

Myopia and Discounting*

Xavier Gabaix and David Laibson

August 3, 2017

Abstract

We assume that perfectly patient agents estimate the value of future events by generating noisy, unbiased simulations and combining those signals with priors to form posteriors. These posterior expectations exhibit *as-if discounting*: agents make choices as if they were maximizing a stream of known utils weighted by a discount function, $D(t)$. This as-if discount function reflects the fact that estimated utils are a combination of signals and priors, so average expectations are optimally shaded toward the mean of the prior distribution, generating behavior that partially mimics the properties of classical time preferences. When the simulation noise has variance that is linear in the event's horizon, the as-if discount function is hyperbolic, $D(t) = 1/(1 + \alpha t)$. Our agents exhibit systematic preference reversals, but have no taste for commitment because they suffer from imperfect foresight, which is not a self-control problem. In our framework, agents that are more skilled at forecasting (e.g., those with more intelligence) exhibit less discounting. Agents with more domain-relevant experience exhibit less discounting. Older agents exhibit less discounting (except those with cognitive decline). Agents who are encouraged to spend more time thinking about an intertemporal tradeoff exhibit less discounting. Agents who are unable to think carefully about an intertemporal tradeoff – e.g., due to cognitive load – exhibit more discounting. In our framework, patience is highly unstable, fluctuating with the accuracy of forecasting.

*xgabaix@fas.harvard.edu, dlaibson@harvard.edu. For very good research assistance we are grateful to Omeed Maghziyan, and for helpful comments to Ned Augenblick, Roland Bénabou, Eric Budish, Stefano Dellavigna, Cary Frydman, Sam Gershman, Yoram Halevy, Alex Imas, Emir Kamenica, Ben Lockwood, Steve Morris, Muriel Niederle, Wolfgang Pesendorfer, Matthew Rabin, Daphna Shohamy, Bruno Strulovici, Dmitry Taubinsky, Harald Uhlig, Michael Woodford, and seminar participants at Berkeley, Chicago, Columbia, Edinburgh, Harvard, Hebrew U., NBER, NYU, SITE, and Stirling. We are grateful to the Pershing Square Fund for Research on the Foundations of Human Behavior for financial support.

1 Introduction

Most people appear to act *as if* they have a strong preference for earlier rewards over later rewards. For the last century economists have usually assumed that this type of behavior reflects (fundamental) time preferences, which economists model with discount factors that multiplicatively weight utils. If the one-period-ahead discount factor is δ , then δ utils experienced now are as valuable as one util experienced next period. If $\delta < 1$, economic agents prefer a current util to a delayed util.

However, such time preferences are only one of many ways to explain the empirical regularity that intertemporal choices are characterized by *declining sensitivity as utils are moved further away in time*. Diminishing sensitivity to future utils is also explained by imperfect information. For example, Böhm-Bawerk (1890) wrote that “we possess inadequate power to imagine and to abstract, or that we are not willing to put forth the necessary effort, but in any event we limn a more or less incomplete picture of our future wants and especially of the remotely distant ones. And then, there are all of those wants that never come to mind at all.” Pigou (1920) similarly observed “that our telescopic faculty is defective, and that we, therefore, see future pleasures, as it were, on a diminished scale. That this is the right explanation is proved by the fact that exactly the same diminution is experienced when, apart from our tendency to forget ungratifying incidents, we contemplate the past.”¹ Pigou believed that our imperfect ability to forecast the future mirrors our imperfect ability to recall the past.

To gain intuition for the role of imperfect forecasting, consider a driver who sees an upcoming pothole and estimates that it is small. A few moments later, she realizes that the pothole is large, but it is too late to avoid hitting it and damaging her car. Striking the large pothole is likely a reflection of imperfect foresight, not procrastination or laziness. In this case, large delayed consequences are misperceived by an imperfectly farsighted driver. We probably wouldn’t infer that the driver didn’t care about the impending impact because it was in the “future.” If the driver had foreseen the consequences, she would have braked earlier. In general, people will not respond optimally to future consequences that they do not fully anticipate.

Likewise, consider a sailor who sees a few clouds forming on the horizon and doesn’t

¹For a review of the history of theories of discounting see Loewenstein (1992).

immediately take the costly action of charting a new course. When her vessel is lashed by a violent storm the next day, it is not clear whether she was lazy the previous night or just mistaken in her forecast about the upcoming weather.

Decision-making is rife with situations in which a current action/inaction causes a stream of current and future consequences, many of which are hard to foresee. If delayed consequences are typically harder to foresee than immediate consequences, then decision-makers will appear to be impatient.

The role of imperfect information is also apparent in the seemingly impatient behavior of non-human animals. When monkeys are given an abstract intertemporal choice task on a computer, they act as if they discount delayed rewards at the rate of 10% *per second*. When the same monkeys are given a temporally analogous foraging task (also presented on a computer screen), the monkeys show very little discounting (Blanchard and Hayden 2015). Animal behavior appears to be impatient in completely novel domains and patient in domains that are evolutionarily relevant. As Blanchard and Hayden (2015) conclude, “Seemingly impulsive behavior in animals is an artifact of their difficulty understanding the structure of intertemporal choice tasks.”

In the current paper, we argue that behavior arising from imperfect foresight is hard to distinguish from behavior arising from time preferences. We study a Bayesian decision-maker with *perfectly* patient time preferences who receives noisy signals about the future. The resulting signal-extraction problem leads the Bayesian agent to behave in a way that is easy to misinterpret as a time preference; we call this seemingly impatient behavior *as-if discounting*. Our analysis shows that lack of foresight generates behavior that has most of the same characteristics of behavior that arises from deep time preferences. In other words, a perfectly patient Bayesian decision-maker who receives noisy signals about the future will behave as if she has time preferences.

Ophthalmic myopia arises when people cannot clearly see distant objects. But myopia also means a “lack of foresight or discernment.”² Such forecasting limitations matter when agents need to judge the value of events that will occur at a temporal distance. In this paper, we show that imperfect foresight – i.e., myopia – generates as-if discounting, even when the actors’ true preferences are perfectly patient. More generally, we show that imperfect foresight makes agents appear to behave more impatiently than implied by their deep time

²Merriam-Webster.

preferences.

Our formal model assumes that agents receive noisy, unbiased signals about future events and combine these signals with their priors to generate posterior beliefs about future events. Our key assumption is that the forecasting noise increases with the horizon of the forecast. We give special attention to the case in which the variance of the forecasting noise rises linearly with the forecasting horizon.

We provide an illustrative example of our framework in Section 2, where we study a binary choice problem: an actor chooses between an early reward and a mutually exclusive later reward. We show that when the variance of forecasting noise rises linearly with the event horizon, Bayesian agents will act as if they are hyperbolic discounters, even though their deep time preferences are perfectly patient.

In Section 3, we describe the broader implications of our framework, and identify predictions that distinguish our framework from time preference models. First, we show that our (perfectly patient) agents exhibit preference reversals of the same kind that are exhibited by agents with hyperbolic discount functions. However, these preference reversals do *not* reflect a self-control problem. The preference reversals arise because the agents obtain less noisy information with the passage of time. Accordingly, our agents do not wish to commit themselves; they act as-if they are naive hyperbolic discounters (Strotz 1956, Akerlof 1991, O’Donoghue and Rabin 1999a, 1999b) rather than sophisticated ones (Laibson 1997).

In the cross-section, our framework implies that agents with greater intelligence exhibit less as-if discounting—their superior forecasting ability enables them to make choices that are more responsive to future utility flows.

In addition, our agents exhibit as-if discounting that is domain specific. They exhibit less as-if discounting (i) when they have more overall life experience, (ii) when they are more experienced in the specific choice domain, (iii) when they have more time to think about an intertemporal choice (e.g., Imas, Khun, and Mironova, 2017), and (iv) when they have more cognitive bandwidth to think about their choice (e.g., Benjamin, Brown, and Shapiro 2013).

In Section 4, we generalize our example by making the action set continuous. We provide sufficient conditions that imply that perfectly patient agents who are imperfect forecasters will act as if they are naive hyperbolic discounters.

In Section 5, we discuss connections between our framework and related literatures on myopia, Bayesian cognition, risk, and discounting. Section 6 concludes.

2 A Basic Case: Binary Choice

Our approach can be explained with a simple example of a binary choice. Consider an agent at time zero, who must choose (irreversibly) between two mutually exclusive rewards: *Early* and *Late*. Reward *Early* would be experienced at date, $t \geq 0$. Reward *Late* would be experienced at date, $t + \tau > t$ (i.e., $\tau > 0$). The agent doesn't know the *true* value of *Early* and *Late*, respectively denoted, u_t and $u_{t+\tau}$. To simplify exposition and without loss of generality, we assume that these utility events are deterministic, though they were originally generated from a prior distribution that we will characterize below. (Note that any non-deterministic, zero-mean component is irrelevant because we are operating in utility space and we assume that our agents have classical expected utility preferences.)

Although the agent doesn't know the value of u_t and $u_{t+\tau}$, the agent can mentally simulate these deterministic rewards and thereby generate unbiased signals of their value:

$$s_t = u_t + \varepsilon_t$$
$$s_{t+\tau} = u_{t+\tau} + \varepsilon_{t+\tau}.$$

In the first equation, u_t is the true value of the *Early* utils and ε_t is the simulation noise. In the second equation, $u_{t+\tau}$ is the true value of the *Late* utils and $\varepsilon_{t+\tau}$ is the associated simulation noise. For tractability, we assume that the simulation noise is Gaussian. To simplify exposition, we assume that the correlation between ε_t and $\varepsilon_{t+\tau}$ is zero.³

2.1 Simulation Noise

We assume that the longer the horizon, the greater the variance of the simulation noise. Intuitively, the further away the event, the harder it is to accurately simulate the event's

³Here we suppose that the decision is made at time 0. If it is made (or revised) at decision time $d \leq t$, our analysis remains the same under the following benchmark case. Suppose that at the decision time, d , the agent contemplating consumption at time $t \geq d$ receives a signal $s_{d,t} = u_t + \sum_{k=d}^t \eta_{k,t}$ where $\eta_{k,t}$ are jointly independent, mean zero Gaussian variables with variance v_{t-k} (so that when $d = 0$, then $\varepsilon_t = \sum_{k=0}^t \eta_{k,t}$). This implies that the signal received at time d is a sufficient statistic for the decision that the agent faces. In particular, the history of signals received before d does not add incremental information to the signal received at d . Hence, our analysis goes through independently of the quality of the memory of our agent: only the current signal matters for the relevant decision.

utility. Because our set-up assumes that $t < t + \tau$, this assumption implies that

$$\text{var}(\varepsilon_t) < \text{var}(\varepsilon_{t+\tau}). \quad (1)$$

We will also sometimes assume that

$$\lim_{t \rightarrow \infty} \text{var}(\varepsilon_t) = \infty,$$

however this property is not necessary for our qualitative results.

