Reference-Dependent Consumption Plans

Botond Kőszegi and Matthew Rabin
Department of Economics
University of California – Berkeley

First Draft: June 2006
This Draft: October 2007

Abstract: We develop a dynamic model in which utility depends on consumption as well as recent changes in beliefs about present and future consumption, and bad news is more painful than good news is pleasant. We assume behavior must meet the rational-consistency condition that the person believes her own plans: given the expectations generated by a plan, in each period the plan maximizes discounted expected utility taking into account that continuation plans must also be consistent with rationality. Applying our model to the acquisition of decision-irrelevant information, we show that (1) to avoid fluctuations in beliefs, a person prefers to receive information clumped together rather than separately, and (2) because updating beliefs has lower utility impact when about outcomes that are further in the future, she prefers learning information earlier to later. Because wealth changes constitute news about the distribution of future consumption, our model predicts ways in which the decisionmaker exhibits reference-dependent attitudes toward current changes in wealth—even though she holds wealth exclusively for future consumption and faces substantial uncertainty about future income. We also apply our theory to a two-period consumption-savings problem. Since raising current consumption above the planned level may be very pleasant, a decisionmaker might overconsume early relative to the optimal committed plan. While the greater aversion to losses from lowering plans for future consumption can deter this tendency when consumption plans are deterministic, uncertain plans mitigate this asymmetry, so overconsumption is exacerbated by current uncertainty. Because a surprise increase in lifetime wealth is most pleasant, and a surprise decrease most unpleasant, if consumed immediately, a person responds asymmetrically to surprises. And since higher consumption reduces the sensation of loss from lower-than-expected consumption, a person prepares for anticipated uncertainty about future income with a first-order increase in current saving.

[0]Acknowledgments: We are grateful to Avinash Dixit, and seminar audiences at University of California, Berkeley, CERGE-EI, University of Edinburgh, London School of Economics, University of California, San Diego, and University of Texas for comments. We thank the National Science Foundation (Grant Award #SES-0648659) for financial support. Kristóf Madarász provided research assistance.

1 Introduction

In this paper, we develop a general dynamic model of reference-dependent utility. Building on ideas and models in Hsee and Tsai (2006), Kimball and Willis (2006), and Kőszegi and Rabin (2006, 2007), we assume that utility depends not only on current consumption, but also on recent changes in beliefs about present and future consumption, and that bad news is more painful than good news is pleasant. We develop various "solution concepts" for predicting behavior and welfare based on the idea that a person must be rational in forming focused beliefs about her future consumption, and derive implications of our model for preferences over receiving information about an exogenous future event, monetary risk preferences, and intertemporal consumption decisions. If news about more imminent consumption is felt more heavily than news about distant consumption, a person prefers to receive the same information sooner rather than later, increases consumption immediately in response to good surprises regarding wealth but delays cuts following bad surprises, and—since surprising herself with extra immediate consumption may be very pleasant—she may overconsume early in life relative to the optimal committed plan. To reduce the impact of losses she may suffer, the decisionmaker prefers to receive bits of information together rather than apart, and—to lower the marginal utility associated with any surprise losses—prepares for future uncertainty by increasing savings.

We present the basic framework in Section 2. In each period $t \in \{0, 1, \ldots, T\}$, the decisionmaker consumes a K-dimensional consumption bundle, c_t. Overall instantaneous utility in a period is the sum of reference-independent "consumption utility" that can be thought of as the classical economic notion of utility, and gain-loss utility that derives from changes in beliefs about each period's consumption in each dimension. The person experiences "contemporaneous" gain-loss utility from any contrast between current consumption and her prior expectations of current consumption, and "prospective" gain-loss utility from changes in her beliefs about future consumption. In all these comparisons, she is loss averse: bad news about current or future consumption is more unpleasant than good news is pleasant. Normalizing the weight $\gamma_{t,t}$ on contemporaneous gain-loss utility to 1, in period t the decisionmaker puts weight $\gamma_{t,\tau} \leq 1$ on prospective gain-loss utility regarding outcomes in period $\tau > t$. In period t the decisionmaker's goal is to maximize the sum of instantaneous utilities starting in period t.

Because preferences depend on the sequence of expectations, our model is complete only when combined with a theory of how expectations are formed. While we follow previous work on belief-based preferences (including Kőszegi and Rabin (2006, 2007) on reference-dependent preferences and Caplin and Leahy (2001) and Kőszegi (2006b) on anticipatory utility) in grounding our model in rational expectations, we propose a framework for thinking about the dynamic formation of rational beliefs and plans without imposing that *initial* beliefs about the future are necessarily correct.[1] Normalizing period 0 as the period in which the first focused plans for the decision are made—which is not necessarily the period in which the first relevant choice must be implemented—we do not impose any restrictions on the beliefs at the beginning of period 0. But we assume that from that point on expectations and plans must be rational. Behavior in periods 0 through T must then correspond to an optimal consistent plan, defined as a plan that in each period maximizes expected reference-dependent utility given past expectations, with the constraint that current plans and expectations for the future must be consistent with a similar maximization in the future. This means that starting in a period $t \geq 1$, executed plans constitute a "personal equilibrium" as defined in various prior papers on beliefs-based preferences—whereby dynamic behavior must be optimal given beliefs generated by the same behavior. If updating initial plans generates no prospective gain-loss utility—perhaps because an initial planning period when beliefs are "unfocused" does not generate sensations of gain and loss—optimal consistent planning accords to the more restrictive notion of a "preferred personal equilibrium," the personal equilibrium that yields a person her highest discounted utility starting in period t.

While our main interest is in predicting behavior when the person can influence consumption outcomes, in Section 3 we explore some implications of the model solely due to our key innovation, loss-averse preferences over changes in beliefs. Specifically, we investigate the decisionmaker's preferences regarding interim information about an exogenous stochastic binary variable to be resolved at some set future date. On the one hand, when prospective gain-loss utility does not

[1] Another strand of models of beliefs-based preferences, such as Akerlof and Dickens (1982) and Brunnermeier and Parker (2005), assumes that agents can *choose* beliefs—even beliefs that are not consistent with rational expectations. Because agents are assumed to maximize their utility given their beliefs once the beliefs are chosen, they face a tradeoff between forming beliefs that make them feel better and ones that help them make good choices.

loom as large as contemporaneous gain-loss utility—so that possible bad news is less painful when outcomes are not imminent—a person likes receiving information sooner rather than later. She dislikes interim information in another sense, however: because receiving information at separate times exposes her to possibly unnecessary bad news due to fluctuations in beliefs, she prefers to receive information as "clumped" together as possible. This also means that *if* she has (voluntarily or involuntarily) just received information, she will prefer to receive more information in the same period.[2]

In Section 4, we show how our model provides foundations for gain-loss utility over wealth. Because changes in wealth provide news about the distribution of future consumption and this news generates immediate gain-loss utility, the decisionmaker attends to gains and losses in wealth even if these are minuscule relative to her lifetime wealth and risk exposure—and even if she holds wealth exclusively for future consumption. This insight also provides a new perspective on some discussions of bracketing in the literature: while treating a risk to wealth in isolation from future decisions and risks is typically considered a mistake, our model says that in some situations it is instead a manifestation of a preference over changes in beliefs.

Our set of results on monetary preferences is the most important example of how our theory helps unify much of the growing body of research on reference-dependent utility more generally. A long-standing literature on "habit formation" assumes that people evaluate consumption in part by comparing it to consumption from the recent past.[3] A large literature building from Knetsch (1989) and Kahneman, Knetsch, and Thaler (1990) instead posits that willingness to pay for durable goods depends on recent ownership status. And the literature building from Kahneman and Tversky's (1979) original model of prospect theory assumes that people care about changes in wealth. While the approach of assuming that people care about changes in beliefs about consumption represents a new modeling approach, it is in fact consistent with much of the evidence motivating these previous approaches. Past consumption is in many circumstances a major determinant of expected

[2] These results deriving predictions from loss aversion about when a decisionmaker is "information loving" and "information averse" complement models such as Caplin and Leahy (2001) and Kőszegi (2006a), which assume a taste or distaste for information as a primitive.

[3] See, for instance, Duesenberry (1952), Ryder and Heal (1973), and Scitovsky (1976), with Bowman, Minehart, and Rabin (1999) more recently combining this approach with Kahneman and Tversky's prospect theory.

future consumption; changes in ownership status of a durable good are presumably pleasant or unpleasant mostly because of beliefs about the good's future usefulness; and although it does not fully capture the complicated psychology of money, in the above ways our model is also consistent with wealth-based theories.

In Section 5, we present our main application, exploring a two-period consumption-savings problem. A consumer with strictly concave consumption utility must split a possibly stochastic amount of lifetime wealth between two periods, where we assume for simplicity that the interest rate is zero. We derive a number of behavioral and welfare results from the assumption that $\gamma \equiv \gamma_{1,2} < 1$, so that a deviation from expected period-1 consumption has a greater effect on utility than a similar change of plans regarding period-2 consumption. The most basic implication of this assumption is an asymmetry in the consumer's response to "surprises"—low-probability changes in her wealth. In the extreme, suppose she had made plans expecting a given deterministic lifetime wealth, but then finds at the beginning of period 1 that her wealth will be a different deterministic level. Because the gain from increasing period-1 consumption is more pleasant than the gain from increasing plans for period-2 consumption, she fully consumes small increases in wealth in period 1. And because the loss from decreasing period-1 consumption is more painful than the loss from decreasing plans for period-2 consumption, she fully absorbs small decreases in wealth in period 2.

As our key set of results in this section, we identify ways in which the consumer's behavior may be suboptimal among the strategies available to her. We replicate in our setting a result Stone (2005) established for deterministic wealth, that expectations-based preferences can generate overconsumption for a completely different reason than present-biased preferences in the sense of Laibson (1997) and O'Donoghue and Rabin (1999), and temptation utility in the sense of Gul and Pesendorfer (2001). To understand the intuition, suppose the consumer had made the ex-ante optimal plan, to consume an equal amount in the two periods. If γ is relatively small, pleasantly surprising herself with extra consumption in period 1 at the cost of lowering consumption plans for period 2 increases utility ex post, so that without commitment the optimal plan is not consistent.

More interestingly, our model differs from the above self-control theories in predicting a strong role of current uncertainty in exacerbating overconsumption: there are γ such that if the consumer faces sufficient uncertainty regarding her wealth, she overconsumes in period 1 for all wealth levels; yet for any wealth level on the support of her wealth distribution, if she knew in advance this would be her wealth level, she would not overconsume. Because the sense of gain from an upward revision in period-1 consumption is often smaller than the necessary sense of loss from the corresponding downward revision of period-2 consumption, a deterministic plan for consumption acts as a bright-line commitment device. When there is uncertainty, however, the implications of a revision in plans are evaluated much less asymmetrically—because the possibility of higher or lower consumption was already incorporated into expectations. If realized wealth is the highest possible, for example, an increase in period-1 consumption relative to plans is evaluated as a gain, and a decrease in period-2 consumption is evaluated largely as a foregone gain.

Finally, we consider environments where the consumer's wealth is possibly stochastic, and uncertainty is resolved in period 2. We identify a novel type of precautionary-savings motive that implies an unambiguously positive first-order response to uncertainty, and that seems more intuitive than a similar motive generated by expected-utility-of-wealth considerations. Because a low realization of consumption utility is evaluated as a painful loss relative to higher possible realizations, whereas a high realization is evaluated merely as a gain relative to lower possible realizations, the consumer dislikes uncertainty in consumption utility. Saving more decreases this uncertainty by decreasing the effect of any given monetary shortfall on consumption utility. Hence, if she knows in period 0 that she will face uncertainty, she plans to save more to prepare for it.

We conclude the paper in Section 6 by noting further natural applications of our model, and especially pointing out some of its limitations.

2 The Model

In this section we present the model and explore some of its general features. There are $T+1$ periods, 0 through T, in each of which a K-dimensional consumption vector $c_t = (c_t^1, \ldots, c_t^K)$ is realized. Because our model pays special attention to the possibility and implications of planning for future decisions, the initial period, period 0, corresponds to the time at which the person first makes plans for the relevant set of decisions. This is not necessarily the period in which the first non-trivial decision is implemented.

The timing within a period t is the following. The decisionmaker starts with fixed beliefs $F_{t-1} = \{F_{t-1,\tau}\}|_{\tau=t}^{T}$ inherited from period $t-1$, where $F_{t-1,\tau} = (F_{t-1,\tau}^1, \ldots, F_{t-1,\tau}^K)$ are the beliefs regarding the K dimensions of consumption in period τ. Then, some uncertainty may be resolved, and the decisionmaker takes an action. Further uncertainty may then be resolved, the decisionmaker forms new beliefs $\{F_{t,\tau}\}_{\tau=t}^{T}$, and c_t is realized. The beliefs $F_{t,t}$ assign probability 1 to c_t.[4]

The decisionmaker's instantaneous utility in each period depends on consumption in that period and the most recent changes in beliefs about contemporaneous and future consumption:

$$u_t = m(c_t) + \sum_{\tau=t}^{T} \gamma_{t,\tau} N(F_{t,\tau}|F_{t-1,\tau}). \tag{1}$$

The term $m(c_t)$ is "consumption utility," which can be thought of as corresponding to classical reference-independent utility. We assume that consumption utility is separable across dimensions, and that the consumption-utility function in dimension k, $m^k(\cdot)$, is differentiable and strictly increasing. The terms $N(F_{t,\tau}|F_{t-1,\tau})$ represent "gain-loss utility," and $\gamma_{\tau,\tau} \geq \gamma_{\tau-1,\tau} \geq \cdots \geq \gamma_{0,\tau} \geq 0$ are weights on these gain-loss utilities. For $t \geq 1$, we normalize $\gamma_{t,t} = 1$. For $\tau > t$, $N(F_{t,\tau}|F_{t-1,\tau})$ is "prospective gain-loss utility," which derives from changes between last period and this period in beliefs regarding future outcomes.[5,6] While notationally, substantively, and psychologically consistent with such prospective gain-loss utility, the functions $N(F_{t,t}|F_{t-1,t})$ can be usefully distinguished as "contemporaneous gain-loss utility." These derive from comparing the consumption outcome in period t to the person's beliefs at the beginning of the period regarding that outcome.

