Theoretical Economics 14 (2019), 173–210 1555-7561/20190173

A theory of personal budgeting

Simone Galperti
Department of Economics, University California, San Diego

Prominent research argues that consumers often use personal budgets to manage self-control problems. This paper analyzes the link between budgeting and self-control problems in consumption–saving decisions. It shows that the use of good-specific budgets depends on the combination of a demand for commitment and the demand for flexibility resulting from uncertainty about *intra*temporal trade-offs between goods. It explains the subtle mechanism that renders budgets useful commitments, their interaction with minimum-savings rules (another widely studied form of commitment), and how budgeting depends on the intensity of self-control problems. This theory matches several empirical findings on personal budgeting.

Keywords. Budget, minimum-savings rule, commitment, flexibility, intratemporal trade-off, uncertainty, present bias.

JEL classification. D23, D82, D86, D91, E62, G31.

1. Introduction

Many studies argue that personal budgeting is a pervasive part of consumer behavior.[1] This practice involves grouping expenses into categories and constraining each with an implicit or explicit cap applied to a specified time period (a week, a month, etc.).[2] While this practice cannot be explained by the classic life-cycle theory of the consumer, it has important consequences. It can account for "mysterious" large differences in wealth accumulation between consumers that time or risk preferences cannot explain (Ameriks

Simone Galperti: sgalperti@ucsd.edu
I thank S. Nageeb Ali, Eddie Dekel, Alexander Frankel, Mark Machina, Meg Meyer, Alessandro Pavan, Carlo Prato, Ron Siegel, Joel Sobel, Ran Spiegler, Charles Sprenger, Bruno Strulovici, Balazs Szentes, Alexis Akira Toda, Joel Watson, the co-editor, and two anonymous referees for useful comments and suggestions. I also thank seminar participants at Boston College, Columbia University, London School of Economics, Oxford University, Penn State University, University of Pennsylvania, University College London, Unuversity of California–Riverside, University of Warwick, and Tel Aviv University. This paper supersedes a previous paper titled "Delegating Resource Allocations in a Multidimensional World." John N. Rehbeck provided excellent research assistance. All remaining errors are mine.

[1]See Bakke (1940), Lee et al. (1962), Thaler and Shefrin (1981), Thaler (1985, 1999), Henderson and Peterson (1992), Baumeister et al. (1994), Heath and Soll (1996), Zelizer (1997), Wertenbroch (2002), Ameriks et al. (2003), Bénabou and Tirole (2004), Antonides et al. (2011), and Beshears et al. (2016).

[2]This paper uses the term "personal budgeting" rather than "mental accounting" because the latter has the much broader meaning of a general process whereby people frame events, outcomes, and decisions. This also includes choice bracketing, narrow framing, and gain–loss utility, which differ from budgeting.

Available at http://econtheory.org. https://doi.org/10.3982/TE2881

et al. 2003). By violating the principle of fungibility of money, it shapes demand differently from satiation and income effects (Heath and Soll 1996). It affects how firms promote their products so as to avoid competing for the same budget (Wertenbroch 2002). It is at the foundation of the economics of commitment devices (Bryan et al. 2010). Almost all existing studies informally suggest that consumers use budgets to manage self-control problems, often caused by present bias, which interfere with their saving goals (Thaler 1999, Ameriks et al. 2003, Antonides et al. 2011).

Despite this consensus, a formal investigation of the link between budgeting and self-control problems seems to be missing. The paper offers such a foundation using a broadly studied aspect of time preferences: present bias. It shows, however, that present bias alone cannot explain budgeting. Present-biased consumers value constraints on future choices. But for budgets to emerge, this preference for commitment has to be combined with a preference for flexibility of a precise but plausible kind, namely that caused by uncertainty about *intra*temporal trade-offs due, for instance, to good-specific taste shocks. The paper also uncovers a tension between good-specific budgets and minimum-savings rules, an often studied form of commitment. This leads to a *negative* relationship between the level of present bias and the use of budgets. These predictions help organize the evidence on budgeting and can guide future empirical studies.

Consider an agent, Ann, who has two selves: a time-consistent self-0 and a present-biased self-1.[3] Both selves have the same per-period consumption utility. In each period, self-1 chooses consumption and savings subject to the usual income constraint. Suppose that (i) consumption involves multiple goods (not a single uniform commodity) and (ii) both selves' preferences depend on a state of the world (capturing taste shocks) that affects not only the rate of substitution between present and future utility, but also the rates of substitution between goods within periods. In each period, before the state realizes, Ann's self-0 can adopt a commitment plan that dictates which income allocations self-1 is allowed to choose. This creates a trade-off between commitment and flexibility. The paper focusses on plans that can freely combine good-specific budgets and an overall limit on consumption expenses via a savings floor. This is in line with its motivation and offers an interesting lower bound on self-0's payoff. Of course, one would want to allow for general forms of commitment, which is, however, much harder in the presence of the foregoing features (i) and (ii) (see Section 4), and is beyond the scope of this paper.[4]

Features (i) and (ii) are the key differences between this paper and Amador et al. (2006), where consumption involves a single commodity and taste shocks affect only the intertemporal utility trade-off. That paper shows that, under very general conditions on the shock distribution, the optimal rule is to impose a savings floor and grant self-1 flexibility otherwise, even if self-0 can choose among arbitrarily general forms of

[3]Dual-self models appear in Thaler and Shefrin (1981), Bénabou and Pycia (2002), Bernheim and Rangel (2004), Benhabib and Bisin (2005), Fudenberg and Levine (2006, 2012), Loewenstein and O'Donoghue (2007), Brocas and Carrillo (2008), Chatterjee and Krishna (2009), and Nageeb (2011).

[4]This paper takes the process of noticing an expense and reporting it to its budget as a defining aspect of budgeting itself. To focus on the issues of interest here, it also assumes that people stick to their commitment plans, as justified in Section 2.

commitment. To establish a benchmark, Section 3.2 shows that their result carries over to a world with multiple goods *if* there is no uncertainty about intratemporal trade-offs between goods. Intuitively, in this case binding good-specific budgets forces self-1 to choose inefficient consumption bundles, which is akin to wasting resources (i.e., "money burning"). Amador et al. (2006) already showed that money burning is generally suboptimal.

Uncertain intratemporal trade-offs change things substantially, as summarized by the main results of the paper. First, if the goods satisfy appropriate substitutability and normality conditions, optimal commitment plans always involve good-specific budgets when present bias is sufficiently *weak*, but only a savings floor when present bias is sufficiently *strong*. Second, fixing a weak bias, for some range of parameters the optimal plans combine budgets with a savings floor, but for another range they rely only on the budgets. By contrast, in Amador et al. (2006) optimal plans always involve a savings floor.[5] The substitutability and normality conditions ensure that the consumption distortions caused by budgets curtail how much self-1 gains in terms of present utility by undersaving, thereby resulting in higher savings. This improvement matters more than those distortions for the time-consistent self-0.

To see the intuition for these results, suppose Ann consumes two goods and is uncertain whether her marginal utility of each good will be high or low. Anticipating her tendency to undersave, she first considers setting a savings floor. This limits overspending if both marginal utilities are high, which makes her want to consume a lot of both goods. If only one marginal utility turns out to be high, however, the floor may not bind; this is especially likely if present bias is weak. In this case, Ann realizes that she will still overspend and this will be mostly driven by the good with high marginal utility. She can then also cap this good with a targeted budget, which raises her savings because now she can overspend only on the good with low marginal utility. By contrast, when present bias is stronger, overspending becomes more severe even for the good with low marginal utility, and the budget leads to a small (if any) rise in savings at the cost of rationing a good with high marginal utility. As a result, Ann prefers to adopt only a savings floor, because it curbs undersaving without distorting consumption.

The results involve some noteworthy subtleties. An agent may adopt budgets that distort consumption spending even though her selves always agree on how to divide every dollar between goods within a period. By contrast, a binding floor distorts only the income division between spending and saving. Perhaps counterintuitively, it is not the case that if a present-biased agent adds budgets to a floor, then a more biased agent should do the same. In addition, agents who use budgets may also set *tighter* floors, as the budgets' distortions lower the value of leaving more income for consumption. Once budgets are allowed, agents with a stronger present bias may adopt a *slacker* floor (in contrast to Proposition 5 in Amador et al. 2006). Section 3 further discusses the results relative to the evidence on budgeting, highlighting findings that other theories struggle to explain.

This paper expands our understanding of consumption–savings behavior under self-control problems and the resulting demand for commitment. Since Thaler and

[5]These properties continues to hold for partially naive agents who incorrectly anticipate their bias.

Shefrin's (1981) and Laibson's (1997) seminal work, the literature has almost always assumed a single, per-period commodity ("money").[6] As this paper shows, that assumption is not innocuous with present-biased consumers (in contrast to the case of time-consistent consumers) and this crucially depends on uncertain intratemporal trade-offs. The literature has focussed on the problem of curbing undersaving and the usefulness of devices like illiquid assets and savings accounts. This paper shows that consumers can do strictly better by (also) adopting good-specific budgets, which opens the door to other commitment devices, such as personal budgeting services.[7] It also suggests which type of consumers will demand which type of devices, which can be used by third-party providers.[8]

To derive the results, the paper uses techniques different from the standard mechanism-design approach. The idea is to exploit the information in the Lagrange multipliers for the constraints that budgets and savings floors add to self-1's optimization problem. Relying on sensitivity-analysis techniques (Luenberger 1969), we can use this information to quantify, after appropriately adjusting for self-1's bias, the marginal benefit for self-0 of modifying a budget or a floor.

Related literature

Existing explanations of personal budgeting are based on Thaler (1985). Using the notions of "transaction utility" and gain–loss utility, he argues that agents treat the consequences of each transaction in isolation. Given this, they can solve their consumption–savings problems by means of transaction-specific budgets, a result that echoes Strotz (1957). In reality, people set budgets for sufficiently long periods so that each covers many transactions. Also, in Thaler's deterministic model, the agents can achieve the same utility with and without budgets, but with uncertainty, they would never set binding budgets. Therefore, they do not exhibit a strict demand for budgets as commitment devices. Finally, transaction and gain–loss utility differ conceptually from self-control problems, which the literature views as the main cause of budgeting. Gain–loss utility can explain other phenomena of mental accounting, such as choice bracketing (Koch and Nafziger 2016), which, however, differs from budgeting.

Other papers in the mechanism-design literature study the trade-offs between commitment and flexibility, usually imposing no restriction on the feasible mechanisms. Papers by Amador et al. (2006) and Halac and Yared (2014) are the closest to the present paper.[9] We borrow their baseline model, but adds multiple consumption goods and uncertainty about intratemporal trade-offs. In so doing, we show how this uncertainty affects the commitment–flexibility trade-off and its solutions. Another difference is that

[6]Brocas and Carrillo (2008) discuss a model with two goods, one of which has ex ante uncertain utility, and self-1 is fully myopic. In this case, the optimal commitment strategy consists of a nonlinear plan that punishes spending on one good by cutting spending on the other, which is not a budgeting plan. Even if one focusses on these plans, self-0 never sets budgets with a fully myopic self-1 (see Proposition 3).

[7]This kind of service is currently offered by firms like Mint, Quicken, and StickK.

[8]In reality, it may be hard to observe each consumer's degree of present bias and offer devices accordingly. Some of the issues that arise in this case are analyzed by Galperti (2015).

[9]See also Athey et al. (2005), Ambrus and Egorov (2013), and Amador and Bagwell (2013).

Halac and Yared (2014) focus on the role of information persistency. In their setting, an optimal commitment plan can distort future choices, even though they cause no conflict between the agent's selves given today's choice. Persistence links self-1's current information and expected utility from future choices, which can be used to relax today's incentive constraints, as in other dynamic mechanism-design problems.[10] Correlation among self-1's pieces of information is not the driver of the present paper's results.

An older literature examined how rationing affects consumer behavior (Howard 1977, Ellis and Naughton 1990, Madden 1991). By setting a savings floor or good-specific budgets, an agent essentially rations his future selves just as the government may ration consumers. In contrast to that literature, here rationing assumes the role of a commitment device. That literature shows that predicting the budgets' effects is far from trivial. Its insights are useful to identify conditions under which budgets can help the agent.

2. The model

Consider an agent, Ann, who lives for two periods. In the first, she chooses a consumption bundle $\mathbf{c} = (c_1, c_2) \in \mathbb{R}^2_+$ and a level of savings $s \in \mathbb{R}_+$. In the second period, consumption involves a single good and, hence, equals s. Ann receives her income, normalized to 1, in the first period.

Ann has self-control problems caused by a conflict between a long-run self-0 and a short-run self-1. Their preferences depend on some taste shocks, represented by the state (θ, r_1, r_2), where $\theta > 0$ and $\mathbf{r} = (r_1, r_2) \in \mathbb{R}^2$. In each period both selves have the same (concave) consumption utility: $u(\mathbf{c}; \mathbf{r})$ in period 1 and $v(s)$ in period 2. In period 1, however, self-0 and self-1 evaluate streams $(\mathbf{c}, s)$ using, respectively, the utility functions

$$\theta u(\mathbf{c}; \mathbf{r}) + v(s) \quad \text{and} \quad \theta u(\mathbf{c}; \mathbf{r}) + \beta v(s).$$

For clarity and tractability, for now assume that

$$u(\mathbf{c}; \mathbf{r}) = u^1(c_1; r_1) + u^2(c_2; r_2) \quad \text{with} \quad \frac{\partial^2 u^i(c_i; r_i)}{\partial c_i \partial r_i} = u^i_{cr}(c_i; r_i) > 0 \quad \text{for } i = 1, 2.$$

Self-1's present bias is captured by $\beta \in (0, 1)$. Self-0 knows β (sophistication); we discuss naiveté later.

A key novelty of this model is that the state affects both *inter*- and *intra*temporal trade-offs. While θ affects only the substitution rate between present and future utility, $\mathbf{r}$ also affects the substitution rates between goods within period 1. Hereafter, let $\boldsymbol{\omega} = (\theta, \mathbf{r})$, let G be its distribution, and let Ω be the state space. Distribution G is allowed to have rich forms of dependence as well as full independence across θ, r_1, and r_2.

Self-0 delegates the consumption–savings choice to self-1 by designing a commitment plan that dictates which choices self-1 is allowed to implement. In the case of budgeting, such a plan involves spending limits on specific consumption categories, denoted by b_i (for budget), or an overall limit on consumption expenditures implemented

[10] See, for example, Courty and Hao (2000), Battaglini (2005), and Pavan et al. (2014).

through a minimum-savings rule f (for floor). Formally, let

$$F = \{(\mathbf{c}, s) \in \mathbb{R}^3_+ : c_1 + c_2 + s \leq 1\}.$$

Think of c_i and s as the share of income allocated to good i and savings. A budgeting plan, B, can then be expressed as

$$B = \{(\mathbf{c}, s) \in F : s \geq f, c_1 \leq b_1, c_2 \leq b_2\},$$

where $f \in [0, 1]$ and $b_i \in [0, 1]$ for $i = 1, 2$. Let $\mathcal{B}$ be the set of all budgeting plans. From the ex ante viewpoint, we call f and b_i *binding* if they bind with strictly positive probability under G. Note that $\mathcal{B}$ defines a specific subclass of commitment plans that—though intuitive and tractable—rule out many other possible ways to restrict self-1's choices. Section 4 discusses some intricacies of allowing for more general plans.

In reality, agents commit to their plans prior to observing all the necessary information for making a decision. This creates a trade-off between commitment and flexibility. In the model, first self-0 commits to a plan B, and then only self-1 observes ω and chooses some $(\mathbf{c}, s)$ from B. Self-0 designs B to maximize her expected payoff from self-1's choices. Note that if self-0 knew ω, the problem would be uninteresting: Setting f at the level of savings that self-0 finds optimal given ω always induces self-1 to choose $\mathbf{c}$ and s that maximize self-0's utility.

