about the strategy future selves—i.e. selves active after time $\max\{t, t'\}$—use. Self t thus believes that if she does not complete the task at time t, the task will be completed at the (random) time

$$\hat{\tau}_t = \min\{s > t \colon y_s > \hat{v}_s\}\,,$$

where $\hat{v}_s$ is the perceived cutoff that selves $t < s$ believe Self s will use. Trivially, for all $s > t$ the stopping time $\hat{\tau}_s$ equals $\hat{\tau}_t$ conditional on not stopping before time $s+1$,

$$\mathbb{P}[\hat{\tau}_s = \hat{\tau}_t \mid \hat{\tau}_t > s] = 1\,.$$

Hence, Self t believes that her *perceived continuation utility* v_t if she does not complete the task at time t is given by

$$v_t = \beta\,\mathbb{E}\left[\delta^{\hat{\tau}_t - t} y_{\hat{\tau}_t}\right]\,.$$

We can rewrite the perceived continuation utility by considering the event that the task is completed in period $t+1$ as well as the complementary event that it is completed later

$$v_t = \beta\;\mathbb{E}\left[\delta^{\hat{\tau}_t - t} y_{\hat{\tau}_t}\right] = \beta\;\mathbb{E}\left[\mathbf{1}_{\hat{\tau}_t = t+1}\delta^{\hat{\tau}_t - t} y_{\hat{\tau}_t} + \mathbf{1}_{\hat{\tau}_t > t+1}\delta^{\hat{\tau}_t - t} y_{\hat{\tau}_t}\right].$$

Because Self t believes the task is completed in period $t+1$ if and only if the benefit is greater than the perceived cutoff $y_{t+1} > \hat{v}_{t+1}$, this equals

$$v_t = \beta\;\mathbb{E}\left[\mathbf{1}_{y_{t+1} > \hat{v}_{t+1}}\delta y_{\hat{\tau}_t} + \mathbf{1}_{y_{t+1} \le \hat{v}_{t+1}}\delta^{\hat{\tau}_t - t} y_{\hat{\tau}_t}\right]\,.$$

Since y_{t+1} is distributed according to F and $\hat{\tau}_t = \hat{\tau}_{t+1}$ conditional on not stopping in period $t+1$, we can use the definition of a Riemann–Stieltjes integral to rewrite the above as

$$v_t = \beta\delta\int_{\hat{v}_{t+1}}^{\infty} z\,dF(z) + F(\hat{v}_{t+1})\,\beta\delta\,\mathbb{E}\left[\delta^{\hat{\tau}_{t+1} - (t+1)} y_{\hat{\tau}_{t+1}}\right].$$

Using the definition of v_{t+1} to rewrite the last summand above, we therefore have that

$$v_t = \beta\delta \int_{\hat{v}_{t+1}}^{\infty} z\, dF(z) + F(\hat{v}_{t+1})\, \delta\, v_{t+1}\,. \tag{5}$$

Here, v_t is the cutoff that Self t actually uses. Prior selves, however, believe that Self t discounts with hyperbolic weight $\hat{\beta}$, so the perceived cutoff $\hat{v}_t$ they think Self t uses solves

$$\hat{v}_t = \hat{\beta}\, \mathbb{E}\left[\delta^{\hat{\tau}_t - t} y_{\hat{\tau}_t}\right] = \left(\hat{\beta}/\beta\right)\, \beta\, \mathbb{E}\left[\delta^{\hat{\tau}_t - t} y_{\hat{\tau}_t}\right] = \left(\hat{\beta}/\beta\right)\, v_t\,.$$

Using this equation to replace $\hat{v}_{t+1}$ in (5) establishes that the continuation values $v_1, \ldots, v_{T-1}$ satisfy the recursive equation

$$v_t = \beta\, \delta \int_{\left(\hat{\beta}/\beta\right)\, v_{t+1}}^{\infty} z\, dF(z) + F(\left(\hat{\beta}/\beta\right)\, v_{t+1})\, \delta\, v_{t+1}\,.$$

That any such pair of cutoff strategies constitutes a perception-perfect equilibrium follows from checking the (perceived) optimality conditions inductively starting from the last period.

□

Define the function $g : \mathbb{R} \to \mathbb{R}$ as

$$g(w) = \hat{\beta}\, \delta \int_{w}^{\infty} z\, dF(z) + F(w)\, \delta\, w\,. \tag{6}$$

As the following lemma formally establishes, g has a number of convenient properties.

Lemma 2. *The function g has the following properties:*

i) For all $t \in \{1, \ldots, T-1\}$, the perceived continuation values satisfy $\left(\hat{\beta}/\beta\right)\, v_t = g\left(\left(\hat{\beta}/\beta\right)\, v_{t+1}\right)$.

ii) $g(w)$ is non-decreasing for $w \geq 0$, is right-continuous, and has only upward jumps.

iii) $g(w) \geq w$ for all $w \leq 0$.

Let $F(0) < 1$. Then g has the following additional properties:

iv) $g(w) > w$ for all $w < 0$.

v) Let $w^\star = \inf\{w \in \mathbb{R} \colon g(w) \leq w\}$.[31] If $w^\star < \infty$ it satisfies $g(w^\star) = w^\star$ and $w^\star \geq 0$.

vi) If $w' \geq 0 \geq w$, then $g(w') \geq g(w)$.

Proof of Lemma 2: i) follows immediately from Lemma 1. To see that ii) holds, observe

[31] We follow the convention that $\inf \emptyset = \infty$.

that we can rewrite g as

$$g(w) = \hat{\beta}\,\delta \int_{w}^{\infty} z\, dF(z) + \hat{\beta} F(w)\,\delta\, w + (1-\hat{\beta})F(w)\,\delta\, w$$
$$= \hat{\beta}\,\delta \int_{-\infty}^{\infty} \max\{z, w\}\, dF(z) + (1-\hat{\beta})F(w)\,\delta\, w\,. \tag{7}$$

Note that both the first and the second summand are non-decreasing for $w \geq 0$, and that the first summand is continuous in w while the second is right-continuous and has only upward jumps as F is a CDF.

To see that *iii*) holds, observe that the integral in the first summand of (7) is bounded from below by w and, thus, for $w \leq 0$

$$g(w) \geq \hat{\beta}\,\delta\, w + (1-\hat{\beta})F(w)\,\delta\, w \geq \hat{\beta}\,\delta\, w + (1-\hat{\beta})\delta\, w = \delta w \geq w\,.$$

To see that *iv*) holds observe that the second inequality is strict whenever $w < 0$ and $F(0) < 1$.

We now show *v*). Suppose $w^\star < \infty$. As $g(w) - w$ and is right-continuous for any decreasing, converging sequence w_k with $g(w_k) - w_k \leq 0$ we get that $g(\lim_{k\to\infty} w_k) - \lim_{k\to\infty} w_k \leq 0$ and thus that $w^\star = \inf\{w \in \mathbb{R}\colon g(w) \leq w\}$ satisfies $g(w^\star) = w^\star$. Furthermore, it follows immediately from *iv*) that the set $\{w \in \mathbb{R}\colon g(w) \leq w\}$ contains only $w \geq 0$, and hence that $w^\star \geq 0$.