We will pay particular attention to the special case of simulation noise that has a variance that is proportional to the simulation horizon:

$$\text{var}(\varepsilon_t) = \sigma_{\varepsilon_t}^2 = \sigma_\varepsilon^2 t \quad (2)$$

$$\text{var}(\varepsilon_{t+\tau}) = \sigma_{\varepsilon_{t+\tau}}^2 = \sigma_\varepsilon^2 (t + \tau). \quad (3)$$

This linearity assumption engenders a specific (hyperbolic) functional form in the analysis that follows. But this linearity assumption is not necessary for our qualitative results. We provide a complete characterization of noise functions below: i.e., necessary and sufficient conditions for the noise function to generate as-if discounting with declining discount rates as the horizon increases. The case of linear variance is a special case in this larger class of noise functions.

2.2 Bayesian Priors and Posteriors

The agents in our model combine Bayesian priors with their signals (s_t and $s_{t+\tau}$) to generate a Bayesian posterior. We model the Bayesian prior over utility events (in whatever class of events we are studying) as a Gaussian density with mean μ and variance σ_u^2 :

$$u \sim \mathcal{N}(\mu, \sigma_u^2). \quad (4)$$

Here μ is the average value in this class of utility events (e.g., visits to Philadelphia), whereas σ_u^2 is the overall variance within the class (e.g., some trips are great–Philadelphia in June–and some trips are much less great–Philadelphia in January).

In the appendix, we derive the agent's Bayesian posterior distribution of u_t , which is generated by combining her prior (4) and her signal s_t :

$$u_t \sim \mathcal{N} \left(\mu + \frac{s_t - \mu}{1 + \frac{\sigma_{\varepsilon_t}^2}{\sigma_u^2}}, \left(1 - \frac{1}{1 + \frac{\sigma_{\varepsilon_t}^2}{\sigma_u^2}} \right) \sigma_u^2 \right). \quad (5)$$

We summarize these results with the following proposition.

Proposition 1 *If the agent generates a mental simulation s_t , then her Bayesian posterior will be*

$$u_t \sim \mathcal{N} (\mu + D(t)(s_t - \mu), (1 - D(t)) \sigma_u^2),$$

where

$$D(t) = \frac{1}{1 + \frac{\sigma_{\varepsilon_t}^2}{\sigma_u^2}}, \quad (6)$$

the variance of her simulation noise is $\sigma_{\varepsilon_t}^2$ and her prior distribution is $u \sim \mathcal{N}(\mu, \sigma_u^2)$.

For reasons that will become apparent below (see Proposition 4 in particular), we refer to $D(t)$ as the agent's as-if discount function. Because we assume that simulation noise, $\sigma_{\varepsilon_t}^2$, is increasing in t , $D(t)$ is decreasing in t , which is a standard property of a discount function. If $\lim_{t \rightarrow \infty} \text{var}(\varepsilon_t) = \infty$, then $\lim_{t \rightarrow \infty} D(t) = 0$, another common property of a discount function. In this case, the posterior expectation of u_t converges to the mean of the prior as the horizon increases. In notation,

$$\lim_{t \rightarrow \infty} \mathbb{E}_0[u_t | s_t] = \mu.$$

It is helpful to integrate posteriors over agents in the economy. We assume that the signals s_t are unbiased, so they are equal to u_t on average. Accordingly, the average forecast of u_t will be

$$\int_{s_t} \mathbb{E}_0[u_t | s_t] dF(s_t | u_t) = \mu + D(t)(u_t - \mu). \quad (7)$$

In general, the mean of the prior will be less extreme than the actual values of u_t . To model this statistical property, consider the illustrative case in which the prior is approximately equal to zero. (We will relax this restriction later.) Under this restriction, the average belief

is

$$\int_{s_t} \mathbb{E}_0[u_t | s_t] dF(s_t | u_t) = D(t)u_t.$$

We now have an expression that looks like a discounted utility framework: $D(t)$ is a decreasing function and it multiplies the actual utility value u_t .

2.3 Hyperbolic As-if Discounting

We explore a benchmark case: noise that is linear in the horizon.

Proposition 2 *When we assume that $\text{var}(\varepsilon_t) = \sigma_{\varepsilon_t}^2 = \sigma_{\varepsilon}^2 t$, we obtain hyperbolic as-if discounting:*

$$D(t) = \frac{1}{1 + \alpha t} \tag{8}$$

where

$$\alpha = \frac{\sigma_{\varepsilon}^2}{\sigma_u^2}, \tag{9}$$

which is the (one-period) noise-to-signal variance ratio.⁴

The discount function, $D(t) = \frac{1}{1 + \alpha t}$, implies an instantaneous discount rate

$$\text{discount rate} = -\frac{D'(t)}{D(t)} = \frac{\frac{\alpha}{(1 + \alpha t)^2}}{\frac{1}{1 + \alpha t}} = \frac{\alpha}{1 + \alpha t}.$$

At horizon 0, the as-if discount rate is α . The as-if discount rate falls with t . As $t \rightarrow \infty$, the as-if discount rate converges to 0.

2.4 An Example When the Mean Prior Is not Zero ($\mu \neq 0$)

As we noted above, actual utility events will tend to be more extreme than priors. To capture this property, we previously set the mean of the prior distribution equal to zero: $\mu = 0$. We now relax this restriction and illustrate the general case with an example in which the mean of the prior distribution is $\mu = 1$. For this example, we assume that the simulation variance

⁴There is a trivial generalization of this proposition. When the variance of the forecasting noise is affine in the time horizon, so that $\text{var}(\varepsilon_t) = \sigma_{\varepsilon_t}^2 = \kappa + \sigma_{\varepsilon}^2 t$, we also obtain hyperbolic as-if discounting:

$$D(t) = \frac{1}{1 + \theta + \alpha t},$$

where $\theta = \frac{\kappa}{\sigma_u^2}$.




Figure 1: Plot of the average perceived value \bar{u}_t , given for three different true utilities u_t ($u_t \in \{-9, 1/2, 11\}$), as a function of the time horizon t . This average perceived value is: $\bar{u}_t = \mu + \frac{u_t - \mu}{1 + \frac{\sigma_\varepsilon^2}{\sigma_u^2} t}$ (see equation (7)). The figure uses $\sigma_\varepsilon^2 / \sigma_u^2 = 0.1$.

is linear in the time horizon and the variance ratio is $\frac{\sigma_\varepsilon^2}{\sigma_u^2} = 0.1$. Figure 1 plots the population level expectations of u_t for three values of u_t (holding the mean of the prior distribution fixed at $\mu = 1$):

$$u_t = \mu + 10 = 11$$

$$u_t = \mu - 1/2 = 1/2$$

$$u_t = \mu - 10 = -9.$$

When the three utility events are in the present ($t = 0$), the three expectations are equal to the true value of each utility event, respectively 11, 1/2, and -9. However, as the three utility events recede into the distant future, the three expectations revert to the mean of the prior, $\mu = 1$. This discounting towards the mean of the prior is hyperbolic because we are assuming linear variance (see subsection 2.3).

The $u_t = 11$ curve is characterized by standard discounting. The further ahead the utility event is shifted, the lower the perceived value of the event. The $u_t = -9$ curve is also characterized by standard discounting on most of its domain. As the event is moved further into the future, its value declines toward zero. However, at $t = 90$, the perceived value crosses the x -axis and continues asymptotically toward $\mu = -1$. Finally, the $u_t = 1/2$

line displays anti-discounting. The further the value is moved into the future, the higher its perceived value, as it asymptotes to the prior mean of $\mu = 1$.

These three lines illustrate the three types of cases that arise in our framework, including the special case of anti-discounting. Note that anti-discounting arises when the true value of u_t lies between 0 and the mean of the prior distribution, μ .⁵

2.5 Another Example: Potholes

Return to the example from the introduction of the paper: a driver who strikes a large pothole that she could have avoided (with perfect foresight). We now map this intuitive example to the notation of our model.

Assume that the driver perceives a choice between (i) adjusting her car's path at an immediate payoff of $u_0 = -2$ when she first imperfectly perceives a *distant* pothole, or (ii) staying on her original course and striking the pothole at a stochastic payoff of u_τ . The Bayesian driver estimates that the expected cost of striking the pothole is

$$\mathbb{E}_0[u_\tau \mid s_\tau] = \mu + \frac{1}{1 + \frac{\sigma_\varepsilon^2}{\sigma_u^2}}(s_\tau - \mu).$$

If almost all potholes are small, then μ is close to zero and σ_u^2 is close to zero. Because potholes are hard to see from a distance, σ_ε^2 is large. Accordingly, her estimate of the damage from the pothole when she first spies it at a distance will be close to zero (even if s_τ is large in magnitude). For example, suppose the true damage from the pothole is $u_\tau = -4$. Assume that $\mu = -1$, $\sigma_u^2 = 1$, and $\sigma_\varepsilon^2 = 9$. Then drivers will on average estimate that striking the pothole will generate a payoff of

$$\mathbb{E}_0[u_\tau \mid s_\tau] = -1 + \frac{1}{1 + \frac{9}{1}}[-4 - (-1)] = -1.3.$$

Accordingly, the driver chooses to stay the course, leading her to strike the pothole with a payoff of -4 when she could have avoided the pothole (with perfect foresight) at a payoff of

⁵There are three reasons why the prior, μ , will often be close to 0. First, good and bad hedonic events are both frequent occurrences, so on average $\mu \approx 0$ is a natural benchmark. Second, there are opportunity costs: even pleasant events can have negative net value, given some opportunity cost (much as in a job search model, where most job offers should be declined as they have a negative net value once the opportunity cost has been incorporated). Third, suppose that the valuation is experienced as a difference relative to a reference point which is the ex ante value of the object: then by construction $\mu = 0$.

-2. Naturally, staying on her original course is a rational choice given the information that she had when she first saw the pothole at a distance.

2.6 Probabilistic Choice

Our framework implies that choice is probabilistic, because agents receive noisy signals about the value of future rewards. In our example, the agent chooses *Early* if and only if

$$D(t)s_t \geq D(t + \tau)s_{t+\tau},$$

where $D(t)$ is the as-if discount function introduced above and s_t and $s_{t+\tau}$ are the (unbiased) signals of the respective values of the *Early* and *Late* rewards.

From the perspective of an observer who knows the values of u_t and $u_{t+\tau}$, the probability that the agent chooses the *Early* reward is

$$\begin{aligned} \mathbb{P}(\text{choose } \textit{Early}) &= \mathbb{P}[D(t)(u_t + \varepsilon_t) \geq D(t + \tau)(u_{t+\tau} + \varepsilon_{t+\tau})] \\ &= \Phi\left(\frac{1}{\Sigma}[D(t)u_t - D(t + \tau)u_{t+\tau}]\right), \end{aligned} \quad (10)$$

where Φ is a Gaussian CDF and Σ is a scaling factor:

$$\Sigma = \sqrt{D(t)^2 \text{var}(\varepsilon_t) + D(t + \tau)^2 \text{var}(\varepsilon_{t+\tau})}. \quad (11)$$

and for simplicity we assume $\mu = 0$ in this subsection. This restriction is innocuous as long as the two μ values are the same (for the earlier and later rewards).⁶

This probabilistic choice function has natural properties. If $t = 0$ (i.e., the *Early* reward is an immediate reward), then,

$$\begin{aligned} \mathbb{P}(\text{choose } \textit{Early}) &= \mathbb{P}[u_0 \geq D(\tau)(u_\tau + \varepsilon_\tau)] \\ &= \Phi\left(\frac{1}{\Sigma}[u_0 - D(\tau)u_\tau]\right). \end{aligned}$$

If we let the time delay between the *Early* reward and the *Late* reward go to infinity (i.e.,

⁶Replace u_t by $u_t - \mu$ as needed.