The goal of the paper is to understand whether and how self-0 sets minimum-savings rules and goods-specific budgets. The problem can be stated as

$$\max_{B \in \mathcal{B}} \mathcal{U}(B) = \int_\Omega \left[\theta u(\mathbf{c}(\omega); \mathbf{r}) + v(s(\omega))\right] dG(\omega) \tag{1}$$

$$\text{s.t.} \quad (\mathbf{c}(\omega), s(\omega)) \in \arg\max_{(\mathbf{c},s)\in B} \theta u(\mathbf{c}; \mathbf{r}) + \beta v(s), \quad \omega \in \Omega. \tag{2}$$

A solution is called an optimal plan.

Technical assumptions

Information distributions Let $\Omega = [\underline{\theta}, \overline{\theta}] \times [\underline{r}_1, \overline{r}_1] \times [\underline{r}_2, \overline{r}_2]$, where $0 < \underline{\theta} < \overline{\theta} < +\infty$ and $0 < \underline{r}_i < \overline{r}_i < +\infty$ for $i = 1, 2$. We only assume that the joint probability distribution G of (θ, r_1, r_2) has full support (that is, $G(O) > 0$ for every open $O \subset \Omega$). The conditions on $\underline{\theta}$ and $\overline{\theta}$ rule out the implausible situation where Ann does not care at all about the present or the future.[11] The conditions on $\underline{r}_i$ and $\overline{r}_i$ have bite only when combined with the properties of u listed next.

Differentiability, monotonicity, and concavity The term v is twice continuously differentiable with $v' > 0$ and $v'' < 0$. For $i = 1, 2$ and $r_i \in [\underline{r}_i, \overline{r}_i]$, $u^i(\cdot; r_i) : \mathbb{R}_+ \to \mathbb{R}$ is twice differentiable with $u^i_c(\cdot; r_i) > 0$ and $u^i_{cc}(\cdot; r_i) < 0$; also, u^i_c and u^i_{cc} are continuous on $(0, 1] \times [\underline{r}_i, \overline{r}_i]$. This implies that u^i_c is bounded below and away from zero; this nonsatiation property seems plausible to the extent that i refers to food, housing, or entertainment and a period corresponds to a week or a month.

[11] A similar assumption appears in Amador et al. (2006), who point out that with unbounded support it may be optimal to grant self-1 full flexibility.

Boundary conditions We have $\lim_{s \to 0} v'(s) = +\infty$ and $\lim_{c \to 0} u^i_c(c; r_i) = +\infty$ for $r_i \in [\underline{r}_i, \overline{r}_i]$ and $i = 1, 2$. This allows us to focus on interior solutions.

Discussion of the model

Nothing significant changes if Ann receives income in both periods and can borrow in period 1 or if the consumption bundle $\mathbf{c}$ involves more than two goods. In fact, the proofs consider the general case of $n \geq 2$ goods. It is straightforward to allow for multiple goods in the second period also.

The two selves' preferences are consistent with the quasi-hyperbolic discounting model of Laibson (1997) and with viewing the agent as a household aggregating its members' preferences, which are time consistent but heterogeneous (Jackson and Yariv 2015). A public finance interpretation of the model is also possible along the lines of Halac and Yared (2014). In each period, a government chooses spending on a list of public goods and services, $\mathbf{c}$, and saving or borrowing, s, subject to the constraint given by the tax revenues. The government may exhibit present bias as a consequence of aggregating the preferences of heterogeneous citizens (Jackson and Yariv 2015) or uncertainty in the political turnover (Aguiar and Amador 2011).

The assumed information structure has some redundancy, as both an increase in θ and an increase in all components of $\mathbf{r}$ render period-1 consumption more valuable. Nonetheless, it is convenient for differentiating uncertainty about intra- and intertemporal trade-offs and for showing that the former is crucial for budgets to arise (Section 3.2). In a nutshell, this is because it allows for situations where overspending is driven by all goods and situations where it is mostly driven by only some good.

To focus on the issues of interest for this paper, it is assumed that Ann sticks to her plans. This is not a minor assumption, of course, but the literature has proposed several mechanisms that can justify it. These include a desire for internal consistency (Festinger 1962), the plans' working as reference points (Heath et al. 1999, Hsiaw 2013), self-reputation mechanisms (Bénabou and Tirole 2004), internal control processes that prevent impulsive processes from breaking ex ante rules (Benhabib and Bisin 2005), and self-enforcement sustained by threats of switching to less desirable equilibria (Bernheim et al. 2015). Perhaps in reality people are able to carry out their plans provided that they are not too stringent or costly ex post. Even in this case, it is worth understanding which forces lead people to find budgets and floors useful despite their ex post inefficiency. For instance, some present-biased agents may not use budgets not because they cannot stick to them, but simply because they do not find them useful. This can also be valuable for third parties that design commitment devices to help people stick to their plans (such as firms like Mint, Quicken, and Stick).

3. Optimal budgeting plans

3.1 *Preliminaries*

First of all, treating self-0's payoff as a function of f, b_1, and b_2, one can easily establish existence of an optimal plan using the maximum theorem.[12]

[12] See Lemma 2 in the Appendix.

It is worth defining two benchmark allocations. For each ω, let $(\mathbf{c}^d(\omega), s^d(\omega))$ be self-1's choice if granted full discretion, namely, the solution to $\max_{(\mathbf{c},s)\in F}\{\theta u(\mathbf{c};\mathbf{r}) + \beta v(s)\}$. Also, let $(\mathbf{c}^p(\omega), s^p(\omega))$ represent what self-0 would like self-1 to choose in ω, which is the solution to $\max_{(\mathbf{c},s)\in F}\{\theta u(\mathbf{c};\mathbf{r}) + v(s)\}$. Call $(\mathbf{c}^d, s^d)$ the *full-discretion allocation* and call $(\mathbf{c}^p, s^p)$ the *first-best (or planned) allocation*. They satisfy the following useful properties.

Remark 1. (i) The mechanisms $(\mathbf{c}^p, s^p)$ and $(\mathbf{c}^d, s^d)$ are continuous in ω.

(ii) Each component of $(\mathbf{c}^p, s^p)$ and $(\mathbf{c}^d, s^d)$ takes values in a closed interval and is bounded away from zero.

(iii) For $i = 1, 2$, c_i^p and c_i^d are strictly increasing in r_i and θ, and decreasing in r_j for $j \neq i$.

(iv) Savings s^p and s^d are strictly decreasing in θ, r_1, and r_2.

(v) For $\omega \in \Omega$, $s^d(\omega) < s^p(\omega)$ and $s^d(\omega)$ is continuous and strictly increasing in β.

(vi) For $\omega \in \Omega$ and $i = 1, 2$, $c_i^d(\omega)$ is continuous and strictly decreasing in β.

Another property worth noting is that all consumption goods are normal for both selves.[13]

For illustration, consider a fully symmetric model with respect to goods 1 and 2. Since self-1 saves whatever he does not consume ($v' > 0$), we can focus on his choices of $\mathbf{c}$ represented in Figure 1. Note that $\mathbf{c}$s on negative 45° lines closer to the origin correspond to a higher s. To understand the shape of $\mathbf{c}^p$, suppose for the moment that θ takes only one value. Start from $(\theta, \underline{r}_1, \underline{r}_2)$, which leads to the highest s. If we raise r_1 up to $\overline{r}_1$, c_1^p increases while c_2^p and s^p decrease, which means that we move along the south part of the dashed line. If we now start from $(\theta, \overline{r}_1, \underline{r}_2)$ and raise r_2 up to $\overline{r}_2$, c_2^p increases while c_1^p and s^p decrease; that is, we move along the east part of the dashed line. Proceeding in this way, we can map the entire dashed line; continuity of $\mathbf{c}^p$ implies that

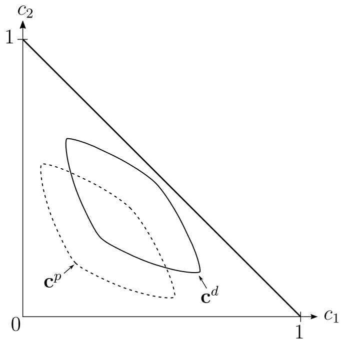

Figure 1. First-best and full-discretion allocations.

[13]This property follows, for instance, from Proposition 1 in Quah (2007).

its range extends inside the boundary in Figure 1. Self-1's systematic undersaving shifts the $\mathbf{c}^d$ region away from the origin; the stronger his bias, the bigger the shift. Figure 1 also highlights the effect of uncertain intratemporal trade-offs. If $\mathbf{r}$ were certain (but not θ), both $\mathbf{c}^p$ and $\mathbf{c}^d$ would collapse to an upward sloping line so that the state calling for minimal savings always coincides with that calling for maximal consumption of *both* goods.

3.2 *A benchmark: Known intratemporal trade-offs*

This section shows that if we remove the uncertainty about intratemporal trade-offs—while keeping that about the intertemporal trade-off and multiple goods—then most of the time optimal plans involve a savings floor but *no* good-specific budgets. This benchmark helps us disentangle the role of uncertain intratemporal trade-offs from the multidimensionality of consumption.

For the sake of the argument, in this section imagine that self-0 observes $\mathbf{r}$ (but not θ) before designing her plan; self-1 continues to observe θ and $\mathbf{r}$. We can then examine the problem defined by (1) and (2) by treating $\mathbf{r}$ as fixed. Let $G_\mathbf{r}$ denote the distribution of θ given $\mathbf{r}$. To state the result, we need some minor conditions about $G_\mathbf{r}$. Assume that $G_\mathbf{r}$ has a strictly positive and continuous density function $g_\mathbf{r}$ on $[\underline{\theta}, \overline{\theta}]$. Define

$$H(\theta) = 1 - G_\mathbf{r}(\theta) - (1-\beta)\theta g_\mathbf{r}(\theta), \quad \theta \in [\underline{\theta}, \overline{\theta}]$$

and

$$\theta^* = \min\left\{\theta \in [\underline{\theta}, \overline{\theta}] : \textstyle\int_{\theta'}^{\overline{\theta}} H(\hat{\theta})\, d\hat{\theta} \leq 0 \text{ for all } \theta' \geq \theta\right\}.$$

Proposition 1. *Suppose H is nonincreasing over $[\underline{\theta}, \theta^*]$. Then a plan B that satisfies $f = s^d(\theta^*)$ and $b_1 = b_2 = 1$ is optimal.*

As Amador et al. (2006) noted, for many distributions—especially those commonly used in applications—H satisfies the above condition for all $\beta \in [0, 1]$. More generally, if $g_\mathbf{r}$ is uniformly bounded away from 0 and changes at a bounded rate, the condition on H holds when β is sufficiently high. Importantly, as we will see, high βs characterize the settings with uncertain intratemporal trade-offs where plans using only f are *not* optimal.

Proposition 1 follows from Amador et al.'s (2006) main result, *once* we establish the following key point: If self-0 knows $\mathbf{r}$, she can focus on commitment plans that regulate only savings and *total* consumption expenses, but not how these are divided between goods. The reason is as follows. Unlike a binding f, which distorts only the income division between spending and saving, a binding b_i also distorts consumption.[14] Thus, by forcing self-1 to consume inefficient bundles, budgets lower the utility she can get from what she does *not* save. To lower this utility, however, another method is simply to

[14] One way to see this is that the positive Lagrange multiplier for the binding $c_i \leq b_i$ introduces a wedge between the goods' marginal utilities.

not let self-1 spend all of $1-s$. The literature called this money burning.[15] Spending a share of $1-s$ efficiently can achieve any utility obtained by spending $1-s$ inefficiently: For all $\mathbf{c}\in\mathbb{R}^2_+$, there exists $y\le c_1+c_2$ that yields $u(\mathbf{c};\mathbf{r})=u^*(y;\mathbf{r})$, where $u^*(y;\mathbf{r})$ is the indirect utility of spending y. Different realizations of the intratemporal trade-offs may affect how self-0 wants to "punish" self-1 for undersaving, holding s fixed. But without that uncertainty, the optimal punishment is unique and can always be achieved with money burning, provided that its amount can flexibly depend on the chosen s. This requires more general forms of commitment than budgeting plans. Formally, let

$$F^{\text{tc}}=\{(y,s)\in\mathbb{R}^2_+ : y+s\le 1\}.$$

Given $D^{\text{tc}}\subset F^{\text{tc}}$, self-1 maximizes $\theta u(\mathbf{c};\mathbf{r})+\beta v(s)$ subject to $c_1+c_2\le y$ and $(y,s)\in D^{\text{tc}}$.

Lemma 1. *Suppose uncertainty affects only the intertemporal utility trade-off. There exists an optimal $D\subset F$ with $\mathcal{U}(D)=\mathcal{U}^*$ if and only if there exists an optimal $D^{\text{tc}}\subset F^{\text{tc}}$ with $\mathcal{U}(D^{\text{tc}})=\mathcal{U}^*$.*

Thus, when only the intertemporal trade-off is uncertain, whether consumption involves one or multiple goods is irrelevant as long as we allow for general commitment plans.

Proposition 1 goes one step further by showing that the number of consumption goods is irrelevant even when self-1 can use only minimum-savings rules. Given Lemma 1, since the constraint $c_1+c_2\le y$ always binds for self-1, the problem becomes

$$\max_{D^{\text{tc}}\subset F^{\text{tc}}}\int_{\underline{\theta}}^{\overline{\theta}}\big[\theta u^*\big(y(\theta);\mathbf{r}\big)+v\big(s(\theta)\big)\big]g_{\mathbf{r}}(\theta)\,d\theta$$

$$\text{s.t.}\quad \big(y(\theta),s(\theta)\big)\in\operatorname*{arg\,max}_{(y,s)\in D^{\text{tc}}}\big\{\theta u^*(y;\mathbf{r})+\beta v(s)\big\},\qquad \theta\in[\underline{\theta},\overline{\theta}].$$

This is isomorphic to the problem studied by Amador et al. (2006). Proposition 1 then follows from their Proposition 3.

3.3 *Main results*

We now return to the model where both inter- and intratemporal trade-offs are uncertain for self-0. The next result is in sharp contrast to the benchmark established before.

Proposition 2. *There exists $\beta^*\in(0,1)$ such that, if $\beta^*<\beta<1$. Then every optimal $B\in\mathcal{B}$ must include binding good-specific budgets.*[16]

The Appendix shows how to derive β^*, which can be significantly smaller than 1. Proposition 2 is silent about whether good-specific budgets are always combined with a floor. Section 3.4 shows that both cases are possible.

Do optimal plans always require good-specific budgets? The answer is no.

[15]Besides Amador et al. (2006), papers that study money burning in delegation problems include: Ambrus and Egorov (2013, 2017) and Amador and Bagwell (2013, 2016).

[16]All proofs are provided in the Appendix.

Proposition 3. *There exists $\beta_* \in (0,1)$ such that if $\beta < \beta_*$, then every optimal $B \in \mathcal{B}$ involves only a binding savings floor.*

The Appendix shows how to calculate β_*, which can be significantly larger than 0 and depends on G only through its support. Using this, we can show that weaker biases suffice to render budgets suboptimal when the uncertainty on intratemporal trade-offs shrinks in the following sense.

Corollary 1. *Consider two agents who have the same utility functions u and v, and their uncertainty has supports $[\underline{\theta}, \overline{\theta}] \times [\underline{r}_1, \overline{r}_1] \times [\underline{r}_2, \overline{r}_2]$ and $[\underline{\theta}, \overline{\theta}] \times [\underline{r}'_1, \overline{r}'_1] \times [\underline{r}'_2, \overline{r}'_2]$. Let β_* and β'_* be the corresponding thresholds in Proposition 3. If $(\underline{r}'_1, \underline{r}'_2) \gneqq (\underline{r}_1, \underline{r}_2)$ and $(\overline{r}'_1, \overline{r}'_2) \lneqq (\overline{r}_1, \overline{r}_2)$, then $\beta'_* > \beta_*$.*

This corollary echoes the benchmark result of Section 3.2: That case corresponds to the limit as $\overline{r}_i - \underline{r}_i \to 0$ for $i = 1, 2$. We saw that under minor conditions, in the limit, essentially $\beta_* = 1$.