To show *vi*), note that for $0 \geq w$ (7) together with $w' \geq 0$ implies that

$$g(w') - g(w) = \hat{\beta}\,\delta \left[\int_{-\infty}^{w} (w' - w) dF(z) + \int_{w}^{w'} (w' - z) dF(z)\right]$$
$$+ (1-\hat{\beta})\delta \left[F(w')w' - F(w)w\right] \geq 0,$$

where the inequality follows from the facts that $w' \geq 0$ and $w \leq 0$. □

Proof of Theorem 1: We begin by establishing the result for the "pure cost" case in which all the agent's payoffs are negative $F(0) = 1$. In this case for $w \leq 0$, we have that

$$g(w) = \hat{\beta}\,\delta \int_{-\infty}^{\infty} \max\{z, w\}\, dF(z) + (1-\hat{\beta})F(w)\,\delta\, w \leq 0\,.$$

As $v_T = \underline{y} \leq 0$ and by Lemma 2, *ii*) $\left(\hat{\beta}/\beta\right)\ v_t = g\left(\left(\hat{\beta}/\beta\right)\ v_{t+1}\right)$, it follows by induction that

$v_t \leq 0$ for all $t \in \{1, \ldots, T-1\}$. By Lemma 2, i) and iii) we have that

$$\left(\hat{\beta}/\beta\right)\ v_t = g\left(\left(\hat{\beta}/\beta\right)\ v_{t+1}\right) \geq \left(\hat{\beta}/\beta\right)\ v_{t+1}\,,$$

which establishes that the subjective continuation values are non-increasing.

We next consider the complementary case where some payoffs are strictly positive $F(0) < 1$. Self T's perceived continuation value is $v_T = \underline{y} \leq 0$. Define $w^\star = \inf\{w \in \mathbb{R} \colon g(w) \leq w\}$. If $w^\star = \infty$ we get that $g(w) > w$ for all w, and thus again by Lemma 2, i) we have that $v_t > v_{t+1}$ for all $t \in \{1, \ldots, T-1\}$, which establishes that the subjective continuation values are non-increasing.

Thus, consider the case where $\omega^\star < \infty$. By Lemma 2, i) and v) we have that

$$w^\star - \left(\hat{\beta}/\beta\right)\ v_t = g(w^\star) - g\left(\left(\hat{\beta}/\beta\right)\ v_{t+1}\right)\,. \tag{8}$$

As $v_T = \underline{y} \leq 0$ and $w^\star \geq 0$ (by Lemma 2, v)), we have that $\left(\hat{\beta}/\beta\right)\ v_T \leq w^\star$. We now proceed by induction to show that this implies that $v_t \leq w^\star$ for all $t \in \{1, \ldots, T\}$. The induction step shows that $v_{t+1} \leq w^\star$ implies $v_t \leq w^\star$. We distinguish two cases: First, $\hat{\beta}/\beta\, v_{t+1} \geq 0$. In this case the monotonicity of g, established in Lemma 2, ii), together with Equation 8 implies that $\operatorname{sgn}(w^\star - \hat{\beta}/\beta\, v_t) = \operatorname{sgn}(w^\star - \hat{\beta}/\beta\, v_{t+1})$ and, thus, the induction step follows $\hat{\beta}/\beta\, v_t \leq w^\star$. Second if $\hat{\beta}/\beta\, v_{t+1} < 0$, then by Lemma 2, vi), $g(w^\star) \geq g(\hat{\beta}/\beta\, v_{t+1})$ and hence it follows from Equation 8 that $\hat{\beta}/\beta\, v_t \leq w^\star$. We conclude that $\hat{\beta}/\beta\, v_t \leq w^\star$ for all $t \in \{1, \ldots, T\}$. Hence, $g\left(\hat{\beta}/\beta\, v_{t+1}\right) \geq \hat{\beta}/\beta\, v_{t+1}$, and thus

$$\hat{\beta}/\beta\, v_{t+1} \leq g(\hat{\beta}/\beta\, v_{t+1}) = \hat{\beta}/\beta\, v_t \Rightarrow v_{t+1} \leq v_t\,.$$

This completes the proof. □

Proof of Theorem 2: Fix a non-decreasing sequence of stopping probabilities $0 < p_1 \leq p_2 \leq \ldots \leq p_T < 1$, $(\delta, \beta) \in (0,1] \times (0,1]$, and a penalty $y_{T+1} = \frac{\underline{y}}{\beta\delta} \in \mathbb{R}$. We will construct a distribution F that implies the stopping probabilities p for a sophisticate.

Pick any constant c satisfying

$$c > \max\left\{0, -(1-\beta)\,\delta\underline{y}\ \frac{1-(\delta\,\frac{1-p_T}{2})^{T-1}}{(1-(\delta\,\frac{1-p_T}{2}))(\delta\,\frac{1-p_T}{2})^{T-1}} - \frac{\underline{y}}{(T-1)(\delta\frac{1-p_T}{2})^{T-1}}\right\}. \tag{9}$$

We denote by $f \in \mathbb{R}_+^{T+2}$ the following non-negative vector

$$f_k = \begin{cases} (1-p_T)/2 & \text{if } k = 0, 1 \\ p_{T-k+2} - p_{T-k+1} & \text{if } k \in \{2, \ldots, T\} \\ p_1 & \text{if } k = T+1 \end{cases} . \tag{10}$$

Note that $f_0 > 0$ as $p_T < 1$ and that $\sum_{k=0}^{T+1} f_k = 1$ and define $F_k = \sum_{j=0}^{k} f_j$. We define the a vector $\pi \in \mathbb{R}^{T+2}$ by $\pi_0 = \underline{y} - c, \pi_1 = \underline{y}$, and for $k \in \{2, \ldots, T\}$,

$$\pi_k = \pi_{k-1} + (1-\beta)\,\delta f_{k-1}\pi_{k-1} + \delta\, F_{k-2}\,(\pi_{k-1} - \pi_{k-2})\ , \tag{11}$$

and

$$\pi_{T+1} = \frac{(1 - \beta\delta(p_2 - p_1))\pi_T - \delta(1-p_2)\pi_{T-1}}{\beta\delta p_1} . \tag{12}$$

As the right-hand-side of (11) and (12) depends only on elements with a smaller index, this equation can be solved forward and admits a unique solution. We note for future reference that (12) is equivalent to

$$\pi_T = \beta\delta f_{T+1}\pi_{T+1} + \beta\delta f_T \pi_T + \delta F_{T-1}\pi_{T-1} . \tag{13}$$

We first establish a serious of auxiliary results about the vector π.

Claim A: The vector π is strictly increasing.