$\tau \rightarrow \infty$), then

$$\lim_{\tau \rightarrow \infty} \mathbb{P}(\text{choose } \textit{Early}) = 1_{u_0 > 0}.$$

This implies that the agent chooses the *Early* reward with probability one if three properties hold: (i) the *Early* reward is available immediately ($t = 0$), (ii) the *Late* reward is available arbitrarily far in the future ($\tau \rightarrow \infty$), and (iii) the *Early* reward is strictly positive ($u_0 > 0$). In other words, the agent behaves as if she places no value on the (infinitely) delayed *Late* reward.

Now assume that the *Early* reward is available with some delay, so that $t > 0$ (i.e., the *Early* reward is not immediate), then

$$\begin{aligned} \lim_{\tau \rightarrow \infty} \mathbb{P}(\text{choose } \textit{Early}) &= \mathbb{P}[u_t + \varepsilon_t > 0] \\ &= \Phi\left(\frac{u_t}{\Sigma}\right). \end{aligned}$$

Accordingly, if the *Late* reward is available arbitrarily far in the future ($\tau \rightarrow \infty$), then the agent chooses the *Early* reward with the same probability that she perceives the *Early* reward to have positive value. Once again, the agent behaves as if she places no value on the (infinitely) delayed *Late* reward.

2.7 Preference Reversals without Commitment

In our setting, an observer who knows the values of u_t and $u_{t+\tau}$ will be able to predict (probabilistic) preference reversals. For example, consider the case of linear variances. In addition, assume that $u_{t+\tau} > u_t > 0$, and

$$u_t > D(\tau)u_{t+\tau}.$$

When the two options are sufficiently far in the future (large t), a majority of agents (if forced to choose) will prefer *Late* over *Early*, because

$$\mathbb{P}(\text{choose } \textit{Early}) = \Phi\left(\frac{1}{\Sigma} [D(t)u_t - D(t + \tau)u_{t+\tau}]\right) < \frac{1}{2}. \quad (12)$$

To see this, note that

$$D(t)u_t - D(t + \tau)u_{t+\tau} = \frac{u_t}{1 + \alpha t} - \frac{u_{t+\tau}}{1 + \alpha(t + \tau)}.$$

For sufficiently large values of t , $u_{t+\tau} > u_t$ implies,

$$\frac{u_t}{1 + \alpha t} - \frac{u_{t+\tau}}{1 + \alpha(t + \tau)} < 0.$$

However, with the passage of time, all agents will eventually choose *Early* because $u_t > D(\tau)u_{t+\tau}$. More precisely, if agents were not forced to choose in advance, but were instead given the chance to choose at time t , all would choose *Early*.

In many economic models, such preference reversals are a sign of dynamic inconsistency in preferences.⁷ That is not the case here. The agents in the current model have imperfect information, not dynamically inconsistent time preferences. Their externally predictable preference reversals are a result of their imperfect information. Accordingly, the agents in our model will not desire to limit their own choice sets. Preference reversals arise from their inference problems, not self-control problems.

2.8 More General Discounting Functions

We can provide necessary and sufficient conditions for the as-if discount function, $D(t)$, to exhibit decreasing impatience. In other words, we can derive necessary and sufficient conditions for the property that the instantaneous as-if discount rate

$$\rho(t) := -\frac{D'(t)}{D(t)}$$

is decreasing in the horizon t .

Proposition 3 *The as-if discount function $D(t)$ exhibits strictly decreasing impatience at*

⁷See McGuire and Kable (2012, 2013) for a setting in which preference reversals arise because of rational learning dynamics. If a delayed reward that was probabilistically expected does not arrive after a period of waiting, the agent infers that the hazard rate of arrival is low and further waiting is not likely to pay off, and therefore reverts to choosing the immediate reward.

time horizon t if and only if

$$\frac{d^2\sigma_{\varepsilon_t}^2}{dt^2} - \frac{\left(\frac{d\sigma_{\varepsilon_t}^2}{dt}/\sigma_u^2\right)^2}{\left(1 + \frac{\sigma_{\varepsilon_t}^2}{\sigma_u^2}\right)} < 0.$$

This proposition is proved in the appendix. Because we assume that $\frac{d\sigma_{\varepsilon_t}^2}{dt} > 0$, this proposition yields an immediate corollary.

Lemma 1 *The as-if discount function $D(t)$ exhibits strictly decreasing impatience if the variance of simulation noise, $\text{var}(\varepsilon_t) = \sigma_{\varepsilon_t}^2$, is a weakly concave function of time.*

Accordingly, our model generates as-if discount rates that decrease as the horizon increases in many cases. We next study a boundary case.

Exponential As-if Discounting Our framework can also be reverse-engineered to generate exponential discounting as a special case. However, this requires assumptions on the variance function that we believe are heroic.

Lemma 2 *The as-if discount function $D(t)$ exhibits a constant discount rate, ρ , if and only if*

$$\sigma_{\varepsilon_t}^2 = [\exp(\rho t) - 1] \sigma_u^2.$$

Accordingly, the discount rate is exponential if and only if the simulation variance, $\sigma_{\varepsilon_t}^2$, rises exponentially. This Lemma is proven by setting

$$D(t) = \frac{1}{1 + \frac{\sigma_{\varepsilon_t}^2}{\sigma_u^2}} = \exp(\rho t).$$

This sort of cognitive discounting is useful because of the tractability it induces (see for instance Gabaix 2016a, 2016b).

3 Implications

We now discuss the key predictions of our model, emphasizing several ways that our model of myopia differs from other models in the intertemporal choice literature. As discussed above, our myopic agent acts as if she is maximizing a utility function with an as-if discount

function, $D(\tau)$, where

$$D(\tau) = \frac{1}{1 + \frac{\sigma_{\varepsilon\tau}^2}{\sigma_u^2}}.$$

When the variance of the forecasting noise is weakly concave in the simulation horizon, the discounting function is characterized by an instantaneous discount rate that falls with the horizon. When the forecasting noise is linear in the simulation horizon, so that $\sigma_{\varepsilon\tau}^2 = \sigma_{\varepsilon}^2\tau$ then the discount function is hyperbolic:

$$D(\tau) = \frac{1}{1 + \frac{\sigma_{\varepsilon}^2}{\sigma_u^2}\tau}.$$

These as-if discounting functions arise because of the imperfect information that the agent has when she generates forecasts. If she were asked to describe her preferences, she would say that she has no time preferences. In other words, she is *trying* to maximize

$$\sum_{\tau=0}^{T-t} u(a_{t+\tau}).$$

Her as-if discounting behavior arises because she doesn't have perfect foresight regarding the future utility flows $u(a_{t+\tau})$.

The next proposition formalizes the sense in which she “appears” to be discounting the future. In this proposition we introduce a “confused social scientist,” who uses the choice data of agents as described in this paper. But he doesn't understand the true structural model. He thinks that the agent chooses accordingly to a true discount function $\widehat{D}(t)$, which he is trying to estimate. To model noise, the scientist uses a Probit model (i.e., a random choice model with Gaussian distributed noise). The next proposition shows that he will estimate a discount function equal to our (normalized) as-if discount function $D(t)/D(0)$.⁸

Proposition 4 *Consider a confused social scientist modelling the agent described in this paper. The scientist fits a Probit model assuming that the agent has discounted utility $\sum_{t=0}^T \widehat{D}(t) \widehat{u}_t$ (with \widehat{u}_t and $\widehat{D}(t)$ unknown, normalizing $\widehat{D}(0) = 1$). His dataset is a collection of probabilities of choosing the Early reward (as in (10)), with at least two different consumptions per date (but potentially a finite set of dates \mathcal{T}), and with choices on lotteries*

⁸The normalization is necessary for the special case in which $D(0)$ is not equal to one. This case arises when $\sigma_{\varepsilon_0}^2$ is not equal to zero (in other words, when utils that are going to be *immediately* experienced are still at least partially noisy in their forecasted value).

with payouts at time 0. This social scientist will estimate a discount function $\widehat{D}(t)$ equal to our (normalized) as-if discount function $D(t)/D(0)$, even though the agent doesn't actually discount utils.

The proof (in the appendix) rests on the fact that the fitted Probit model predicts exactly the probability distribution of choices made by the agent, given by equation (12).⁹

3.1 Absence of Commitment

The agents in this model have a forecasting problem, not a self-control problem. Therefore they are never willing to reduce their choice set (unless they are paid to do so). This absence of a willingness to pay for commitment may explain the lack of commitment technologies in markets. In real markets there is little commitment for commitment's sake.¹⁰ Personal trainers and website blocking apps are frequently mentioned exceptions, but such technologies are not commonly used.

By contrast, economists have been able to elicit commitment in experiments (see Cohen, Ericson, Laibson, and White forthcoming for a review). However, most of these experiments elicit only a weak taste for a commitment and little or no willingness to pay for commitment (e.g., Augenblick, Niederle, and Sprenger 2015, Sadoff, Samek, and Sprenger 2016).

Our myopia model predicts that agents will exhibit as-if hyperbolic discounting with preference reversals *and* no willingness to pay for commitment. In this sense, our model reproduces the predictions of the standard hyperbolic discounting model with naive beliefs (see O'Donoghue and Rabin 1999a, 1999b, 2001, Laibson 2015, Ericson forthcoming). However, it also generates further implications, to which we now turn.

3.2 Intelligence is Associated with Less As-if Discounting

Our model predicts that agents with less forecasting noise will exhibit less as-if discounting. Because of this mechanism, agents that are more intelligent will exhibit less as-if discounting.¹¹

⁹In addition, the scientist recovers the correct utility function, up to an affine transformation.

¹⁰However, there is a great deal of ancillary commitment, like mortgage contracts, which create a forced savings system as a by-product of a stream of loan/principal repayments.

¹¹The underlying assumption is that more intelligent agents simulate the future with less noise—for instance because they generate more simulations. If they run n simulations, the variance of the average simulation will decrease by a factor of $1/n$, so it will be lower.

To see this formally, let H represent human capital and assume that the variance of forecasting noise is declining in human capital:

$$\frac{d\sigma_\varepsilon^2(H)}{dH} < 0.$$

The as-if discount rate is given by

$$-\frac{D'(t)}{D(t)} = \frac{\alpha}{1 + \alpha t},$$

where

$$\alpha = \frac{\sigma_\varepsilon^2(H)}{\sigma_u^2}.$$

The as-if discount rate is increasing in $\sigma_\varepsilon^2(H)$, so as-if discounting is decreasing in human capital, H .

The available evidence supports this prediction. Measured discount rates are negatively correlated with scores on IQ tests: see Benjamin, Brown, and Shapiro (2013), Burk et al. (2009), Shamosh and Gray (2008). Indeed, such effects also arise across species. Tobin and Logue (1994) show that patience increases as the study population switches from pigeons, to rats, to humans.