One subtlety of the model is that f, b_1, and b_2 can bind simultaneously, thereby affecting self-1's choices in possibly complex ways. To handle this, the proof proceeds in several steps, which are sketched here to also uncover the intuitions for the results.

The first step is to consider how self-0 would use f in isolation. In this case, the best f lies strictly between the highest and lowest first-best savings, $\overline{s}^p$ and $\underline{s}^p$, and rises as β falls. Though similar, this step is not a corollary of Amador et al.'s (2006) results and uses different techniques. Intuitively, self-0 never finds any $s < \underline{s}^p$ justifiable, and f never distorts the chosen $\mathbf{c}$ because the two selves have the same consumption utility u. Consequently, self-0 always sets $f \geq \underline{s}^p$. Setting $f = \underline{s}^p$ cannot be optimal, as raising f a bit causes a second-order loss when $s^p(\omega) = \underline{s}^p$, but a first-order gain when $s^p(\omega) > \underline{s}^p$ and f binds. A similar logic explains why $f < \overline{s}^p$. Thus, f has to balance the benefit of curbing undersaving and the cost of causing oversaving. A lower β raises the optimal f because the benefit of raising f grows if self-1 tends to undersave more, but the cost stays the same: When self-0 wants $s < f$, self-1 does too and f binds for *any* β. To obtain these properties, the proof shows that the derivative of self-0's payoff in f exists, has a simple form, and is decreasing in β. This uses the fact that we can focus on the states where f binds for self-1 and so $s = f$, which allows us to immediately infer the effect of varying f on self-0's savings utility, v. Since both selves share u, the effect on self-0's consumption utility can be inferred from self-1's *indirect utility* from spending $1 - f$ via Lagrangian sensitivity analysis, which links this effect to the marginal utility of any good at the chosen $\mathbf{c}$. These effects are shown to matter for a set of states with strictly positive probability using the continuity of self-1's choices in f and ω and the full support of G.

The second step is to consider how self-0 would use b_i in isolation. It turns out that capping even only one good dominates granting self-1 full discretion. To see why, start from the level of b_i where it starts to bind (i.e., $b_i = \overline{c}^d_i = \max_\omega c^d_i(\omega)$). Lowering b_i creates a benefit and a cost for self-0 when b_i binds. The cost is that it distorts consumption, but this is initially a second-order cost, because the full discretion $\mathbf{c}^d$ is efficient in the sense of equalizing marginal utilities between goods. The benefit is that b_i curbs undersaving, which is of first-order importance for self-0. Overall b_i should then benefit self-0, but there is a subtlety: Self-1 should not reallocate income to the unrestricted c_j much faster than to s, which is not obvious and need not be true. This key property holds for the additively separable u, but also more generally (see Section 4). Once this is established, the proof uses the fact that both selves share u to show that self-0's payoff can be written as her savings utility scaled by $(1-\beta)$ plus self-1's total payoff subject to b_i. Lagrangian sensitivity analysis on the latter pins down the second-order negative effects of b_i. The former part directly quantifies the first-order positive effects of b_i. These effects again matter for a set of states with strictly positive probability for the same reasons as before.

These points highlight the general mechanism whereby multidimensional consumption can help to curb the consequences of present bias. A budget b_i incentivizes self-1 to save more because it forces him to choose inefficient bundles—not just to spend less on c_i, which he could fully shift to c_j—and this inefficiency limits the present utility self-1 can gain by undersaving.

The above two steps are combined to obtain Proposition 2. Intuitively, when present bias is weak, an optimal plan must use budgets because they help improve savings when doing so via f would require it to be too tight. Consider Figure 2, which focusses on consumption choices and reports the full-discretion and first-best allocations from Figure 1. Graphically, b_1 defines a vertical line allowing only c_1s to its left, b_2 defines a horizontal line allowing only c_2s below it, and f defines a line with slope -1 allowing only $\mathbf{c}$s below it. Thus, in Figure 2(b) for instance, self-1 must choose $\mathbf{c}$ within the diamond-shaped region to the southwest of the solid lines f, b_1, and b_2. Uncertain intratemporal trade-offs imply that the states where both selves want to spend the most on c_1 or c_2 (which map to the light-shaded areas) are *not* the states where they want to save the least (which map to the dark-shaded areas). Indeed, by Remark 1, for $i = 1, 2$ and $k = p, d$,

$$\overline{c}_i^k = c_i^k(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}) > c_i^k(\overline{\theta}, \overline{r}_i, \overline{r}_{-i}) \quad \text{and} \quad \underline{s}^k = s^k(\overline{\theta}, \overline{r}_i, \overline{r}_{-i}) < s^k(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}).$$

We saw that if self-0 can use only f, she relaxes it as β rises; that is, the f line moves farther away from the origin. Consequently, f constrains choices in the dark-shaded area, but at some point stops affecting them in the light-shaded areas: compare panel (a)

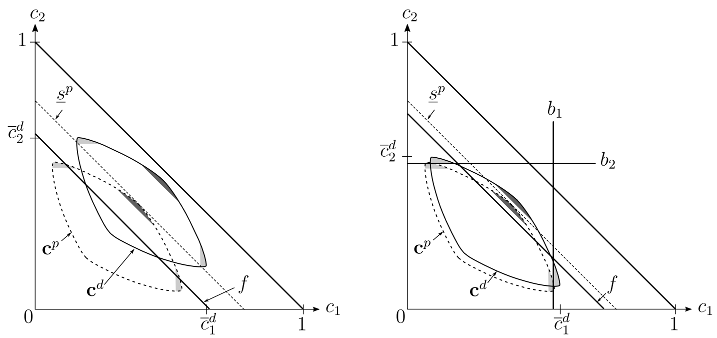

(a) Strong Bias (b) Weak Bias

FIGURE 2. Optimal budgeting plans: Intuition.

to (b). To curb undersaving in these states, self-0 prefers not to use f, but can add budgets that bind when f does not (as in (b)). As noted, such budgets benefit self-0.

The proof of Proposition 3 first shows that it is never optimal to let self-1 save $s < \underline{s}^p$. Intuitively, if B allows this, raising f up to $\underline{s}^p$ uniformly improves self-0's savings utility; moreover, since all goods are normal, the resulting lower spendable income renders any budget in B less likely to bind and distort $\mathbf{c}$, which again benefits self-0. Now note that when β is sufficiently small, self-1 always wants to save $s < \underline{s}^p$. Hence, when b_i forces self-1 to spend less on c_i, he shifts all the money to c_j, but not s. Since budgets distort $\mathbf{c}$, if they do not raise s, they cannot benefit self-0 and, hence, be part of optimal plans.

It is now easy to see why shrinking uncertainty about intratemporal trade-offs expands the set of strong biases for which optimal plans use only f (Corollary 1). Budgets are useful to curb undersaving in states with large asymmetry in the goods' marginal utilities (recall Figure 2). If this asymmetry shrinks, so does the scope for budgets to be useful.

As should be expected, the strongest bias for which optimal plans use good-specific budgets depends on the details of the context. This does not change the main takeaway. As β falls, for every B, it raises the probability of the states where self-1 is constrained by B's actual lower bound on savings at some $\tilde{s} \geq \underline{s}^p$. Since in these states binding budgets only distort $\mathbf{c}$, their appeal falls accordingly. How self-0 balances the distortions in those states with the budgets' benefits in other states ultimately depends on their distribution G. Nonetheless, since self-0 can always set $f = \tilde{s}$, for βs below some $\hat{\beta} \geq \beta_*$, every optimal plan must use only f.

Finally, one may think that these results are driven by the fact that budgets and the savings floor are substitute tools in the sense that, everything else equal, optimal plans set a slacker f if they can also use b_1 and b_2. Their interaction is actually more subtle. By curbing self-1's undersaving, b_1 and b_2 lower the return of tightening f, which can result in a slacker f. At the same time, b_1 and b_2 also lower the return of loosening f because they prevent self-1 from consuming efficiently the extra spendable income, which can result in a tighter f. Therefore, when agents can also use budgets, their savings floor need not vary monotonically with β. This is another difference from Amador et al. (2006), who predict that f decreases in β. Clearly, all constraints are eventually removed as $\beta \to 1$, yet this need not occur monotonically along the way.

3.4 *Optimal plans can involve no minimum-savings rule*

This section shows that with multiple goods there exist both settings where self-0 combines budgets with a savings floor and settings where she uses only the budgets. In contrast, in the case of a single good, optimal plans always involve a binding floor.[17] We focus on the symmetric model $u^1(c; r) = u^2(c; r) = r \ln(c)$, $\underline{r}_1 = \underline{r}_2 = \underline{r} > 0$, $\overline{r}_1 = \overline{r}_2 = \overline{r} > \underline{r}$, and $v(s) = \ln(s)$.[18]

[17] See Propositions 2 and 10 in Amador et al. (2006).

[18] The function $\ln(\cdot)$ violates the continuity and differentiability assumptions of Section 2 at 0, but this is irrelevant for the analysis.

Proposition 4. *There exist full-support distributions G such that f, b_1, and b_2 are all binding for every optimal $B \in \mathcal{B}$. There also exist full-support distributions G' such that, for every optimal $B \in \mathcal{B}$, b_1 and b_2 are binding, but f never binds.*

While the result holds for full-support distributions, its intuition can be best explained by considering a three-state case. Let $\omega^0 = (\overline{\theta}, \overline{r}_1, \overline{r}_2)$, $\omega^1 = (\underline{\theta}, \overline{r}_1, \underline{r}_2)$, and $\omega^2 = (\underline{\theta}, \underline{r}_1, \overline{r}_2)$ with respective probabilities g, $\frac{1}{2}(1-g)$, and $\frac{1}{2}(1-g)$. Remark 1 and symmetry imply that

$$s^d(\omega^0) < s^d(\omega^1) = s^d(\omega^2), \qquad c_1^d(\omega^2) = c_2^d(\omega^1) < c_1^d(\omega^1) = c_2^d(\omega^2), \qquad c_1^d(\omega^0) = c_2^d(\omega^0);$$

similar properties hold for $(\mathbf{c}^p, s^p)$. By continuity, there exists $\beta < 1$ sufficiently high that $s^d(\omega^1) = s^d(\omega^2) > s^p(\omega^0)$; hereafter, fix such a β. There exists $\underline{\theta}$ sufficiently close to $\overline{\theta}$ that $c_1^p(\omega^1) > c_1^p(\omega^0)$ and $c_2^p(\omega^2) > c_2^p(\omega^0)$. Figure 3(a) represents this situation, focussing again on consumption. Concretely, imagine that Ann has two friends, Becky and Cindy. In a given week, Ann may go out with Becky (ω^1), Cindy (ω^2), or both together (ω^0). Ann likes shopping for clothes with Becky and trying new restaurants with Cindy. When out with both, she enjoys both activities even more. Finally, Ann anticipates that once in the store or the restaurant, she will tend to spend too much.

One can show that if Ann deems going out with both friends sufficiently likely (i.e., $g > g^*$ for some $g^* \in (0,1)$), then she wants to set a binding f as well as budgets for both goods. In fact, the optimal B satisfies $f = s^p(\omega^0)$, $b_1 = c_1^p(\omega^1)$, and $b_2 = c_2^p(\omega^2)$.[19] The intuition is this. If Ann was sure to go out with one friend at a time, she could set f so as to eliminate splurging in ω^1 and ω^2: this f corresponds to the dotted line in Figure 3(a). However, this f will be too stringent if she ends up going out with both friends. Since this is very likely, Ann prefers $f = s^p(\omega^0)$. She knows that this f will not bind when she is out with only one friend. But for this case she can curb overspending using b_1 and b_2; also, here she can do so without affecting her choice in ω^0.

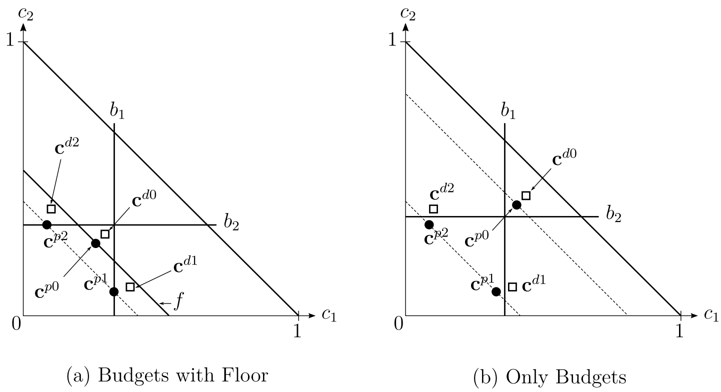

Figure 3. Three-state example ($\mathbf{c}^{di} = \mathbf{c}^d(\omega^i)$ and $\mathbf{c}^{pi} = \mathbf{c}^p(\omega^i)$).

[19]This claim is shown as part of the constructive proof of Proposition 4. The specific levels of b_1 and b_2 are just a by-product of logarithmic payoffs.

A simple change of this three-state setting suffices to explain why optimal plans can involve only good-specific budgets. Fix $g > g^*$ and all the other parameters except $\overline{\theta}$. If we increase $\overline{\theta}$, both selves want to consume more in ω^0. This eventually leads to a situation as in Figure 3(b), where $c_1^p(\omega^0) > c_1^p(\omega^1)$ and $c_2^p(\omega^0) > c_2^p(\omega^2)$. In this case, Ann wants to keep b_1 and b_2, but drop f. In fact, her optimal B satisfies $b_1 = b_2$ and $c_i^p(\omega^i) < b_i < c_i^p(\omega^0)$ for every $i = 1, 2$, but $f = 0$.[20] Figure 3(b) helps with the intuition. Now Ann is willing to spend even more in ω^0. Therefore, the budgets she would set to curb splurging in ω^1 and ω^2 start to bind also in ω^0. As a result, she wants to relax them. She realizes, however, that her first-best spending on clothes and food in ω^0 is just above those budgets. Relaxing them a bit will allow her to curb splurging in ω^1 and ω^2 as well as ω^0. Since these budgets already push savings above the first best in ω^0, Ann cannot benefit by adding a binding f.

In short, a weakly present-biased agent may use only good-specific budgets for the following reason. To curb undersaving in states with large asymmetry in consumption marginal utilities, she may prefer to use the budgets rather than a savings floor, which would have to be too stringent. Together the budgets then impose a cap on total spending. If this already ensures sufficiently high savings in states where present consumption is very valuable overall, then any binding floor will have to cause additional oversaving and this inefficiency can exceed the floor's commitment benefits.

4. Discussion

Theory and evidence

How does this theory relate to the evidence on budgeting and other explanations thereof? One finding is that people may set budgets on "unobjectionable goods like sports tickets and blue jeans" (Heath and Soll 1996) or housing, food, and even charitable giving (Thaler 1985, 1999). This is difficult to explain with an alternative theory arguing that people set budgets for the goods they find tempting ("vice goods"), although this can be true in some cases. By contrast, present bias *combined with* uncertain intratemporal trade-offs can lead people to set budgets on "unobjectionable goods," as doing so helps them manage their overall tendency to overspend better than with just a savings floor. Thus, it is not the tempting nature of a good that matters.

Another plausible theory is that budgeting is a technique to simplify the complex matter of household finance (Simon 1997, Johnson 1984). This theory is complementary to the one in this paper, but again struggles to explain some evidence. For instance, it is not clear why computational complexity would lead people to systematically set budgets that seem too strict and to cause underconsumption, as found by Heath and Soll (1996). By contrast, present-biased people do optimally set budgets that systematically bind and thus exhibit those properties.

[20] Again, this claim is shown as part of the constructive proof of Proposition 4. Note that, although this three-state example is intuitive, it takes some work to rule out the possibility of multiple, perhaps asymmetric, optimal plans featuring different properties from those in the proposition.