Proof of Claim A: We will show that $\pi_k - \pi_{k-1} > 0$ by induction for $k \in \{1, \ldots, T\}$. $\pi_0 < \pi_1$ by construction as $c > 0$. We next do the induction step and assume that $\pi_0 < \pi_1 < \ldots < \pi_{k-1}$. Since for $k \geq 2$ one has $\pi_{k-1} \geq \pi_1 = \underline{y}$ and $F_{k-2} \geq f_0$, (11) implies that

$$(\pi_k - \pi_{k-1}) \geq (1-\beta)\,\delta f_{k-1}\underline{y} + \delta\, f_0\,(\pi_{k-1} - \pi_{k-2}) .$$

Since for $\underline{y} \geq 0$, we have $(1-\beta)\,\delta f_{k-1}\underline{y} \geq 0$, it follows that π is strictly increasing if $\underline{y} \geq 0$. Thus, assume for the rest of the proof that $\underline{y} < 0$. In this case we get that

$$(\pi_k - \pi_{k-1}) \geq \alpha + \gamma\,(\pi_{k-1} - \pi_{k-2})\ , \tag{14}$$

where $\alpha = (1-\beta)\,\delta \underline{y} < 0$ and $\gamma = \delta\, f_0 \in (0,1)$. This implies that

$$(\pi_k - \pi_{k-1}) \geq \alpha \sum_{j=0}^{k-2} \gamma^j + \gamma^{k-1}(\pi_1 - \pi_0) = \alpha \frac{1-\gamma^{k-1}}{1-\gamma} + \gamma^{k-1} c \geq \alpha \frac{1-\gamma^{T-1}}{1-\gamma} + \gamma^{T-1} c$$
$$= \gamma^{T-1}\left(c - |\alpha| \frac{1-\gamma^{T-1}}{(1-\gamma)\gamma^{T-1}}\right) > 0\,. \tag{15}$$

The last inequality here follows from our choice of c satisfying (9). We thus have shown that $\pi_0 < \pi_1 < \ldots < \pi_T$. We are left to show that $\pi_T < \pi_{T+1}$.

Our choice of c furthermore ensures that $\pi_T > 0$, since

$$\pi_T = \pi_1 + \sum_{l=2}^{T} (\pi_l - \pi_{l-1}) \geq \underline{y} + (T-1)\left[\gamma^{T-1}\left(c - |\alpha| \frac{1-\gamma^{T-1}}{(1-\gamma)\gamma^{T-1}}\right)\right] > 0,$$

where the first inequality follows from (15) and the second from our choice of c.

Using $\pi_T > 0$, if $\pi_{T+1} \leq \pi_T$, we have that (13) implies

$$\pi_T = \beta\delta(f_T \pi_T + f_{T+1}\pi_{T+1}) + \delta \pi_{T-1} F_{T-1} \leq \pi_T \beta\delta f_T + \pi_T \beta \delta f_{T+1} + \pi_{T-1}\delta F_{T-1}$$
$$\Leftrightarrow 1 \leq \beta\delta f_T + \beta\delta f_{T+1} + \frac{\pi_{T-1}}{\pi_T}\delta F_{T-1}\,.$$

As $f_T + f_{T+1} + F_{T-1} = 1$, $F_{T-1} \geq f_0 > 0$, and $\pi_{T-1} < \pi_T$, this is a contradiction and completes the proof of Claim A.

Claim B: For $k \in \{2, \ldots, T\}$ the vector π satisfies the equation

$$\pi_k = \beta\,\delta \sum_{j=k}^{T+1} f_j \pi_j + \delta\, F_{k-1}\pi_{k-1}\,. \tag{16}$$

Proof of Claim B: We proof the claim by induction starting from T. For $k = T$ the claim is satisfied by (13). Now we establish that if (16) is satisfied for $k > 3$ it is also satisfied for

$k-1$. By (11) we have that for $k \geq 2$

$$\begin{aligned}\pi_{k-1} &= \pi_k - (1-\beta)\,\delta f_{k-1}\pi_{k-1} - \delta\, F_{k-2}\,(\pi_{k-1} - \pi_{k-2}) \\ &= \beta\,\delta \sum_{j=k}^{T+1} f_j \pi_j + \delta\, F_{k-1}\pi_{k-1} - (1-\beta)\,\delta f_{k-1}\pi_{k-1} - \delta\, F_{k-2}\,(\pi_{k-1} - \pi_{k-2}) \\ &= \beta\,\delta \sum_{j=k-1}^{T+1} f_j \pi_j + \delta\, F_{k-2}\pi_{k-1} - \delta\, F_{k-2}\,(\pi_{k-1} - \pi_{k-2}) = \beta\,\delta \sum_{j=k-1}^{T+1} f_j \pi_j + \delta\, F_{k-2}\pi_{k-2}\,.\end{aligned}$$

This completes the induction step and thus the proof of Claim B.

We define the distribution $F : \mathbb{R} \to [0,1]$ that assigns to the points $\pi_0, \ldots, \pi_{T+1}$ the probabilities $f_0, \ldots, f_{T+1}$

$$F(x) = \sum_{k=0}^{T+1} f_k\, \mathbf{1}_{\pi_k \leq x}\,. \tag{17}$$

Claim C: The continuation values $v_1, \ldots, v_T$ induced by the distribution F are given by $v_t = \pi_{T+1-t}$ for $t \in \{1, \ldots, T\}$.

Proof of Claim C: Using $\hat{\beta} = \beta$ in Lemma 1, the perceived continuation values are the unique solution to the equation

$$v_t = \begin{cases} \beta\,\delta \int_{v_{t+1}}^{\infty} z\, dF(z) + F(v_{t+1})\,\delta\, v_{t+1} & \text{for } t \in \{1, \ldots, T-1\} \\ \underline{y} & \text{for } t = T \end{cases}\,. \tag{18}$$

As $\pi_1 = \underline{y}$, we have that $\pi_{T+1-t} = v_t$ for $t = T$. For $t < T$, (18) simplifies to

$$v_t = \beta\delta \int_{v_{t+1}}^{\infty} z dF(z) + F(v_{t+1})\delta v_{t+1} = \beta\,\delta \sum_{j=0}^{T+1} \mathbf{1}_{v_{t+1} < \pi_j} f_j \pi_j + \delta\, \left(\sum_{j=0}^{T+1} \mathbf{1}_{v_{t+1} \geq \pi_j} f_j \right) v_{t+1}\,.$$

Plugging in $v_t = \pi_{T+1-t} = \pi_k$ on the left and right-hand-side and using that π is strictly increasing by Claim A, the above equation simplifies to the condition that for $k \in \{2, \ldots, T\}$

$$\pi_k = \beta\,\delta \sum_{j=1}^{T+1} \mathbf{1}_{\pi_{k-1}<\pi_j} f_j \pi_j + \delta \left(\sum_{j=0}^{T+1} \mathbf{1}_{\pi_{k-1}\geq\pi_j} f_j\right) \pi_{k-1}$$
$$= \beta\,\delta \sum_{j=k}^{T+1} f_j \pi_j + \delta \left(\sum_{j=0}^{k-1} f_j\right) \pi_{k-1}\,.$$

We established that this equation is satisfied in Claim B and thus completed the proof of Claim C.