3.3 Myopia Is Domain Specific

These comparative statics on cognitive function generate a wider set of predictions when forecasting ability varies across domains. For example, our framework predicts that agents with more domain-relevant experience, and hence better within-domain forecasting ability, will exhibit less discounting. Read, Frederick and Scholten (2013) report that people exhibit more patience when an intertemporal choice is posed as an investment rather than a (seemingly novel) money-now-vs-money-later decision. Relatedly, recall our earlier discussion of the monkey experiments reported by Blanchard and Hayden (2015): when an intertemporal choice is presented as a reward-now-vs-reward-later decision, monkeys choose far more impatiently than they do when a foraging problem is used to frame the intertemporal tradeoffs.

Likewise, our framework predicts that older agents – who generally have more life experience and consequently better forecasting skills – will exhibit less discounting. This prediction

is supported by Green, Fry, and Myerson (1994). Relatedly, Addessi et al. (2014) show that replacing one-for-one representations of future reward with more abstract one-for-many representations of the same future rewards, leads capuchin monkeys and (human) children to exhibit more impatience. In contrast, adults, who have more experience using abstract symbols, do *not* behave more impatiently when one-for-one representations of future reward are replaced with one-for-many representations. The Addessi et al. (2014) experimental evidence implies that childhood impatience is due, at least in part, to children’s less developed ability to cognitively represent (abstract) future rewards. Our framework also predicts that people who experience cognitive decline (e.g., due to normal aging) will exhibit more discounting, which is supported by evidence from James et al. (2015).

Our framework predicts that agents who are unable to think carefully about an intertemporal tradeoff – e.g., due to a cognitive load manipulation or the effects of alcohol – will exhibit more discounting. Steele and Josephs (1990), Shiv and Fedorikhin (1999), Hinson, Jameson, and Whitney (2003), and Benjamin, Brown, and Shapiro (2013) document such effects. This prediction is closely related to the theory of cognitive scarcity: see Spears (2012), Mullainathan and Shafir (2013), and Schilbach, Schofield, and Mullainathan (2016).

Our framework predicts that agents who are encouraged to spend more time thinking about a future tradeoff will exhibit less discounting. Imas, Kuhn, and Mironova (2017) robustly measure such an effect experimentally. In their experiment, some subjects decide at time 0 how to divide an effort task between time 0 and time t . Other subjects are given a preceding hour to decide how to divide the effort task between time 0 and time t . Subjects in the latter condition choose more patiently: their measured discount rate is 16 percentage points lower. The authors also find additional evidence that the additional decision time has little impact on decisions outside the domain of the task allocation and that, in the treatment that mandated a longer decision time, subjects measured to have higher information processing capacities are relatively more patient.

Our framework also predicts that rewards delivered in future periods that are cognitively well-simulated will exhibit less discounting. Peters and Büchel (2010) exogenously manipulate the salience of various future periods and find that higher salience from imagery of future reward periods increases the value of rewards delivered during those periods.

Finally, our framework predicts that discounting behavior will only be weakly correlated across domains because discounting is not a domain general preference, but rather the

result of imperfect forecasting that will naturally vary across domains. Chapman (1996) and Chabris et al. (2008) document the low level of correlation in discount rates that are measured in different decision-making domains.

4 Extension to a Continuous Action

Until now we have studied the case in which the agent has two mutually exclusive actions: choose an *Early* reward or a *Late* reward. We now generalize the action space to a continuum. We then provide sufficient conditions that enable us to apply our framework to a general, multi-period intertemporal choice problem.

The upshot of this section is that the economics of the binary action case still goes through, though with more complex mathematics.

4.1 Modelling How Agents Observe with Noise a Whole Utility Function

Suppose that an action a leads to a true payoff $u(a)$. However, the agent observes this noisily: we suppose that the agent observes the “noised up” version of the utility function:

$$\mathbf{s} = (s(a))_{a \in \mathcal{A}} \tag{13}$$

of the whole function $u = (u(a))_{a \in \mathcal{A}}$, where $\mathcal{A} = [\underline{a}, \bar{a}]$ is the support of the action, which is assumed to contain 0 (this is just a normalization). This noised-up version is assumed to take the form:

$$s(a) = u(a) + \sigma_{\varepsilon_t} W(a) + \chi \sigma_{\varepsilon_t} \eta_0 \tag{14}$$

for all $a \in \mathcal{A}$. There is a continuous noise $W(a)$, modelled as standard Brownian motion with $W(0) = 0$ except that W is “two-sided”, i.e. runs to the left and right of 0.^{12,13} The noise is modelled as proportional to σ_{ε_t} when the utility is seen from a distance t . For instance, the linear case is $\sigma_{\varepsilon_t} = \sigma_{\varepsilon} \sqrt{t}$. The term $\chi \sigma_{\varepsilon_t} \eta_0$ ensures that the value at $a = 0$ is also perceived with noise ($\sigma_{\eta_0} = 1$, χ is a parameter).

¹²Formally, $W(x)_{x \geq 0}$ and $W(-x)_{x \geq 0}$ are independent Brownian motions.

¹³See Callender and Hummel (2014) for a recent model using inference on Brownian paths, though with a signal structure different from ours.

Given this perceived curve \mathbf{s} , what’s his posterior about $u(a)$? We will see the under the “right” assumptions (to be specified soon), we simply have

$$\mathbb{E}[u(a) \mid \mathbf{s}] = D(t) s(a)$$

with the same $D(t)$ as in the binary case. This means that the representative agent just dampens the true function.

4.2 Assumptions for Our Result

Here we detail the assumptions we use for the results. The reader may wish to skip to the result itself, in the next subsection 4.3.

Assumption 1 (Wiener decomposition) *We suppose that function $v(\cdot) := \frac{u(\cdot)-u(0)}{\sigma_u}$ is perceived as drawn from the Wiener measure, and $u(0)$ is drawn as $u(0) \sim N(0, \chi^2 \sigma_u^2)$ independent of v . We call*

$$D(t) = \frac{1}{1 + \frac{\sigma_{\varepsilon_t}^2}{\sigma_u^2}} \tag{15}$$

where σ_{ε_t} is as in (14).

Let us state this assumption in more user-friendly language. The value of $u(0)$ is also seen as random—and we index its randomness by χ . The rest of the function u , outside the intercept, is also random. To specify this, we set $v(a) := \frac{u(a)-u(0)}{\sigma_u}$, which is u normalized to have 0 intercept and standardized size (so that $\mathbb{E}[v(1)^2] = 1$). We view v a a “random function” drawn from a distribution. For simplicity, we consider that it’s drawn from the simplest distribution of random functions – the so-called Wiener measure (Brownian motions are typical instances of such functions).¹⁴ Basically, the assumption is that the component of $du(a)$ are drawn as i.i.d. normal increments, like a Brownian motion, with square width $\sigma_u^2 da$. Note that this refers to the distribution assumed by the agent when he performs his Bayesian inference, not necessarily the true distribution.

Section 7.2 of the appendix proposes a variant, Assumption 2, with polynomial utility, that uses more elementary mathematics, at the cost of heavier notations and proofs.

¹⁴We could imagine a number of variants, e.g. u'' would be drawn from this distribution; or, to keep u concave, we could have $\ln(-u'')$ drawn from that distribution. This becomes quickly more mathematically involved, so we leave this to a separate investigation and focus on what we view as the simplest case.

4.3 Perceived Utility Function Given the True Utility

We can now derive the utility perceived by the agent, given she agent sees the whole noised-up function \mathbf{s} (equation (14)).

Proposition 5 (Perceived utility for a continuous utility function) *Make Assumption 1 or 2. Then, the perceived utility is proportional to the signal:*

$$\mathbb{E}[u(a) | \mathbf{s}] = D(t) s(a)$$

where $D(t) = \frac{1}{1 + \frac{\sigma_{\varepsilon}^2 t}{\sigma_{\bar{u}}^2}}$. As a result, the average perceived utility $\bar{u}(a)$, defined as:

$$\bar{u}(a) := \mathbb{E}[\mathbb{E}[u(a) | \mathbf{s}] | \mathbf{u}]$$

satisfies:

$$\bar{u}(a) = D(t) u(a) \tag{16}$$

This means that the average perceived utility is $D(t) u(a)$ rather than plainly $u(a)$, exactly like in the simple two-action (consume or don't consume) case.

4.4 The Representative Agent Perspective

4.4.1 Assumptions for a Tractable Generalization

To cleanly study the dynamic problem, we assume the following (in addition to the assumptions of Proposition 5).¹⁵

- A1. The agent treats the noise at all simulation horizons as uncorrelated.
- A2. The agent has Gaussian priors with 0 mean (and no correlation between $u_t, u_{t+\tau}$).
- A3. The agent acts as if she won't learn new simulation information in the future.¹⁶

The notion of “average behavior” is potentially messy with non-linear utilities. Hence, we find it useful to define the following form of “representative agent” version of the model.

¹⁵There are many alternative ways to generate variances that are linear in the forecasting horizon, including new signals that contain all of the information of the old signals.

¹⁶This is the assumption of the “anticipated utility” framework of Kreps (1998) used also by Cogley and Sargent (2008).

We study the equilibrium path in which all simulation noise happens to be realized as zero (but the agent doesn't know this). In our illustrative example, this corresponds to $\varepsilon_t = 0$. For instance, we had $s_t = u_t + \varepsilon_t$ and $\mathbb{E}[u_t | s_t] = D(t)s_t$. The representative agent draws noise $\varepsilon_t = 0$, so for the representative agent, $\mathbb{E}[u_t | s_t] = D(t)u_t$.

Proposition 6 (Dynamic choices of the representative agent) *Assume that the agent has dynamically consistent preferences*

$$\sum_{t=0}^T u(a_t).$$

Then A1-A3 imply that at each time period $t \in \{0, \dots, T\}$ the representative agent acts as if she is trying to maximize

$$\sum_{\tau=0}^{T-t} D(\tau)u(a_{t+\tau})$$

where

$$D(\tau) = \frac{1}{1 + \frac{\sigma_{\varepsilon\tau}^2}{\sigma_u^2}}.$$

Corollary 1 *Assume that simulation variance is linear in the horizon of the simulation: $\sigma_{\varepsilon\tau}^2 = \tau\sigma_{\varepsilon}^2$. Then, at each time period $t \in \{0, \dots, T\}$ the representative agent acts as if she is trying to maximize*

$$\sum_{\tau=0}^{T-t} D(\tau)u(a_{t+\tau}),$$

where

$$D(\tau) = \frac{1}{1 + \alpha\tau}$$

$$\alpha = \frac{\sigma_{\varepsilon}^2}{\sigma_u^2}.$$

Proposition 6 shows that our basic results extend to arbitrary utility functions with continuous actions.

5 Literatures on Related Mechanisms

We now review other lines of research that are related to this paper and on which this paper builds. We review three literatures: models of myopia, Bayesian foundations of imperfect

and costly cognition, and risk-based models of as-if discounting (including risk-based models with probability distortions).

5.1 Myopia

Political economists, psychologists, and other social scientists have long posited that impatient behavior was due in part to imperfect foresight. These ideas were informally described by political economists, including Böhm-Bawerk (1890), and economists, including Pigou (1920), both of whom are quoted in the introduction of this paper.