The prediction that only weakly biased agents should use good-specific budgets is consistent with some findings in Antonides et al. (2011). In their sample, people who exhibit a "short-term time orientation" (which according to their description is consistent with strong present bias) are *less* likely to use budgets than people who exhibit a "long-term time orientation" (a weak bias). Unfortunately, Antonides et al. (2011) do not measure how stringent budgets or floors are in relation to present bias. For that matter, we saw that this relation need not be monotonic. As an alternative explanation, strongly biased agents may not use budgets because they are less sophisticated or able to commit. The anticipated bias (not the true one) is what matters for self-0's problem, however. Therefore, by Proposition 2 *under*estimating that bias—not entirely, of course—may actually render it *more* likely that self-0 finds budgets beneficial. If this same agent were instead sophisticated, he might *not* adopt any budget by Proposition 3. We also saw that once a strongly biased agent can use a savings floor, the reason why budgets do not work for him is not that he cannot honor them: Even if he could, they would strictly lower his utility.

Relaxing separability

The message of the paper generalizes to settings where utility is not separable across goods. Continue to assume that $u(\mathbf{c};\mathbf{r})$ is strictly concave in $\mathbf{c}$ and twice differentiable with continuous $u_{c_i}(\mathbf{c};\mathbf{r}) > 0$ and $u_{c_ic_j}(\mathbf{c};\mathbf{r})$ in both arguments for all i and j. We saw that budgets help curb self-1's undersaving if (a) they increase savings and (b) there exist states that call for high consumption of some good, but not of all goods. Property (b) holds if some good is a sufficiently strong substitute of all other goods. Property (a) holds if the capped good is a Hicks substitute of savings (Howard 1977); in general, such a good always exists (Madden 1991, Theorem 2). As noted, however, a budget has to curb undersaving faster than it exacerbates overspending on other goods for it to benefit self-0. Given space constraints, these properties are stated directly in terms of allocations.

Condition 1. Both $(\mathbf{c}^p, s^p)$ and $(\mathbf{c}^d, s^d)$ are interior for every ω. Both s^p and s^d are strictly decreasing in θ and r_i for $i = 1, 2$. There exists some good j that satisfies the following statements: (i) c_j^p and c_j^d are strictly increasing in θ and r_j, and decreasing in r_i for $i \neq j$; (ii) there exists $\varepsilon > 0$ such that, for every $b_j < \max_\omega c_j^d(\omega)$, self-1's optimal $(\mathbf{c}^*, s^*)$ subject to plans involving only b_j satisfies $s^*(\omega) - s^d(\omega) \geq \varepsilon[c_j^d(\omega) - c_j^*(\omega)]$ for all $\omega \in \Omega$.

Appendix A.6 presents an example that satisfies Condition 1.

To state the result, consider a more general class of budgeting plans, denoted by $\overline{\mathcal{B}}$, which allow us to also set good-specific floors and a savings cap,

$$\overline{B} = \{(\mathbf{c}, s) \in F : f_0 \leq s \leq b_0, f_1 \leq c_1 \leq b_1, f_2 \leq c_2 \leq b_2\},$$

where $f_i, b_i \in [0, 1]$ satisfy $f_i \leq b_i$ for $i = 0, 1, 2$ and $f_0 + f_1 + f_2 \leq 1$.

Proposition 5. *Under Condition 1, there exists $\beta^* \in (0, 1)$ such that, if $\beta^* < \beta < 1$. Then every optimal $\overline{B} \in \overline{\mathcal{B}}$ must use distorting good-specific restrictions.*

The proof is omitted, because using Condition 1, one can adapt the proof of Proposition 2 to show that plans using only f_0 are strictly dominated for sufficiently high β. Since setting a binding b_0 is never optimal, the result follows.

Do optimal good-specific restrictions always take the form of budgets? The answer depends on the substitutability and complementarity between goods and between each good and savings, which can be affected by the restrictions themselves. A sufficient condition for optimal plans to never use f_1 and f_2 is that all goods are Hicks substitutes and collectively sufficiently normal (see Ellis and Naughton 1990 for a formal statement of this property). Given this, by Theorems 3 and 4 of Madden (1991) two goods remain substitutes independently of which goods are restricted, and Ellis and Naughton's (1990) analysis implies that, given any f_1 and f_2, relaxing them raises s. Hence, since f_1 and f_2 distort consumption, they strictly harm self-0. Optimal plans use only b_1 and b_2 for the example in Appendix A.6.

General mechanisms

One may wonder what the best among all conceivable plans (not just those in $\mathcal{B}$) looks like and whether it belongs to $\mathcal{B}$. These are important questions, but also hard in the presence of multidimensional consumption and uncertainty. The main challenges come from the income constraint and the complexity of the incentive constraints, which as usual cannot be reduced to only the local ones. Here one can try to apply the insights from multidimensional screening (Rochet and Stole 2003), but substantive differences remain. First, screening problems allow for transfers. Here one can view the utility from savings as a transfer and use Rochet and Choné's (1998) approach to simplify the incentive constraints and self-0's objective. But the state-wise income constraint (the second difference from screening) cannot be simplified. General techniques exist for handling such constraints (Luenberger 1969), but unlike in the case of unidimensional consumption, here they do not go far.

5. Concluding remarks

This paper provides a theoretical analysis of the link between self-control problems and personal budgeting using a parsimonious consumption–savings model with a present-biased agent. Unlike minimum-savings rules, good-specific spending caps help to curtail overspending because they cause inefficiencies in consumption that lower the return from undersaving, thereby counteracting present bias. Consequently, good-specific budgets are no free lunch and are used only by agents who are weakly biased and uncertain about their intratemporal trade-offs between goods. Those who are strongly biased or do not face such uncertainty prefer to rely exclusively on a minimum-savings rule.

This theory offers insights into the subtle forces underlying a widely observed phenomenon, which has far-reaching consequences for consumer behavior and welfare by affecting demand differently from satiation and income effects, and by significantly contributing to households' wealth accumulation. The theory matches existing empirical findings, such as that people often set budgets for goods normally not viewed as tempting and only those who exhibit weak present bias seem to use budgets. The theory also suggests new directions for enriching the sparse evidence on budgeting by demonstrating its dependence on uncertain intratemporal trade-offs, and for designing commitment devices whose functions are targeted to the right type of present-biased agents.

Appendix

A.1 *Technical lemmas*

Lemma 2. *There exists B that maximizes $\mathcal{U}(B)$ over $\mathcal{B}$.*

Proof. Each $B \in \mathcal{B}$ can be viewed as an element $(f, \mathbf{b})$ of the compact set $[0, 1]^{n+1}$. Thus, we can think that self-0 chooses $(f, \mathbf{b}) \in [0, 1]^{n+1}$.

Given any such $(f, \mathbf{b})$, let $(\mathbf{c}(\omega|f, \mathbf{b}), s(\omega|f, \mathbf{b}))$ be self-1's optimal allocation in state ω from the compact set $B_{f,\mathbf{b}}$ defined by $(f, \mathbf{b})$. Since $B_{f,\mathbf{b}}$ is convex (Theorem 2.1 in Rockafellar 1997), $(\mathbf{c}(\omega|f, \mathbf{b}), s(\omega|f, \mathbf{b}))$ is unique for every $\omega \in \Omega$ by strict concavity of self-1's utility function. Clearly, the correspondence that for each $(f, \mathbf{b}) \in [0, 1]^{n+1}$ maps to $B_{f,\mathbf{b}}$ is nonempty, compact valued, and continuous. It follows from the maximum theorem that $(\mathbf{c}(\omega|\cdot, \cdot), s(\omega|\cdot, \cdot))$ is continuous for every $\omega \in \Omega$.

We can now show that self-0's payoff is continuous in $(f, \mathbf{b})$. For each $(f, \mathbf{b}) \in [0, 1]^{n+1}$, let

$$\mathcal{U}(f, \mathbf{b}) = \int_{\Omega} \big[\theta u\big(\mathbf{c}(\omega|f, \mathbf{b}); \mathbf{r}\big) + v(s(\omega|f, \mathbf{b})\big]\, dG(\omega).$$

Since the integrand is continuous in $(f, \mathbf{b})$ for every $\omega \in \Omega$ and is uniformly bounded over $B_{(f,\mathbf{b})}$, Lebesgue's dominated convergence theorem implies the claimed property of $\mathcal{U}(\cdot, \cdot)$.

A second application of the maximum theorem gives the result. □

Lemma 3. *Fix $i \in \{1, \dots, n\}$ and consider $B \in \mathcal{B}$ with $b_j = 1$ for all $j \neq i$. For any ω, if $b_i < c_i^d(\omega)$, self-1 chooses $s > s^d(\omega)$ and $c_j > c_j^d(\omega)$ for all $j \neq i$.*

Proof. Let $i = 1$ and $b_1 \in (0, c_1^d(\omega))$. Consider self-1's problem in state ω to maximize $\theta u(\mathbf{c}, \mathbf{r}) + \beta v(s)$ for $(\mathbf{c}, s) \in F$ subject to $c_1 \leq b_1$. The first-order conditions of its Lagrangian are $\beta v'(s(\omega)) = \mu(\omega)$, $\theta u_c^1(c_1(\omega); r_1) = \mu(\omega) + \lambda_1(\omega)$, and $\theta u_c^i(c_i(\omega); r_i) = \mu(\omega)$ for all $i \neq 1$, where $\mu(\omega) \geq 0$ and $\lambda_1(\omega) \geq 0$ are the Lagrange multipliers for $\sum_{i=1}^n c_i \leq 1$ and $c_1 \leq b_1$.

Suppose $s(\omega) \leq s^d(\omega)$. Since $c_1(\omega) = b_1 < c_1^d(\omega)$ and $s(\omega) + \sum_j c_j(\omega) = s^d(\omega) + \sum_j c_j^d(\omega) = 1$ by strong monotonicity of preferences, $c_j(\omega) > c_j^d(\omega)$ for some $j \neq 0, 1$. By strict concavity of u^j and v, $\theta u_c^j(c_j(\omega); r_j) < \theta u_c^j(c_j^d(\omega); r_j) = \beta v'(s^d(\omega)) \leq \beta v'(s(\omega))$. This violates the first-order conditions for $\mathbf{c}(\omega)$. So we must have $s(\omega) > s^d(\omega)$. This in turn implies that $\theta u_c^j(c_j(\omega); r_j) = \beta v'(s(\omega)) < \beta v'(s^d(\omega)) = \theta u_c^j(c_j^d(\omega); r_j)$ for $j \neq i$. By concavity, $c_j(\omega) > c_j^d(\omega)$ for $j \neq 1$. □

For $k = p, d$, let $\underline{s}^k = \min_\omega s^k(\omega)$ and $\overline{s}^k(\omega) = \max_\omega s^k(\omega)$. Focussing on $f \in [\underline{s}^d, \overline{s}^p]$,[21] denote by B_f the corresponding policy in $\mathcal{B}$.

Lemma 4. *Define* $\overline{\Omega}(f) = \{\omega \in \Omega : s^d(\omega) \le f\}$ *and let* $\mathbf{c}^f(\omega)$ *be the maximizer of* $u(\mathbf{c}; \mathbf{r})$ *subject to* $\sum_{i=1}^n c_i \le 1 - f$. *Then* $\mathcal{U}(B_f)$ *is differentiable in* f *over* $[\underline{s}^d, \overline{s}^p]$ *with*

$$\frac{d}{df}\mathcal{U}(B_f) = \int_{\overline{\Omega}(f)} \left[v'(f) - \theta u_{c_i}\big(\mathbf{c}^f(\omega); \mathbf{r}\big)\right] dG \quad \textit{for any } i = 1, \dots, n.$$

Proof. Given f and any ω, define

$$\tilde{u}(f; \omega) \equiv u\big(\mathbf{c}^f(\omega); \mathbf{r}\big) = \max_{\{\mathbf{c} \in \mathbb{R}^n_+ : \sum_{i=1}^n c_i \le 1 - f\}} u(\mathbf{c}; \mathbf{r})$$

and $\tilde{U}(f; \omega) = \theta \tilde{u}(f; \omega) + v(f)$. Since $u(\cdot; \mathbf{r})$ is strictly concave in $\mathbf{c}$, so is $\tilde{u}(\cdot; \mathbf{r})$ in f by standard arguments. Hence, $\tilde{U}(\cdot; \omega)$ is also strictly concave in f. Now consider $\tilde{U}'(f; \omega)$. Whenever it is defined, $\tilde{U}'(f; \omega) = \theta \tilde{u}'(f; \omega) + v'(f)$. By first-order conditions of the Lagrangian defining $\tilde{u}(f; \mathbf{r})$, we have $u_{c_i}(\mathbf{c}^f(\omega); \mathbf{r}) = \lambda(\omega; f)$ for $i = 1, \dots, n$, where $\lambda(\omega; f)$ is the Lagrange multiplier for $\sum_{i=1}^n c_i \le 1 - f$. Since $\mathbf{c}^f(\omega)$ is continuous in f for every ω, so is $\lambda(\omega; f)$. By Theorem 1 of Luenberger 1969, p. 222) $\lambda(\omega; f')(f'' - f') \le \tilde{u}(f'; \mathbf{r}) - \tilde{u}(f''; \mathbf{r}) \le \lambda(\omega; f'')(f'' - f')$ for every $f', f'' \in (0, 1)$. Continuity of $\lambda(\omega; \cdot)$ implies that $\tilde{u}'(f; \mathbf{r})$ exists for every $f \in (0, 1)$ and $\tilde{u}'(f; \mathbf{r}) = -\lambda(\omega; f) = -u_{c_i}(\mathbf{c}^f(\omega); \mathbf{r})$. Therefore,

$$\tilde{U}'(f; \omega) = v'(f) - \theta u_{c_i}\big(\mathbf{c}^f(\omega); \mathbf{r}\big), \quad \omega \in \Omega. \tag{3}$$

For any f, denote by $(\mathbf{c}^f, s^f)$ self-1's behavior as a function of ω under B_f. By the maximum theorem, $(\mathbf{c}^f(\omega), s^f(\omega))$ is continuous in both f and ω. Since, fixing any s, both selves would choose the same $\mathbf{c}$ in every ω, by definition

$$\Psi(f) \equiv \mathcal{U}(B_f) = \int_\Omega \tilde{U}\big(s^f(\omega); \omega\big)\, dG.$$

Consider any $f > \hat{f}$ and recall that $\overline{\Omega}(f) = \{\omega : s^d(\omega) \le f\}$. Then

$$\begin{aligned}\Psi(f) - \Psi(\hat{f}) &= \int_{\overline{\Omega}(f)} \left[\tilde{U}(f; \omega) - \tilde{U}\big(s^{\hat{f}}(\omega); \omega\big)\right] dG \\ &= \int_{\overline{\Omega}(f) \cap (\overline{\Omega}(\hat{f}))^c} \left[\tilde{U}(f; \omega) - \tilde{U}\big(s^{\hat{f}}(\omega); \omega\big)\right] dG + \int_{\overline{\Omega}(\hat{f})} \left[\tilde{U}(f; \omega) - \tilde{U}(\hat{f}; \omega)\right] dG;\end{aligned}$$

the first equality holds because $s^f(\omega) = s^{\hat{f}}(\omega)$ for $\omega \notin \overline{\Omega}(f)$ and $s^f(\omega) = f$ for $\omega \in \overline{\Omega}(f)$. Divide both sides by $f - \hat{f}$ and consider the limit as $f \downarrow \hat{f}$. First, for all ω, we have

$$\lim_{f \downarrow \hat{f}} \frac{\tilde{U}(f; \omega) - \tilde{U}(\hat{f}; \omega)}{f - \hat{f}} = \tilde{U}'(\hat{f}; \omega).$$

[21] Any other f is dominated by one in this range.