We finally show that the distribution F leads to the stopping probabilities p. We have that the continuation probability in period $t \in \{1, \ldots, T\}$ equals

$$F(v_t) = \sum_{j=0}^{T+1} \mathbf{1}_{\pi_j \leq v_t} f_j = \sum_{j=0}^{T+1} \mathbf{1}_{\pi_j \leq \pi_{T+1-t}} f_j = \sum_{j=0}^{T+1-t} f_j$$
$$= (1-p_T) + \sum_{j=2}^{T+1-t} (p_{T+2-j} - p_{T+1-j}) = (1-p_t)\,.$$

This completes the proof of Theorem 1. □

Lemma 3. *Suppose $\delta < 1$ and the agent believes to be time-consistent $\hat{\beta} = 1$.*

i) For every distribution F with $F(\underline{y}) > 0$ and $\underline{y} < 0$, the continuation values are strictly decreasing $v_1 > v_2 > \ldots > v_T$.

ii) A first-order stochastic dominance increase in the payoff distribution F increases the vector of perceived continuation values point-wise.

Proof of Lemma 3: *i)***:** Suppose towards a contradiction that the continuation values are not strictly decreasing. Since they are weakly decreasing by Theorem 1, we thus have $v_{t'-1} = v_{t'}$ for some $t' \in \{2, \ldots, T\}$. By Lemma 2 *i)*, we have that $v_t/\beta = g(v_{t+1}/\beta)$ for all $t \in \{1, \ldots, T-1\}$, where, by (7),

$$g(x) = \delta \int_{-\infty}^{\infty} \max\left\{z, x\right\} dF(z)\,.$$

Thus, $v_t = v_{t'}$ for all $t \leq t'$.

We next prove that $v_{t'+1} = v_{t'}$. Denote by $\underline{m} = \min\left(\operatorname{supp} F\right)$ the left end-point of the support of F. By assumption $\underline{m} \leq \underline{y} < 0$. As $\underline{y} < 0$ and $F(\underline{y}) > 0$, we get that

$0 < F(\underline{y}) \leq F(0)$. By Lemma 2 *iv*), any fixed point of g is non-negative, so that $v'_t \geq 0$. We have that

$$g\left(\frac{v_{t'}}{\beta}\right) = \frac{v_{t'}}{\beta} = g\left(\frac{v_{t'+1}}{\beta}\right). \tag{19}$$

As g is strictly increasing for $x \geq \underline{m}$ and $\frac{v_{t'}}{\beta} \geq 0 \geq \underline{m}$, (19) implies that $v_{t'+1} = v_{t'}$. By induction, $v_s = v_t$ for all $s, t \in \{1, \ldots, T\}$. As $v_T = \underline{y}$, this implies that $v_t = \underline{y}$ for all t. This is a contradiction since we established that $v_{t'} \geq 0$ and $\underline{y} < 0$ by assumption.

We now show *ii*): Let v be the continuation values associated with F and $\tilde{v}$ the continuation values associated with $\tilde{F} \prec_{FOSD} F$. We want to show that $v_t \geq \tilde{v}_t$ for every $t \in \{1, \ldots, T\}$. We show the result by backward induction over T. The start of the induction is that $v_T = \tilde{v}_T = \underline{y}$. To complete the induction step, we show that $v_{t+1} \geq \tilde{v}_{t+1}$ implies $v_t \geq \tilde{v}_t$

$$\begin{aligned} {}^{v_t}\!/_{\beta} &= \delta \int_{-\infty}^{\infty} \max\left\{z, {}^{v_{t+1}}\!/_{\beta}\right\} dF(z) \geq \delta \int_{-\infty}^{\infty} \max\left\{z, {}^{\tilde{v}_{t+1}}\!/_{\beta}\right\} dF(z) \\ &\geq \delta \int_{-\infty}^{\infty} \max\left\{z, {}^{\tilde{v}_{t+1}}\!/_{\beta}\right\} d\tilde{F}(z) = {}^{\tilde{v}_t}\!/_{\beta}. \qquad \square \end{aligned}$$

Proof of Theorem 3: Throughout the proof fix an arbitrary non-decreasing sequence of stopping probabilities $0 < p_1 \leq \ldots \leq p_T < 1$, a discount factor $\delta < 1, \beta \in (0, 1]$ and a continuation payoff $\underline{y} = \delta\beta v_{T+1} < 0$ in period T. We will show that there exists a payoff distribution that leads to the stopping probabilities p for a naive agent with the time preference (δ, β).

Let $G_{a,b}$ denote the uniform CDF on $[a, b]$ for $a < b$

$$G_{a,b}(x) = \begin{cases} 0 & \text{for } x < a \\ \frac{x-a}{b-a} & \text{for } x \in [a, b] \\ 1 & \text{for } x > b \end{cases}$$

and a Dirac measure $G_{a,a}(x) = \mathbf{1}_{a \leq x}$ for $a = b$.

Fix two arbitrary constants $c_1, c_2 > 0$. For every non-increasing sequence $v_1 \geq \ldots \geq v_{T-1}$ with $v_{T-1} \geq \underline{y}$, define the function $F(\cdot; v)$ as the weighted sum of the CDFs of $T + 1$ uniform distributions on the intervals $[\pi_k(v), \pi_{k+1}(v)]$ for $k \in \{0, \ldots T\}$ as

$$F(x; v) = \sum_{k=0}^{T} f_k \, G_{\pi_k(v), \pi_{k+1}(v)}(x). \tag{20}$$

We set the endpoints of the intervals $[\pi_k(v), \pi_{k+1}(v)]$

$$\pi_k(v) = \begin{cases} \underline{y} - c_1 & \text{if } k = 0 \\ \underline{y} & \text{if } k = 1 \\ v_{T-k+1} & \text{if } k \in \{2, \ldots, T\} \\ v_1 + c_2 & \text{if } k = T+1 \end{cases}, \tag{21}$$

and the probabilities f_k assigned to each interval as

$$f_k = \begin{cases} 1 - p_T & \text{if } k = 0 \\ p_{T-k+1} - p_{T-k} & \text{if } k \in \{1, \ldots, T-1\} \\ p_1 & \text{if } k = T \end{cases}. \tag{22}$$

Note that $f_k \geq 0$, that for $k < T$

$$\sum_{j=0}^{k} f_j = 1 - p_{T-k}, \tag{23}$$

and that $\sum_{j=0}^{T} f_j = 1$. For every v, the function $F(\cdot; v)$ is non-decreasing and non-negative as the CDF G is non-decreasing and non-negative. It thus follows that F is a well defined CDF whose support satisfies $\operatorname{supp} F(\cdot; v) \subseteq [\pi_0, \pi_{T+1}] = [\underline{y} - c_1, v_1 + c_2]$.