These informal explanations have been joined by formal, mathematical definitions, models, and analyses of that incorporate various formulations of myopia. For example, Brown and Lewis (1981) provide an axiomatic definition of myopia. Feldstein (1985) evaluates the optimality of social security under the assumptions of myopia and partial myopia (modelled as a low discount factor in a two-period decision problem). Jéhiel (1995) studies two-player games in which players have limited forecasting horizons. Spears (2012) generate a forecasting horizon that is endogenous because forward-looking calculations are costly. Gabaix, Laibson, Moloche, and Weinberg (2006) report experimental evidence that supports a model in which agents choose an endogenous forecasting horizon at which the cognitive cost and estimated utility benefit of marginally increasing the forecasting horizon are equalized. This optimal forecasting framework generates a complex option value problem with respect to information acquisition (see also Fudenberg, Strack, and Strzalecki 2016).¹⁷

In the current paper, we assume that the agent has noisy signals about the future, which engenders Bayesian forecasts that have “myopic” properties: i.e., declining sensitivity to future events. When the noise is linear in the forecasting horizon, the as-if discounting takes a simple hyperbolic form. Accordingly, our paper introduces a tractable microfoundation for myopia.

¹⁷Neuroscientists have long hypothesized that perceptions of consequences are neurally represented as noisy unbiased signals of those consequences. For example, this is the foundational assumption of the large literature on the drift-diffusion model (e.g., Shadlen and Shohamy 2016). This framework includes a noise term analogous to our noise term. By implication, more noise in the drift-diffusion model would imply more discounting in our model.

5.2 Bayesian Models of Attention and Cognition

The current paper assumes that agents are Bayesian, which adopts the approach of early decision-theory pioneers like Raiffa and Schlaifer (1961). There is a growing body of literature (in economics, cognitive psychology, and neuroscience) that studies the effects of noisy perception and Bayesian inference, and uses this combination to explain seemingly suboptimal behaviors. One of the earliest examples is the work of Commons, Woodford, and Ducheny (1982), and Commons, Woodford, and Trudeau (1991) who use this approach to generate a theory of hyperbolic memory recall. In their framework, the noisy signals are memories, whereas the noisy signals in our model are simulations of the future. The literature on attention allocation assumes that agents have limited information, which is mathematically equivalent to the assumption that agents have noisy signals about the state of the world. Geanakoplos and Milgrom (1991), Sims (2003), Kamenica (2008), Woodford (2009), Gabaix (2014), Schwartzstein (2014), Hanna, Mullainathan, and Schwartzstein (2014), Allcott and Taubinsky (2015), Matejka, Steiner, and Stewart (forthcoming), Natenzon (2016), Taubinsky and Rees-Jones (2016a, 2016b), study agents who allocate their limited attentional bandwidth to the activities that they believe are the most valuable.¹⁸ Steiner and Stewart (2016) and Khaw, Li, and Woodford (2017) study an environment in which agents react to the noise in their probability perceptions by (optimally) distorting their perceived probabilities in a way that mimics the probability mapping in prospect theory (Kahneman and Tversky 1979).

Our paper adopts the approach that unifies the work above: noisy signals plus Bayesian inference jointly produce as-if behavior that appears to be imperfectly rational. Specifically, in our case, this combination generates as-if hyperbolic discounting.

5.3 Risk-Based Models of As-if Discounting

It has long been recognized that time preferences engender the same kind of behavior that is associated with risk or mortality (e.g., Yaari 1965). For example, if promised future rewards may be permanently withdrawn or lost at a constant hazard rate, ρ , then a perfectly patient decision-maker should be indifferent between 1 util at time zero and $\exp(\rho\tau)$ utils at a time τ .

¹⁸Another strand of the literature uses non-Bayesian rules to govern attention and salience (e.g. Bordalo, Gennaioli, and Shleifer 2012, 2013), though it might be probably be given some quasi-Bayesian interpretation.

In this example, risk induces a perfectly patient agent to appear to discount the future with exponential discount rate ρ .

This type of risk-based discounting can also produce hyperboloid discount functions under specific assumptions about a non-constant hazard rate (see Sozou 1998, Azfar 1999, Weitzman 2001, Halevy 2005, Dasgupta and Maskin 2005, Fernández-Villaverde and Mukherji 2006, Halevy 2014, 2015). For instance, Azfar, Sozou, and Weitzman all assume that the hazard rate that governs the withdrawal of rewards is itself drawn from a distribution and has a value that can only be inferred from the observed data. This assumption produces preferences that are characterized by a declining discount rate as the horizon increases—the more time that passes without a withdrawal, the more likely that one of the low hazard rates is the hazard rate that was drawn from the distribution at the start of time, implying a lower effective discount rate at longer horizons. Risk can also produce hyperboloid discount functions because of probability transformations that are characterized by a certainty effect, whereby a certain present reward is discretely more valuable than an even slightly uncertain delayed reward (see the non-expected utility frameworks in Prelec and Loewenstein 1991, Quiggin and Horowitz 1995, Keren and Roelofsma 1995, Weber and Chapman 2005, Halevy 2008, Epper, Fehr-Duda, and Bruhin 2011, Baucells and Heukamp 2012, Andreoni and Sprenger 2012, Epper and Fehr-Duda 2015, Chakraborty 2017).

Our model works off a related but different risk mechanism than those listed above. The uncertainty in our model is due to noise that is generated by the forecaster herself. For example, our mechanism predicts that an expert would exhibit little as-if discounting in her domain of expertise (she forecasts the future with little or no noise) while a non-expert would exhibit substantial as-if discounting in the *same* domain (she forecasts the future with relatively more noise than the expert). Likewise, our framework predicts that cognitive load should increase as-if discounting because it reduces an agent’s ability to forecast accurately. Accordingly, our noise-based discounting mechanism is not propagated by external risk (like mortality or the likelihood of default), but rather by noise associated with the limited forecasting ability of the decision maker.

Finally, our framework is consistent with Bayesian decision-making and expected utility theory. Accordingly, our agent will not be dynamically inconsistent and will not pay for commitment. In our framework, preference reversals reflect classical information acquisition, not weakness of will.

Our key assumption is that the agent has (unbiased) noise in her signals about the future. This noise leads our agent to optimally down-weight her simulations of the future and therefore place more weight on her priors. Consequently, she ends up being (rationally) imperfectly responsive to future contingencies and therefore behaves as if she discounts the future. As her expertise and experience improves (over her lifetime, or as she gains domain-specific knowledge), she shifts her behavior and acts as-if she has become more patient.

5.4 True Discounting vs. Extrinsic Risk vs. Myopia

We have now summarized three mechanisms that induce declining sensitivity to delayed utility flows: (i) true discounting arising from deep time preferences (e.g., exponential discounting, present bias, or some other time preference function); (ii) extrinsic risk (e.g., mortality, default, or some other source of risk); and (iii) *myopia* arising from forecasting noise (the focus of the current paper).

The reader may wonder how one can pull these mechanisms apart empirically. Extrinsic risk is the easiest to distinguish from the other two. Extrinsic risk can be measured directly (e.g., by measuring sources of extrinsic risk), but it can also be measured indirectly by studying agents' beliefs. For example, a researcher could elicit an agent's subjective probability that she will fail to receive a promised payment at a future date. Subjects might report that they trust an immediate payment more than a delayed payment (e.g., payable in a year's time). Failing to trust a delayed payment is an example of (perceived) extrinsic risk. Extrinsic risk is associated with "flat" learning dynamics in the sense that experience/expertise may either reduce or increase perceptions of extrinsic risk. For example, a merchant may discover through experience that a customer is reliable (unreliable), thereby leading the merchant to increase (decrease) credit to the customer. Likewise, an experimental subject may learn to trust (distrust) an experimenter, leading the subject to increase (decrease) her willingness to choose larger/later rewards over smaller/immediate rewards.

True time discounting is easy to conflate empirically with myopia. For example, we have shown that (true) hyperbolic discounting with naiveté is observationally equivalent to myopia with linear simulation noise (see Propositions 2 and 4).

Despite this similarity, true time discounting and myopia induce very different learning dynamics. With true time discounting, learning/experience/expertise induce *no* change in

the underlying time preferences. In the presence of myopia, learning/experience/expertise induce *less* as-if discounting because the simulation noise falls and/or because the priors become more refined and pull away from zero. For example, consider a household/subject who is thinking about some future opportunity. The more the household thinks about the future decision, the more responsive the household will be to future tradeoffs and the less as-if discounting the household will exhibit (e.g., Imas, Khun, and Mironova, 2017).

Learning can also be used to identify the existence of true time discounting that is dynamically inconsistent (and therefore induces a self-control problem, such as present-bias). If time preferences are dynamically inconsistent, learning will generate a preference for commitment (e.g., a 50-year-old who plays too much computer chess may finally realize that he needs to delete the app from his iPad to stop himself from playing too much). With myopia, by contrast, learning does not generate a preference for commitment. The importance of this distinction may be played out in policy design, where, for example, both the preference for commitment and the true time discounting rate impact the effectiveness of tax payments delays in improving the welfare of lower income households (Lockwood 2016).

In summary, it is possible to empirically distinguish between true time preferences and myopia by studying learning dynamics. With true time preferences, learning generates no change in the time preferences and, if the time preferences are dynamically inconsistent, learning engenders a taste for commitment. With myopia, learning generates less (as-if) discounting and no taste for commitment.

6 Conclusion

We assume that *perfectly patient* agents estimate the value of future events by generating noisy, unbiased simulations of those events. Our agents combine these noisy signals with their priors, thereby forming posterior utility expectations. We show that these expectations exhibit a property that we call *as-if discounting*. Specifically, the agent makes choices as if she were maximizing a stream of known utils weighted by an as-if discount function, $D(t)$. This as-if discount function adjusts for the fact that future utils are not actually known by the agent and must be estimated with noisy signals and priors. This estimation shades the estimated utils toward the mean of the prior distribution, creating behavior that largely mimics the effect of classical time preferences.

When the simulation noise has a variance that is linear in the event's horizon, the as-if discount function is hyperbolic:

$$D(t) = \frac{1}{1 + \alpha t},$$

where α is the ratio of the variance of (per-period) simulation noise to the variance of events in the agent's prior distribution.

Our model generates several predictions that match the known empirical evidence. Our agents exhibit systematic preference reversals. Our agents have no intrinsic taste for commitment, because they suffer from an imperfect forecasting problem, not a self-control problem. Our agents will exhibit comparative statics with respect to cognitive function: people who are more skilled at forecasting (e.g., those with greater intelligence) will exhibit less discounting.

Our framework predicts many domain-specific discounting effects. Agents with more domain-relevant experience will exhibit less discounting. Older agents will exhibit less discounting (except those with cognitive decline, who will exhibit more discounting). Agents who are encouraged to spend more time thinking about a future tradeoff will exhibit less discounting. Finally, agents who are unable to think carefully about an intertemporal tradeoff—e.g., due to a cognitive load manipulation—will exhibit more discounting.

Our framework predicts that discounting is a highly variable and plastic phenomenon that arises from imperfect forecasting of future rewards or costs. Our model provides a complementary alternative to the classical assumption that discounting arises from a deep preference for known rewards (costs) to be moved earlier (later) in time.