Since $\tilde{U}(\cdot;\omega)$ is concave,

$$\left|\frac{\tilde{U}(f;\omega)-\tilde{U}(\hat{f};\omega)}{f-\hat{f}}\right| \le \max\{|\tilde{U}'(f;\omega)|, |\tilde{U}'(\hat{f};\omega)|\}.$$

Since $\tilde{U}'(f;\omega)$ is continuous in ω and f as illustrated by (3), $|\tilde{U}'(f;\omega)|$ is bounded by some $M<+\infty$ for $(f,\omega)\in[\underline{s}^d,\overline{s}^p]\times\Omega$. Therefore, by Lebesgue's bounded convergence theorem,

$$\lim_{f\downarrow\hat{f}}\int_{\overline{\Omega}(\hat{f})}\frac{\tilde{U}(f;\omega)-\tilde{U}(\hat{f};\omega)}{f-\hat{f}}\,dG=\int_{\overline{\Omega}(\hat{f})}\tilde{U}'(\hat{f};\omega)\,dG.$$

Consider now the second part of the limit. Again, by concavity of $\tilde{U}(\cdot;\omega)$ and since $s^f(\omega)\in[\underline{s}^d,\overline{s}^p]$ for $f\in[\underline{s}^d,\overline{s}^p]$, we have

$$\left|\frac{\tilde{U}(f;\omega)-\tilde{U}(s^{\hat{f}}(\omega);\omega)}{f-s^{\hat{f}}(\omega)}\right|\le M.$$

Therefore,

$$\begin{aligned}\left|\int_{\overline{\Omega}(f)\cap(\overline{\Omega}(\hat{f}))^c}\frac{\tilde{U}(f;\omega)-\tilde{U}(s^{\hat{f}}(\omega);\omega)}{f-\hat{f}}\,dG\right| &\le \int_{\overline{\Omega}(f)\cap(\overline{\Omega}(\hat{f}))^c}\left|\frac{\tilde{U}(f;\omega)-\tilde{U}(s^{\hat{f}}(\omega);\omega)}{f-\hat{f}}\right|dG\\ &\le \int_{\overline{\Omega}(f)\cap(\overline{\Omega}(\hat{f}))^c}\left|\frac{\tilde{U}(f;\omega)-\tilde{U}(s^{\hat{f}}(\omega);\omega)}{f-s^{\hat{f}}(\omega)}\right|dG\\ &\le M\int_{\overline{\Omega}(f)\cap(\overline{\Omega}(\hat{f}))^c}dG.\end{aligned}$$

Observe that $\overline{\Omega}(f)\cap(\overline{\Omega}(\hat{f}))^c=\{\omega:\hat{f}<s^{\hat{f}}(\omega)\le f\}$, which converges to an empty set as $f\downarrow\hat{f}$. Since then the second part of the limit converges to zero as $f\downarrow\hat{f}$, for every $\hat{f}\in[\underline{s}^d,\overline{s}^p)$,

$$\Psi'(\hat{f}+)=\int_{\overline{\Omega}(\hat{f})}\tilde{U}'(\hat{f};\omega)\,dG.$$

A similar argument implies that $\Psi'(\hat{f}-)=\int_{\overline{\Omega}(\hat{f})}\tilde{U}'(\hat{f};\omega)\,dG$ for every $\hat{f}\in(\underline{s}^d,\overline{s}^p]$. Hence, $\Psi(f)$ is differentiable over $[\underline{s}^d,\overline{s}^p]$. □

A.2 *Proof of Lemma 1*

Fix $\mathbf{r}$. Recall the definition of $\mathcal{U}(D)$ in (1) of the main text and that $(\mathbf{c}(\theta),s(\theta))$ represents self-1's optimal choice in state θ. There exists $D\subset F$ such that $\mathcal{U}(D)\ge\mathcal{U}(D')$ for all $D'\subset F$ if and only if there exist functions $\boldsymbol{\chi}:[\underline{\theta},\overline{\theta}]\to\mathbb{R}^n_+$ and $t:[\underline{\theta},\overline{\theta}]\to\mathbb{R}_+$ that satisfy the following two conditions:

Condition 1. For all $\theta,\theta'\in[\underline{\theta},\overline{\theta}]$,

$$\theta u(\chi(\theta);\mathbf{r})+\beta v(t(\theta))\ge\theta u(\chi(\theta');\mathbf{r})+\beta v(t(\theta'))$$

and

$$\sum_{i=1}^{n} \chi_i(\theta) + t(\theta) \leq 1.$$

CONDITION 2. The pair (χ, t) maximizes

$$\int_{\underline{\theta}}^{\overline{\theta}} [\theta u(\chi(\theta); \mathbf{r}) + v(t(\theta))] g_{\mathbf{r}}(\theta)\, d\theta.$$

By contrast, there exists $D^{\mathrm{tc}} \subset F^{\mathrm{tc}}$ such that $\mathcal{U}(D^{\mathrm{tc}}) \geq \mathcal{U}(\hat{D}^{\mathrm{tc}})$ for all $\hat{D}^{\mathrm{tc}} \subset F^{\mathrm{tc}}$ if and only if there exist functions $\varphi : [\underline{\theta}, \overline{\theta}] \to \mathbb{R}_+$ and $\tau : [\underline{\theta}, \overline{\theta}] \to \mathbb{R}_+$ that satisfy the following two conditions:

CONDITION 1′. For all $\theta, \theta' \in [\underline{\theta}, \overline{\theta}]$,

$$\theta u^*(\varphi(\theta); \mathbf{r}) + \beta v(\tau(\theta)) \geq \theta u^*(\varphi(\theta'); \mathbf{r}) + \beta v(\tau(\theta')),$$

where $u^*(y; \mathbf{r}) = \max_{\{\mathbf{c}' \in \mathbb{R}^n_+ : \sum_{i=1}^n c'_i \leq y\}} u(\mathbf{c}'; \mathbf{r})$ and

$$\varphi(\theta) + \tau(\theta) \leq 1.$$

CONDITION 2′. The pair (φ, τ) maximizes

$$\int_{\underline{\theta}}^{\overline{\theta}} [\theta u^*(\varphi(\theta); \mathbf{r}) + v(\tau(\theta))] g_{\mathbf{r}}(\theta)\, d\theta.$$

Suppose (χ, t) that satisfies Conditions 1 and 2. Then, by our discussion on money burning before the statement of Lemma 1, there exists a function $\varphi : [\underline{\theta}, \overline{\theta}] \to \mathbb{R}_+$ such that $u^*(\varphi(\theta); \mathbf{r}) = u(\chi(\theta); \mathbf{r})$ and $\varphi(\theta) \leq \sum_{i=1}^n \chi_i(\theta)$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$. Hence, letting $\tau \equiv t$, we have that (φ, τ) satisfies both Conditions 1′ and 2′.

Suppose (φ, τ) satisfy Conditions 1′ and 2′. For every $\theta \in [\underline{\theta}, \overline{\theta}]$, let

$$\chi(\theta) = \underset{\{\mathbf{c} \in \mathbb{R}^n_+ : \sum_{i=1}^n c_i \leq \varphi(\theta)\}}{\arg\max} u(\mathbf{c}; \mathbf{r}).$$

Then, by definition, $u(\chi(\theta); \mathbf{r}) = u^*(\varphi(\theta); \mathbf{r})$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$. Letting $t \equiv \tau$, we have that (χ, t) satisfies both Conditions 1 and 2.

A.3 *Proof of Proposition 2*

The proof uses the following three lemmas.

LEMMA 5. *When self-*0 *can set only* f*, every optimal* f *satisfies* $\underline{s}^p < f < \overline{s}^p$.

PROOF. We will show that $\Psi'(f) > 0$ for all $f \in (\underline{s}^d, \overline{s}^p]$ and $\Psi'(f-) < 0$ for $f = \overline{s}^p$. Recall that $(\mathbf{c}^f, s^f)$ is continuous in f for every ω and, therefore, so is $\Psi(f)$. These observations imply that every optimal f^* is in $(\underline{s}^p, \overline{s}^p)$.

For any $f \in (\underline{s}^d, \overline{s}^p]$, define $\Omega^+(f) = \{\omega : s^p(\omega) > f\}$ and $\Omega^-(f) = \{\omega : s^p(\omega) \leq f\}$. For $\omega \in \Omega^+(f)$, consider the fictitious problem maximizing $\theta u(\mathbf{c}; \mathbf{r}) + v(s)$ for $(\mathbf{c}, s) \in \mathbb{R}_+^{n+1}$ subject to $s + \sum_i c_i \leq 1$ and $s \leq f$. Letting $\mu(\omega)$ and $\phi^+(\omega)$ be the corresponding Lagrange multipliers, the first-order conditions are[22] $v'(s) = \mu(\omega) + \phi^+(\omega)$ and $\theta u_{c_i}(\mathbf{c}; \mathbf{r}) = \mu(\omega)$ for all i. Clearly, $s = f$ and $\phi^+(\omega) > 0$ for $\omega \in \Omega^+(f)$. Also, conditional on $s = f$, both selves would choose the same $\mathbf{c}$ in state ω, which is, therefore, $\mathbf{c}^f(\omega)$. Using (3), it follows that, for every i,

$$\phi^+(\omega) = v'(f) - \theta u_{c_i}\big(\mathbf{c}^f(\omega); \omega\big) = \tilde{U}'(f; \omega), \quad \omega \in \Omega^+(f). \tag{4}$$

For $\omega \in \Omega^-(f)$, consider the fictitious problem of maximizing $\theta u(\mathbf{c}; \mathbf{r}) + v(s)$ for $(\mathbf{c}, s) \in \mathbb{R}_+^{n+1}$ subject to $s + \sum_i c_i \leq 1$ and $s \geq f$. Letting $\mu(\omega)$ and $\phi^-(\omega)$ be the corresponding Lagrange multipliers, the first-order conditions are $v'(s) = \mu(\omega) - \phi^-(\omega)$ and $\theta u_{c_i}(\mathbf{c}; \mathbf{r}) = \mu(\omega)$ for all i. Clearly, $s = f$ and $\phi^-(\omega) \geq 0$ for $\omega \in \Omega^-(f)$. Also, conditional on $s = f$, both selves would choose the same $\mathbf{c}$ in state ω, which is, therefore, $\mathbf{c}^f(\omega)$. Using (3), it follows that, for every i,

$$\phi^-(\omega) = \theta u_{c_i}\big(\mathbf{c}^f(\omega); \mathbf{r}\big) - v'(f) = -\tilde{U}'(f; \omega), \quad \omega \in \Omega^-(f).$$

Consider any $f \in (\underline{s}^d, \overline{s}^p]$. Recall that $\overline{\Omega}(f) = \{\omega : s^d(\omega) \leq f\}$. Using Lemma 4, we have

$$\Psi'(f) = \int_{\overline{\Omega}(f) \cap \Omega^+(f)} \tilde{U}'(f; \omega)\, dG + \int_{\overline{\Omega}(f) \cap \Omega^-(f)} \tilde{U}'(f; \omega)\, dG = \int_{\overline{\Omega}(f) \cap \Omega^+(f)} \phi^+(\omega)\, dG,$$

where the last equality follows because either $\Omega^-(f) = \varnothing$ or $\phi^-(\omega) = 0$ for $\omega \in \Omega^-(f)$. The function $\phi^+(\omega)$ is strictly positive over $\overline{\Omega}(f) \cap \Omega^+(f)$. We need to show that $G(\overline{\Omega}(f) \cap \Omega^+(f)) > 0$, which implies $\Psi'(f) > 0$. This is immediate if $f \in (\underline{s}^d, \underline{s}^p)$, because $\Omega^+(f) = \Omega$. Consider $f = \underline{s}^p$. Clearly, $\overline{\Omega}(\underline{s}^p) \cap \Omega^+(\underline{s}^p)$ contains the open set $\overline{\Omega}^\circ(\underline{s}^p) \cap \Omega^+(\underline{s}^p) = \{\omega : s^d(\omega) < \underline{s}^p < s^p(\omega)\}$. If this set is nonempty, we are done because G has full support. Both $\overline{\Omega}^\circ(\underline{s}^p)$ and $\Omega^+(\underline{s}^p)$ are nonempty. Suppose that $\Omega^+(\underline{s}^p) \cap \overline{\Omega}^\circ(\underline{s}^p) = \varnothing$. Then, for every $\omega \in \Omega^+(\underline{s}^p)$, we have $s^d(\omega) \geq \underline{s}^p$ and that $\overline{\Omega}^\circ(\underline{s}^p) \subset \Omega^-(\underline{s}^p) = \{\omega : s^p(\omega) = \underline{s}^p\}$. Now, consider $\hat{\omega} \in \overline{\Omega}^\circ(\underline{s}^p)$ and any sequence $\{\omega_n\}$ in $\Omega^+(\underline{s}^p)$ converging to $\hat{\omega}$. We have that $\lim_{\omega_n \to \hat{\omega}} \inf s^d(\omega_n) \geq \underline{s}^p > s^d(\hat{\boldsymbol{\omega}})$. But this violates the continuity of s^d—a contradiction.

Now consider $f = \overline{s}^p$. Again using Lemma 4, we have

$$\Psi'\big(\overline{s}^p-\big) = \int_{\overline{\Omega}(\overline{s}^p)} \tilde{U}'\big(\overline{s}^p; \omega\big)\, dG = \int_\Omega \tilde{U}'\big(\overline{s}^p; \omega\big)\, dG = -\int_\Omega \phi^-(\omega)\, dG,$$

where $\phi^-(\omega) > 0$ for all ω such that $s^p(\omega) < \overline{s}^p$. Therefore, $\Psi'(\overline{s}^p-) < 0$.[23] □

[22] Here—and in the other proofs—the complementary slackness conditions are omitted for simplicity.

[23] It is easy to see that the optimal f satisfies $f \leq \overline{s}^p$. Suppose $f \in (\overline{s}^p, 1)$. Then, for all ω, self-1 chooses $s(\omega) = f$ and $\mathbf{c}(\omega) = \mathbf{c}^f(\omega)$. Take any $f' \in (\overline{s}^p, f)$. Then, for every ω, $f' = \zeta(\omega)f + (1 - \zeta(\omega))s^p(\omega)$ for some $\zeta(\omega) \in (0, 1)$. Therefore, for every ω, $\tilde{U}(f'; \omega) > \tilde{U}(f; \omega)$ because $\tilde{U}(s^p(\omega); \omega) > \tilde{U}(f; \omega)$ and $\tilde{U}(\cdot; \omega)$ is strictly concave. It follows that self-0's payoff is strictly larger under f' than under f.

Let $E(\beta)$ be the set of optimal floors and let $u^*(y;\mathbf{r})$ be the indirect utility of spending $y\in[0,1]$.

Lemma 6. *The set $E(\beta)$ is decreasing in β in the strong set order.*[24] *The largest optimal f converges monotonically to $\underline{s}^p$ as $\beta\uparrow 1$. There exists $\underline{\beta}>0$ such that $E(\beta)=\{\overline{f}\}$ for all $\beta\le\underline{\beta}$, where $\overline{f}$ satisfies $\overline{f}<\overline{s}^p$ and*

$$\mathcal{U}(B_{\overline{f}})=\max_{f\in[\underline{s}^p,\overline{s}^p]}\int_\Omega\big[\theta u^*(1-f;\mathbf{r})+v(f)\big]\,dG.$$

Proof. Fix $f\in[\underline{s}^d,\overline{s}^p]$. The set $\overline{\Omega}(f)$ in Lemma 4 depends on β via $(\mathbf{c}^d,s^d)$. By standard arguments, if $\beta<\beta'<1$, then $s^d(\omega;\beta)<s^d(\omega;\beta')$ for every ω and, hence, $\overline{\Omega}(f;\beta')\subset\overline{\Omega}(f;\beta)$. Also, $\Omega^-(f)\subset\overline{\Omega}(f;\beta)$ for every $\beta<1$ because $s^d(\omega;\beta)<s^p(\omega)$ for every ω. So, if $\beta<\beta'<1$,

$$\Psi'(f;\beta)-\Psi'\big(f;\beta'\big)=\int_{(\overline{\Omega}(f;\beta)\setminus\overline{\Omega}(f;\beta'))\cap\Omega^+(f)}\phi^+(\omega)\,dG\ge 0,$$

where the inequality uses (4). By standard results, $E(\beta)$ decreases in the strong set order.