[4] While the above within-period timing is sufficient for all the applications in this paper, the model of preferences is fully compatible with other decision-making structures.

[5] While it would be more realistic to assume that comparisons to past beliefs matter beyond a one-period lag, it seems most essential intuitions of such preferences can be captured in our simplified specification by positing that multiple decisions are made within a period.

[6] By including anticipation of consumption in a given period in instantaneous utility in all prior periods, our formulation may appear to be "multiple-counting" consumption. Repeated changes in beliefs about the same future period can indeed have substantial utility effects—generating some interesting results below about aversion to incessant arrival of information. But it is *only* when beliefs are changing incessantly that perception of one period's consumption has a huge effect on prior periods' utility. In fact, when consumption in a given period implements plans that were always held, our specification below implies that prospective gain-loss utility regarding that period is zero in all prior periods.

The weights $\{\gamma_{t,\tau}\}$ are central in determining the relative importance of news as a function of how far in advance of consumption the news is received. When $\gamma_{t,\tau} = 0$ for all $\tau > t$, changed expectations of and plans about the future do not affect a person's current well-being. By assuming that it is ultimately consumption that generates sensations of gain and loss, previous beliefs-based theories of reference-dependent utility in Stone (2005) and Kőszegi and Rabin (2006, 2007) have implicitly made exactly this parameter restriction, and as such have ignored considerations that are central in this paper. When, by contrast, $\gamma_{t,\tau} = 1$ for all $\tau \geq t$, getting news about an outcome before it happens has no less resonance with a person than learning it at the time. The in-between case we consider, where $0 < \gamma_{t,\tau} < 1$ with $\gamma_{t,\tau}$ increasing in t, is that the impact of news about a period is greater the closer is the period. Although we find this last case most plausible, the role of temporal distance in the impact of learning news is (as far as we know) largely unknown. Our results identify ways that this affects behavior, so they can be used to guide further empirical investigation in the area.[7]

Our formulation follows some previous research suggesting that changes in beliefs are carriers of utility. For instance, Loewenstein (1988) finds that people are willing to pay more to avoid delaying the delivery of a durable good when they expected to receive it immediately than to speed up its delivery when they expected to receive it later. This is consistent with our assumption that lowering consumption utility in the interim period below expected is a painful loss. More specifically related to our specification of prospective gain-loss utility, Hsee and Tsai (2006) find that the "news utility" from learning about near-term consumption can be stronger than the utility from consumption itself, and the model by Kimball and Willis (2006) assumes that news about future utility is one component of current happiness. The premise of our model is that "news" is not only an additional component of reference-dependent utility, but it is the central component. While this means that in formal terms our model differs markedly from most other theories of reference-dependent utility, it is in fact consistent with previous models and intuitions. Because past consumption is often a reliable indicator of future consumption, in many environments our

[7] With similarly little evidence to go on, we conjecture that hyperbolic decay in $\gamma_{t,\tau}$ as $\tau - t$ increases is more plausible than exponential decay: bad news about tomorrow may be considerably less painful than bad news about today, while bad news about 117 days from now may resonate much the same as bad news about 116 days from now.

model makes similar predictions to theories of habit formation and reference-dependent utility starting from Ryder and Heal (1973). As we show below, the use of money in most classical applications of prospect theory can be understood as people treating money as news about future consumption. Similarly, concern for ownership of a durable good in models of the endowment effect is naturally interpretable as concern for the news getting an object conveys about future consumption. In this light, the wealth and ownership-based models of reference dependence can be seen as simply not harping on details of how outcomes generate gain-loss utility rather than as reflecting a hypothesis that it really is the receipt of objects or wealth per se that people care about. Our formal elaboration that *does* harp on such details helps provide some unifying foundations for these approaches, while—more importantly—also establishing cases where the more explicit focus on beliefs about consumption modifies and extends insights.

We now turn to specifying the gain-loss utility function N. While our definition is notationally quite cumbersome, the basic idea is that the decisionmaker makes a rank-dependent comparison between her previous beliefs $F^k_{t-1,\tau}(\cdot)$ and new beliefs $F^k_{t,\tau}(\cdot)$: she compares the worst outcomes under $F^k_{t,\tau}(\cdot)$ to the worst outcomes under $F^k_{t-1,\tau}(\cdot)$, the best outcomes under $F^k_{t,\tau}(\cdot)$ to the best outcomes under $F^k_{t-1,\tau}(\cdot)$, and so on, and experiences sensations of gain or loss for each of these comparisons. Suppose, for example, that she had believed she has a 50-50 chance of spending either 0 or 100 minutes with Johnny Depp. If she now believes that she has a 50-50 chance of spending either 20 or 120 minutes in Johnny's company, she experiences a gain of 20 minutes. If she now believes that she has a 50-50 chance of having either 20 or 80 minutes with Johnny, she experiences "mixed feelings" of a gain from comparing 20 minutes to 0 and a loss from comparing 80 minutes to 100. If she now believes she has a 60-40 chance at either 0 or 100 minutes, that feels like a 10% chance of losing 100 minutes. And if her beliefs were correct and uncertainty is resolved in period t, with probability 1/2 she compares a deterministic 0 minutes to the fifty-fifty 0/100 lottery—which feels like a one-half chance of losing 100 minutes with Johnny—and with probability 1/2 she compares a deterministic 100 minutes to the lottery—which feels like a one-half chance of gaining 100 minutes with Johnny.

Formally, for any distribution F over $\mathbb{R}$ and any $p \in (0,1)$, let $c_F(p)$ be the consumption level at percentile p, defined implicitly by the conditions that $F(c_F(p)) \geq p$ and $F(c) < p$ for all $c < c_F(p)$. We define gain-loss utility from the change in beliefs in dimension k as

$$N^k(F^k_{t,\tau}|F^k_{t-1,\tau}) = \int_0^1 \mu\left(m^k(c_{F^k_{t,\tau}}(p)) - m^k(c_{F^k_{t-1,\tau}}(p))\right) dp,$$

where $\mu(\cdot)$ is a "universal gain-loss utility function" with the following properties:

*A*0. $\mu(x)$ is continuous for all x, twice differentiable for $x \neq 0$, and $\mu(0) = 0$.

*A*1. $\mu(x)$ is strictly increasing.

*A*2. If $y > x \geq 0$, then $\mu(y) + \mu(-y) < \mu(x) + \mu(-x)$.

*A*3. $\mu''(x) \leq 0$ for $x > 0$ and $\mu''(x) \geq 0$ for $x < 0$.

*A*4. $\frac{\mu'_-(0)}{\mu'_+(0)} \equiv \lambda > 1$, where $\mu'_+(0) \equiv \lim_{x\to 0} \mu'(|x|)$ and $\mu'_-(0) \equiv \lim_{x\to 0} \mu'(-|x|)$.

Properties *A*0-*A*4, first stated by Bowman, Minehart, and Rabin (1999), correspond to Kahneman and Tversky's (1979) explicit or implicit assumptions about their "value function" defined on the difference between an outcome and the reference point. Loss aversion is captured by *A*2 for large stakes and *A*4 for small stakes, and diminishing sensitivity is captured by *A*3. While the inequalities in *A*3 are most realistically considered strict, to characterize the implications of loss aversion without diminishing sensitivity as a force on behavior, we define a subcase of *A*3:

$A3'$. For all $x \neq 0, \mu''(x) = 0$.

When we apply $A3'$ below, we will parameterize μ as $\mu'_+(0) = \eta$ and $\mu'_-(0) = \lambda\eta > \eta$, so that η can be interpreted as the weight attached to gain-loss utility, and λ as the coefficient of loss aversion.[8]

Finally, we assume that total gain-loss utility in period t is simply the sum of gain-loss utilities in each dimension: $N(F_{t,\tau}|F_{t-1,\tau}) = \sum_{k=1}^{K} N^k(F^k_{t,\tau}|F^k_{t-1,\tau})$.

Having specified per-period utilities, we assume that in period t the person wishes to maximize

[8] Most applications of reference-dependent utility rely on the—clearly correct—assumption that $\lambda > 1$, and we also make this assumption. Nevertheless, it is worth noting that our theory also generates some interesting results with $\lambda = 1$. In the model of Section 5, for instance, the decisionmaker overconsumes and responds asymmetrically to surprises regarding wealth even when $\lambda = 1$. In fact, if $\lambda = 1$, our model in Section 5 is observationally equivalent to the $\beta - \delta$ model presented in Laibson (1997), with $\beta = (1 + \gamma\eta)/(1 + \eta)$ and $\delta = 1$.

the sum of utilities,[9]

$$U^t \equiv \sum_{\tau=t}^{T} u_\tau. \tag{2}$$

For notational as well as substantive reasons, our formulation of total utility assumes no discounting. Our theory to some extent allows for uncertainty-based "heuristic discounting," where the decisionmaker puts lower weight on a future date because she assigns positive probability that something renders her modeled decisions for that date irrelevant. But whereas heuristic discounting can capture any such uncertainty in classical utility-maximization models because it does not affect predictions for the modeled contingency, the same is not generally true in our model: optimal planning for the modeled contingency could depend on the person's beliefs about what happens in unmodeled states of the world. Although unlikely to be of any calibrational interest, to the extent that there is time-consistent "hedonic discounting"—whereby a person simply cares less about how happy she is in the future—this will not substantively affect most of our results.[10]

While the above defines the person's utility as a function of consumption and beliefs, a full model must also specify how those beliefs are formed. Our formulation ignores sundry types of prevalent errors in belief formation, and as a conceptually useful starting point is premised on consistency with rationality: the person correctly anticipates the implications of her own plans, and cannot make plans she knows she will not carry through. Since she first makes focused plans in period 0, this means that the beliefs F_0 through F_T must be rational. But we do not impose any restrictions on initial beliefs, F_{-1}, so that we do not make any assumptions about the reasonableness of a person's initial unfocused beliefs.

More precisely, our theory requires that $F_0, \ldots, F_T$ be determined by an optimal consistent plan (OCP), a plan that in each period maximizes expected reference-dependent utility given the expectations generated by the plan (or, in period 0, F_{-1}, with the constraint that future behavior and plans must be similarly optimal. To formalize this while still suppressing a fair amount of

[9] Welfare statements below, including examples of suboptimal behavior, are made with respect to these time-consistent preferences.

[10] For some of our results assuming that $\gamma_{t,\tau} < 1$, this is true only so long as the person's discount factor is not too small relative to the rate with which $\gamma_{t,\tau}$ decreases as the outcome moves further into the future. It is more likely that incorporating a time-inconsistent taste for immediate gratification would substantially change some of our results, but we have not worked out how.

cumbersome notation, let $\{D_t\}_{t=0}^T$ be the set of feasible plans beginning each period t, which depend on both prior choices and exogenous stochastic events.[11] For given initial beliefs F_{-1}, an optimal consistent plan is defined as follows:

Definition 1. Define the sets $\{D_t^*\}_{t=1}^T$ in the following backward-recursive way. A plan $d_t \in D_t$ is in D_t^* if, given the expectations generated by d_t, in any contingency, (i) it prescribes a continuation plan in D_{t+1}^* that maximizes the expectation of U^t; and (ii) it prescribes an action in period t that maximizes the expectation of U^t, assuming that future plans are made according to (i). A plan $d_0 \in D_0$ is an *optimal consistent plan (OCP)* if it satisfies the above given the expectations F_{-1}.

To summarize, our model predicts that a person will play an OCP, which makes a (generically unique) prediction based on the primitives of the decision environment from the moment of first focus to full execution, the preferences, and the initial beliefs.[12] Practical applications of the model, however, are often likely to require predictions using more limited information about the person's choice and characterization of the problem. Suppose that we "look in" at a person in period t, knowing her utility function and choice problem for all periods t and on—including any probabilistic beliefs she has about the choice problem in the morning of period t. What are the predictions we can make as a function of possible ancillary assumptions about her prior decisionmaking environment and preferences?

The least restrictive predictions we might make are the following:

Definition 2. A plan $d_t \in D_t$ is *reference-dependent rational* (RDR) if there exists a decision problem, weights $\gamma_{0,\tau}, \ldots, \gamma_{t-1,\tau}$ for all τ with $t \leq \tau \leq T$, and initial beliefs F_{-1} such that in OCP the decision problem D_t is on the support in period t and generates continuation plan d_t.

[11] Our notation and definition suppresses that (i) the beliefs F_t depend on the available information and are generated by the entire state-contingent plan; (ii) actions can depend on the entire history, including past beliefs; and (iii) expected utility is taken over the entire sequence of outcomes and beliefs.

[12] Our model is a single-person but multi-self special case of a *dynamic psychological game* as introduced by Battigalli and Dufwenberg (2005) for multi-player games, and our definition of OCP is theoretically very similar to the solutions concepts they propose. Although we have not fully explored whether our ideas can be coded in their very rich framework, it seems that there are two important differences between our model and theirs. First, whereas we assume that a person can always affect her future selves' preferences merely by changing plans (so long as those plans are credible), in Battigalli and Dufwenberg's model this is only possible through taking different actions. This difference is natural since a person is likely to be able to choose her favorite plan for her *own* behavior but not for *others'* behavior. Second, as discussed above, we do not impose any restrictions on the reasonableness of a person's initial beliefs.

The set of RDR plans is the set of plans that we would deem possible if we knew the decision-maker's situation and preferences from the time she wakes up on date t, but knew nothing about her initial beliefs, her earlier choice sets and behaviors, and so on. This would be the set of behaviors consistent with our model when researchers are fully agnostic about prior preferences and any features of the environment that would not typically be included in the specification of a classical model. Indeed, once assumptions are made about the structures of $\gamma_{t',\tau}$ for all $t \leq t' \leq \tau \leq T$, the shape of μ, and the consumption dimensions, RDR is the set of predictions consistent with our model using solely the person's consumption utility and decision problem starting in period t—which is exactly the information used in conventional analysis.