References

- Addessi, Elsa, Francesca Bellagamba, Alexia Delfino, Francesca De Petrillo, Valentina Focaroli, Luigi Macchitella, Valentina Maggiorelli, et al. 2014. “Waiting by Mistake: Symbolic Representation of Rewards Modulates Intertemporal Choice in Capuchin Monkeys, Preschool Children and Adult Humans.” *Cognition* 130 (3): 428–41.
- Akerlof, George A. 1991. “Procrastination and Obedience.” *American Economic Review* 81 (2): 1–19.
- Allcott, Hunt, and Dmitry Taubinsky. “Evaluating Behaviorally-Motivated Policy: Experimental Evidence from the Lightbulb Market.” *American Economic Review* 105 (8): 2501–38.
- Andreoni, James, and Charles Sprenger. 2012. “Risk Preferences Are Not Time Preferences.” *American Economic Review* 102 (7): 3357–76.
- Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2015. “Working Over Time: Dynamic Inconsistency in Real Effort Tasks.” *Quarterly Journal of Economics* 130 (3): 1067–115.
- Azfar, Omar. 1999. “Rationalizing Hyperbolic Discounting.” *Journal of Economic Behavior & Organization* 38 (2): 245–52.
- Baucells, Manel, and Franz H. Heukamp. 2012. “Probability and Time Trade-Off.” *Management Science* 58 (4): 831–42.
- Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro. 2013. “Who is ‘Behavioral’? Cognitive Ability and Anomalous Preferences.” *Journal of the European Economic Association*, 11 (6): 1231–55.
- Blanchard, Tommy C., and Benjamin Y. Hayden. 2015. “Monkeys Are More Patient in a Foraging Task than in a Standard Intertemporal Choice Task.” *PLoS ONE* 10 (2): e0117057.
- Bohm-Bawerk, Eugen von. 1890. *Capital and Interest: A Critical History of Economical Theory*. London: Macmillan.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. “Salience Theory of Choice Under Risk.” *The Quarterly Journal of Economics* 127 (3): 1243–85.
- Bordalo, Pedro, Gennaioli, Nicola and Andrei Shleifer. 2013. “Salience and Consumer Choice.” *Journal of Political Economy* 121(5): 803-43.

Brown, Donald J., and Lucinda M. Lewis. 1981. "Myopic Economic Agents." *Econometrica* 49 (2): 359-68.

Burks, Stephen V., Jeffrey P. Carpenter, Lorenz Goette, and Aldo Rustichini. 2009. "Cognitive Skills Affect Economic Preferences, Strategic Behavior, and Job Attachment." *Proceedings of the National Academy of Sciences* 106 (19): 7745-50.

Callander, Steven, and Patrick Hummel. 2014. "Preemptive Policy Experimentation." *Econometrica* 82 (4): 1509-28.

Chabris, Christopher F., David Laibson, Carrie L. Morris, Jonathon P. Schuldt, and Dmitry Taubinsky. 2008. "Individual Laboratory-Measured Discount Rates Predict Field Behavior." *Journal of Risk and Uncertainty* 37 (2-3): 237-69.

Chakraborty, Anujit. 2017. "Present Bias." Working paper, University of British Columbia.

Chapman, Gretchen B. 1996. "Temporal Discounting and Utility for Health and Money." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22 (3): 771-91.

Cogley, Timothy, and Thomas J. Sargent. 2008. "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making." *International Economic Review* 49 (1): 185-221.

Cohen, Jonathan, Keith M. Ericson, David Laibson, and John M. White. Forthcoming. "Measuring Time Preferences." *Journal of Economic Literature*.

Commons, Michael L., Michael Woodford, and John R. Ducheny. 1982. "How Reinforcers Are Aggregated in Reinforcement-Density Discrimination and Preference Experiments." In *Quantitative Analyses of Behavior: Vol. 2, Matching and Maximizing Accounts*, edited by Michael L. Commons, Richard J. Herrnstein, and Howard Rachlin, 25-78. Cambridge, MA: Ballinger.

Commons, Michael L., Michael Woodford, and Edward J. Trudeau. 1991. "How Each Reinforcer Contributes to Value: 'Noise' Must Reduce Reinforcer Value Hyperbolically." In *Signal Detection: Mechanisms, Models, and Applications*, edited by Michael L. Commons, John A. Nevin, and Michael C. Davison, 139-68. Hillsdale, NJ: Lawrence Erlbaum Associates.

Dasgupta, Partha, and Eric Maskin. 2005. "Uncertainty and Hyperbolic Discounting." *American Economic Review* 95(4): 1290-9.

Epper, Thomas, and Helga Fehr-Duda. 2015. "The Missing Link: Unifying Risk Taking and Time Discounting." Working Paper no. 096, University of Zurich.

Epper, Thomas, Helga Fehr-Duda, and Adrian Bruhin. 2011. “Viewing the Future Through a Warped Lens: Why Uncertainty Generates Hyperbolic Discounting.” *Journal of Risk and Uncertainty* 43 (3): 169–203.

Ericson, Keith M. Forthcoming. “On the Interaction of Memory and Procrastination: Implications for Reminders, Deadlines, and Empirical Estimation.” *Journal of the European Economic Association*.

Feldstein, Martin. 1985. “The Optimal Level of Social Security Benefits.” *The Quarterly Journal of Economics* 100 (2): 303–20.

Fernández-Villaverde, Jesús, and Arijit Mukherji. 2006. “Can We Really Observe Hyperbolic Discounting?” Working paper, University of Pennsylvania.

Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki. “Stochastic Choice and Optimal Sequential Sampling.” Working paper, Harvard University.

Gabaix, Xavier. 2014. “A Sparsity-Based Model of Bounded Rationality.” *The Quarterly Journal of Economics* 129 (4): 1661–710.

Gabaix, Xavier. 2016a. “Behavioral Macroeconomics via Sparse Dynamic Programming.” Working Paper no. 21848, NBER, Cambridge, MA.

Gabaix, Xavier. 2016b. “A Behavioral New Keynesian Model.” Working Paper no. 22954, NBER, Cambridge, MA.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model.” *American Economic Review* 96 (4): 1043–68.

Geanakoplos, John, and Paul Milgrom. 1991. “A Theory of Hierarchies Based on Limited Managerial Attention.” *Journal of the Japanese and International Economies* 5 (3): 205–25.

Green, Leonard, Astrid F. Fry and Joel Myerson. 1994. “Discounting and Delayed Rewards: A Life-Span Comparison. ” *Psychological Science* 5 (1): 33–6.

Halevy, Yoram. 2005. “Diminishing Impatience: Disentangling Time Preference from Uncertain Lifetime.” Working paper, University of British Columbia.

Halevy, Yoram. 2008. “Strotz Meets Allais: Diminishing Impatience and the Certainty Effect.” *American Economic Review* 98 (3): 1145–62.

Halevy, Yoram. 2014. “Some Comments on the Use of Monetary and Primary Rewards in the Measurement of Time Preferences.” Technical report, University of British Columbia.

Halevy, Yoram. 2015. “Time Consistency: Stationarity and Time Invariance.” *Econo-*

metrica 83 (1): 335–52.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein. 2014. “Learning Through Noticing: Theory and Evidence from a Field Experiment.” *The Quarterly Journal of Economics* 129 (3): 1311–53.

Hinson, John M., Tina L. Jameson, and Paul Whitney. 2003. “Impulsive Decision Making and Working Memory.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29 (2): 298–306.

Imas, Alex, Michael A. Kuhn, and Vera Mironova. 2017. “Waiting to Choose.” Working Paper, Carnegie Mellon University.

James, Bryan D., Patricia A. Boyle, Lei Yu, S. Duke Han, and David A. Bennett. 2015. “Cognitive Decline is Associated with Risk Aversion and Temporal Discounting in Older Adults Without Dementia.” *PLoS ONE* 10 (4): e0121900.

Jéhiel, Phillippe. 1995. “Limited Horizon Forecast in Repeated Alternate Games.” *Journal of Economic Theory* 67 (2): 497–519.

Kahneman, Daniel, and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision Under Risk.” *Econometrica* 47 (2): 263–91.

Kamenica, Emir. 2008. “Contextual Inference in Markets: On the Informational Content of Product Lines.” *American Economic Review* 98 (5): 2127–49.

Khaw, Mel Win, Ziang Li, and Michael Woodford. 2017. “Risk Aversion as a Perceptual Bias.” Working paper, Columbia University.

Keren, Gideon, and Peter Roelofsma. 1995. “Immediacy and Certainty in Intertemporal Choice.” *Organizational Behavior and Human Decision Processes* 63 (3): 287–97.

Kreps, David M. 1998. “Anticipated Utility and Dynamic Choice.” In *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997*. Edited by D. Jacobs, E. Kalai and M. Kamien. Econometric Society Monographs, 242–74. Cambridge, UK: Cambridge University.

Laibson, David. 1997. “Golden Eggs and Hyperbolic Discounting.” *The Quarterly Journal of Economics* 112 (2): 443–77.

Laibson, David. 2015. “Why Don’t Present-Biased Agents Make Commitments?.” *American Economic Review* 105 (5): 267–72.

Lockwood, Benjamin. 2016. “Optimal Income Taxation with Present Bias.” Working paper, University of Pennsylvania.

Loewenstein, George. 1992. “The Fall and Rise of Psychological Explanation in the Economics of Intertemporal Choice.” In *Choice over Time*, edited by George Loewenstein and Jon Elster, 3–34. New York: Russell Sage.

Matejka, Filip, Jakub Steiner and Colin Stewart. Forthcoming. “Rational Inattention Dynamics: Inertia and Delay in Decision-Making.” *Econometrica*.

McGuire, Joseph T., and Joseph W. Kable. 2012. “Decision Makers Calibrate Behavioral Persistence on the Basis of Time-Interval Experience.” *Cognition* 124 (2): 216–26.

McGuire, Joseph T., and Joseph W. Kable. 2013. “Rational Temporal Predictions Can Underlie Apparent Failures to Delay Gratification.” *Psychological Review* 120 (2): 395–410.

Mullainathan, Sendhil, and Eldar Shafir, 2013. *Scarcity: Why Having Too Little Means So Much*. New York: Henry Holt & Company.

Natenzon, Paulo. 2016. “Preference reversal or limited sampling? Maybe túngara frogs are rational after all.” Working paper, Washington University in St. Louis.

O’Donoghue, Ted, and Matthew Rabin. 1999a. “Doing It Now or Later.” *American Economic Review* (1999): 103–24.

O’Donoghue, Ted, and Matthew Rabin. 1999b. “Incentives for Procrastinators.” *The Quarterly Journal of Economics* 114 (3): 769–816.

O’Donoghue, Ted, and Matthew Rabin. 2001. “Choice and Procrastination.” *The Quarterly Journal of Economics* 116 (1): 121–60.

Peters, Jan, and Christian Büchel. 2010. “Episodic Future Thinking Reduces Reward Delay Discounting Through an Enhancement of Prefrontal-Mediotemporal Interactions.” *Neuron* 66 (1): 138–48.

Pigou, Arthur C. 1920. *The Economics of Welfare*. London: Macmillan.

Prelec, Drazen, and George Loewenstein. 1991. “Decision Making over Time and Under Uncertainty: A Common Approach.” *Management science*, 37 (7): 770–86.