Consider now the continuation values w induced by $F(\cdot; v)$. By Lemma 1, they can be computed by solving the equation

$$\frac{w_t}{\beta} = \delta \int_{-\infty}^{\infty} \max\left\{z, \frac{w_{t+1}}{\beta}\right\} dF(z; v) \quad \text{for } t \in \{1, \ldots, T-1\}, \tag{24}$$

with $w_T = \underline{y}$. Denote by $L : \mathbb{R}^{T-1} \to \mathbb{R}^{T-1}$ the function mapping $(v_1, \ldots, v_{T-1})$ to $(w_1, \ldots, w_{T-1})$ using (24). By Theorem 1, $w = L(v)$ is non-increasing. As w is non-increasing and $w_T = \underline{y}$, it follows that $(Lv)_t \geq \underline{y}$ for all $t \in \{1, \ldots, T\}$. Furthermore, as $\operatorname{supp} F(\cdot; v) \subseteq [\underline{y} - c_1, v_1 + c_2]$

$$\begin{aligned} w_1 &= \beta\delta \int_{-\infty}^{\infty} \max\left\{z, \frac{w_2}{\beta}\right\} dF(z; v) \leq \beta\delta \int_{-\infty}^{\infty} \max\left\{v_1 + c_2, \frac{w_1}{\beta}\right\} dF(z; v) \\ &= \delta\beta \max\left\{(v_1 + c_2), \frac{w_1}{\beta}\right\}. \end{aligned}$$

We distinguish two cases: $w_1 > 0$ and $w_1 \leq 0$. If $w_1 > 0$ we have that $w_1 > \delta w_1$ and for all v_1 such that $v_1 \leq \frac{\delta}{1-\delta}c_2$ we have in addition

$$w_1 \leq \delta\beta \max\left\{(v_1 + c_2), \frac{w_1}{\beta}\right\} = \delta\beta(v_1 + c_2) \leq \delta(v_1 + c_2) \leq \frac{\delta}{1-\delta}c_2\,.$$

If $w_1 \leq 0$ we have that for all $v_1 \leq \frac{\delta}{1-\delta}c_2$

$$w_1 \leq \delta\beta \max\left\{(v_1 + c_2), \frac{w_1}{\beta}\right\} \leq \max\{\delta\beta(v_1 + c_2), 0\} \leq \frac{\delta}{1-\delta}c_2\,.$$

In either case, we have that $\underline{y} \leq w_t \leq \frac{\delta}{1-\delta}c_2$ for all $t \in \{1, \ldots, T-1\}$. Consequently, $L : \mathbb{R}^{T-1} \to \mathbb{R}^{T-1}$ maps $M \subset \mathbb{R}^{T-1}$ into itself, where M is the set of non-increasing sequences contained in $[\underline{y}, \frac{\delta}{1-\delta}c_2]^{T-1}$, i.e.

$$M = \left\{m \in \left[\underline{y}, \frac{\delta}{1-\delta}c_2\right]^{T-1} : m_1 \geq m_2 \geq \ldots \geq m_{T-1}\right\}.$$

We note that M combined with the pointwise order $\geqq$ forms a complete bounded lattice, as the point-wise maximum and minimum over any set of non-increasing sequences is non-increasing.

We note that $v \mapsto \pi(v)$ is monotone in the pointwise order. Furthermore, $G_{a,b}(x) \geq G_{a',b'}(x)$ for all $(a, b) \leqq (a', b')$ and all $x \in \mathbb{R}$. By (20) this implies that $F(x; v) \geq F(x, v')$ for all $v \leqq v'$ and all $x \in \mathbb{R}$, which means that $v \mapsto F(\cdot; v)$ is monotone in first-order stochastic dominance (FOSD). By Lemma 3 *ii*), increasing the distribution of payoffs in FOSD will (weakly) increase the perceived continuation values. This implies that the operator $L : M \to M$ is monotone with respect to the pointwise order. As L is a monotone operator, i.e. $L(v) \geqq L(w)$ if $v \geqq w$, it admits at least one fixed point on the complete lattice M by Tarski's fixed point theorem. We pick an arbitrary fixed point of L and denote it by $\omega^\star$. By construction the fixed point $\omega^\star$ is such that the payoff distribution $F(\cdot; \omega^\star)$ will lead to the sequence of continuation values $\omega^\star$.

We next argue that the distribution $F(\cdot; w^\star)$ induces the stopping probabilities p and thus solves our problem. First, we note that $F(\underline{y}; \omega^\star) = 1 - p_T > 0$ and that $\underline{y} < 0$ by assumption. By Lemma 3 *i*), it follows that the continuation values $w^\star$ induced by $F(\cdot; \omega^\star)$ must be strictly decreasing $w_1^\star > w_2^\star > \ldots > w_{T-1}^\star$. As $w^\star$ is the continuation value associated with $F(\cdot; \omega^\star)$, the agent stops in period $t \in \{1, \ldots, T\}$ if and only if $y_t \geq w_t^\star$, which happens with probability

$$\begin{aligned}\mathbb{P}[y > w_t^\star] = 1 - F(w_t^\star; w^\star) &= 1 - \sum_{k=0}^{T} f_k \, G_{\pi_k(w^\star),\pi_{k+1}(w^\star)}(w_t^\star) = 1 - \sum_{k=0}^{T} f_k \, \mathbf{1}_{\pi_{k+1}(w^\star)\le w_t^\star} \\ &= 1 - \sum_{k=1}^{T-1} f_k \, \mathbf{1}_{w_{T-k}^\star \le w_t^\star} - f_0 \mathbf{1}_{\underline{y}\le w_t^\star} - f_T \mathbf{1}_{w_1^\star + c_2 \le w_t^\star} \\ &= 1 - \sum_{k=0}^{T-t} f_k \; = 1 - (1-p_t) = p_t\,;\end{aligned}$$

where we used (23) in the second to last equality. Thus, $F(\cdot;\omega^\star)$ leads to the stopping probabilities p, which completes the proof. □

Lemma 4. *Whenever* (4) *admits a solution for a plausible dataset, there exists a solution F that consists of exactly $T+1$ mass points located at $(\pi_0, \dots, \pi_T)$ that satisfy*

$$\pi_0 \le v_T < \pi_1 \le v_{T-1} < \ldots \le \pi_{T-1} \le v_1 < \pi_T,$$

with associated probabilities $f_k = \mathbb{P}[y = \pi_k]$ given by

$$f_k = \begin{cases} 1 - p_T & \textit{if } k = 0 \\ p_{T-k+1} - p_{T-k} & \textit{if } k \in \{1, \dots, T-1\} \\ p_1 & \textit{if } k = T \end{cases}.$$

Proof of Lemma 4: Let the pair u, G solve 4. From now one, fix u. Let $\mathbb{E}_G$ denote the expectation taken with respect to the cumulative distribution function G, and $\mathbb{P}_G$ the probability mass with respect to G.