We illustrate RDR and later concepts using the most important application in Kőszegi and Rabin (2006): shoe purchases. Suppose that in period t, a consumer might face a wide range of prices for shoes. Although RDR does not place restrictions on what the person had expected coming into period t, we know she will buy for very low prices and not buy for very high prices. More generally, we can say that the consumer uses a cutoff strategy, buying below a reservation price and not buying above it. But the reservation price can be anywhere within a range. If the consumer had expected not to buy, for example, she will experience paying as a loss and getting the shoes as merely a gain, so her reservation price will be low; if she had expected to buy, but at a low price, her reservation price would be moderate; and if she had expected to buy the shoes at a high price, her reservation price will be high.

If we are willing to make assumptions about the person's beliefs in period $t-1$ regarding the impending decision problem, a stronger solution concept is appropriate. In the case of shoe purchases, for instance, we would want to specify the person's previous beliefs about the probabilistic price distribution she would face. The solution concept to apply in this case is:

Definition 3. A plan $d_t \in D_t$ is a *personal equilibrium* (PE) if $d_t \in D_t^*$.

We and others have previously used the term "personal equilibrium" as a solution concept defining reasonable planned behavior when people have beliefs-based preferences, and we use the same term here in a new context because our new definition shares the essential characteristics of the earlier ones. Namely, we assume that the person not only maximizes her utility given her beliefs, but also that her beliefs must turn out to be probabilistically correct given her own behavior and exogenous happenings. More precisely, the set of personal equilibria is the set of credible plans in the sense that if the person believed at the end of period $t-1$ that she would carry through the plan, she would indeed want to carry it through.[13] Despite the term "equilibrium", this amounts only to a rational-consistency requirement: the person can only make plans she believes she will follow through.

The restrictions PE places on possible prior beliefs relative to RDR can be illustrated using our example of shoe purchases above. As we have argued, the person always buys the shoes for very low prices. PE requires that she realize this and incorporate it into her expectations, which in turn means that when the probability of a low price is substantial, PE requires that she experience a loss from not getting the shoes, ruling out relatively low reservation prices.

While any PE plan is credible, different PE plans may generate different ex-ante expected utilities starting in period t. Our third, most restrictive solution concept is that the person chooses the (typically unique) PE yielding the highest utility:

Definition 4. A plan $d_t \in D_t$ is a *preferred personal equilibrium* (PPE) if it is a PE, and it maximizes expected utility starting in period t among PE plans.

Because it yields the highest expected utility beginning in period t among plans she knows she will carry through, it may appear that a person would always plan to play the PPE. Indeed, we have invoked exactly this intuition in our previous research to motivate the use of PPE. But the model here says that this justification is not always appropriate because it ignores the possibility of prospective gain-loss utility: if one has to *change* one's plans—and incur possible disappointments in some dimensions of future consumption—to get things "right," choosing the PPE may be unattractive.

Yet PPE is the appropriate solution concept whenever two conditions are satisfied. First, it must be the case that the decisionmaker had known about the distribution of choice sets in a period t' when her γ's for outcomes starting in period t are zero (that is, $\gamma_{t',\tau} = 0$ for all $\tau \geq t$). This is

[13] Notice that although D_t^* is introduced as part of the definition of OCP (Definition 1) beginning in period 0, its definition requires only the person's decision problem starting in period t.

likely to be the case if she had known about her decision problem very early, when she cared little about changing plans. In addition, while we do not formalize issues of focusing in this paper, it seems that updating unfocused beliefs (in period $t' = 0$) generates little or no immediate gain-loss utility. In either case, the person can update her plans in period t' at no cost, and so would like to form PPE plans. A second requirement is that having formed a PPE plan, the person does not want to switch to another PE plan before period t. This condition is less likely to be binding, because if the person does not like to play the other PE in period t, it will typically not be pleasant to switch to it.[14]

Our new model also allows us to clarify the circumstances under which each of the reduced-form concepts we have introduced and applied in previous papers (Kőszegi and Rabin 2006,2007)—which use the same terms PE and PPE—might be appropriate. To link those single-decision, single-outcome models to the current one, suppose that all non-trivial consumption outcomes occur in period $t \geq 1$, and all uncertainty about this outcome is resolved in the same period. The person selects the lottery determining this outcome from the choice set D in a given period $t' \leq t$. Because consumption outcomes are realized and uncertainty is resolved in period t, the decisionmaker's utility collapses to a simple form. Let F be the lottery of outcomes the person receives in period t, and G the previous period's expectations regarding those outcomes. The decisionmaker's utility in period t is then

$$U(F|G) = \sum_{k=1}^{K} E_{F^k} \left[m^k(c^k) + N^k(c^k|G^k) \right].$$

The following proposition states how to check for PE and PPE in this situation.

Proposition 1. *Suppose all outcomes occur and all uncertainty is resolved in period t, and the outcome lottery is the result of a single choice from the choice set D. If the choice is made in period t, then (i) $F \in D$ is a PE distribution of outcomes if and only if $U(F|F) \geq U(F'|F)$ for all $F' \in D$; (ii) $F \in D$ is a PPE distribution of outcomes if and only if it is a PE, and there is no PE distribution $F'' \in D$ such that $U(F''|F'') > U(F|F)$; and (iii) if A3$'$ holds and there is a $t' < t$ such that $\gamma_{t',t} = 0$, then OCP induces a PPE starting in period t. If the choice from D is made in period*

[14] The condition is met, for example, for all of the types of situations we consider in Section 5 of this paper, even if the relevant decision is made in some period $t > 1$. When A3$'$ holds, the condition is also met in the single-decision setting of Proposition 1 below.

$t' < t$ and $\gamma_{t',t} = 0$, $F \in D$ is an OCP distribution of outcomes if and only if $U(F|F) \geq U(F'|F')$ for all $F' \in D$.

Our old notion of PE is a situation where the stochastic outcome of a person's utility-maximizing choice given expectations is identical to the expectations (Kőszegi and Rabin 2006, Definition 1). Lemma 1 says that if the person makes a single decision in period t with consumption also occurring in period t, our current notion of PE is equivalent to the old notion. In the same situation, our old notion of PPE—the PE that maximizes ex-ante expected utility (Kőszegi and Rabin 2006, Definition 2)—is equivalent to the current notion. Furthermore, if A3′ holds and there is a period in which the person can update her beliefs at no hedonic cost, then OCP indeed implies PPE because the two conditions discussed above are met.

In Kőszegi and Rabin (2007), we introduced the solution concept *choice-acclimating personal equilibrium* (CPE) for situations in which a person commits to a lottery to be resolved in the future (Definition 3). A lottery is a CPE if it maximizes expected reference-dependent utility given that it determines both the reference lottery and the outcome lottery. Lemma 1 says that this is necessarily appropriate only if the relevant γ in the period of choice is zero. Similarly to PPE, our previous justification for CPE—that when choosing ahead of time the person commits to the lottery that maximizes expected utility at the time of resolution—ignored prospective gain-loss utility: it could be that changing plans to the option that maximizes expected utility in period t is painful, and hence the person will not do it.

3 Information Preferences

In this section, we explore some significant implications of the most basic premise of our model—loss-averse preferences over changes in beliefs—for how a person feels about information regarding fixed but unknown future consumption. As a complement to models—such as Caplin and Leahy (2001, 2004), Caplin and Eliaz (2003), and Kőszegi (2006a)—that assume directly a taste or distaste for information, we *derive* such tastes from the same preferences towards good and bad news that (as we interpret is) underlie prospect theory, and emphasize how a person's like or dislike of information may depend on features of the information and the environment. Our results follow mostly from two key principles: that people prefer to get information clumped together rather than apart, and that—if prospective gain-loss utility weakens with time lag to the outcome—people prefer to get information sooner rather than later.

To isolate the implications of loss aversion, we assume that $\mu(\cdot)$ satisfies A3$'$, so that $\mu(x) = \eta x$ for $x \geq 0$ and $\mu(x) = \eta\lambda x$ for $x < 0$. Because the general statements below are notationally somewhat cumbersome, we illustrate the intuition for many of our results in a simple example. Suppose $T = 2$, there is no consumption in period 1, and there is one dimension of consumption in period 2, with $m(c_2) = c_2$. There are two equiprobable possible consumption levels, $c_2 = 0$ and $c_2 = 1$, and the decisionmaker has no control over this outcome. She may, however, receive information about c_2 in period 1. Specifically, she may observe a signal $s \in \{0, 1\}$, where the signal is accurate ($s = c_2$) with probability $q > 1/2$. We investigate how observing the early signal affects the decisionmaker's expected utility as a function of q and the strength of her concern for prospective gain-loss utility, $\gamma \equiv \gamma_{1,2}$.

If the person observes the signal, her expected gain-loss utility is[15]

$$-\frac{1}{2}\gamma\eta(\lambda - 1)\left(q - \frac{1}{2}\right) - q(1-q)\eta(\lambda - 1). \tag{3}$$

The first term captures expected prospective gain-loss utility in period 1. After observing the early signal, the decisionmaker will either believe the high outcome $c_2 = 1$ happens with probability q—leading to a gain of $q - 1/2$ as compared to her prior—or she will believe it happens with probability $1 - q$—leading to a loss of $q - 1/2$ as compared to her prior. Because the loss is more heavily felt, expected prospective gain-loss utility is negative. The second term in Expression 3 captures expected gain-loss utility in period 2. With probability 1/2, the person leaves period 1 believing consumption will be high with probability q. In that case, with probability q she later learns that $c_2 = 1$—leading to a gain of $1 - q$—and with probability $1 - q$ she learns that $c_2 = 0$—leading to a loss of q. The expected utility from this possibility is therefore $-q(1-q)\eta(\lambda - 1)/2$. Similar considerations apply if she leaves period 1 assigning probability $1 - q$ to the high outcome.

If the decisionmaker does not observe the signal, her expected gain-loss utility is $-\eta(\lambda - 1)/4$.

[15] Since consumption utility is independent of the arrival of information, we focus on gain-loss utility.

Hence, observing the signal generates strictly more expected utility than not observing it if and only if

$$\gamma < 2\left(q - \frac{1}{2}\right). \tag{4}$$

Suppose first that $q = 1$, so that the signal provides the same information the decisionmaker would otherwise learn in period 2. By Inequality 4, if $\gamma < 1$, she strictly prefers to receive the information early, and if $\gamma = 1$, she is indifferent. Intuitively, given that there is an equal chance of information moving beliefs up and down, loss aversion implies that the decisionmaker finds it unpleasant in expected terms to learn information. If $\gamma < 1$, the sense of loss for non-immediate outcomes is not as large, so the person is better off receiving the information early. If $\gamma = 1$, the sense of loss is exactly as aversive in period 1 as in period 2, so the person is indifferent to the timing of information.

In addition to liking early information, the person dislikes partial information. For the key intuition, suppose that $\gamma = 1$. Then, for any $q < 1$ she prefers not to receive the signal. Although the decisionmaker does not care when she experiences a given change in beliefs, piecemeal information exposes her to fluctuations in her beliefs: there is a possibility, for instance, that she receives a positive signal and her hopes rise, but then she is all the more disappointed by getting the worse of the two outcomes. Since the pleasure from receiving positive news is smaller than the pain from finding out that the news was incorrect, such fluctuations in beliefs decrease expected utility.

When $\gamma < 1$ and $q < 1$, the decisionmaker faces a tradeoff between her taste for early information and her distaste for partial information. For "weak" news, the latter effect dominates: for any $\gamma > 0$, the decisionmaker dislikes sufficiently weak early information. Intuitively, a very weak signal causes a small but first-order immediate change in beliefs, which loss aversion makes aversive in expected terms. While this early information decreases the expected change in beliefs in the future, it does so only by a second-order effect: it moves beliefs a small amount towards what eventual beliefs will be with slightly higher probability than it moves beliefs away from what eventual beliefs will be.

We now generalize these intuitions. Suppose that consumption occurs solely in period T, but that a person is potentially receiving information about this consumption in periods 1 through $T-1$.

We call any non-identity mapping from all possible prior beliefs to posteriors a "signal." Consider a sequence σ of signals about consumption, $s_1, s_2, ..., s_J$, and let $t(s_j|\sigma)$ denote the time that signal s_j is received under σ, with $t(s_j) \leq t(s_{j+1})$ for all j. This specification fits many information-acquisition scenarios, and we do not assume, for instance, that signals are independent. Also note that in applying the propositions below, it is useful that "two" signals arriving in the same period can be labeled as either one signal or two. We assume, however, that consumption in period T is binary and the signals are discrete, and as above the two possible outcomes are $c_T = 0$ and $c_T = 1$; we conjecture that the gist of our results are true more generally, but have not found an appropriate way to characterize and prove them.

To be able to compare information structures, we say that an information structure σ' is (t_a, t_b, j)-equivalent to σ if 1) σ and σ' involve the same sequence of signals, 2) in both σ and σ' exactly the two signals s_j and s_{j+1} arrive between periods t_a and $t_b > t_a$ (inclusive), and 3) for all $i \neq j, j+1$, $t(s_i|\sigma') = t(s_i|\sigma)$. That is, σ' and σ differ solely in the timing of the two signals s_j and s_{j+1}. Under a given set of preferences, let $U(\sigma)$ and $U(\sigma')$ be the discounted expected utilities for the two information structures.

Formalizing the intuition above that receiving all information at the beginning rather than in pieces is welfare-increasing, Proposition 2 says that collapsing two signals into one, so long as that does not delay the signals, always strictly improves welfare:

Proposition 2. *Suppose that σ' is (t_a, t_b, j)-equivalent to σ with $t(s_{j+1}|\sigma') = t(s_j|\sigma') \leq t(s_j|\sigma) < t(s_{j+1}|\sigma)$. Then $U(\sigma') > U(\sigma)$ for any $\gamma_{t,T} > 0$ non-decreasing in t.*

By iteratively applying this proposition, it is clear that any change in information structure that collapses different signals without delaying any of them will raise utility for the person.