Quiggin, John, and John Horowitz. 1995. “Time and Risk.” *Journal of Risk and Uncertainty* 10 (1): 37–55.

Raiffa, Howard, and Robert Schlaifer. 1961. *Applied Statistical Decision Theory*. Cambridge, MA: Harvard University and MIT Press.

Read, Daniel, Shane Frederick, and Marc Scholten. 2013. “DRIFT: An Analysis of Outcome Framing in Intertemporal Choice.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 39 (2): 573–88.

Sadoff, Sally, Anya Samek, and Charles Sprenger. 2016. “Dynamic Inconsistency in Food Choice: Experimental Evidence from a Food Desert.” Working paper, University of California, San Diego.

Schilbach, Frank, Heather Schofield, and Sendhil Mullainathan. 2016. “The Psychological Lives of the Poor.” *American Economic Review: Papers and Proceedings* 106 (5): 435–40.

Schwartzstein, Joshua. 2014. “Selective Attention and Learning.” *Journal of the European Economic Association* 12 (6): 1423–52.

Shadlen, Michael N., and Daphna Shohamy. 2016. “Decision making and sequential sampling from memory.” *Neuron* 90 (5): 927–939.

Shamosh, Noah A., and Jeremy R. Gray. 2008. “Delay Discounting and Intelligence: A Meta-Analysis.” *Intelligence* 36 (4): 289–305.

Shiv, Baba, and Alexander Fedorikhin. 1999. “Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making.” *Journal of Consumer Research* 26 (3): 278–92.

Sims, Christopher A. 2003. “Implications of Rational Inattention.” *Journal of Monetary Economics* 50 (3): 665–690.

Sozou, Peter D. 1998. “On Hyperbolic Discounting and Uncertain Hazard Rates.” *Proceedings of the Royal Society B: Biological Sciences* 265 (1409): 2015–20.

Spears, Dean. 2012. “Cognitive Limits, Apparent Impatience, and Monthly Consumption Cycles: Theory and Evidence from the South Africa Pension.” Working paper.

Steele, Claude M., and Robert A. Josephs. 1990. “Alcohol Myopia: Its Prized and Dangerous Effects.” *American Psychologist* 45 (8): 921–33.

Steiner, Jakub, and Colin Stewart. 2016. “Perceiving Prospects Properly.” *American Economic Review* 106 (7): 1601–31.

Strotz, Robert H. 1956. “Myopia and Inconsistency in Dynamic Utility Maximization.” *Review of Economic Studies* 23 (3): 165–80.

Taubinsky, Dmitry, and Alex Rees-Jones. 2016a. “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment.” Working Paper no. 22545, NBER, Cambridge, MA.

Taubinsky, Dmitry, and Alex Rees-Jones. 2016b. “Heuristic Perceptions of the Income Tax: Evidence and Implications.” Working Paper no. 22884, NBER, Cambridge, MA.

Tobin, Henry, and Alexandra W. Logue. 1994. "Self-Control across Species (*Columba Livia*, *Homo Sapiens*, and *Rattus Norvegicus*)." *Journal of Comparative Psychology* 108 (2): 126–33.

Weber, Bethany J., and Gretchen B. Chapman. 2005. "The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediacy Effect? Does Delay Eliminate the Certainty Effect?" *Organizational Behavior and Human Decision Processes* 96 (2): 104–18.

Weitzman, Martin. 2001. "Gamma Discounting." *American Economic Review*. 91 (1): 260–71.

Woodford, Michael. 2009. "Information-Constrained State-Dependent Pricing." *Journal of Monetary Economics* 56 (S): 100–24.

Yaari, Menahem E. 1965. "Uncertain Lifetime, Life Insurance, and the Theory of Consumer." *Review of Economic Studies*, 32 (1): 37–50.

7 Appendix: Proofs and Complements

7.1 Omitted Proofs

Proof of Proposition 1 This proof is very elementary, but for completeness we provide its calculations. We normalize $\mu = 0$ without loss of generality (for instance, by considering $u'_t = u_t - \mu$ and $s'_t = s_t - \mu$). It is well-known that $u_t \mid s_t$ is Gaussian distributed, and can be represented:

$$u_t = \lambda s_t + \eta_t \tag{17}$$

for some λ , and some Gaussian variable η_t independent of s_t , so that $\mathbb{E}[s_t \eta_t] = 0$. Multiplying (17) by s_t on both sides and taking the expectations gives: $\mathbb{E}[u_t s_t] = \lambda \mathbb{E}[s_t^2]$, i.e.

$$\begin{aligned} \lambda &= \frac{\mathbb{E}[u_t s_t]}{\mathbb{E}[s_t^2]} = \frac{\mathbb{E}[u_t (u_t + \varepsilon_t)]}{\mathbb{E}[(u_t + \varepsilon_t)^2]} = \frac{\mathbb{E}[u_t^2]}{\mathbb{E}[u_t^2 + \varepsilon_t^2]} \text{ as } \mathbb{E}[u_t \varepsilon_t] = 0 \\ &= \frac{\sigma_u^2}{\sigma_u^2 + \sigma_{\varepsilon_t}^2} = \frac{1}{1 + \frac{\sigma_{\varepsilon_t}^2}{\sigma_u^2}} = D(t). \end{aligned}$$

Next, taking the variance of both sides of (17), we have

$$\sigma_u^2 = \lambda^2 \sigma_s^2 + \text{var}(\eta_t)$$

as $\text{cov}(s_t, \eta_t) = 0$ and with $\sigma_s^2 = \sigma_u^2 + \sigma_{\varepsilon_t}^2$. So, using $\lambda \sigma_s^2 = \sigma_u^2$,

$$\text{var}(\eta_t) = \sigma_u^2 - \lambda^2 \sigma_s^2 = \sigma_u^2 - \lambda \sigma_u^2 = (1 - \lambda) \sigma_u^2.$$

Hence, $u_t \mid s_t \sim \mathcal{N}(\lambda s_t, (1 - \lambda) \sigma_u^2)$, as announced.

Proof of Proposition 4 In this proof, we normalize μ to 0 (again, this is innocuous, replacing $u(c)$ by $u(c) - \mu$ as necessary). Because the scientist has access to the lotteries on the goods at time 0, he can recover the agent's cardinal utility up to affine transformation, i.e. he can fit $\hat{u}(c) = K(u(c) + b)$ for all c , with constants K and b to be determined. Then, moving to choices for more general dates, the scientist needs to fit the probability of choosing *Early* for all dates $t, t + \tau$, which was derived in (10), i.e. he needs to make sure that his

fitted values (all denoted with hats) satisfy:

$$\Phi \left(\frac{1}{\widehat{\Sigma}_{t,t+\tau}} \left[\widehat{D}(t)\widehat{u}(c_t) - \widehat{D}(t+\tau)\widehat{u}(c_{t+\tau}) \right] \right) = \Phi \left(\frac{1}{\Sigma_{t,t+\tau}} [D(t)u(c_t) - D(t+\tau)u(c_{t+\tau})] \right)$$

Both the left- and right-hand sides of this equation represent the probability of choosing good c_t at t over good $c_{t+\tau}$ at $t+\tau$. On the left, this probability is expressed in the scientist's model (with $\widehat{\Sigma}_{t,t+\tau}$ the standard deviation of the Probit noise), and on the right, this probability is represented by the true data generating process (see (10)). In other terms, the scientist's fitted values must insure that for all consumptions and dates in his dataset:

$$\frac{1}{\widehat{\Sigma}_{t,t+\tau}} \left[\widehat{D}(t)\widehat{u}(c_t) - \widehat{D}(t+\tau)\widehat{u}(c_{t+\tau}) \right] = \frac{1}{\Sigma_{t,t+\tau}} [D(t)u(c_t) - D(t+\tau)u(c_{t+\tau})]. \quad (18)$$

It is clear that a possible way to make (18) hold exactly is to have: $\widehat{D}(t) = \frac{D(t)}{D(0)}$, $\widehat{u}(c) = Ku(c)$ and $\widehat{\Sigma}_{t,t+\tau} = \frac{K\Sigma_{t,t+\tau}}{D(0)}$. Then, the two sides of (18) are equal.

This almost proves the proposition, but not quite: we also have to verify that this is the only possible fit (up to the usual arbitrary scaling factor $K > 0$). To finish the last step and prove uniqueness, take (18) for two different values of c_t , keeping the other terms constant, and subtract them. This implies $\frac{K\widehat{D}(t)}{\widehat{\Sigma}_{t,t+\tau}} = \frac{D(t)}{\Sigma_{t,t+\tau}}$, i.e., $\frac{K\Sigma_{t,t+\tau}}{\widehat{\Sigma}_{t,t+\tau}} = \frac{D(t)}{\widehat{D}(t)}$. The same reasoning applied to two different values of $c_{t+\tau}$ gives $\frac{K\Sigma_{t,t+\tau}}{\widehat{\Sigma}_{t,t+\tau}} = \frac{D(t+\tau)}{\widehat{D}(t+\tau)}$, i.e.

$$\frac{K\Sigma_{t,t+\tau}}{\widehat{\Sigma}_{t,t+\tau}} = \frac{D(t)}{\widehat{D}(t)} = \frac{D(t+\tau)}{\widehat{D}(t+\tau)}. \quad (19)$$

Using (19) for $t = 0$ gives: $D(0) = \frac{D(\tau)}{\widehat{D}(\tau)}$ (indeed, recall that we normalized $\widehat{D}(0) = 1$), i.e., $\widehat{D}(\tau) = \frac{D(\tau)}{D(0)}$. This holds for all dates τ . Then (19) implies that for all dates $t, t+\tau$, we have $\frac{K\Sigma_{t,t+\tau}}{\widehat{\Sigma}_{t,t+\tau}} = D(0)$. Finally, (18) then implies that $b = 0$. This concludes the proof of uniqueness, up to the usual multiplicative factor $K > 0$.

Proof of Proposition 5 It is a corollary of Proposition 7 (for Assumption 1) and Proposition 8 (for Assumption 2) below.

Proof of Proposition 6 Given our assumptions, the agent at time t will want to maximize

$$\max_{(a_{t+\tau})_{\tau \geq 0}} \mathbb{E} \left[\sum_{\tau=0}^{T-t} u(a_{t+\tau}) \mid \mathbf{s} \right] = \max_{(a_{t+\tau})_{\tau \geq 0}} \sum_{\tau=0}^{T-t} D(\tau) s_{t+\tau}(a_{t+\tau})$$

where $\mathbf{s} = (s_t(y), \dots, s_{t+\tau}(y))_{y \in \mathcal{A}}$. Assumption A1-A3 allows us to remove expected values. For our representative agent, we have $s_{t+\tau}(a) = u_{t+\tau}(a)$. Hence, this representative agent maximizes at time t :

$$\max_{(a_{t+\tau})_{\tau \geq 0}} \sum_{\tau=0}^{T-t} D(\tau) u(a_{t+\tau}).$$

7.2 Complements to the Continuous Actions Case

Here are some complements to Section 4. To simplify the notations, we set $\sigma = \sigma_{\varepsilon_t}$.

7.2.1 Result for the Wiener case

Proposition 7 (Bayesian updating with functions) *Under Assumption 1, we have*

$$\mathbb{E}[u(a) \mid \mathbf{s}] = \lambda s(a)$$

with $\lambda = \frac{1}{1 + \sigma_{\varepsilon_t}^2 / \sigma_u^2}$. This means that we can do Bayesian updating on this space of functions.