We now specify a distribution F that has the properties specified in the Lemma. The $T+1$ mass points $(\pi_0, \dots, \pi_T)$ are located at

$$\pi_k = \begin{cases} \mathbb{E}_G[y | y \le v_T] & \text{if } k = 0 \\ \mathbb{E}_G[y | v_{T-k+1} < y \le v_{T-k}] & \text{if } k \in \{1, \dots, T-1\} \\ \mathbb{E}_G[y | v_1 < y] & \text{if } k = T \end{cases}.$$

and their probability mass is given by f_k as specified in the Lemma. Observe that by construction, we have

$$\pi_0 \leq v_T < \pi_1 \leq v_{T-1} < \ldots \leq \pi_{T-1} \leq v_1 < \pi_T .$$

Since G solves 4 and $1 - F(v_t) = p_t$ for all $t \in \{1, \ldots, T\}$ by construction, one has

$$1 - F(v_t) = 1 - G(v_t) \ \ \forall t \in \{1, \ldots, T\}.$$

Furthermore,

$$\begin{aligned}
\int_{v_{t+1}}^{\infty} z\, dG(z) &= \sum_{k=T-t}^{T-1} \mathbb{E}_G[y | v_{T-k+1} < y \leq v_{T-k}] \mathbb{P}_G[y | v_{T-k+1} < y \leq v_{T-k}] + \mathbb{E}_G[y | v_1 < y] \mathbb{P}_G[y | v_1 < y] \\
&= \sum_{k=T-t}^{T} f_k \pi_k \\
&= \int_{v_{t+1}}^{\infty} z\, dF(z) \, .
\end{aligned}$$

Thus, since u, G solve 4 so do u, F. □

Proof of Theorem 4: Lemma 4 implies for a plausible dataset that (4) admits a solution if and only if there exists $\pi \in \mathbb{R}^{T+1}, f \in \Delta^{T+1}$ and a monotone function u such that

$$v_t = u(m_t) \qquad \forall t \in \{1, \ldots, T\} \, , \tag{25}$$

$$\pi_0 \leq v_T < \pi_1 \leq v_{T-1} < \ldots \leq \pi_{T-1} \leq v_1 < \pi_T , \tag{26}$$

$$\sum_{k=T-t}^{T} \pi_k f_k = \frac{\delta^{-1} v_t - (1 - p_{t+1}) v_{t+1}}{\beta} \qquad \forall t \in \{1, \ldots, T-1\} \, , \tag{27}$$

$$\sum_{k=T-t+1}^{T} f_k = p_t \ \ , \qquad \forall t \in \{1, \ldots, T\} \, . \tag{28}$$

Equation 28 is equivalent to $f_T = p_1, f_0 = 1 - p_T$ and for all $t \in \{2, \ldots, T\}$

$$p_t - p_{t-1} = \sum_{k=T-t+1}^{T} f_k - \sum_{k=T-t+2}^{T} f_k = f_{T-t+1} \, ,$$

and thus completely determines f. From now on we thus consider f as given.

Equation 27 for $t = 1$ is equivalent to

$$\pi_{T-1} f_{T-1} + \pi_T f_T = \frac{\delta^{-1} v_1 - (1 - p_2) v_2}{\beta} .$$

We note that there exists π satisfying the above equation and (26) if and only if

$$v_2 f_{T-1} + v_1 f_T < \frac{\delta^{-1} v_1 - (1 - p_2) v_2}{\beta} . \tag{29}$$

That this is necessary follows as (26) provides a lower bound on π_{T-1} and π_T. Since, $f_T = p_1 > 0$, this is also sufficient as you can always chose π_T arbitrarily large. Rearranging for β and plugging in f yields

$$\beta < \frac{\delta^{-1} v_1 - (1 - p_2) v_2}{v_2(p_2 - p_1) + v_1 p_1} . \tag{30}$$

Next, we consider (27) for $t \in \{2, \dots, T-1\}$. Subtracting (27) evaluated at $t-1$ from (27) evaluated at t yields

$$\pi_{T-t} f_{T-t} = \sum_{k=T-t}^{T} \pi_k f_k - \sum_{k=T-t+1}^{T} \pi_k f_k = \frac{\delta^{-1} v_t - (1 - p_{t+1}) v_{t+1}}{\beta} - \frac{\delta^{-1} v_{t-1} - (1 - p_t) v_t}{\beta} ,$$

which is equivalent to

$$\pi_{T-t} = \frac{v_{t+1}(p_{t+1} - p_t) - \delta^{-1}(v_{t-1} - v_t) + (1 - p_t)(v_t - v_{t+1})}{\beta(p_{t+1} - p_t)}.$$

The above equation admits a solution satisfying (26) if and only if for all $t \in \{2, \dots, T-1\}$, $v_{t+1} < \pi_{T-t} \leq v_t$. Rewriting using the definition of $a(\delta, t)$ from the statement of the theorem, 27 admits a solution satisfying (26) if for all $t \in \{2, \dots, T-1\}$ both $v_{t+1}\beta < v_{t+1} a(\delta, t)$ and $v_t \beta \geq v_{t+1} a(\delta, t)$, and in addition

$$\beta < \frac{\delta^{-1} v_1 - (1 - p_2) v_2}{v_2(p_2 - p_1) + v_1 p_1} . \tag{31}$$

This completes the proof. □

References

Abbring, Jaap H. and Øystein Daljord, "Identifying the Discount Factor in Dynamic Discrete Choice Models," September 2019.

Altmann, Steffen, Christian Traxler, and Philipp Weinschenk, "Deadline and Cognitive Limitations," November 2019. Working Paper.

Augenblick, Ned and Matthew Rabin, "An Experiment on Time Preference and Misprediction in Unpleasant Tasks," *Review of Economic Studies*, 2019, *86* (3), 941–975.

_ , Muriel Niederle, and Charles Sprenger, "Working over Time: Dynamic Inconsistency in Real Effort Tasks," *Quarterly Journal of Economics*, 2015, *130* (3), 1067–1115.

Bajari, Patrick, Chenghuan Sean Chu, Denis Nekipelov, and Minjung Park, "Identification and Semiparametric Estimation of a Finite Horizon Dynamic Discrete Choice Model with a Terminating Action," *Quantitative Marketing and Economics*, 2016, *14*, 271–323.

Bergemann, Dirk and Stephen Morris, "Robust Predictions in Games with Incomplete Information," *Econometrica*, 2013, *81* (4), 1251–1308.

_ and _ , "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 2016, *11* (2), 487–522.

Bisin, Alberto and Kyle Hyndman, "Present-Bias, Procrastination and Deadlines in a Field Experiment," June 2018. Working Paper.

Brown, Jeffrey R. and Alessandro Previtero, "Saving for Retirement, Annuities and Procrastination," 2018. Working Paper.

Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick, "Optimal Defaults and Active Decisions," *Quarterly Journal of Economics*, 2009, *124* (4), 1639–1674.

Chou, Cheng, "Identification and Linear Estimation of General Dynamic Programming Discrete Choice Models," May 2016.

Christensen, Timothy and Benjamin Connault, "Counterfactual Sensitivity and Robustness," March 2019. Working Paper.