Define $\overline{f}(\beta)=\max\{f:f\in E(\beta)\}$. Since $\overline{f}(\beta)\ge\underline{s}^p$ for all β and $\overline{f}(\cdot)$ is decreasing, $\lim_{\beta\uparrow 1}\overline{f}(\beta)$ exists; denote it by $\overline{f}(1-)\ge\underline{s}^p$. Clearly, $\overline{f}(1)=\underline{s}^p$. Now suppose that $\overline{f}(1-)>\overline{f}(1)$. By a similar argument, for any $f>\underline{s}^p$, $\lim_{\beta\uparrow 1}\Psi'(f;\beta)$ exists and equals $-\int_{\Omega^-(f)}\phi(\omega)\,dG<0$. Therefore, for β close enough to 1, $\overline{f}(\beta)\ge\overline{f}(1-)$ cannot be optimal—a contradiction that implies $\overline{f}(1-)=\overline{f}(1)$.

Note that $\overline{s}^d(\beta)=\max_\Omega s^d(\mathbf{s};\beta)$ falls monotonically to 0 as $\beta\downarrow 0$. Let $\underline{\beta}=\max\{\beta\in[0,1]:\overline{s}^d(\beta)\le\underline{s}^p\}$, which is strictly positive because $\underline{s}^p>0$. Then $\overline{\Omega}(f)=\Omega$ for all $\beta\le\underline{\beta}$ and $f\in[\underline{s}^p,\overline{s}^p]$, and, hence,

$$\Psi(f;\beta)=\int_\Omega\big[\theta u\big(\mathbf{c}^f(\omega);\mathbf{r}\big)+v(f)\big]\,dG. \tag{5}$$

From the proof of Lemma 4, $u(\mathbf{c}^f(\omega);\mathbf{r})=\tilde{u}(f;\omega)$ is strictly concave in f for all ω. Thus, the maximizer of (5) is unique. From the proof of Lemma 5, the derivative of (5) is negative at $\overline{s}^p$ and, hence, $\overline{f}<\overline{s}^p$. □

We now show that self-0 benefits from using only b_i.

Lemma 7. *Fix i and consider plans B_{b_i} with $b_j=1$ for all $j\neq i$ and $f=0$. There exists $b_i<\max_\omega c_i^d(\omega)\equiv\overline{c}_i^d$ such that $\mathcal{U}(B_{b_i})>\mathcal{U}(F)$.*

Proof. Fix $i=1$ and any $b_1\in(0,\overline{c}_1^d]$. Let $(\mathbf{c}^{b_1},s^{b_1})$ describe self-1's choices under B_{b_1} and let

$$\Phi(b_1)=\int_\Omega\big[\theta u\big(\mathbf{c}^{b_1}(\omega);\mathbf{r}\big)+v\big(s^{b_1}(\omega)\big)\big]\,dG.$$

[24] Given two sets E and E' in $\mathbb{R}$, $E\ge E'$ in the strong set order if, for every $f\in E$ and $f'\in E'$, $\min\{f,f'\}\in E'$ and $\max\{f,f'\}\in E$ (Milgrom and Shannon 1994).

Let $\Omega(b_1)=\{\omega : c_1^d(\omega)>b_1\}$. Since c_1^d is continuous, $\Omega(b_1)$ is nonempty and open if $b_1<\overline{c}_1^d$, and, hence, $G(\Omega(b_1))>0$. We have

$$\begin{aligned}\Phi(b_1)-\Phi(\overline{c}_1^d) &= \int_{\Omega(b_1)} \big\{\big[\theta u(\mathbf{c}^{b_1}(\omega);\mathbf{r})+v(s^{b_1}(\omega))\big]-\big[\theta u(\mathbf{c}^d(\omega);\mathbf{r})+v(s^d(\omega))\big]\big\}\,dG \\ &= (1-\beta)\int_{\Omega(b_1)}\big[v(s^{b_1}(\omega))-v(s^d(\omega))\big]\,dG \\ &\quad + \int_{\Omega(b_1)}\big[\tilde V(c_1^{b_1}(\omega);\omega)-\tilde V(c_1^d(\omega);\omega)\big]\,dG,\end{aligned}$$

where

$$\tilde V(\hat b_1;\omega)=\max_{\{(\mathbf{c},s)\in\mathbb{R}_+^{n+1}:\sum_{j=1}^n c_j\le 1, c_1\le \hat b_1\}}\big\{\theta u(\mathbf{c};\mathbf{r})+\beta v(s)\big\}.$$

Clearly, $\tilde V(c_1^d(\omega);\omega)\ge \tilde V(b_1;\omega)$ for all ω. From the first-order conditions of the Lagrangian defining $\tilde V(\hat b_1;\omega)$, we have $\lambda_1(\omega;\hat b_1)=\theta u_c^1(c_1^{\hat b_1}(\omega);r_1)-\beta v'(s^{\hat b_1}(\omega))$, where $\lambda_1(\omega;\hat b_1)$ is the Lagrange multiplier for $c_1\le \hat b_1$. Since $(\mathbf{c}^{\hat b_1}(\omega),s^{\hat b_1}(\omega))$ is continuous in $\hat b_1$ and ω, so is $\lambda_1(\omega;\hat b_1)$. Again by Theorem 1 of Luenberger 1969, p. 222), $\tilde V'(\hat b_1;\omega)$ exists for all $\hat b_1$ and equals $\lambda_1(\omega;\hat b_1)$. It follows that $\tilde V'(c_1^d(\omega);\omega)=0$ for all ω by the definition of $(\mathbf{c}^d,s^d)$. By the mean value theorem (MVT), $\tilde V(c_1^{b_1}(\omega);\omega)-\tilde V(c_1^d(\omega);\omega)=\tilde V'(\chi(\omega);\omega)(c_1^{b_1}(\omega)-c_1^d(\omega))$ and $v(s^{b_1}(\omega))-v(s^d(\omega))=v'(\xi(\omega))(s^{b_1}(\omega)-s^d(\omega))$, where $\chi(\omega)\in[c_1^{b_1}(\omega),c_1^d(\omega)]$ and $\xi(\omega)\in[s^d(\omega),s^{b_1}(\omega)]$.

Let $b_1^\varepsilon=\overline{c}_1^d-\varepsilon$ for some $\varepsilon>0$. Fix $\omega\in\Omega(b_1^\varepsilon)$ and, for now, drop the dependence on ω. Recall that $s^{b_1^\varepsilon}+\sum_i c_i^{b_1^\varepsilon}=s^d+\sum_i c_i^d=1$. Since $s^{b_1^\varepsilon}>s^d$ for $\varepsilon>0$ (Lemma 3), we can write

$$-\frac{c_1^{b_1^\varepsilon}-c_1^d}{s^{b_1^\varepsilon}-s^d}=1+\sum_{j\ne 1}\frac{c_j^{b_1^\varepsilon}-c_j^d}{s^{b_1^\varepsilon}-s^d}. \tag{6}$$

Now, for any b_1^ε, the first-order condition $\beta v'(s)-\theta u_c^j(c_j;r_j)=0$ must hold for every $j\ne 1$. Therefore, again by the MVT, for all $j\ne 1$,

$$c_j^{b_1^\varepsilon}-c_j^d=\frac{\beta\big[v'(s^{b_1^\varepsilon})-v'(s^d)\big]}{\theta u_{cc}^j(\zeta_j;r_j)} \tag{7}$$

for some $\zeta\in[c_j^d,c_j^{b_1^\varepsilon}]$. Now, since v'' is continuous, $v'(y)-v'(\hat y)\ge \underline{v}''[y-\hat y]$ for every $y>\hat y\ge \underline{s}^d$, where $\underline{v}''=\min_{\xi\in[\underline{s}^d,1]}v''(\xi)<0$. Therefore, using (6) and (7),

$$\begin{aligned}-\frac{c_1^{b_1^\varepsilon}-c_1^d}{s^{b_1^\varepsilon}-s^d} &= 1+\frac{1}{s^{b_1^\varepsilon}-s^d}\sum_{j\ne 1}\frac{\beta}{\theta u_{cc}^j(\zeta_j;r_j)}\big[v'(s^{b_1^\varepsilon})-v'(s^d)\big] \\ &\le 1+\frac{1}{s^{b_1^\varepsilon}-s^d}\sum_{j\ne 1}\frac{\beta\underline{v}''}{\theta u_{cc}^j(\zeta_j;r_j)}\big[s^{b_1^\varepsilon}-s^d\big]\le 1+\frac{\beta\underline{v}''}{\underline{\theta}}\sum_{j\ne 1}\frac{1}{\overline{u}_{cc}^j},\end{aligned}$$

where the first inequality uses $u^j_{cc} < 0$ and $\overline{u}^j_{cc} = \max_{\xi \in [\underline{s}^d, 1], r_j \in [\underline{r}_j, \overline{r}_j]} u^j_{cc}(\xi; r_j) < 0$. Letting $K = [1 + \frac{\beta \underline{v}''}{\underline{\theta}} \sum_{j \neq 1} \frac{1}{\overline{u}^j_{cc}}]^{-1}$, it follows that $s^{b^\varepsilon_1}(\omega) - s^d(\omega) \geq K[c^d_1(\omega) - c^{b^\varepsilon_1}_1(\omega)]$ for every $\omega \in \Omega(b^\varepsilon_1)$.

These observations imply that $\Phi(b^\varepsilon_1) - \Phi(\overline{c}^d_1)$ is bounded below by

$$\int_{\Omega(b^\varepsilon_1)} \left[K(1-\beta)v'\big(\xi(\omega)\big) - \tilde{V}'\big(\chi(\omega); \omega\big)\right]\big(c^d_1(\omega) - b^\varepsilon_1\big)\, dG. \tag{8}$$

Since v' is continuous and strictly positive everywhere and $\xi(\omega) \in [\underline{s}^d, 1]$ with $\underline{s}^d > 0$ for all $\omega \in \Omega(b^\varepsilon_1)$, there exists a finite $\kappa > 0$ such that $v'(\xi(\omega)) \geq \kappa$ for all $\omega \in \Omega(b^\varepsilon_1)$.

Next let $\overline{\Omega}(b^\varepsilon_1) = \{\omega : c^d_1(\omega) \geq b^\varepsilon_1\}$. By continuity of c^d_1, $\overline{\Omega}(\cdot)$ is a compact-valued and continuous correspondence. Note that $\tilde{V}'(\chi(\omega); \omega) = \tilde{V}'(c^d_1(\omega); \omega) = 0$ if $c^d_1(\omega) = b^\varepsilon_1$. We have

$$\sup_{\omega \in \Omega(b^\varepsilon_1)} \tilde{V}'\big(\chi(\omega); \omega\big) = \sup_{\omega \in \overline{\Omega}(b^\varepsilon_1)} \tilde{V}'\big(\chi(\omega); \omega\big) \leq \max_{b^\varepsilon_1 \leq \zeta \leq \overline{c}^d_1, \omega \in \overline{\Omega}(b^\varepsilon_1)} \tilde{V}'(\zeta; \omega) \equiv \kappa\big(b^\varepsilon_1\big).$$

Clearly, $\kappa(b^\varepsilon_1) \geq 0$ for every $\varepsilon > 0$, $\kappa(b^\varepsilon_1) \leq \kappa(b^{\varepsilon'}_1)$ for $\varepsilon' > \varepsilon > 0$, and $\lim_{\varepsilon \to 0} \kappa(b^\varepsilon_1) = 0$ because $\kappa(\cdot)$ is continuous. Therefore, there exists $\varepsilon^* > 0$ such that $\kappa(b^{\varepsilon^*}_1) < \kappa(1-\beta)K$. It follows that for all $\varepsilon \in (0, \varepsilon^*]$, expression (8) is strictly positive and, hence, $\Phi(b^{\varepsilon^*}_1) > \Phi(\overline{c}^d_1)$. □

We can now complete the proof. By Lemma 6, $\overline{f}(\beta)$ falls monotonically to $\underline{s}^p$ when $\beta \uparrow 1$. Also, for every $i = 1, \ldots, n$, $s^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}; \beta)$ rises monotonically to $s^p(\overline{\theta}, \overline{r}_i, \underline{r}_{-i})$ as $\beta \uparrow 1$. By Remark 1, $s^p(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}) > \underline{s}^p$. Given this, define

$$\beta^* = \inf\Big\{\beta \in (0,1) : \overline{f}(\beta) < \max_i s^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}; \beta)\Big\}.$$

Clearly, $\beta^* < 1$ and, for every $\beta > \beta^*$, we have $s^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}; \beta) > \overline{f}(\beta)$ for at least some $i = 1, \ldots, n$. Hereafter, fix $\beta > \beta^*$ and any i that satisfies this last condition.

For $\varepsilon \geq 0$, consider $b^\varepsilon_i = \overline{c}^d_i - \varepsilon$ as in Lemma 7, where $\overline{c}^d_i = c^d_i(\overline{\theta}, \overline{r}_i, \underline{r}_{-i})$ by Remark 1. Let $\Phi(b^\varepsilon_i, \overline{f}(\beta))$ be self-0's payoff from adding b^ε_i to the existing $\overline{f}(\beta)$. We will show that there exists $\varepsilon > 0$ such that $\Phi(b^\varepsilon_i, \overline{f}(\beta)) > \Phi(b^0_i, \overline{f}(\beta)) = \mathcal{U}(B_{\overline{f}(\beta)})$. To do so, for any $\varepsilon \geq 0$, let $(\mathbf{c}^\varepsilon, s^\varepsilon)$ be self-1's choice function under $(b^\varepsilon_i, \overline{f}(\beta))$ and $\Omega(b^\varepsilon_i) = \{\omega \in \Omega : c^0_i(\omega) > b^\varepsilon_i\}$. Then

$$\begin{aligned}\Phi\big(b^\varepsilon_i, \overline{f}(\beta)\big) - \Phi\big(b^0_i, \overline{f}(\beta)\big) = \int_{\Omega(b^\varepsilon_i)} \big\{&\big[\theta u\big(\mathbf{c}^\varepsilon(\omega); \mathbf{r}\big) + v\big(s^\varepsilon(\omega)\big)\big] \\ &- \big[\theta u\big(\mathbf{c}^0(\omega); \mathbf{r}\big) + v\big(s^0(\omega)\big)\big]\big\}\, dG.\end{aligned}$$

Note that if there exists $\overline{\varepsilon} > 0$ such that $(\mathbf{c}^\varepsilon(\omega), s^\varepsilon(\omega)) = (\mathbf{c}^d(\omega), s^d(\omega))$ for all $\omega \in \Omega(b^\varepsilon_i)$ and $0 < \varepsilon < \overline{\varepsilon}$, then for such εs the previous difference equals $\Phi(b^\varepsilon_i) - \Phi(\overline{c}^d_i)$ in the proof of Lemma 7. By the conclusion of that proof, there exists $\varepsilon^{**} \in (0, \overline{\varepsilon})$ such that $\Phi(b^{\varepsilon^{**}}_i, \overline{f}(\beta)) > \Phi(b^0_i, \overline{f}(\beta))$.