Our second proposition generalizes the point that receiving information earlier increases welfare if $\gamma < 1$, and does not affect welfare if $\gamma = 1$:

Proposition 3. *Suppose that σ' is (t_a, t_b, j)-equivalent to σ with $t(s_j|\sigma') < t(s_j|\sigma)$, $t(s_{j+1}|\sigma') \leq t(s_{j+1}|\sigma)$, and $t(s_{j+1}|\sigma') = t(s_j|\sigma')$ if and only if $t(s_{j+1}|\sigma) = t(s_j|\sigma)$. Then $U(\sigma') > U(\sigma)$ when $\gamma_{t,T}$ is strictly increasing in t and $U(\sigma') = U(\sigma)$ when $\gamma_{t,T} = 1$ for all t.*

This proposition says that, so long as it does not change the order of signals (including which signals arrive simultaneously), getting signals earlier rather than later is better if prospective gain-loss utility weakens with time to the outcome, and is equally good if prospective gain-loss utility does not weaken with time to the outcome. Combined with the previous proposition, the former point implies that when prospective gain-loss utility is always as strong as contemporaneous, all the person cares about is to avoid the "dribbling in" of information—collapsing signals, even when this involves delaying the signals, is always a good thing. The two propositions also imply that learning everything right away is always at least as good as any other information structure.

For the final two propositions in this section, we assume that the signals s_1 through s_J are independent conditional on c_T. To state the next proposition, we call a signal *always informative* if for any realization s of the signal, $\Pr[s|c_T = 1] \neq \Pr[s|c_T = 0]$. The proposition establishes that, no matter how weak is prospective gain-loss utility, a sufficiently weak signal always harms welfare by a little. Although the proposition could be (more clumsily) stated for any small signal, for simplicity we consider symmetric binary signals.

Proposition 4. *Choose any* $\gamma_{t,T} > 0, j, t_a, t_b, \{s_i\}_{i \neq j}$, *and* (t_a, t_b, j)*-equivalent* σ *and* σ' *with* $t(s_j|\sigma) < t(s_{j+1}|\sigma) = t(s_j|\sigma') = t(s_{j+1}|\sigma')$*. Suppose* s_j *is a binary signal with accuracy* $1/2 + \epsilon$ *and* s_{j+1} *is always informative. Then, if* ϵ *is sufficiently small,* $U(\sigma') > U(\sigma)$.

But while our model implies a dislike of *isolated* inaccurate information, the logic of the model also predicts a very different attitude towards additional weak information when a person has just been or is about to be exposed to other information: in this case, she is eager to receive more information immediately. To state this result, say that a signal is of size $\epsilon > 0$ if the maximum change in the probability assigned to the high outcome resulting from the signal is ϵ. Proposition 5 says that a person always strictly prefers information of a size smaller than her recent change in beliefs, even if $\gamma = 1$, the information is very inaccurate, and receiving it means splitting it away from a later signal:

Proposition 5. *Let* σ *and* σ' *be* $(t_a, t_b, 1)$*-equivalent with* $t(s_1|\sigma') = t_a < t(s_1|\sigma)$ *and* $t(s_2|\sigma') = t(s_2|\sigma)$*. Suppose the decisionmaker has received information in period* t_a *that changed her subjective probability of getting* $c_T = 1$ *from* p *to* p'*. If* s_1 *is of size less than* $|p - p'|$*,* $U(\sigma') > U(\sigma)$.

For an intuition, suppose, say, that the person has just received information that has decreased her beliefs by a given amount. Any further information that is of smaller size will not move her beliefs back to their original level. Hence, since any news will only change the degree of loss she suffers, but not whether she suffers a loss or a gain, further positive news are evaluated just as strongly as further negative news. Therefore, these news have zero immediate expected utility. Yet these news decrease expected future fluctuations in beliefs, increasing expected utility.

An implication of Proposition 4 is that the decisionmaker hates to receive information in the form of a sequence of small news. This prediction of our model may be a problematic one: for instance, many investors follow the performance of their stock portfolios on a day-to-day basis, and many sports fans follow the online ticker during a game rather than look only at the final score. This kind of behavior seems partly to be due to motivations, such as a curiosity-induced inability to avoid available information or the pure enjoyment of following a sports game, our model ignores. But even in our setting, Proposition 5 qualifies when exactly a decisionmaker dislikes small news. If an investor can avoid all information regarding her retirement wealth, for example, she prefers not to monitor her prospects too closely. But if she receives some unavoidable information, she will immediately look for additional information herself. Even if the first of these predictions turns out to be innaccurate because of alternative motivations not included in our model, the comparative-static prediction—that the more information a person receives in a given period, the more willing she is to receive additional information—is likely to be robust to other motives.

A simple extension of our results in this section can be used to formalize and extend a key point by Kimball and Willis (2006). They argue that since news about consumption affects immediate happiness, changes in happiness following news can be used to infer a person's consumption utility. In our model, not only happiness, but information-seeking *behavior* can be used to identify consumption utility. While we have for simplicity set $m(0) = 0$ and $m(1) = 1$, in each of the above instances it is clear that the decisionmaker's like or dislike of information is proportional to $m(1) - m(0)$. Hence, how much the person is willing to pay to receive or avoid information about a particular outcome reveals how much she cares about that outcome, even if she never makes any choice that affects the probability of the outcome. This means, in principle, that if we have identified the nature of a person's prospective and contemporaneous gain-loss utilities, we may be able to use "revealed" preference over information to identify her consumption utility for outcomes over which she has no control.

4 Monetary Preferences

Many previous models of reference-dependent utility assume—presumably as a shortcut—that individuals care directly about receiving money or experiencing changes in wealth. In this section, we show how our model can provide consumption-based foundations for such preferences, and use these foundations to derive some patterns in monetary preferences. To motivate some insights, we begin with a puzzle regarding the psychology of money that has been noted by researchers, and was elaborated most clearly by Barberis, Huang, and Thaler (2006) and Kőszegi and Rabin (2007). Because the typical person holds wealth primarily for future consumption and faces substantial uncertainty regarding future wealth—so that modest changes in current wealth are unlikely to determine whether she ends up above or below her reference point in the future—it would seem that even a loss-averse person would do much better maximizing expected value over modest stakes rather than sticking with the complicated pattern of reference-dependent behavior she does exhibit. Hence, a non-neutral attitude toward small risks seems to require "narrow bracketing"—ignoring that the current risk will be integrated with substantial other risk—as well as substituting a complicated suboptimal pattern of behavior for a simple near-optimal one.

By defining gain-loss utility over changes in beliefs, our approach provides a new perspective on these issues. While some people clearly do bracket narrowly in ways that our model fails to explain, we predict that people may care about small changes in wealth even if they recognize that the changes contribute negligible risk to the consumption ultimately determined by their wealth. The basic reason is simple: gains and losses in money are news about future consumption, and this news generates immediate prospective gain-loss utility that looks similar to gain-loss utility over wealth. But our specification also predicts how the timing of news about a risk affects a person's attitude toward that risk. This means that—elaborating on some of the points about timing made in Kőszegi and Rabin (2007)—our model provides a unifying framework for determining whether and when various existing reduced-form models of money apply. The model also predicts when rationally accounting for background risk would in fact eliminate any significant aversion to moderate amounts of additional risk, rendering narrow bracketing a mistake.

We develop our formal results for situations where the timing of consumption is fixed and a large part of the consumption uncertainty is exogenous, isolating how the person treats risks to future consumption.[16] Suppose $T = 2$, $K = 1$, non-trivial consumption only occurs in period 2, and $m(c_2) = c_2$. In line with our arguments that updating unfocused beliefs has no hedonic consequences, we assume that $\gamma_{0,1} = \gamma_{0,2} = 0$. The decisionmaker's consumption in period 2 is the sum of two components: "background risk" with a fixed and known distribution, and the outcome resulting from her choice from the set D of independent small risks to period-2 consumption. The decisionmaker might have known since period 0 that she would make a choice from D, or—in what can be thought of as a "surprise" situation—she might have believed she would only be facing the background risk, and then find in period 1 or 2 that she must choose from D.[17] In addition, the decisionmaker may have to implement her choice in period 0, 1, or 2, and the uncertainty in the lotteries in D might be realized in period 1 or 2. Appendix A characterizes behavior formally as a function of all these features of the environment, and here we discuss intuitively some key features of the results.

A stark case illustrating some implications of the model is when the person is confronted with D, implements her decision, and learns the outcome all in period 1. Then, relative to her previous expectation of not facing any risk other than the background risk, the outcome shifts the distribution of future consumption exactly by the realized outcome of the chosen lottery, and hence generates a gain or loss equal to that realized outcome. This means that *independently of the background risk*, the decisionmaker chooses from the set D as postulated by prospect theory, maximizing reference-dependent utility from money receipts with a reference point of zero. In this sense, our model says

[16] We speculate at the end of the section how endogenous consumption choices of the type we explore in Section 5 would affect loss aversion over money. But because the path of future consumption in dynamic settings is quite complicated, we do not know to what extent many other results extend to a more general model.

[17] This latter environment is the limiting case of situations where the decisionmaker had expected to face only the background risk with probability $1 - \epsilon$ and to also choose from D with probability ϵ, and the low-probability event is realized.

that some forms of "narrow bracketing" are not necessarily errors. Whatever the background risk, losing \$10 conveys the bad news that one will have less to consume in one's lifetime, and in our theory this fully and rationally generates a sense of loss.

Although the exact form of behavior is a little more complicated, the decisionmaker also attends to the gains and losses resulting from her choice if she learns D and implements her decision in period 0, and uncertainty is resolved in period 1. In this case, she anticipates that the resolution of uncertainty will generate utility-inducing news in period 1, so whether or not there is background risk she chooses from D according to something like our CPE concept: she maximizes expected reference-dependent utility taking into account that both the reference lottery and the outcome lottery are determined by her choice from D. The complication is that the way in which she evaluates outcomes relative to the reference lottery depends on the background risk. For some forms of background risk, the decisionmaker evaluates each outcome relative to all possibilities in the reference lottery, as proposed by Kőszegi and Rabin (2006, 2007). But for other forms of background risk, she evaluates an outcome relative to the mean of the reference lottery, as in the disappointment-aversion models of Loomes and Sugden (1982) and Bell (1985). In either case, she is first-order risk averse with respect to the options in D.

Similar considerations arise if the decisionmaker learns D in period 0, implements her decision in period 1, and uncertainty in D is also resolved in period 1. Then, both a deviation of her own behavior from expected, and the resolution of uncertainty in period 1, generate news that shifts the distribution of period-2 consumption and induces immediate gain-loss utility. Hence, when choosing from D in period 1, the person attends to gains and losses relative to her expectations regarding her choice. This means that she behaves according to something like our static PPE concept as applied to D: she makes the best plan she knows she will carry through. As above, however, how she evaluates an outcome relative to the reference lottery depends on the nature of the background risk.

In all these situations, the uncertainty in D is resolved in period 1, and in various ways the decisionmaker cares about gains and losses in money. The common thread is clear: although the decisionmaker fully understands the fungibility of money and that realizations of wealth in period 1 have no immediate consumption implications, since she cares about the changes in beliefs induced, she attends to those realizations.

But because our theory posits that only temporally isolated news generate separate gain-loss utility, it typically predicts approximate risk neutrality when uncertainty in D is resolved in period 2 and is integrated with large background uncertainty.[18] Our theory therefore fails to explain important instances of narrow bracketing that have been observed. For instance, Tversky and Kahneman (1981) and Rabin and Weizsäcker (2007) found that subjects narrowly bracket two separate pairwise choices even when they know the two chosen lotteries would be played out and reported at the same time. But while inconsistent with the fully rational framework we develop, these mistakes seem more interpretable when viewed through the lens of our model than with either classical reference-independent preferences or other reference-dependent theories. Since we learn about the consequences of many or most of our decisions separately, and "narrow bracketing" in these situations is *not* a mistake in our model, mistaken narrow bracketing may naturally result when people heuristically treat even simultaneous decisions as if they would be informed about the outcomes in isolation.

By dint of providing consumption-based foundations for loss aversion in money, our model establishes an endogenous relationship between loss aversion over goods and loss aversion over money. If $\gamma_{t,\tau} < 1$ for $\tau > t$, a person will be less loss averse over money than over immediately consumed goods, as suggested for example by Novemsky and Kahneman (2005). This contrasts with our single-decision model in Kőszegi and Rabin (2006), where we have assumed equal loss aversion across dimensions, including money. While underlying gain-loss utility in our dynamic setting is still a single gain-loss utility function $\mu(\cdot)$, the dynamic aspect introduces differences in loss aversion according to differences in the timing of consumption.

Incorporating an endogenous choice regarding the timing of consumption can further weaken loss aversion over money, and may also affect other results in this section. As we show in the next section, when $\gamma < 1$ a person may adjust consumption to increases in wealth immediately but to decreases in wealth only later. Because $\gamma < 1$, this decreases her sensitivity to losses but not gains,

[18] As we explain in the appendix, the only exception can happen when the decisionmaker learns D in period 1, because in that case the news of learning D can itself generate risk preferences.

weakening and in some cases eliminating (but never reversing) loss aversion.

5 Wealth and Consumption in Intertemporal Choice

In this section we explore the pattern of consumption in simple two-period consumption-savings decisions with and without uncertainty regarding wealth. We elaborate on a new form of overconsumption relative to the optimal committed plan that was first identified by Stone (2005), and demonstrate that a stochastic plan for period-1 consumption, and hence ex-ante uncertainty in wealth that is resolved in period 1, facilitates overconsumption. Our theory predicts an asymmetric response to surprises regarding wealth. Finally, we provide a novel and intuitive explanation for precautionary savings, while also noting that whether a person increases or decreases savings in response to future uncertainty typically depends on when she learns about the uncertainty.