Chung, Doug J., Thomas Steenburgh, and K. Sudhir, "Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans," *Marketing Science*, 2014, *33* (2), 165–187.

Daljord, Øystein, Denis Nekipelov, and Minjung Park, "Comments on "identification and semiparametric estimation of a finite horizon dynamic discrete choice model with a terminating action"," *Quantitative Marketing and Economics*, 2019, *17* (4), 439–449.

DellaVigna, Stefano and Ulrike Malmendier, "Paying Not to Go to the Gym," *American Economic Review*, 2006, *96* (3), 694–719.

Dixit, Avinash K. and Robert S. Pindyck, *Investment under Uncertainty*, Princeton, NJ: Princeton University Press, 1994.

Dziewulski, Paweł, "Revealed Time Preference," *Games and Economic Behavior*, 2018, *112*, 67–77.

Echenique, Federico, Taisuke Imai, and Kota Saito, "Testable Implications of Models of Intertemporal Choice: Exponential Discounting and its Generalizations," *American Economic Journal: Microeconomics*, forthcoming.

Ericson, Keith M. Marzilli, "On the Interaction of Memory and Procrastination: Implications for Reminders, Deadlines, and Empirical Estimation," *Journal of the European Economic Association*, July 2017, *15* (3), 692–719.

_ **and David I. Laibson**, "Intertemporal Choice," in B. Douglas Bernheim, Stefano DellaVigna, and David I. Laibson, eds., *Handbook of Behavioral Economics - Foundations and Applications 2*, Elsevier, 2019.

Frakes, Michael D. and Melissa F. Wasserman, "Procrastination in the Workplace: Evidence from the U.S. Patent Office," 2016. NBER Working Paper #22987.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 2002, *40* (2), 351–401.

Fudenberg, Drew and David K Levine, "A Dual-self Model of Impulse Control," *American Economic Review*, 2006, *96* (5), 1449–1476.

Heffetz, Ori, Ted O'Donoghue, and Henry S. Schneider, "Forgetting and Heterogeneity in Task Delay: Evidence from New York City Parking-Ticket Recepients," December 2016. NBER Working Paper #23012.

Laibson, David, "Why don't Present-Biased Agents make Commitments?," *American Economic Review Papers and Proceedings*, 2015, *105* (5), 267–272.

Laibson, David I., "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

Levy, Matthew and Pasquale Schiraldi, "Identification of Intertemporal Preferences in History-dependent Dynamic Discrete Choice Models," 2020.

Magnac, Therry and David Thesmar, "Identifying Dynamic Discrete Decision Processes," *Econometrica*, March 2002, *70* (2), 801–816.

Mahajan, Aprajit, Christian Michel, and Alessandro Tarozzi, "Identification of Time-Inconstistent Models: The Case of Insecticide Treated Nets," March 2020. Working Paper.

Martinez, Seung-Keun, Stephan Meier, and Charles Sprenger, "Procrastination in the Field: Evidence from Tax Filing," 2017. Working Paper.

Norets, Andrij and Xun Tang, "Seminparametric Inference in Dynamic Binary Choice Models," *Review of Economic Studies*, 2014, *81* (3), 1229–1262.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 1999, *89* (1), 103–124.

_ **and** _ , "Choice and Procrastination," *Quarterly Journal of Economics*, 2001, *116* (1), 121–160.

Oliveira, Henrique De and Rohit Lamba, "Rationalizing Dynamic Choices," 2019. Working Paper.

Ramsey, Frank P., "A Mathematical Theory of Saving," *Economic Journal*, 1928, *38*, 543–559.

Rust, John, "Structural Estimation of Markov Decision Processes," in Robert F. Engle and Daniel L. McFadden, eds., *Handbook of Econometrics*, Vol. IV 1994, pp. 3082–3139.

Slemrod, Joel, Charles Christian, Rebeccas London, and Jonathan A. Parker, "April 15 Syndrome," *Economic Inquiry*, October 1997, *35*, 695–709.

Spiegler, Ran, *Bounded Rationality and Industrial Organization*, Oxford University Press, 2011.

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 1956, *23*, 165–180.

Wald, Abraham, "Sequential Tests of Statistical Hypotheses," *Annals of Mathematical Statistics*, June 1945, *16* (2), 117–186.

Weisbrod, Burton A., "Collective-Consumption Services of Individual-Consumption Goods," *Quarterly Journal of Economics*, August 1964, *78* (3), 471–477.

Yao, Song, Carl F Mela, Jeongwen Chiang, and Yuxin Chen, "Determining Consumers' Discount Rates with Field Studies," *Journal of Marketing Research*, 2012, *49* (6), 822–841.

Supplementary Appendix

This supplementary appendix contains variations of Example 2 in the main body of the paper, which further illustrate the importance of functional form assumptions for the analyst's estimates. Table SA1 provides the corresponding log-likelihood estimates when the analyst does not know (and thus estimates) the mean and standard deviation of the payoff distribution. As in the case with known mean and variance, the analyst misestimates β to be substantially below 1.[32] Table SA2 and Table SA3 illustrate that also with 30 or 60 periods, β is incorrectly estimated to be substantially below 1. Table SA4 provides estimates for a variant of Example 2 in which the true distribution is logistic. In this variant, the functional form assumption determines whether β is over- or underestimated.[33] Table SA5 (and SA6) provide estimates for β analogous to Example 2 if the agent is truly present-biased and naive with $\beta = 0.9$, and the analyst does (or does not) know the mean and variance of the distribution. In either case, the analyst significantly underestimates β. Furthermore, Figure SA1 illustrates that eventually as the number of periods the analyst observes increases, her estimates move further and further away from truth.

Parametric Family	β	*Mean*	*Std. Deviation*	*Log-Likelihood*
Uniform Naive	1.	-1.86762	5.78115	-1.59186
Uniform Sophisticate	1.	-2.04179	1.87369	-1.59186
Normal Naive	0.822972	0.0942045	3.47898	-1.59187
Normal Sophisticate	0.826388	0.0978794	3.10058	-1.59187
Extreme Value Naive	0.807256	-2.05785	2.37227	-1.59186
Extreme Value Sophisticate	0.830535	-1.84762	1.85227	-1.59187
Logistic Naive	0.763135	0.193664	9.44528	-1.59187
Logistic Sophisticate	0.768789	0.105082	4.10288	-1.59188
Laplace Naive	0.640929	0.206991	8.82003	-1.59199
Laplace Sophisticate	0.650699	0.0614326	2.24342	-1.59204

Table SA1: Log-likelihood estimates of β and the mean and standard deviation for Example 2 if the analyst does not know the mean and standard deviation of the payoff distribution.

[32] If the analyst selects the model with the highest log-likelihood, for example, she concludes that the agent is naive time-inconsistent with $\beta = 0.807$ and that the shocks follow an extreme value distribution.

[33] Independently of whether she supposes the agent is naive or sophisticated, (when not imposing $\beta \leq 1$ a priori) she estimates β to be 0.96 for the Laplace distribution and 1.19 for the extreme value distribution.