Proof. Take the increments:

$$ds(a) = du(a) + \sigma dW(a)$$

The key observation is that the $ds(a)$'s are all Gaussians innovations, independent of the value of the functions at other points $y \neq a$. So, by the formulation for Gaussian updating we used before:

$$\mathbb{E}[du(a) \mid ds(a)] = \lambda ds(a)$$

with $\lambda = \frac{1}{1 + \sigma^2 / \sigma_u^2}$. Next, because the $du(a)$ and $dW(a)$ are independent,

$$\mathbb{E}[du(a) \mid \mathbf{s}] = \mathbb{E}[du(a) \mid ds(a)] = \lambda ds(a). \quad (20)$$

Next, the behavior at 0 needs a special treatment. Because $s(0) = u(0) + \chi \sigma \eta_0$,

$\mathbb{E}[u(0) | s(0)] = \lambda_0 s(0)$, with $\lambda_0 = \frac{1}{1 + \frac{\text{var}(\chi\sigma\eta_0)}{\text{var}(u_0)}} = \frac{1}{1 + \frac{\chi^2\sigma^2}{\chi^2\sigma_u^2}} = \lambda$. Then, $\mathbb{E}[u(0) | s(0)] = \lambda s(0)$, and by independence:

$$\mathbb{E}[u(0) | \mathbf{s}] = \lambda s(0). \quad (21)$$

Hence, integrating from 0 to a , we get

$$\begin{aligned} \mathbb{E}[u(a) | \mathbf{s}] &= \mathbb{E}\left[u(0) + \int_{y=0}^a du(y) | \mathbf{s}\right] = \mathbb{E}[u(0) | \mathbf{s}] + \int_{y=0}^a \mathbb{E}[du(y) | \mathbf{s}] \\ &= \lambda s(0) + \int_0^a \lambda ds(y) \\ &= \lambda s(a). \end{aligned}$$

7.2.2 Polynomial utility

Here we provide assumptions that are a little more elementary, but apply only when the utility function $u(a)$ is a polynomial in a . For instance, we want to capture that $u(a) = b_0 + b_1a + b_2a^2$ with unknown coefficients b_i , that the agent wants to learn from noisy signals.

Assumptions for the polynomial utility case We shall use the Legendre $P_i(a)$ polynomials as a basis, as they are more convenient than the plain monomials a^i . We have for instance:¹⁹

$$P_0(a) = 1, \quad P_1(a) = a, \quad P_2(a) = \frac{1}{2}(3a^2 - 1), \quad P_3(a) = \frac{1}{2}(5a^3 - 3a).$$

Using the inner product on $L^2([-1, 1])$:

$$\langle f | g \rangle := \int_{-1}^1 f(a) g(a) da, \quad (22)$$

we have the standard result: $\langle P_i | P_j \rangle = \frac{1}{i + \frac{1}{2}} 1_{i=j}$. So we define q_i to be a rescaled version of the standard Legendre polynomial:

$$q_i(a) := \sqrt{i + \frac{1}{2}} P_i(a), \quad (23)$$

¹⁹More generally we have $P_i(a) = \frac{1}{2^i i!} \frac{d^i}{da^i} [(a^2 - 1)^i]$ by Rodrigues' formula.

so that

$$\langle q_i | q_j \rangle = 1_{i=j}. \quad (24)$$

Polynomial q_i has degree i , and the q_i 's form an orthogonal basis for polynomial functions.

We can now state our assumption.

Assumption 2 (Utility function as drawn from a random distribution on polynomial basis)

We decompose the true utility function $u(a)$ as:

$$u(a) = \sum_{i=-1}^{\infty} f_i Q_i(a) \quad (25)$$

where $Q_{-1}(a) \equiv 1$ and for $i \geq 0$, $Q_i(a) = \int_0^a q_i(y) dy$, where $q_i(y)$ is the i -th normalized Legendre polynomial (23). We assume that a finite subset I such that coefficients $\{f_i\}_{i \in I}$ are nonzero and that the f_i for $i \in I$ are i.i.d. and follow a $N(0, \sigma_u^2)$ distribution. Also assume $\sigma_{f_{-1}}^2 = \chi^2 \sigma_u^2$.

We note that, in the limit where all coefficients are non-zero, we get the ‘‘Wiener’’ case.

Result We prove a more general proposition.

Proposition 8 Suppose that coefficients f_i are drawn from the Gaussian $N(0, \sigma_{f_i}^2)$, and jointly Gaussian and uncorrelated. Then, the posterior $\mathbb{E}[u(a) | \mathbf{s}] := \mathbb{E}\left[u(a) | (s(y))_{y \in [-1,1]}\right]$ is:

$$\mathbb{E}[u(a) | \mathbf{s}] = \sum_{i=-1}^{\infty} \mathbb{E}[f_i | \mathbf{s}] Q_i(a)$$

where, for $i \geq 1$

$$\begin{aligned} \mathbb{E}[f_i | \mathbf{s}] &= \lambda_i \langle q_i | d\mathbf{s} \rangle = \lambda_i \int_{a=-1}^1 q_i(a) ds(a) \\ \lambda_i &= 1 / (1 + \sigma^2 / \text{var}(f_i)) \end{aligned}$$

while $\mathbb{E}[f_{-1} | \mathbf{s}] = \lambda_{-1} s(0)$ with $\lambda_{-1} = 1 / (1 + \chi^2 \sigma^2 / \text{var}(f_{-1}))$. This implies that the average posterior is:

$$\bar{u}(a) := \mathbb{E}[[u(a) | \mathbf{s}] | f] = \sum_{i=-1}^{\infty} \lambda_i f_i Q_i(a). \quad (26)$$

In particular, take the case of flat priors of Assumption 2, and call $\lambda = 1 / \left(1 + \frac{\sigma^2}{\sigma_u^2}\right)$. Then,

$$\bar{u}(a) = \lambda u(a) \quad (27)$$

i.e. we obtain uniform dampening.

Proof of Proposition 8 Suppose that we have a function $u(a)$, and we observe, as in (14),

$$s(a) = u(a) + \sigma W(a) + \chi \sigma \eta_0 \quad (28)$$

where $W(a)$ is a Brownian motion and $\tilde{\eta}_0 = \chi \sigma \eta_0$ is a Gaussian variable of mean zero. Differentiate:

$$\begin{aligned} ds(a) &= u'(a) da + \sigma dW(a) \\ u'(a) &= \sum_{j=-1}^{\infty} f_j Q'_j(a) = \sum_{j=0}^{\infty} f_j q_j(a). \end{aligned}$$

Hence:

$$ds(a) = \sum_{j=0}^{\infty} f_j q_j(a) da + \sigma dW(a).$$

The agent wants to infer u given s , i.e. f given ds (we consider the intercept $u(0)$ at the end). Multiplying the previous equation by $q_i(a)$ and integrating between -1 and 1 gives:

$$\begin{aligned} S_i &:= \langle q_i \mid ds \rangle \\ &= \sum_j f_j \langle q_i \mid q_j \rangle + \sigma \langle q_i \mid dW \rangle \\ &= f_i + \sigma \langle q_i \mid dW \rangle \end{aligned} \quad (29)$$

because of (24).

Hence we can write the signal $S_i := \langle q_i \mid ds \rangle$ as

$$S_i = f_i + \sigma \varepsilon_i \quad (30)$$

with $\varepsilon_i := \langle q_i \mid dW \rangle = \int_{-1}^1 q_i(a) dW_a$ satisfies $\mathbb{E}[\varepsilon_i] = 0$. In addition

$$\begin{aligned} \mathbb{E}[\varepsilon_i \varepsilon_j] &= \mathbb{E} \left[\left(\int q_i(a) dW_a \right) \left(\int q_j(a) dW_a \right) \right] = \int q_i(a) q_j(a) da \\ &= \mathbf{1}_{i=j}. \end{aligned}$$

Hence, the signal-extraction problem $\mathbb{E}[f_i \mid \mathbf{s}]$ is quite simple, as only S_i is informative about f_i : $\mathbb{E}[f_i \mid \mathbf{s}] = \mathbb{E}[f_i \mid S_i]$. Given (30),

$$\mathbb{E}[f_i \mid \mathbf{s}] = \lambda_i S_i \tag{31}$$

$$\lambda_i = 1 / (1 + \sigma^2 / \text{var}(f_i)). \tag{32}$$

Hence, we have

$$\begin{aligned} \mathbb{E}[u'(a) \mid \mathbf{s}] &= \sum_{i=0}^{\infty} \mathbb{E}[f_i \mid \mathbf{s}] q_i(a) \\ &= \sum_{i=0}^{\infty} \mathbb{E}[f_i \mid \mathbf{s}] Q'_i(a) \end{aligned}$$

We next study the intercept in (28), $u(0)$. Given $s(0) = u(0) + \chi \sigma \eta_0$ and $u(0) = f_{-1}$,

$$\mathbb{E}[u(0) \mid \mathbf{s}] = \mathbb{E}[f_{-1} \mid s(0)] = \lambda_{-1} s(0) = \lambda_{-1} S_{-1}$$

where $S_{-1} := s(0)$ and $\lambda_{-1} = 1 / (1 + \chi^2 \sigma^2 / \text{var}(f_{-1}))$. Integrating,

$$\begin{aligned} \mathbb{E}[u(a) \mid \mathbf{s}] &= \mathbb{E}[u(0) \mid \mathbf{s}] + \mathbb{E} \left[\int_0^a u'(b) db \mid \mathbf{s} \right] \\ &= \lambda_{-1} S_{-1} + \sum_{i=0}^{\infty} \lambda_i S_i Q_i(a) = \sum_{i=-1}^{\infty} \lambda_i S_i Q_i(a). \end{aligned}$$

In addition, the average perception is:

$$\begin{aligned} \bar{u}(a) &:= \mathbb{E}[[u(a) \mid \mathbf{s}] \mid \mathbf{u}] = \sum_{i=-1}^{\infty} \lambda_i \mathbb{E}[S_i \mid f] Q_i(a) \\ &= \sum_{i=-1}^{\infty} \lambda_i f_i Q_i(a) \end{aligned} \tag{33}$$

If we assume a “flat” prior of Assumption 2, where $\text{var}(f_i)$ is independent of i (if $\text{var}(f_i) > 0$), we have for $i \geq 0$

$$\lambda_i = \lambda = \frac{1}{1 + \frac{\sigma^2}{\text{var}(f_i)}} = \frac{1}{1 + \frac{\sigma^2}{\sigma_u^2}}.$$

Furthermore, as $\sigma_{f_{-1}}^2 = \chi^2 \sigma_u^2$,

$$\lambda_{-1} = \frac{1}{1 + \frac{\chi^2 \sigma^2}{\text{var}(f_{-1})}} = \lambda.$$

Hence, $\lambda_i = \lambda$ for all $i \geq -1$, and (33) implies:

$$\bar{u}(a) = \sum_{i=-1}^{\infty} \lambda_i f_i Q_i(a) = \lambda \sum_{i=-1}^{\infty} f_i Q_i(a) = \lambda u(a).$$