Suppose a consumer needs to decide how to allocate consumption spending between periods 1 and 2, given an intertemporal budget constraint $c_1 + c_2 = W$. Her consumption utility $m(\cdot)$ is strictly increasing and strictly concave.[19] As in Section 3, for notational simplicity we let $\gamma \equiv \gamma_{1,2}$. We begin by identifying the restrictions our model imposes when the consumer knows her wealth level when choosing consumption in period 1, but we do not know anything about what she might have believed about her future consumption in period 0.

Proposition 6. *Suppose the consumer knows her wealth W in period 1. Then, the consumption path $(c_1, W - c_1)$ is RDR if and only if*

$$\frac{1+\gamma\eta}{1+\eta\lambda} \leq \frac{m'(c_1)}{m'(W-c_1)} \leq \frac{1+\gamma\eta\lambda}{1+\eta}.$$

[19] With reference-independent utility, it is generally appropriate to use an indirect utility function defined over spending as a reduced form for the solution to a decisionmaker's full optimization over multiple goods. The same is true in our model (under Assumption $A3'$) in deterministic settings, but not, for instance, when there is uncertainty about prices: because a person's expectations can affect her preferences, a single indirect utility function may not capture her preferences across all different situations. We have little sense for the general implications or the calibrational significance of such examples, however. To derive results in the simplest possible setting, we use a single-dimensional utility function.

The highest possible consumption in period 1 happens if the consumer experiences changes in c_1 as a loss and changes in c_2 as a gain, and the lowest possible consumption happens if the opposite is the case. This gives the bounds in Proposition 6.

In the rest of this section, we derive much more specific results on the allocation of consumption over time applying the PPE solution concept. Recall that this concept is implied by OCP if revising initial (unfocused) beliefs has no hedonic consequences. The sharper predictions relative to Proposition 6 result from using PPE in combination with assumptions about the decisionmaker's information in period 0, and some important results identify exactly how behavior depends on that information.

Our first and biggest goal is to explore, in settings with and without uncertainty, whether the decisionmaker chooses the consumption path that maximizes her ex-ante utility among the strategies available to her.[20] Suppose first that W is deterministic. Then, choosing $c_1 = c_2 = W/2$ maximizes both ex-ante consumption utility and ex-ante expected gain-loss utility (which is zero for deterministic plans and negative for non-deterministic plans), and so is the ex-ante optimal committed strategy. But this strategy may not be consistent. If the consumer had planned $c_1 = W/2$, then her period-1 utility for $c_1 \geq W/2$ is

$$m(c_1) + \eta(m(c_1) - m(W/2)) - \gamma\eta\lambda(m(W/2) - m(W - c_1)) + m(W - c_1). \tag{5}$$

The first and last terms constitute consumption utility in periods 1 and 2, respectively. The second term is the period-1 contemporaneous gain from consuming above plans in that period, and the third term is the period-1 prospective loss from having to plan lower future consumption as a result. Whatever the person consumes in period 1, she forms new expectations that determine her reference point in period 2, so there is no gain-loss utility in that period.

The derivative of Expression 5 with respect to c_1 evaluated at $W/2$ is

$$(1 + \eta)m'(W/2) - (1 + \gamma\eta\lambda)m'(W/2) = \eta(1 - \gamma\lambda)m'(W/2). \tag{6}$$

Hence, if $\gamma < 1/\lambda$—if the consumer cares much more about contemporaneous gain-loss utility than

[20] As we have mentioned above, we do so under the assumption that preferences are time consistent. We do not believe that our main results, especially regarding the role of uncertainty in overconsumption, would be qualitatively different if a source of overconsumption was a time-inconsistent taste for immediate gratification.

about prospective gain-loss utility—deviating from the ex-ante optimal plan increases utility ex post, so that this plan is not consistent, and any consistent plan must have $c_1 > W/2$. This over-consumption relative to the ex-ante optimal committed level is behaviorally very similar to recent models of hyperbolic discounting and present bias (Laibson 1997, O'Donoghue and Rabin 1999) and temptation disutility (Gul and Pesendorfer 2001). But in our model behavior is suboptimal for a completely different reason: whereas the decisionmaker takes the reference point as given in period 1, her awareness in period 0 that she would deviate raises her reference point for period 1, and consequently lowers her ex-ante utility in that period.

That expectations-based loss aversion can generate overconsumption is the extension to our setting of the same point made by Stone (2005). Stone assumes that deviations of current consumption from previous expectations induce gain-loss sensations, but there are no gain-loss sensations from the implied deviation of future consumption from previous expectations. Translated into our model, this means that Stone (2005) implicitly assumed $\gamma = 0$. Our theory says that Stone's implicit assumption is not without loss of generality: the weight on prospective gain-loss utility is crucial in determining whether overconsumption occurs. If $\gamma \geq 1/\lambda$, the derivative in Expression 6 is negative, so increasing c_1 above c_2 does not increase utility ex post. As emphasized above, when the decisionmaker deviates from the optimal consumption plan by increasing c_1, she experiences a contemporaneous gain and a prospective loss. Because of loss aversion, the prospective loss tends to act as an internal commitment device that discourages her from deviating. Summarizing:

Proposition 7. *Suppose that wealth W is deterministic. If $\gamma \geq 1/\lambda$, the PPE consumption path is $(W/2, W/2)$. If $\gamma < 1/\lambda$, the PPE consumption path (c_1^*, c_2^*) satisfies*

$$(1+\eta)m'(c_1^*) = (1+\gamma\eta\lambda)m'(c_2^*). \tag{7}$$

We now show that uncertainty that is resolved in period 1 exacerbates the tendency to overconsume. Suppose that wealth W is distributed according to the continuous distribution $F(\cdot)$, and that the uncertainty regarding wealth is resolved in period 1. We posit that the consumer makes a PE plan that in both periods calls for strictly increasing consumption as a function of total wealth;

a sufficiently large amount of wealth uncertainty will force the consumer to make such plans. Let these consumption functions be $c_1(W)$ and $c_2(W)$, respectively. Suppose that the realized wealth level is W, and the consumer is considering whether to change consumption from the planned level of $c_1(W)$. Since the probability that c_1 was going to be lower than $c_1(W)$ is $F(W)$, an increase in c_1 is evaluated as a gain in proportion to $F(W)$, and as a decreased loss in proportion to $1-F(W)$. Hence, the marginal utility from an increase in c_1 is $m'(c_1(W))[1+F(W)\eta+(1-F(W))\eta\lambda]$. The same deviation also lowers c_2, lowering prospective gain-loss utility in period 1. As with period-1 consumption, the probability that c_2 was going to be lower than $c_2(W)$ is $F(W)$, so that a decrease in c_2 is evaluated as a foregone prospective gain in proportion to $F(W)$, and as a prospective loss in proportion to $1-F(W)$. Because the consumer learns her wealth in period 1, there is no gain-loss utility in period 2. Hence, the total marginal disutility from a decrease in c_2 is $m'(c_2(W))[1+F(W)\gamma\eta+(1-F(W))\gamma\eta\lambda]$. In order for $c_1(W)$ to be consistent, therefore, it must be that

$$0=m'(c_1(W))[1+F(W)\eta+(1-F(W))\eta\lambda]-m'(c_2(W))[1+F(W)\gamma\eta+(1-F(W))\gamma\eta\lambda]. \quad (8)$$

Equation 8 implies that if $\gamma<1$, then $c_1(W)>c_2(W)$ for all W. Intuitively, the consumer cares more about surprises regarding c_1 than about surprises regarding c_2, so she is willing to surprise herself with an increase in consumption utility in period 1 if the cost is a similar or smaller decrease in consumption utility in period 2. Hence, no continuous plan with $c_1(W)\leq c_2(W)$ can be consistent. Although the ex-ante optimal committed plan is itself not generally to set $c_1(W)=c_2(W)$, this logic also leads $c_1(W)$ to be higher than optimal:

Proposition 8. *Suppose wealth is distributed continuously, the uncertainty regarding wealth is resolved in period 1, and PE plans $c_1^*(W), c_2^*(W)$ are strictly increasing in W. If $\gamma<1$, then $c_1^*(W)>c_2^*(W)$, and decreasing $c_1^*(W)$ in any neighborhood on the support of W would strictly increase ex-ante expected utility.*

The above analysis implies a major difference between deterministic plans and stochastic plans. For $1/\lambda\leq\gamma<1$, in fact, there is a striking result: if a wealth distribution induces the consumer to make stochastic plans, she overconsumes for *all* wealth realizations, but if she knew her wealth level in advance, she would not overconsume for *any* wealth realization. A prescribed consumption level that has high probability (or is deterministic) acts as a bright line that discourages deviations in either direction: the sense of gain from increasing consumption in a period is much smaller than the sense of loss from the corresponding decrease in consumption in the other period. But when a prescribed consumption level has low (or zero) probability, these two sensations are evaluated much more symmetrically because the possibility of higher or lower consumption was already incorporated into expectations. Hence, uncertainty in wealth exacerbates overconsumption.

Note that if the consumer learns her wealth in period 2 instead of period 1, uncertainty does not similarly undermine self-control. In that case, a plan involves a deterministic consumption level in period 1, so that if the consumer deviates by increasing c_1, she experiences that as a contemporaneous gain combined with a prospective loss due to a downward shift in the distribution of period-2 consumption. In fact, as we show below, uncertainty that is resolved in period 2 leads the consumer to *decrease* consumption in period 1.

The next result shows that when $\gamma < 1$, the consumer responds asymmetrically to surprises about wealth. To illustrate this most simply, suppose that she had been expecting to have lifetime wealth W with probability one, but at the beginning of period 1 learns (to her surprise) that her wealth is actually $W + y$. This is the limiting case of environments where the decisionmaker assigns a small probability to having wealth $W + y$.

Proposition 9. *Suppose that the consumer had expected a deterministic wealth level W, made PPE plans to consume c_1^*, c_2^* in the two periods, and in period 1 learns that her wealth is $W + y$. If $\gamma < 1$, there are constants $\underline{y} < 0$ and $\overline{y} > 0$ with the following properties:*

(i) If $0 \leq y \leq \overline{y}$, the resulting consumption pattern is $c_1^* + y, c_2^*$. If $y > \overline{y}$, consumption satisfies $c_1 > c_1^*, c_2 > c_2^*$ and is given by

$$(1 + \eta)u'(c_1) = (1 + \gamma\eta)u'(c_2).$$

(ii) If $0 \geq y \geq \underline{y}$, the resulting consumption pattern is $c_1^*, c_2^* + y$. If $y < \underline{y}$, consumption satisfies

$c_1 < c_1^*, c_2 < c_2^*$ *and is given by*

$$(1+\eta\lambda)u'(c_1) = (1+\gamma\eta\lambda)u'(c_2).$$

Proposition 9 says that (i) the person consumes sufficiently small windfalls entirely in period 1, leaving period-2 consumption unchanged; and (ii) she does not cut immediate consumption in response to small bad surprises, absorbing the entire shock in period 2. Larger increases or decreases in wealth are split between the two periods. Intuitively, $\gamma < 1$ implies that in period 1, consuming a windfall gain immediately is more pleasant than planning for higher future consumption, and unexpectedly lowering consumption immediately is more painful than planning for lower future consumption. Hence, the consumer immediately takes advantage of surprise improvements in circumstances, so as to be able to enjoy the pleasant surprise most; and she puts off absorbing negative news, so as to feel the unpleasant surprise least.

Finally, we analyze how a person responds to future uncertainty. Suppose that wealth is stochastic, and the uncertainty regarding wealth is resolved in period 2. Specifically, suppose that y is a mean-zero non-deterministic lottery, $s > 0$ is a scalar, and wealth is equal to $W_0 + s \cdot y$. We establish properties of the PPE level of savings as a function of the scale s of the risk. Although our formal results are more general, the main point can be illustrated easily by assuming that y is a binary lottery that gives 1 or -1 with probability 1/2 each. This means that wealth takes on two possible values, $W_0 + s$ and $W_0 - s$, with equal probabilities. To abstract from issues of overconsumption, for the illustration suppose that γ is sufficiently high for the optimal consumption path to be a PPE. Let $c_2 = W_0 + s - c_1$ and $c_2' = W_0 - s - c_1$ be the two possible consumption levels in period 2. Then, the consumer's expected utility in period 2 is

$$\frac{1}{2}m(c_2) + \frac{1}{2}m(c_2') - \frac{1}{4}\eta(\lambda-1)[m(c_2) - m(c_2')]. \tag{9}$$

The first two terms represent expected consumption utility, while the last term represents expected contemporaneous gain-loss utility in period 2. Consuming c_2', which has a one-half chance of occurring, feels like a loss relative to c_2, which the consumer also expected to occur with probability

1/2. The expected disutility from this loss is $\eta\lambda(m(c_2) - m(c_2'))/4$. Similarly, c_2 feels like a gain relative to the possibility of getting c_2', generating utility $\eta(m(c_2) - m(c_2'))/4$. Because of loss aversion, the overall impact of uncertainty on gain-loss utility is negative. Since the consumer learns no new information in period 1, her prospective gain-loss utility in that period is zero.

Using Expression 9, the optimal consumption path satisfies

$$m'(c_1) = \frac{1}{2}m'(c_2) + \frac{1}{2}m'(c_2') + \frac{1}{4}\eta(\lambda - 1)[m'(c_2') - m'(c_2)]. \tag{10}$$

Equation 10 extends the standard Euler equation to loss-averse preferences, and collapses to the standard Euler equation for either $\eta = 0$ or $\lambda = 1$. The non-standard term is the last term on the right-hand side, which reflects a loss-aversion-based incentive to increase savings. Intuitively, future risk exposes the decisionmaker to a sensation of loss should realized consumption utility be lower than other possible realizations. In order to decrease the pain from this loss, the consumer saves more to decrease the impact of any given wealth shock on consumption utility.