Thus we need only to prove the existence of $\overline{\varepsilon}$. Let $\overline{\Omega}(\overline{f}(\beta)) = \{\omega \in \Omega : s^d(\omega) \leq \overline{f}(\beta)\}$, which is compact by continuity of s^d. Define $\tilde{c}_i = \max_{\overline{\Omega}(\overline{f}(\beta))} c^0_i(\omega)$. Since $s^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}) >$

$\overline{f}(\beta)$, $(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}) \notin \overline{\Omega}(\overline{f}(\beta))$ and, hence, $c_i^0(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}) = c_i^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i})$, where $c_i^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}) = \overline{c}_i^d$ by Remark 1. We must also have $\tilde{c}_i < \overline{c}_i^d$: Indeed, for all $\omega \in \overline{\Omega}(\overline{f}(\beta))$, optimality requires

$$\theta u_c^i\big(c_i(\omega); r_i\big) = \beta v'\big(\overline{f}(\beta)\big) + \lambda_0(\omega) > \beta v'\big(s^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i})\big) = \overline{\theta} u_c^i\big(\overline{c}_i^d; \overline{r}_i\big),$$

where $\lambda_0(\omega) \geq 0$ is the Lagrange multiplier for $s \geq \overline{f}(\beta)$. If ω is such that $c_i^0(\omega) > \tilde{c}_i$, then $\omega \notin \overline{\Omega}(\overline{f}(\beta))$—otherwise it would contradict the definition of $\tilde{c}_i$—and, hence, $\mathbf{c}^0(\omega) = \mathbf{c}^d(\omega)$. Let $\overline{\varepsilon} = \overline{c}_i^d - \tilde{c}_i > 0$. By construction for $\varepsilon \in (0, \overline{\varepsilon})$, $c_i^0(\omega) > b_i^\varepsilon$ implies $\mathbf{c}^0(\omega) = \mathbf{c}^d(\omega)$, as desired.

A.4 *Proofs of Proposition 3 and Corollary 1*

Lemma 8. *For every $\beta \in (0,1)$, if $B \in \mathcal{B}$ is optimal, then*

$$\max\left\{f, 1 - \sum_{i=1}^n b_i\right\} \geq \min_{\omega \in \Omega} s^p(\omega).$$

Proof. Define $\sigma = \max\{f, 1 - \sum_{i=1}^n b_i\}$. Since $s(\omega) + \sum_{i=1}^n c_i(\omega) = 1$, $s(\omega) \geq \sigma$ for all ω. Without loss of generality, we can let $\sigma = \min_\Omega s(\omega)$: If $\min_\Omega s(\omega) > \sigma$, we could raise f to the level $\min_\Omega s(\omega)$ and nothing would change.

Now fix $\beta \in (0,1)$. Suppose B' is optimal, but $\sigma' < \underline{s}^p$. Consider $B'' \in \mathcal{B}$ equal to B', except for $f'' = \underline{s}^p$. Since B' is convex and compact, the ensuing allocation $(\mathbf{c}', s')$ is continuous in ω. So $\Omega(\underline{s}^p) = \{\omega \in \Omega : s'(\omega) < \underline{s}^p\}$ contains an open subset and $G(\Omega(\underline{s}^p)) > 0$. Consider $\omega \in \Omega(\underline{s}^p)$ and the problem of maximizing $\theta u(\mathbf{c}; \mathbf{r}) + v(s)$ for $(\mathbf{c}, s) \in \mathbb{R}_+^{n+1}$ subject to $c_i \leq b_i'$ for $i = 1, \dots, n$ and $s \leq f$. For any $f < \underline{s}^p$, the latter must bind because, by the logic of Lemma 3, self-0 would want to save at least $s^p(\omega) \geq \underline{s}^p$ if facing only $c_i \leq b_i'$ for $i = 1, \dots, n$. Therefore, self-0's payoff from this fictitious problem is strictly increasing in f for $f \leq \underline{s}^p$. When self-1 faces B'', the constraint $s \geq \underline{s}^p$ must bind, so his allocation $(\mathbf{c}''(\omega), \underline{s}^p)$ solves $\max u(\mathbf{c}; \mathbf{r})$ subject to $\mathbf{c} \in \mathbb{R}_+^n$, $c_i \leq b_i'$, and $\sum_{i=1}^n c_i \leq 1 - \underline{s}^p$. This allocation coincides with self-0's allocation in the fictitious problem with $f = \underline{s}^p$. Hence, in ω, $(\mathbf{c}''(\omega), \underline{s}^p)$ is strictly better for self-0 than $(\mathbf{c}''(\omega), s'(\omega))$. For all $\omega \in \Omega(\underline{s}^p)$, self-0's payoff is then strictly larger under B'' than B'. Since for $\omega \notin \Omega(\underline{s}^p)$, self-1's allocation is unchanged, $\mathcal{U}(B'') > \mathcal{U}(B')$—a contradiction. □

Given Lemma 8, we now complete the proof of Proposition 3. We first show that there exists $\beta_{**} > 0$ such that, if $\beta < \beta_{**}$, then for any $B \in \mathcal{B}$ with $\sigma \geq \underline{s}^p$, the resulting allocation $(\mathbf{c}, s)$ satisfies $s(\omega) = \sigma$ for all $\omega \in \Omega$. It is enough to show that $s(\underline{\theta}, \underline{\mathbf{r}}) = \overline{s} = \max_\Omega s(\omega)$ must equal σ. By strict concavity of v, $v'(\overline{s}) \leq v'(\underline{s}^p) < +\infty$ because $\underline{s}^p > 0$. By considering the Lagrangian of self-1's problem in $\underline{\omega} = (\underline{\omega}, \underline{\mathbf{r}})$, we have that $(\mathbf{c}(\underline{\omega}), s(\underline{\omega}))$ must satisfy $\beta v'(\overline{s}) + \phi_0(\underline{\omega}) + \gamma_i(\underline{\omega}) = \underline{\theta} u_c^i(c_i(\underline{\omega}); \underline{r}_i)$ for all $i = 1, \dots, n$, where $\phi_0(\underline{\omega}) \geq 0$ and $\gamma_i(\underline{\omega}) \geq 0$ are the Lagrange multipliers for $s \geq f$ and $c_i \leq b_i$. For every $i = 1, \dots, n$, since $c_i(\underline{\omega}) \leq 1$ and $u^i(\cdot; \underline{r}_i)$ is strictly concave, $u_c^i(c_i(\underline{\omega}); \underline{r}_i) \geq u_c^i(1; \underline{r}_i) > 0$. Now let

$$\beta_{**} = \min_i \frac{\underline{\theta} u_c^i(1; \underline{r}_i)}{v'\big(\underline{s}^p\big)} > 0. \tag{9}$$

Then, for every $\beta < \beta_{**}$, we have $\beta v'(s(\underline{\omega})) < \underline{\theta} u_c^i(c_i(\underline{\omega}); \underline{r}_i)$ for all $i = 1, \ldots, n$. Therefore, $\phi_0(\underline{\omega}) + \gamma_i(\underline{\omega}) > 0$ for all $i = 1, \ldots, n$. Hence, either $\phi_0(\underline{\omega}) > 0$, in which case $\overline{s} = f = \sigma$, or $\gamma_i(\underline{\omega}) > 0$ for all $i = 1, \ldots, n$, in which case $\overline{s} = 1 - \sum_{i=1}^n c_i(\underline{\omega}) = 1 - \sum_{i=1}^n b_i = \sigma$.

Finally, let $\beta < \beta_* = \min\{\underline{\beta}, \beta_{**}\}$, where $\underline{\beta} > 0$ was defined in Lemma 6. Let $B^\beta \in \mathcal{B}$ be an optimal plan for β. By Lemma 8, $\sigma^\beta \geq \underline{s}^p$. The previous result then implies that $\mathcal{U}(B^\beta) = \int_\Omega [\theta u(\mathbf{c}(\omega); \mathbf{r}) + v(\sigma^\beta)]\, dG$. Hence,

$$\mathcal{U}(B^\beta) \leq \int_\Omega \big[\theta u\big(\mathbf{c}^{\sigma^\beta}(\omega); \mathbf{r}\big) + v\big(\sigma^\beta\big)\big]\, dG \leq \int_\Omega \big[\theta u\big(\mathbf{c}^{\overline{f}}(\omega); \mathbf{r}\big) + v(\overline{f})\big]\, dG = \mathcal{U}(B_{\overline{f}}),$$

where the first inequality holds since $u(\mathbf{c}(\omega); \mathbf{r}) \leq \max_{\{\mathbf{c} \in \mathbb{R}_+^n : \sum_{i=1}^n c_i \leq \sigma^\beta\}} u(\mathbf{c}; \mathbf{r}) = u(\mathbf{c}^{\sigma^\beta}(\omega); \mathbf{r})$ for all $\omega \in \Omega$ and from the definition of $\overline{f}$ in Lemma 6. If B^β involves budgets that bind for a set Ω' with $G(\Omega') > 0$, then $u(\mathbf{c}(\omega); \mathbf{r}) < u(\mathbf{c}^{\sigma^\beta}(\omega); \mathbf{r})$ for all $\omega \in \Omega'$ and $\mathcal{U}(B^\beta) < \mathcal{U}(B_{\overline{f}})$. Therefore, optimal plans can only use f.

Finally, let $\underline{\mathbf{r}}'$, $\underline{\mathbf{r}}$, $\overline{\mathbf{r}}'$, and $\overline{\mathbf{r}}$ satisfy the properties in Corollary 1. Then $\underline{s}^{p\prime} = s^p(\overline{\theta}, \overline{\mathbf{r}}') \geq s^p(\overline{\theta}, \overline{\mathbf{r}}) = \underline{s}^p$ with strict inequality if $\overline{\mathbf{r}} \neq \overline{\mathbf{r}}'$ (Remark 1). Similarly, for $\beta \in (0, 1)$, $\overline{s}^{d\prime}(\beta) = s^d(\underline{\theta}, \underline{\mathbf{r}}'; \beta) \leq s^d(\underline{\theta}, \underline{\mathbf{r}}; \beta) = \overline{s}^d(\beta)$ with strict inequality if $\underline{\mathbf{r}}' \neq \underline{\mathbf{r}}$. The definition of β_{**} in (9), the strict concavity of v, and $\underline{r}_i' \geq \underline{r}_i$ imply that $\beta_{**}' > \beta_{**}$. The definition of $\underline{\beta}$ in the proof of Lemma 6 and the fact that $\overline{s}^d$ is strictly increasing in β imply that $\underline{\beta}' > \underline{\beta}$. Therefore, $\beta_*' > \beta_*$.

A.5 *Proof of Proposition 4*

The proof is constructive: It first establishes the claimed properties for the illustrative three-state setting in Section 3.4, which are then used for the general result.

Part I Let $\omega^0 = (\overline{\theta}, \overline{r}_1, \overline{r}_2)$, $\omega^1 = (\underline{\theta}, \overline{r}_1, \underline{r}_2)$, and $\omega^2 = (\underline{\theta}, \underline{r}_1, \overline{r}_2)$ with respective probabilities g, $\frac{1-g}{2}$, and $\frac{1-g}{2}$. Choose $\beta < 1$ so that $s^d(\omega^1) = s^d(\omega^2) > s^p(\omega^0)$ and $\underline{\theta}$ so that $c_1^p(\omega^1) > c_1^p(\omega^0)$ and $c_2^p(\omega^2) > c_2^p(\omega^0)$.

Lemma 9. *There exists $g^* \in (0, 1)$ such that if $g > g^*$, then the optimal $B \in \mathcal{B}$ satisfies $f = s^p(\omega^0)$, $b_1 = c_1^p(\omega^1)$, and $b_2 = c_2^p(\omega^2)$.*

Existence of an optimal B follows from an argument similar to the proof of Lemma 2. Claims 1–5 characterize its properties.

Claim 1. There exists $g^* \in (0, 1)$ such that if $g > g^*$ and self-0 can use only f, she sets $f = s^p(\omega^0)$.

Proof. We can focus on $f \in [s^d(\omega^1), s^p(\omega^1)] \cup \{s^p(\omega^0)\}$. For simplicity, let $U(\mathbf{c}, s; \omega) = \theta u(\mathbf{c}; \mathbf{r}) + v(s)$. If $f = s^p(\omega^0)$, by symmetry, self-0's payoff is $gU(\mathbf{c}^p(\omega^0), s^p(\omega^0); \omega^0) + (1-g)U(\mathbf{c}^d(\omega^1), s^d(\omega^1); \omega^1)$; if instead $f \in [s^d(\omega^1), s^p(\omega^1)]$, self-0's payoff is $gU(\mathbf{c}^f(\omega^0), f; \omega^0) + (1-g)U(\mathbf{c}^f(\omega^1), f; \omega^1)$, where $\mathbf{c}^f(\omega)$ is defined in Lemma 4. Thus, $f = s^p(\omega^0)$ identifies the best B that uses only f if

$$\frac{g}{1-g} > \max_{f \in [s^d(\omega^1), s^p(\omega^1)]} \frac{U\big(\mathbf{c}^f(\omega^1), f; \omega^1\big) - U\big(\mathbf{c}^d(\omega^1), s^d(\omega^1); \omega^1\big)}{U\big(\mathbf{c}^p(\omega^0), s^p(\omega^0); \omega^0\big) - U\big(\mathbf{c}^f(\omega^0), f; \omega^0\big)} \geq 0.$$

The term on the right-hand side is well defined; also, for all $f \in [s^d(\omega^1), s^p(\omega^1)]$ we have $U(\mathbf{c}^p(\omega^1), s^p(\omega^1); \omega^1) \geq U(\mathbf{c}^f(\omega^1), f; \omega^1) \geq U(\mathbf{c}^d(\omega^1), s^d(\omega^1); \omega^1)$ and $U(\mathbf{c}^p(\omega^0), s^p(\omega^0); \omega^0) > U(\mathbf{c}^f(\omega^0), f; \omega^0)$ because $s^d(\omega^1) > s^p(\omega^0)$. □

Hereafter, assume that $g > g^*$.

Claim 2. Fix $i \in \{1, 2\}$. Suppose self-0 knows that the state is ω^i and can only use b_i. Then $b_i = c_i^p(\omega^i)$ is optimal.

Proof. Let $i = 1$. The argument in the proof of Lemma 7 implies that it is optimal to set $b_1 < c_1^d(\omega^1)$. To find the optimal b_1, consider first self-1's problem to maximize $\underline{\theta}[\bar{r}\ln(c_1) + \underline{r}\ln(c_2)] + \beta\ln(s)$ subject to $s + c_1 + c_2 \leq 1$ and $c_1 \leq b_1$. Since both constraints must bind, this becomes $\max_{s\in[0,1]}\{\underline{\theta r}\ln(1 - b_1 - s) + \beta\ln(s)\}$. The solution is characterized by the first-order conditions, which lead to

$$s(b_1) = \frac{\beta}{\underline{\theta r} + \beta}(1 - b_1) \quad \text{and} \quad c_2(b_1) = \frac{\underline{\theta r}}{\underline{\theta r} + \beta}(1 - b_1).$$

Given this, self-0's payoff in ω^1 as a function of b_1 becomes (up to a constant)

$$\underline{\theta}\big[\bar{r}\ln(b_1) + \underline{r}\ln(1 - b_1)\big] + \ln(1 - b_1). \tag{10}$$

The optimal b_1 is again characterized by the first-order condition, which leads to

$$b_1 = \frac{\underline{\theta}\bar{r}}{1 + \underline{\theta}[\underline{r} + \bar{r}]}. \tag{11}$$

Finally, $c_1^p(\omega^1)$ results from maximizing $\underline{\theta}[\bar{r}\ln(c_1) + \underline{r}\ln(c_2)] + \ln(s)$ subject to $s + c_1 + c_2 \leq 1$. Substituting $s = 1 - c_1 - c_2$ and combining the first-order conditions yields

$$c_1^p(\omega^1) = \frac{\underline{\theta}\bar{r}}{1 + \underline{\theta}[\underline{r} + \bar{r}]}.$$ □

Claim 3. Fix $i \in \{1, 2\}$. Suppose self-0 knows that the state is ω^i. Then she strictly prefers to use only b_i rather than only b_{-i}.