Parametric Family	β	*Log-Likelihood*
Uniform Naive	1.	-3.29153
Uniform Sophisticate	1.	-3.29153
Normal Naive	0.871612	-3.29198
Normal Sophisticate	0.88423	-3.29228
Extreme Value Naive	0.765061	-3.29383
Extreme Value Sophisticate	0.792468	-3.29483
Logistic Naive	0.814908	-3.29203
Logistic Sophisticate	0.836259	-3.29254
Laplace Naive	0.758422	-3.29317
Laplace Sophisticate	0.787311	-3.29418

Table SA2: Log-likelihood estimates of β for the payoff distribution and parameters specified in Example 2 if the analyst knows the mean and standard deviation of the payoff distribution with $T = 30$ periods.

Parametric Family	β	*Log-Likelihood*
Uniform Naive	1.	-3.95505
Uniform Sophisticate	1.	-3.95505
Normal Naive	0.889306	-3.95576
Normal Sophisticate	0.903474	-3.95624
Extreme Value Naive	0.801094	-3.95715
Extreme Value Sophisticate	0.8301	-3.95833
Logistic Naive	0.835118	-3.95584
Logistic Sophisticate	0.85936	-3.9566
Laplace Naive	0.794377	-3.95701
Laplace Sophisticate	0.824827	-3.95823

Table SA3: Log-likelihood estimates of β for the payoff distribution and parameters specified in Example 2 if the analyst knows the mean and standard deviation of the payoff distribution with $T = 30$ periods.

Parametric Family	β	*Log-Likelihood*
Uniform Naive	1.1051	-1.61023
Uniform Sophisticate	1.10823	-1.61029
Normal Naive	1.02514	-1.60953
Normal Sophisticate	1.0253	-1.60953
Extreme Value Naive	1.1942	-1.61034
Extreme Value Sophisticate	1.19231	-1.61008
Logistic Naive	1.	-1.60944
Logistic Sophisticate	1.	-1.60944
Laplace Naive	0.959755	-1.61017
Laplace Sophisticate	0.960106	-1.61016

Table SA4: Log-likelihood estimates of β if the true distribution is Logistic and has the same mean and standard deviation as in Example 2. We suppose the analyst knows the mean and standard deviation of the payoff distribution, and that $T = 5$ periods.

Parametric Family	β	*Log-Likelihood*
Uniform Naive	0.9	-1.57692
Uniform Sophisticate	0.900684	-1.57692
Normal Naive	0.725994	-1.57692
Normal Sophisticate	0.730595	-1.57693
Extreme Value Naive	0.467228	-1.58092
Extreme Value Sophisticate	0.477292	-1.58106
Logistic Naive	0.670309	-1.57692
Logistic Sophisticate	0.676695	-1.57693
Laplace Naive	0.545986	-1.57699
Laplace Sophisticate	0.555965	-1.57705

Table SA5: Log-likelihood estimates of β for the mean and standard deviation from Example 2 if the agent is naive and $\beta = 0.9$, the true distribution is Uniform, and the analyst knows the mean and standard deviation of the payoff distribution with $T = 5$ periods.

Parametric Family	β	*Mean*	*Std. Deviation*	*Log-Likelihood*
Uniform Naive	0.899999	-0.0000121032	3.08835	-1.57692
Uniform Sophisticate	0.901039	0.00221368	0.838862	-1.57692
Normal Naive	0.729808	0.0281063	2.91605	-1.57692
Normal Sophisticate	0.736594	0.0731089	4.76987	-1.57692
Extreme Value Naive	0.706168	-0.347689	0.621169	-1.57692
Extreme Value Sophisticate	0.633785	0.144273	0.652626	-1.60398
Logistic Naive	0.6741	0.0166023	2.176	-1.57692
Logistic Sophisticate	0.683439	0.0773394	5.63958	-1.57693
Laplace Naive	0.55626	0.017136	1.21714	-1.57698
Laplace Sophisticate	0.569426	0.0941048	5.09827	-1.57703

Table SA6: Log-likelihood estimates of β, the mean, and standard deviation if the agent is naive and $\beta = 0.9$, the true distribution is Uniform with parameters as in Example 2, and the analyst does not know the mean and standard deviation of the payoff distribution with $T = 5$ periods.

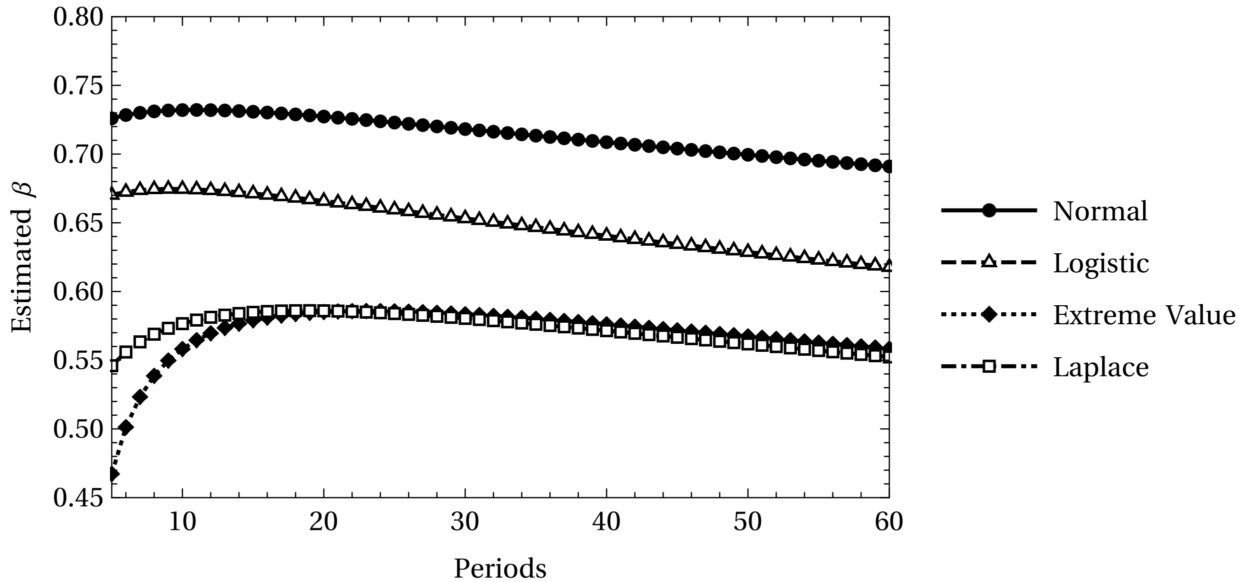

Figure SA1: Estimates of β in Example 2 when the agent is naive and time-inconsistent with $\beta = 0.9, \hat{\beta} = 1, \delta = 1$ for different number of periods T under different parametric assumptions. The analyst knows that $\delta = 1, \hat{\beta} = 1$, as well as the mean and standard deviation of the shock distribution, and estimates β. As the analyst observes the behavior in more and more periods, the estimated value of β eventually moves further away from the true value of 0.9.