While the classical reference-independent model is also consistent with a precautionary-savings motive, we view our model as providing a potentially more intuitive reason for it. Both accounts predict precautionary savings in response to increased wealth uncertainty only if an increase in expected future consumption decreases the cost of future consumption uncertainty. But whereas the classical account predicts that increased consumption meliorates the disutility of uncertainty only insofar as higher spending makes consumption utility more linear, our model predicts the melioration happens simply because more spending lowers the marginal utility of consumption. For the same reason, loss aversion as we have modeled it also provides a more robust explanation for precautionary savings. To make this case formally, consider the Taylor-expansion approximation of the right-hand side of Equation 10 around $s = 0$:

$$m'(c_2) + \frac{1}{2}m'''(c_2)s^2 + \frac{1}{2}\eta(\lambda - 1)(-m''(c_2))s.$$

When there is no loss aversion ($\eta = 0$ or $\lambda = 1$), uncertainty has a second-order effect on savings, and whether this effect is positive or negative depends on the third derivative of the utility function. With loss aversion, the effect is first-order and unambiguously positive for any strictly concave consumption-utility function.

This feature, in fact, generalizes to other situations:[21]

Proposition 10. *Suppose wealth is equal to $W_0 + s \cdot y$, where y is a non-deterministic mean-zero lottery that is resolved in period 2. For any strictly concave $m(\cdot)$ and any $\eta > 0$, $\lambda > 1$, $\gamma \geq 0$, the PPE consumption rule satisfies $dc_1/ds|_{s=0} < 0$.*

While our framework predicts a novel reason for precautionary savings, it also implies that the decisionmaker's response to future uncertainty depends on whether she had known about the uncertainty before the savings decision. Suppose she enters period 1 having made plans anticipating a deterministic wealth level W, but then learns that wealth is uncertain and will be resolved in period 2. Because building precautionary savings would require her to decrease period-1 consumption below the expected level, for sufficiently low amounts of uncertainty she does not decrease consumption. In fact, in some situations she increases consumption.[22] Intuitively, when wealth is deterministic, the sensation of loss that would result from lowering planned consumption for period 2 can keep a person from overconsuming. If she learns that her wealth may be higher, however, she evaluates a decrease in period-2 consumption partly as a foregone gain from the possible windfall. Hence, she does not find the decrease in future consumption so aversive.

The predictions of our model differ in several ways from prominent existing dynamic specifications of reference-dependent utility and habit formation, such as Ryder and Heal (1973), Becker and Murphy (1988), and Campbell and Cochrane (1999), where the reference point for evaluating a consumption outcome is lagged consumption. The most striking difference is that in our setting increasing consumption today reduces tomorrow's reference point by reducing expectations

[21] Technically speaking, the unambiguous prediction that uncertainty increases savings is only true for small amounts of risk. For large amounts of risk, the consumer's consumption utility can dominate gain-loss utility, and in that case—as in the standard model—savings in general depends on the third derivative of $m(\cdot)$. Since $m(\cdot)$ represents a global utility function, however, a small risk in our model can still be very large in practical terms.

[22] For example, suppose that $2/(1+\lambda) > \gamma \geq 1/\lambda$—so that the consumer plans to set $c_1 = c_2 = W/2$—and she learns that her wealth has a fifty-fifty chance of being either $W + \epsilon$ or $W - \epsilon$. For a sufficiently small ϵ, she responds by increasing consumption by ϵ, fully consuming the better realization of the uncertainty. To see this, note that the consumer's utility from the perspective of period 1 when $W/2 < c_1 < W/2 + \epsilon$ is:

$$\begin{aligned} & m(c_1) + \eta(m(c_1) - m(W/2)) + \frac{1}{2}m(W + \epsilon - c_1) + \frac{1}{2}m(W - \epsilon - c_1) \\ + \quad & \frac{1}{2}\gamma\eta(m(W + \epsilon - c_1) - m(W/2)) - \frac{1}{2}\gamma\eta\lambda(m(W/2) - m(W - \epsilon - c_1)) - \frac{1}{4}\eta(\lambda - 1)(m(W + \epsilon - c_1) - m(W - \epsilon - c_1)). \end{aligned}$$

It is easy to show that for a sufficiently low ϵ, the derivative of this expression is positive.

of future consumption, whereas in lagged-consumption-based models the same raises tomorrow's reference point by raising the habitual level of consumption. As a result, our model (unlike these others) does not predict the kind of preference for increasing consumption profiles that was found in surveys by Loewenstein and Sicherman (1991) and Loewenstein and Prelec (1993). Such habit formation is realistic and could be added to consumption utility in our model, but we have not done so. Even if we did, our model (like other models of time-inconsistent preferences) calls into question whether a one-shot expression of preference over a profile is what will manifest itself in a dynamic setting. In addition, our model generates many realistic predictions that lagged-consumption-based theories do not. For instance, because in these other theories a person typically plans for a strictly increasing consumption profile, her consumption is not at the reference point in most periods, so that she does not respond to news regarding wealth in the asymmetric way that our model predicts. In addition, none of these models predict overconsumption, a role of prior uncertainty on behavior, or a preference over the timing of decision-irrelevant information.

6 Conclusion

In addition to unifying existing theories and intuitions about reference-dependent utility and making new predictions in the settings explored in Sections 3-5, we hope the model developed can be applied in many other economic settings. Because our model predicts that deterministic plans are easier to stick to than stochastic plans, it may help provide some insights into the relationship between self control, budgeting, and mental accounting. Intuitively, because deterministic plans accentuate the power of loss aversion to deter the temptation to generate pleasant surprises, the model helps explain why someone might make rigid plans in uncertain environments where the rigidity would seem to be costly. Also, because the theory predicts that expectations to acquire a good are painful to give up, it may capture the common intuition that as an auction progresses, bidders who seem to have a chance of winning get excited and bid higher. And by emphasizing the importance of plans and beliefs that occur in a person's head before any observable decisions are made, our model may provide a foundation for the relationship between decisions and contemplation.

But there are several important ways in which our model is both an incomplete and an inaccurate model of reference-dependent preferences. As discussed in Kőszegi and Rabin (2006, 2007), perhaps the greatest weakness of our theory—as well as other theories of reference-dependent utility put forward—is that it takes as a primitive the set of decisions and risks a person is focusing on. In this paper, for example, our theory predicts that a person is more likely to overconsume if she had been expecting a lot of uncertainty to be resolved just prior to her choice than if she had been expecting little. To the extent that it is often hard to know what people have been thinking before a particular decision, this kind of prediction may be hard to test.

Our theory also takes as a primitive input the set of dimensions on which gain-loss utility is evaluated separately. We discuss this issue in detail in the context of our static model in Kőszegi and Rabin (2004), but a dynamic setting introduces further complications. In particular, our formulation of gain-loss utility has at least one unattractive feature. As an example, it implies that if a person revises her consumption in 101 periods down by one apple and her consumption in 100 periods up by one apple, the former is evaluated as a loss and the latter merely as a gain, generating negative gain-loss utility. It seems plausible, however, that the latter gain fully compensates for the former loss, so that the person experiences no gain-loss utility. That is, consumption 100 periods from now may be on the same psychological dimension as consumption 101 periods from now. Finally, psychological evidence in the context of adaptation and other domains indicates that people often underestimate the extent to which changes in circumstances will change their preferences (Kahneman 1991, Loewenstein, O'Donoghue, and Rabin 2003, for example). In the context of expectations-based reference dependence, this could mean that a decisionmaker underestimates how much changes in expectations will change how she will feel in the future. Any results that are driven by a motive to manage future gain-loss utility by changing current expectations—e.g. the preference for early information in Section 3—are then likely to be weaker.

Appendix A: An Elaboration of Monetary Preferences

This appendix characterizes monetary preferences in the setting of Section 4. Beyond saying that a person may asymmetrically attend to small gains and losses in wealth even when the background risk is large, our theory can identify exactly how she treats small risks as a function of her decisionmaking environment. In fact, in each of the cases we consider in this section, the person's attitude toward small gambles can be characterized by a "reduced-form" solution concept that depends only on the choice set D of small risks.

For our characterization, we first define two reduced-form utility functions for comparing a lottery the individual faces to a reference lottery. We will apply these utility functions to lotteries from D. Denote the mean of a lottery G by $\overline{G}$. For a function $v : \mathbb{R} \rightarrow \mathbb{R}$ satisfying assumptions $A0$ through $A4$, let

$$U(F|G) = \iint v(c-r)dG(r)dF(c), \text{ and} \tag{11}$$

$$V(F|G) = \int v(c-\overline{G})dF(c). \tag{12}$$

The functions U and V are different ways of evaluating the outcome lottery F relative to the reference lottery G. Under the utility function U, each outcome of F is evaluated relative to all possible outcomes under G, and the overall evaluation is an average of these gain-loss sensations. This is the utility function we postulated in our static models of reference-dependent utility (Kőszegi and Rabin 2006,2007). Under V, each outcome of F is compared only to the mean of the reference lottery.

Based on the above utility functions, we define reduced-form solution concepts for a decision-maker's choice from the choice set D.

Definition 5. $F \in D$ is a *risk-neutral choice* (RN) if for all $F' \in D$, $\overline{F} \geq \overline{F}'$. $F \in D$ is a *prospect-theory choice* (PT) if there is a U of the form 11 such that $U(F|0) \geq U(F'|0)$ for all $F' \in D$. $F \in D$ is a *choice-acclimating personal equilibrium* (CPE) if there is a U of the form 11 such that $U(F|F) \geq U(F'|F')$ for all $F' \in D$. $F \in D$ is a *static preferred personal equilibrium* (SPPE) if there is a U of the form 11 such that $U(F|F) \geq U(F'|F)$ for all $F' \in D$, and for any $F'' \in D$ satisfying the same condition, $U(F|F) \geq U(F''|F'')$. $F \in D$ is a *Bell-Loomes-Sugden choice* (BLS) if there is a V of the form 12 such that $V(F|F) \geq V(F'|F')$ for all $F' \in D$. $F \in D$ is an *expected-value preferred personal equilibrium* (EVPPE) if there is a V of the form 12 such that $V(F|F) \geq V(F'|F)$ for all $F' \in D$, and for any $F'' \in D$ satisfying the same condition, $V(F|F) \geq V(F''|F'')$.

Any of the above concepts is *strict* if the weak inequalities in the definition are replaced with strict inequalities for all lotteries not equal to F.

Risk neutrality simply means that the decisionmaker chooses the lottery with the highest mean. As we have emphasized in Section 4, previous theories of reference dependence imply that a person approaches risk neutrality for large amounts of background risk. A person accords to prospect theory if she maximizes the expectation of a reference-dependent utility function given a reference point of zero. A choice is a CPE if it maximizes the expectation of U given that it determines both the reference lottery and the outcome lottery. As mentioned in Section 2, this is the solution concept we introduced in Kőszegi and Rabin (2007) for situations when the outcome of a lottery is resolved long after the decision. Similarly, a lottery in D is BLS if it maximizes the expectation of V given that it determines both the outcome lottery and the reference lottery. This solution concept is analogous to the disappointment-aversion models of Loomes and Sugden (1982) and Bell (1985). SPPE, first defined in Kőszegi and Rabin (2006) as PPE, is the person's favorite plan among plans she knows she will carry through if her reference point is determined by an expectation to carry through her plan and her utility function is U. Finally, EVPPE is the analogue of PPE when the person's utility function is V instead of U.

Because it gives rise to qualitatively different behavior, we consider two extreme types of background risk. One type is a uniform distribution on an interval $[x_l, x_h]$. The other is a bounded discrete distribution in which the distance between any two atoms is at least twice the absolute value of the largest possible realization of any lottery in D. While such background risk is somewhat artificial, it is of interest for two reasons. First, it captures in extreme form a situation where much of the risk a person faces comes in large increments—e.g. whether she gets a promotion—and the additional risk at hand is small relative to these increments. Second, because a decisionmaker in our model dislikes risks to consumption, she may endogenously choose consumption distributions that are discrete.

Our main interest is in determining how the decisionmaker treats small risks as the background risk to her consumption becomes arbitrarily large; that is, as $x_h - x_l \to \infty$ in the case of uniform background risk and the weight on the most probable atom approaches zero in the case of discrete

Learn	Implement	Resolved	Discrete	Uniform
0	0,1,2	2	RN	RN
0	0	1	CPE*	BLS
0	1	1	SPPE*	EVPPE
1	1	1	PT*	PT*
1	1	2	PT	RN
1,2	2	2	RN	RN

Table 1: Limiting Behavior for Large Background Risk
In the cases denoted by *, the result is true for any background risk of the given form.

background risk. In particular, we will identify to which of the solution concepts in Definition 5 behavior approaches as the background risk becomes large. But because of possible indifferences between options in D, we have to be careful (and unfortunately somewhat clumsy) in defining what this convergence means:

Definition 6. The concept X=RN, PT, CPE, SPPE, BLS, or EVPPE introduced in Definition 5 represents the decisionmaker's limiting behavior as the background risk becomes large if there is a utility function of the given form such that (i) if $F \in D$ satisfies the strict version of X, then F is an OCP choice for a sufficiently large background risk; and (ii) if $F' \in D$ does not satisfy the weak version of X, then F' is not an OCP choice for a sufficiently large background risk.

Table 1 summarizes the decisionmaker's limiting behavior as a function of when she learns about D, when she implements her choice, when the uncertainty from the chosen lottery is resolved, and the type of background risk. For the cases denoted with an asterisk (*), the result is true not only in the limit, but for *any* lottery of the given type.

As we have emphasized in Section 4, when the uncertainty in the chosen small lottery is resolved in period 1, this resolution generates immediate prospective gain-loss utility, so the person attends asymmetrically to gains and losses coming from the small risk. The difference between discrete and uniform background distributions arises due to a difference in how the person evaluates news with the two types of background risk. With a discrete distribution, the additional small risk replaces each atom of the background risk with a small "local lottery," so when the person receives new about the small lottery, her prospective gain-loss utility is determined by a rank-dependent comparison between her old beliefs and new beliefs *regarding the small lottery.* With a uniform distribution, the small lottery does not change the distribution of outcomes over most of the support of period-2 consumption—the distribution remains uniform other than near the edges. As a result, when the person receives news, she only cares about the comparison of her new *mean* beliefs and old *mean* beliefs regarding the small risk. Hence, while the two types of background risk generate conceptually similar reduced-form solution concepts, these concepts are based on the utility function U in the case of a discrete distribution but on the utility function V in the case of a uniform distribution.

As we have also discussed in Section 4, when the uncertainty in the small lottery is resolved in period 2, the resolution does not generate separate gain-loss utility, so the person integrates the lottery with the background risk and is typically risk-neutral to it. The only exception to this logic occurs when the person learns about D in period 1 and makes an immediate choice. In this case, her choice from D generates immediate news relative to her previous expectations of not having a choice, and hence induces prospective gain-loss utility. Nevertheless, when the background risk is uniform, she only cares about the mean of the chosen lottery, so despite prospective gain-loss utility in period 1 she is still risk neutral. But when the background risk is discrete, she compares the full distribution of the chosen lottery to her previous expectation of zero, so she behaves as postulated by prospect theory.

Due to the many cases and possibilities, our model may appear to be an overly complicated and unwieldy theory of monetary preferences. We feel, however, that the model reflects part of the complicated psychology of money rather than unnecessary formal complexity. In fact, in almost all cases in Table 1, our theory reduces to a previous theory of monetary preferences—which researchers presumably introduced because they believed it was realistic and there was evidence in its favor. Furthermore, our theory not only says that all these previous theories may be right in some circumstances, it predicts exactly in *which* circumstances each of them is appropriate.

Appendix B: Proofs

Proof of Proposition 1. All parts except for (iii) are obvious from the definition of OCP. To prove part (iii), it is sufficient to prove that if the person makes a PPE plan in period t', she does not switch away from this plan in a subsequent period; if this is the case, the person will clearly make a PPE plan. Suppose, then, that F is a PPE plan starting in period t. We prove by contradiction. Suppose that in period τ satisfying $t > \tau > t'$, the person strictly prefers to switch to plan $G \neq F$. Since she can only credibly switch to a PE plan, and there is no PE plan that she strictly prefers to F, to be preferred the switching itself must generate positive prospective gain-loss utility: $N(G|F) > 0$. Since the gains of G relative to F count less than the losses, this immediately implies that average consumption utility under G is strictly greater than under F: $E_G[m(c)] > E_F[m(c)]$. Next, we prove that for any dimension k, $N^k(G^k|F^k) \leq E_{G^k}[N^k(c^k|F^k)] - E_{F^k}[N^k(c^k|F^k)]$. Combined with the fact that consumption utility is higher under G, this means that in period t the person wants to switch away from a plan to choose F, so that F is not a PE, a contradiction.

To prove our claim, take any $p \in [0, 1]$. We will show that

$$\mu(m^k(c_{G^k}(p)) - m^k(c_{F^k}(p))) \leq N^k(c_{G^k}(p)|F^k) - N^k(c_{F^k}(p)|F^k) \tag{13}$$

by considering two cases. First, if $c_{G^k}(p) \geq c_{F^k}(p)$, then

$$\mu(m^k(c_{G^k}(p)) - m^k(c_{F^k}(p))) = \eta(m^k(c_{G^k}(p)) - m^k(c_{F^k}(p))) \leq N^k(c_{G^k}(p)|F^k) - N^k(c_{F^k}(p)|F^k)$$

because the difference on the right-hand side is a mixture of gains and avoided losses rather than just a pure gain. Second, similarly, if $c_{G^k}(p) \leq c_{F^k}(p)$, then

$$\mu(m^k(c_{G^k}(p)) - m^k(c_{F^k}(p))) = -\eta\lambda(m^k(c_{G^k}(p)) - m^k(c_{F^k}(p))) \leq N^k(c_{G^k}(p)|F^k) - N^k(c_{F^k}(p)|F^k)$$

because the right-hand side is a mixture of eliminated gains and losses rather than a pure loss.

Finally, integrating Inequality 13 over p gives exactly the desired inequality. □

Proof of Proposition 2. The two sequences of signals generate the same expected utility up to s_{j-1} being received and after s_{j+1} is received, so we compare expected utilities for the two in-between signals. Notice that because the outcome is binary and μ is linear, for any updating of the probability of $c_T = 1$ from p_{t-1} to p_t generates gain-loss utility equal to $N_t(p_t|p_{t-1}) = \mu(p_t - p_{t-1})$.

Consider any realized signals $s_1, \ldots, s_{j-1}$. For simplicity, we introduce the following notation for the purposes of this proof and that of Proposition 3. Let p be the decisionmaker's posterior on $c_T = 1$ after $s_1, \ldots, s_{j-1}$. Let $p(s_j)$ and $p(s_j, s_{j+1})$ denote the updated probabilities after one or two additional signals. Finally, let $\gamma_1 = \gamma_{t(s_j|\sigma'),T} = \gamma_{t(s_{j+1}|\sigma'),T}, \gamma_2 = \gamma_{t(s_j|\sigma),T}, \gamma_3 = \gamma_{t(s_{j+1}|\sigma),T}$. Since $\gamma_{t,T}$ is non-decreasing in t, we have $\gamma_1 \leq \gamma_2 \leq \gamma_3$. Now expected utility with signal structure σ' is

$$\begin{aligned}\gamma_1 E_{s_j,s_{j+1}}[\mu(p(s_j, s_{j+1}) - p)] &= \gamma_1 E_{s_j,s_{j+1}}[\mu(p(s_j) - p + p(s_j, s_{j+1}) - p(s_j))] \\ &\geq \gamma_1 E_{s_j,s_{j+1}}[\mu(p(s_j) - p) + \mu(p(s_j, s_{j+1}) - p(s_j))],\end{aligned}$$

where the last inequality holds because μ is steeper for losses than for gains. Furthermore, because these signals are non-trivial, there are realizations of $s_1, \ldots, s_{j-1}$ such that the above inequality is strict. Now we can further rewrite the above as

$$\gamma_1 E_{s_j}[\mu(p(s_j) - p)] + E_{s_j}[\gamma_1 E_{s_{j+1}}[\mu(p(s_j, s_{j+1}) - p(s_j))|s_j]] \geq$$

$$\gamma_2 E_{s_j}[\mu(p(s_j) - p)] + E_{s_j}[\gamma_3 E_{s_{j+1}}[\mu(p(s_j, s_{j+1}) - p(s_j))|s_j]]$$

since all these expectations are negative and $\gamma_1 \leq \gamma_2 \leq \gamma_3$. But the right-hand side above is exactly the decisionmaker's expected utility under σ. □

Proof of Proposition 3. We use a similar method of proof to that of Proposition 2. We prove for the case $t(s_{j+1}|\sigma') > t(s_j|\sigma')$ and $t(s_{j+1}|\sigma) > t(s_j|\sigma)$. The other case, when $t(s_{j+1}|\sigma') = t(s_j|\sigma')$ and $t(s_{j+1}|\sigma) = t(s_j|\sigma)$, is similar. Let $\gamma_1 = \gamma_{t(s_j|\sigma),T}, \gamma_2 = \gamma_{t(s_{j+1}|\sigma),T}$ and $\gamma'_1 = \gamma_{t(s_j|\sigma'),T}, \gamma'_2 = \gamma_{t(s_{j+1}|\sigma'),T}$. Then, using the notation of Proposition 2, the expected gain-loss utility from these two signals with information structure σ' is

$$\gamma'_1 E_{s_j}[\mu(p(s_j) - p)] + E_{s_j}[\gamma'_2 E_{s_{j+1}}[\mu(p(s_j, s_{j+1}) - p(s_j))|s_j]],$$

whereas with information structure σ it is

$$\gamma_1 E_{s_j}[\mu(p(s_j) - p)] + E_{s_j}[\gamma_2 E_{s_{j+1}}[\mu(p(s_j, s_{j+1}) - p(s_j))|s_j]].$$

Since all these expectations are non-positive, and they are strictly negative for some realizations of $s_1, \ldots, s_{j-1}$, the proposition immediately follows. □

Proof of Proposition 4. Suppose that given the signals she has observed so far, the decisionmaker's posterior of $c_T = 1$ is p. Consider a signal realization $\tilde{s}$. For this proposition and the following one, we make use of the following lemma:

Lemma 1. *The amount by which $\tilde{s}$ moves the decisionmaker's subjective probability of $c_T = 1$ is differentiable and strictly concave in p.*

Proof of Lemma 1. For notational simplicity, let $q = \Pr[\tilde{s}|c_T = 1]$ and $q' = \Pr[\tilde{s}|c_T = 0]$. By Bayes' rule, upon observing $\tilde{s}$ the decisionmaker updates her beliefs that $c_T = 1$ to

$$\frac{pq}{pq + (1-p)q'},$$

so the change in her beliefs is

$$\frac{pq}{pq + (1-p)q'} - p = (q - q')\frac{1}{\frac{q}{1-p} + \frac{q'}{p}}.$$

The denominator in the final expression is clearly positive, differentiable, and strictly convex in p. The reciprocal of such a function is always differentiable and strictly concave. This completes the proof of the lemma.

To prove proposition, we consider any beliefs $p \in (0,1)$ the decisionmaker holds after signals $s_1, \ldots, s_{j-1}$. By assumption, the probability that the decisionmaker ends up with such beliefs after $j-1$ signals is positive. Similarly to the proof of Proposition 2, let $p(s_j), p(s_{j+1}), p(s_j, s_{j+1})$ denote updated probabilities after the given signals, and let $\gamma_1 = \gamma_{t(s_j|\sigma),T}, \gamma_2 = \gamma_{t(s_{j+1}|\sigma),T}$. Notice that since s_{j+1} is always informative and has only finitely many possible realizations, for a sufficiently small ϵ no realization of s_j turns good news from s_{j+1} into bad news or vice versa. This means that by virtue of the law of iterated expectations, for a sufficiently small ϵ we have $E_{s_j,s_{j+1}}[|p(s_j, s_{j+1}) - p|] = E_{s_{j+1}}[|p(s_{j+1}) - p|]$. Hence, under σ' the expected utility from signals s_j and s_{j+1} is $-1/2 \cdot \gamma_2 \eta(\lambda - 1)E_{s_{j+1}}[|p(s_{j+1}) - p|]$ where this expression reflects the utility of the expected change in beliefs which is negative since losses are felt heavier than gains. The expected utility under σ is instead

$$-\frac{\gamma_1}{2}\eta(\lambda - 1)E_{s_j}[|p(s_j) - p|] - \frac{\gamma_2}{2}\eta(\lambda - 1)E_{s_j}[E_{s_{j+1}}[|p(s_j, s_{j+1}) - p(s_j)||s_j]].$$

By Lemma 1, the difference between $E_{s_j}[E_{s_{j+1}}[|p(s_j, s_{j+1}) - p(s_j)||s_j]]$ and $E_{s_{j+1}}[|p(s_{j+1}) - p|]$ is second-order in ϵ. But $E_{s_j}[|p(s_j) - p|]$ is first-order in ϵ, so for a sufficiently low ϵ the information structure σ yields strictly lower expected utility. □

Proof of Proposition 5. Notice that under σ', the signal s_1 has zero expected utility impact in period t_a because its size is less than $|p' - p|$. Hence, the result is immediate if $t(s_1|\sigma) < t(s_2|\sigma)$ since in this case the utility impact of s_1 is negative under σ and the impact of s_2 is the same under σ and σ'. Now suppose $t(s_1|\sigma) = t(s_2|\sigma)$, and let $\gamma = \gamma_{t(s_1|\sigma),T} = \gamma_{t(s_2|\sigma),T} = \gamma_{t(s_2|\sigma'),T}$. Using the same notation as in the previous propositions, notice that by virtue of the law of iterated expectation and the convexity of the absolute value function

$$E_{s_1,s_2}[|p(s_1, s_2) - p|] \geq E_{s_2}[|p(s_2) - p|].$$

Notice that $-\gamma\eta(\lambda - 1)/2$ times the term on the left-hand side of the above inequality is the decisionmaker's expected utility from receiving the two signals under σ. By Lemma 1, the right-hand side of the above inequality is strictly greater than

$$E_{s_1}[E_{s_2}[|p(s_1, s_2) - p(s_1)||s_1]].$$

which is $-\gamma\eta(\lambda - 1)/2$ times the decisionmaker's expected utility under σ'. This completes the proof. □

Proof of Proposition 6. First we show that if the consumption path is RDR then the inequalities in the proposition hold. To do so, it is sufficient to prove that the inequalities hold for any F_0. Since a consumption choice is made only in period 1, to check whether a plan is consistent it in turn suffices to check whether the plan is consistent with self 1's behavior. Furthermore, since m is strictly concave, for any given expectations self 1's utility function is strictly concave, so it suffices to consider only local deviations by self 1 from period 0's expectations.

Note that the marginal utility from increasing period-1 consumption above expected is at least

$$(1 + \eta)m'(c_1) - (1 + \gamma\eta\lambda)m'(W - c_1), \tag{14}$$

so deviating by increasing consumption is optimal unless

$$\frac{m'(c_1)}{m'(W-c_1)} \leq \frac{1+\gamma\eta\lambda}{1+\eta}.$$

The marginal utility from decreasing consumption below planned is at least

$$-(1+\eta\lambda)m'(c_1)+(1+\gamma\eta)m'(W-c_1), \tag{15}$$

so that deviating in this direction is optimal unless

$$\frac{m'(c_1)}{m'(W-c_1)} \geq \frac{1+\gamma\eta}{1+\eta\lambda}.$$

To see that any consumption path satisfying the inequalities is an RDR, it is sufficient to prove that any such consumption path is a PE. To check this, suppose that period-0 beliefs put unit mass on the consumption path $(c_1, W-c_1)$. Then, the marginal utility from increasing period-1 consumption is given exactly by Equation 14, and the marginal utility from decreasing period-1 consumption is given exactly by 15. Hence, if the inequalities in the proposition are satisfied, neither local deviation can increase self 1's utility. Given that self 1's utility function is concave, the same also holds non-locally. □

Proof of Proposition 7. Consider $\gamma \geq 1/\lambda$ first. Given the strict concavity of consumption utility, the ex-ante optimal plan is $c_1 = c_2 = W/2$. We show that this plan is consistent. We have established in the text that self 1 does not want to deviate by locally increasing period-1 consumption, and it is easy to show that she does not want to deviate by locally decreasing consumption. Furthermore, since self 1's utility function is concave, this means that self 1 does not want to deviate from a plan to consume equally.

We use an extension of the same argument to derive the PPE for $\gamma < 1/\lambda$. In this case, self 0 makes plans to follow the smoothest consumption path from which self 1 will not deviate. That is, self 0 plans the lowest consumption level in period 1 that is consistent with period-1 behavior. We show that this is given by the Equation 7. Again, since self 1's utility function is concave, it is sufficient to consider local deviations. Furthermore, for plans with $c_1 \geq c_2$ self 1 would clearly not deviate by decreasing period-1 consumption, so we consider only local increases in period-1 consumption. With plans to consume the amounts (c_1, c_2) in the two periods, a period-1 change of plans to slightly increase immediate consumption induces contemporaneous gain-loss utility of $\eta m'(c_1)$ and prospective gain-loss utility of $-\gamma\eta\lambda m'(c_2)$. For any $c_1 < c_1^*$, the net utility is positive, but it is zero for $c_1 = c_1^*$. □

Proof of Proposition 8. The marginal effect of changing $c_1^*(W)$ in a neighborhood of W on *ex-ante* expected utility is

$$\begin{aligned} & f(W)m'(c_1^*(W))[\underbrace{1+F(W)\eta+(1-F(W))\eta\lambda}_{\text{direct effect}}-\underbrace{(1-F(W))\eta-F(W)\eta\lambda}_{\text{indirect effect}}] \\ - \quad & f(W)m'(c_2^*(W))[\overbrace{1+F(W)\gamma\eta+(1-F(W))\gamma\eta\lambda}-\overbrace{(1-F(W))\gamma\eta-F(W)\gamma\eta\lambda}]. \end{aligned} \tag{16}$$

Increasing $c_1(W)$ above $c_1^*(W)$ has two effects on ex-ante expected utility. It has a *direct* effect through changing utility—both consumption utility and gain-loss utility—when realized wealth is W. This effect corresponds exactly to the right-hand of Equation 8. But changing c_1 also affects ex-ante utility *indirectly* through changing gain-loss utility when realized wealth is not W: it increases losses for lower wealth realizations and decreases gains for higher wealth realizations. Similar considerations hold for the ex-ante effect of decreasing $c_2(W)$ on prospective-gain loss utility. Expression (16) can be rewritten as

$$f(W)m'(c_1(W))[1+(1-2F(W))\eta(\lambda-1)] - f(W)m'(c_2(W))[1+(1-2F(W))\gamma\eta(\lambda-1)]. \tag{17}$$

Because

$$\frac{1+F(W)\eta+(1-F(W))\eta\lambda}{1+F(W)\gamma\eta+(1-F(W))\gamma\eta\lambda} > \frac{1+(1-2F(W))\eta(\lambda-1)}{1+(1-2F(W))\gamma\eta(\lambda-1)}$$

for any $\gamma < 1$, whenever Equation (8) holds Expression (17) is negative. □

Proof of Proposition 9. By Proposition 7, when $\gamma \geq 1/\lambda$, then $m'(c_1^*) = m'(c_2^*)$, and when $\gamma < 1/\lambda$, then $(1+\eta)m'(c_1^*) = (1+\gamma\eta\lambda)m'(c_2^*)$. Using that $\gamma < 1$ and $\lambda > 1$, this implies that in either case $(1+\eta)m'(c_1^*) > (1+\gamma\eta)m'(c_2^*)$. This immediately implies the statements regarding wealth increases. By similar considerations, $(1+\eta\lambda)m'(c_1^*) > (1+\gamma\eta\lambda)m'(c_2^*)$, implying the statements regarding wealth decreases. □

Proof of Proposition 10. We prove that the derivative of the marginal utility of increasing saving with respect to s is positive. This implies that $dc_1/ds|_{s=0} < 0$ both when $\gamma > 1/\lambda$ (because the ex-ante optimal plan, which the person follows, features a lower c_1) and when $\gamma \leq 1/\lambda$ (because a higher marginal utility in period 2 means that a lower c_1 becomes consistent).

Let F be the cumulative distribution function of the random variable y. Similarly to the argument in the text, the expected utility in period 2 is

$$\int m(c_2 + sy)dF(y) + \iint \mu(m(c_2 + sy) - m(c_2 + sy'))dF(y')dF(y)$$

$$= \int m(c_2 + sy)dF(y) - \frac{1}{2}\eta(\lambda - 1)\iint (m(c_2 + s\max\{y, y'\}) - m(c_2 + s\min\{y, y'\}))dF(y')dF(y).$$

The marginal utility from increasing savings is therefore

$$\int m'(c_2 + sy)dF(y) + \frac{1}{2}\eta(\lambda - 1)\iint (m'(c_2 + s\min\{y, y'\}) - m'(c_2 + s\max\{y, y'\}))dF(y')dF(y).$$

The derivative of the above expression with respect to s evaluated at $s = 0$ is

$$\frac{1}{2}\eta(\lambda - 1)(-m''(c_2))\iint |y - y'|dF(y')dF(y).$$

This completes the proof. □

Proof of Table 1. We prove each case in turn; in the notation below, the three numbers refer to when the decision is learned, implemented, and resolved, respectively.

0,2,2; 1,2,2; 2,2,2: By Part 1 of Proposition 2 in Kőszegi and Rabin (2007), whenever the person implements the decision in period 2 her limiting behavior is risk neutral.

0,0,2: This is a direct implication of Part 1 of Proposition 6 in Kőszegi and Rabin (2007).

To prove our statements for the rest of the cases, we solve for how the person feels about a period-1 change in beliefs. We start with discrete distributions. To state the following lemma, let Y be the highest absolute value of any realization of any lottery in D.

Lemma 2. *Suppose the background risk B is of the discrete type. Then, for any lotteries F, G distributed on $[-Y, Y]$ we have*

$$N(B + F|B + G) = N(F|G).$$

Proof of Lemma 2. For any atom x of B, let $B_-(x)$ be the probability of strictly smaller realizations of x, and $b(x)$ the weight on x. (Hence, $B_-(x) = B(x) - b(x)$). Since any realization of F has absolute value of at most Y, and the distance between any two atoms of B is at least $2Y$, for any atom x of B and any $y \in [-Y, Y]$ we have

$$(B + F)(x + y) = B_-(x) + b(x)F(y).$$

For any $p \in [0, 1]$, let x_p be the atom of B that satisfies $B_-(x_p) \leq p < B(x_p)$. Using the above equality, we have

$$c_{B+F}(p) = c_F\left(\frac{p - B_-(x_p)}{b(x_p)}\right).$$

This implies the statement of the lemma.

We prove a similar lemma for the case of uniform background risk.

Lemma 3. *Suppose the background risk B is of the uniform type. Then*

(i) For any constant y, $N(B + y|B) = N(y|0)$.

(ii) For any lotteries F, G distributed on $[-Y, Y]$,

$$\lim_{x_h - x_l \to \infty} N(B + F|B + G) = \mu[\overline{F} - \overline{G}].$$

Proof of Lemma 3. Part (i) is obvious. To prove Part (ii), we start by noting that for any F distributed on $[-Y, Y]$ and any $c \in (x_l + Y, x_h - Y)$ we have

$$(B + F)(c) = \int \frac{c - (x_l + y)}{x_h - x_l} dF(y) = \frac{c - x_l - \overline{F}}{x_h - x_l} = B(c - \overline{F}).$$

Hence, for any p satisfying $p \in ((B + F)(x_l + Y), (B + F)(x_h - Y))$ and $p \in ((B + G)(x_l + Y), (B + G)(x_h - Y))$, we have $c_{B+F}(p) - c_{B+G}(p) = \overline{F} - \overline{G}$. Since $\lim_{x_h - x_l \to \infty}(B + F)(x_l + Y) = \lim_{x_h - x_l \to \infty}(B+G)(x_l+Y) = 0$, and $\lim_{x_h - x_l \to \infty}(B+F)(x_h-Y) = \lim_{x_h - x_l \to \infty}(B+G)(x_h-Y) = 1$, this implies the statement of the lemma.

Now we return to proving the rest of the cases in Table 1.

0,1,2: Since in the case 0,0,2 the decisionmaker becomes risk neutral in the limit, to prove that she will become risk neutral in this case it is sufficient to prove that if she had planned in period 0 to make a risk neutral choice, switching away from it in period 1 does not generate positive prospective gain-loss utility. This is obvious from Lemmas 2 and 3.

0,0,1; 0,1,1: Obvious from the definition of OCP and Lemma 2 and Lemma 3, Part (ii).

1,1,1: Obvious from the definition of OCP and Lemma 2 and Lemma 3, Part (i).

1,1,2: Suppose for a second that $\gamma_{1,2} = 0$, so that there is no prospective gain-loss utility in period 1. Then, by Part 1 of Proposition 6 in Kőszegi and Rabin (2007), the decisionmaker approaches risk-neutrality as the background risk becomes large. By Part (ii) of Lemma 3, with uniform background risk this does not change when prospective gain-loss utility is added in period 1. Now consider discrete background risk. By Lemma 2, for any lottery $F \in D$, $N(B + F|B) = N(F|0) = U(F|0)$ for some U satisfying Equation 11. Given that period-2 utility approaches risk neutrality, this implies that the person chooses according to prospect theory in period 1.

This completes the proof of all statements in the table. □

References

Akerlof, G., and W. T. Dickens (1982): "The Economic Consequences of Cognitive Dissonance," *American Economic Review*, 72(3), 307–319.

Barberis, N., M. Huang, and R. Thaler (2006): "Individual Preferences, Monetary Gambles and Stock Market Participation: A Case for Narrow Framing," *American Economic Review*, 96(4), 1069–1090.

Battigalli, P., and M. Dufwenberg (2005): "Dynamic Psychological Games," Innocenzo Gasparini Institute Working Paper #287.

Becker, G. S., and K. M. Murphy (1988): "A Theory of Rational Addiction," *Journal of Political Economy*, 96(4), 675–700.

Bell, D. E. (1985): "Disappointment in Decision Making Under Uncertainty," *Operations Research*, 33(1), 1–27.

Bowman, D., D. Minehart, and M. Rabin (1999): "Loss Aversion in a Consumption-Savings Model," *Journal of Economic Behavior and Organization*, 38(2), 155–178.

Brunnermeier, M., and J. Parker (2005): "Optimal Expectations," *American Economic Review*, 95(4), 1092–1118.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

Caplin, A., and K. Eliaz (2003): "Aids Policy and Psychology: A Mechanism-Design Approach," *RAND Journal of Economics*, 34(4), 631–646.

Caplin, A., and J. Leahy (2001): "Psychological Expected Utility and Anticipatory Feelings," *Quarterly Journal of Economics*, 116(1), 55–80.

——— (2004): "The Supply of Information by a Concerned Expert," *Economic Journal*, 114(497), 487–505.

Duesenberry, J. S. (1952): *Income, Saving and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, MA.

Gul, F., and W. Pesendorfer (2001): "Temptation and Self-Control," *Econometrica*, 69(6), 1403–1435.

Hsee, C. K., and C. I. Tsai (2006): "Hedonomics in Consumer Behavior," Chapter to appear in *Handbook of Consumer Psychology.*

Kahneman, D. (1991): "Judgement and Decision-Making: A Personal View," *Psychological Science*, II, 142–145.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1990): "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy*, 98(6), 1325–1348.

Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47(2), 263–291.

Kőszegi, B. (2006a): "Ego Utility, Overconfidence, and Task Choice," *Journal of the European Economic Association*, 4(4), 673–707.

——— (2006b): "Emotional Agency," *Quarterly Journal of Economics*, 121(1), 121–155.

Kőszegi, B., and M. Rabin (2004): "A Model of Reference-Dependent Preferences," Economics Department, University of California, Berkeley, Working Paper E04-337.

——— (2006): "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 121(4), 1133–1166.

——— (2007): "Reference-Dependent Risk Attitudes," *American Economic Review*, forthcoming.

Kimball, M., and R. Willis (2006): "Utility and Happiness," Working Paper, University of Michigan.

Knetsch, J. L. (1989): "The Endowment Effect and Evidence of Nonreversible Indifference Curves," *American Economic Review*, 79(5), 1277–84.

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

Loewenstein, G., T. O'Donoghue, and M. Rabin (2003): "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics*, 118(4), 81–123.

Loewenstein, G., and D. Prelec (1993): "Preferences for sequences of outcomes.," *Psychological Review*, 100(1), 91–108.

Loewenstein, G. F. (1988): "Frames of Mind in Intertemporal Choice," *Management Science*, 34(2), 200–214.

Loewenstein, G. F., and N. Sicherman (1991): "Do Workers Prefer Increasing Wage Profiles?," *Journal of Labor Economics*, 9(1), 67–84.

Loomes, G., and R. Sugden (1982): "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal*, 92(368), 805–824.

Novemsky, N., and D. Kahneman (2005): "The Boundaries of Loss Aversion," *Journal of Marketing Research*, 42(2), 119–128.

O'Donoghue, T., and M. Rabin (1999): "Doing It Now or Later," *American Economic Review*, 89(1), 103–124.

Rabin, M., and G. Weizsäcker (2007): "Narrow Framing and Dominated Choices," Working Paper.

Ryder, Harl E., J., and G. M. Heal (1973): "Optimum Growth with Intertemporally Dependent Preferences," *Review of Economic Studies*, 40(1), 1–33.

Scitovsky, T. (1976): *The Joyless Economy: The Psychology of Human Satisfaction*. Oxford University Press, US.

Stone, R. (2005): "Loss Aversion and Self-Control," Working Paper.

Tversky, A., and D. Kahneman (1981): "The Framing of Decisions and the Psychology of Choice," *Science*, 211, 453–458.