Proof. Let $i = 1$. Similar calculations to the proof of Claim 2 imply that if self-0 can impose only b_2, she sets

$$b_2 = \frac{\underline{\theta r}}{1 + \underline{\theta}[\underline{r} + \bar{r}]}. \tag{12}$$

We want to argue that self-0's payoff in ω^1 is strictly larger if she uses only b_1 as in (11) than if she uses only b_2 as in (12). Substituting the choices implied by b_1 and b_2 into self-0's utility function, one can show that b_1 in (11) is strictly better than b_2 in (12) if and only if

$$(1 + \underline{\theta}\bar{r})\ln(\beta + \underline{\theta}\bar{r}) - (1 + \underline{\theta r})\ln(\beta + \underline{\theta r}) > (1 + \underline{\theta}\bar{r})\ln(1 + \underline{\theta}\bar{r}) - (1 + \underline{\theta r})\ln(1 + \underline{\theta r}).$$

To show that this holds, consider $\varphi(\beta, r) = (1 + \underline{\theta} r)\ln(\beta + \underline{\theta} r)$ for $0 < \beta < 1$ and $r > 0$. Since

$$\varphi_{\beta r}(\beta, r) = \frac{\partial}{\partial r}\left(\frac{1+\underline{\theta} r}{\beta + \underline{\theta} r}\right) = \frac{\underline{\theta}(\beta - 1)}{(\beta + \underline{\theta} r)^2} < 0,$$

$\varphi(\beta, \overline{r}) - \varphi(\beta, \underline{r})$ is strictly decreasing in β. Continuity gives the result. □

Claim 4. If B is optimal, then f can bind at most in ω^0.

Proof. If f binds in all ωs, then B is weakly dominated by a policy using only f and no budgets, as they distort $\mathbf{c}$ without raising s. Given $g > g^*$, by Claim 1 the latter plan is strictly dominated by one using only $f = s^p(\omega^0)$. Clearly, if f binds in ω^1 and ω^2, it must bind in ω^0.

Now suppose that f binds only in ω^0 and another state, say, ω^1. There are two cases:

Case 1: b_1 does not bind in ω^2. Then removing b_1 leads to a weakly superior policy in which f binds only in ω^0 and ω^1. Given $g > g^*$, however, the gain from raising f above $s^p(\omega^0)$ to improve self-1's allocation only in ω^1 does not justify the loss created in ω^0. Therefore, B is again strictly dominated by the policy obtained if we remove b_1 and set $f = s^p(\omega^0)$.

Case 2: b_1 binds also in ω^2. This implies that f has to bind in all ωs. Indeed, since b_1 binds in both ω^1 and ω^2, self-1 chooses $c_1 = b_1$ in both states; also, since good 2 is more valuable in ω^2 than in ω^1, self-1 wants to allocate more income to good 2 than to s relative to ω^1, and so f also binds in ω^2. But we know that such a policy is strictly dominated by using only $f = s^p(\omega^0)$. □

Claim 5. If B involves binding budgets, then b_i can bind at most in ω^i for $i = 1, 2$.

Proof. Without loss, consider b_1. Suppose first that b_1 binds in all states, which implies that $c_1(\omega^i) = b_1$ for all $i = 0, 1, 2$. There are five cases to consider:

Case 1: Neither b_2 nor f binds in any state. Since $\overline{\theta} > \underline{\theta}$, $c_2(\omega^0) > c_2(\omega^2)$. The plan cannot be optimal because, given b_1, self-0 would be strictly better off by adding an f that binds only in ω^0: Even if b_1 were binding for self-0 in ω^0, she would strictly prefer $c_2 < c_2(\omega^0)$ of good 2.

Case 2: b_2 binds in all states. Then $c_2(\omega^i) = b_2$ and $s(\omega^i) = 1 - b_1 - b_2$ for all $i = 0, 1, 2$. This plan is strictly dominated by one that imposes only $f = 1 - b_1 - b_2$ (because budgets are distorting), which is, in turn, strictly dominated by the plan with only $f = s^p(\omega^0)$ given $g > g^*$.

Case 3: b_2 binds in no state. Then, as in Case 1, for B to be optima, f must bind at least in ω^0 and only in ω^0 by Claim 4. Since, by assumption, b_1 binds in all states, it must be that $b_1 < c_1^p(\omega^1)$. Indeed, if $b_1 \geq c_1^p(\omega^1)$, the optimal f equals $s^p(\omega^0)$; since, by assumption, $c_1^p(\omega^0) < c_1^p(\omega^1)$, b_1 cannot bind in ω^0. It follows that with regard to ω^0 and ω^1, self-0 would be strictly better off replacing b_1 and f with $\hat{b}_1 = c_1^p(\omega^1)$ and $\hat{f} = s^p(\omega^0)$. With regard to ω^2, self-0 would be better off by replacing b_2 with $\hat{b}_2 = c_2^p(\omega^2)$: By Claim 3, even if b_1 were perfectly tailored for ω^2, it would be strictly dominated in ω^2 by $\hat{b}_2$.

Case 4: b_2 binds only in ω^0. Since self-1's choices satisfy $c_2(\omega^0) > c_2(\omega^2)$ if the plan used only b_1, it follows that self-0 can obtain in all ωs the same allocations induced by B if she uses an f that binds only in ω^0. Such a plan, however, is again strictly dominated as in Case 3.

Case 5: b_2 binds in ω^0 and in ω^2. Since self-1's choices satisfy $c_2(\omega^0) > c_2(\omega^2)$ if the plan used only b_1, self-0 could again obtain the same allocation in all ωs with an f that binds only in ω^0 and ω^2. By Claim 4, however, such a plan cannot be optimal.

Now suppose b_1 binds in only two states. If these states are ω^1 and ω^0, by the same argument as in Case 3 above, self-0 is strictly better off by replacing b_1 and f with $\hat{b}_1 = c_1^p(\omega^1)$ and $\hat{f} = s^p(\omega^0)$ as well as b_2 with $\hat{b}_2 = c_2^p(\omega^2)$. If b_1 binds in ω^1 and ω^2, it must also bind in ω^0, which is the case we considered before. Indeed, if b_1 binds in ω^2, then it will also bind at the fictitious state $(\underline{\theta}, \overline{r}, \overline{r})$ and, hence, in ω^0, where both goods are more valuable. Last, b_1 cannot bind only in ω^0 and ω^2: It would have to bind also in ω^1, since in ω^1, good 1 is more valuable than in ω^2.

Finally, suppose b_1 binds in only one ω. We have just argued that if b_1 binds in ω^2, it must also bind in ω^1. Thus, we have to rule out the case where b_1 binds only in ω^0. This is possible only if b_2 also binds in ω^0, inducing self-1 to overconsume good 1. However, such a b_2 must also bind in ω^2, but we just showed that a b_i cannot bind in more than one ω. □

Part II The three-state setting can be modified so that the optimal B uses only b_1 and b_2. Fix $g > g^*$ and the other parameters, except $\overline{\theta}$. Raise $\overline{\theta}$ to $\overline{\theta}'$ so that $c_i^p(\omega^0) > c_i^p(\omega^1)$ for $i = 1, 2$.

Lemma 10. *There exists $\overline{\theta}'$ such that in the optimal $B \in \mathcal{B}$ both b_1 and b_2 bind, but f never binds. In particular, the optimal B satisfies $b_1 = b_2$ and $c_i^p(\omega^i) < b_i < c_i^p(\omega^0)$ for $i = 1, 2$.*

Start from $\overline{\theta}$ that implies $c_i^p(\omega^1) > c_i^p(\omega^0)$ for $i = 1, 2$ and, hence, leads to the optimal B in Lemma 9. If we raise $\overline{\theta}$, $c_1^p(\omega^0)$ and $c_2^p(\omega^0)$ rise continuously while always satisfying $c_1^p(\omega^0) = c_2^p(\omega^0)$. There exists a unique $\overline{\theta}^\dagger$ such that when $\overline{\theta} = \overline{\theta}^\dagger$, $c_i^p(\omega^1) = c_i^p(\omega^0)$ for $i = 1, 2$. For $\overline{\theta} \le \overline{\theta}^\dagger$, the optimal B remains $b_1 = c_1^p(\omega^1)$, $b_2 = c_2^p(\omega^2)$, and $f = s^p(\omega^0)$, where the latter falls continuously as $\overline{\theta}$ rises toward $\overline{\theta}^\dagger$.

Now let $\mathcal{B}(\overline{\theta}) \subset \mathcal{B}$ be the set of optimal Bs as a function of $\overline{\theta}$. By the previous argument, $\mathcal{B}(\overline{\theta})$ is singleton for $\overline{\theta} \le \overline{\theta}^\dagger$. Define the distance between any B and B' as the Euclidean distance between (f, b_1, b_2) describing B and (f', b_1', b_2') describing B'. By the maximum theorem, $\mathcal{B}(\overline{\theta})$ is upper hemicontinuous in $\overline{\theta}$.[25] By choosing $\overline{\theta} > \overline{\theta}^\dagger$ sufficiently close to $\overline{\theta}^\dagger$, we can make the distance between $B(\overline{\theta}^\dagger)$ and every $B \in \mathcal{B}(\overline{\theta})$ arbitrarily small. There exists $\varepsilon > 0$ such that if $\overline{\theta} \in (\overline{\theta}^\dagger, \overline{\theta}^\dagger + \varepsilon)$ for every $B \in \mathcal{B}(\overline{\theta})$, we have (a) $b_i(\overline{\theta}) < c_i^d(\omega^i)$ for $i = 1, 2$ and (b) $f(\overline{\theta})$ cannot bind in either ω^1 or ω^2. To see (b), note

[25]Although self-0's and self-1's utility functions are not continuous at the boundary of $\mathbb{R}_+^3$ due to their logarithmic form, this is irrelevant because it is never optimal to choose B that forces 0 allocation to some dimension. Formally, there exists $\varepsilon > 0$ such that if we required $f \le 1 - \varepsilon$ and $b_i \ge \varepsilon$ for all $i = 1, 2$, we would never affect self-0's problem.

that $\mathcal{B}(\overline{\theta}^\dagger)$ contains the plan defined by $b_i(\overline{\theta}^\dagger) = c_i^p(\omega^i)$ for $i = 1, 2$ and $f(\overline{\theta}^\dagger) = s^p(\omega^0)$, where $f(\overline{\theta}^\dagger) = 1 - b_1(\overline{\theta}^\dagger) - b_2(\overline{\theta}^\dagger)$ and, hence, f is redundant. Thus, $\mathcal{B}(\overline{\theta})$ contains no plan with $f(\overline{\theta}) > 1 - b_1(\overline{\theta}) - b_2(\overline{\theta})$, as such plans are strictly dominated for the same argument that rules them out in the proof of Lemma 9. Since the largest value of $f(\overline{\theta})$ must be close to $f(\overline{\theta}^\dagger)$ for $\overline{\theta} \in (\overline{\theta}^\dagger, \overline{\theta}^\dagger + \varepsilon)$, $f(\overline{\theta})$ cannot bind in ω^1 and ω^2 as well.

Fix $\overline{\theta} \in (\overline{\theta}^\dagger, \overline{\theta}^\dagger + \varepsilon)$. Claims 6–9 characterize the properties of every $B \in \mathcal{B}(\overline{\theta})$.

Claim 6. For every $B \in \mathcal{B}(\overline{\theta})$, $b_1(\overline{\theta})$ and $b_2(\overline{\theta})$ must bind in ω^0; that is, $c_i(\omega^0) = b_i(\overline{\theta})$ for $i = 1, 2$. Given this, $s(\omega^0) = 1 - b_1(\overline{\theta}) - b_2(\overline{\theta})$ and, hence, f can be removed.

Proof. Note that self-0's objective in ω^i as a function of b_i is strictly concave and decreasing for $b_i > c_i^p(\omega^i)$ (see (10)). Thus, if, for example, $b_1(\overline{\theta})$ is not binding for self-1 in ω^0 (that is, $b_1(\overline{\theta}) > c_1(\omega^0)$), self-0 can lower b_1 without affecting self-1's choice in ω^0 and ω^2, and can strictly improve her payoff in ω^1. Hence, the initial plan would not be optimal. □

Claim 7. We have $b_1(\overline{\theta}) = b_2(\overline{\theta})$ for every $B \in \mathcal{B}(\overline{\theta})$.

Proof. Without loss, suppose $b_1(\overline{\theta}) > b_2(\overline{\theta})$. Note that $b_2(\overline{\theta}) < c_2^d(\omega^0)$ because, otherwise, we would have $b_1(\overline{\theta}) > c_1^d(\omega^0) = c_2^d(\omega^0)$, contradicting the previous claim. Consider the alternative with $b_1^\varepsilon = b_1(\overline{\theta}) - \varepsilon$ and $b_2^\varepsilon = b_2(\overline{\theta}) + \varepsilon$ for $\varepsilon > 0$. For ε small, b_1^ε and b_2^ε continue to bind in ω^0, so $1 - b_1^\varepsilon - b_2^\varepsilon = s(\omega^0)$. In ω^0, self-0's payoff is higher, because given $s(\omega^0)$, the chosen $\mathbf{c}$ is closer to being symmetric and, hence, to the best one according to self-0's preference. Due to symmetry and the strict concavity in self-0's payoff induced by b_i in ω^i for $i = 1, 2$ (see (10)), the decrease in the her payoff in ω^2 resulting from the slacker b_2 is more than compensated by the increase in ω^1 resulting from the tighter b_1. Hence, overall, self-0's payoff is strictly larger with $(b_1^\varepsilon, b_2^\varepsilon)$ than with $(b_1(\overline{\theta}), b_2(\overline{\theta}))$, contradicting the optimality of the latter plan. □

Claim 8. We have $1 - b_1(\overline{\theta}) - b_2(\overline{\theta}) > s^p(\omega^0)$ for every $B \in \mathcal{B}(\overline{\theta})$.

Proof. If $1 - b_1(\overline{\theta}) - b_2(\overline{\theta}) < s^p(\omega^0)$, self-0 can set $f = s^p(\omega^0)$ and get a strictly higher payoff in ω^0 without affecting self-1's choices in ω^1 and ω^2. If $1 - b_1(\overline{\theta}) - b_2(\overline{\theta}) = s^p(\omega^0)$, then $b_i = c_i^p(\omega^0)$ for $i = 1, 2$ and so $(\mathbf{c}(\omega^0), s(\omega^0)) = (\mathbf{c}^p(\omega^0), s^p(\omega^0))$. Therefore, it would be possible to lower both $b_1(\overline{\theta})$ and $b_2(\overline{\theta})$ by the same small ε, so as to induce a first-order gain in self-0's payoff in ω^1 and ω^2 because $b_i(\overline{\theta}) > c_i^p(\omega^i)$ for $i = 1, 2$, while causing only a second-order loss in ω^0. □

Claim 9. All $B \in \mathcal{B}(\overline{\theta})$ have the same b_1 and b_2 and satisfy the properties in Lemma 10.

Proof. Let $b_1 = b_2 = b$. Self-0's payoff in ω^1 and ω^2 is given by (10) up to a constant: $\underline{\theta}[\overline{r}\ln(b) + \underline{r}\ln(1-b)] + \ln(1-b)$. Her payoff in ω^0 is given, up to a constant, by $2\underline{\theta}\overline{r}\ln(b) + \ln(1-2b)$. Therefore, self-0's expected payoff from b is strictly concave. To see that $c_i^p(\omega^i) > b_i > c_i^p(\omega^0)$ for $i = 1, 2$, consider the following situation. A $b_i > c_i^p(\omega^i)$ is strictly dominated by $b_i = c_i^p(\omega^i)$ for every i, as this is the optimal level of b_i in ω^i. Consequently, we must have $b_i < c_i^p(\omega^i)$, because, by assumption, $1 - c_1^p(\omega^1) - c_2^p(\omega^2) > s^p(\omega^0)$ for $\overline{\theta} > \overline{\theta}^\dagger$, and so reducing b_i below $c_i^p(\omega^i)$ by the same small amount for $i = 1, 2$ causes a first-order gain in ω^0 and a second-order loss in ω^1 and ω^2. □

Part III Let G^{fb} be a distribution over $(\omega^0, \omega^1, \omega^2)$ that leads to Lemma 9 and let $\overline{G}$ be the uniform over $[\underline{\theta}, \overline{\theta}] \times [\underline{r}, \overline{r}]^2$. Let G^{b} be a distribution that leads to Lemma 10 and let $\overline{G}'$ be the uniform over $[\underline{\theta}, \overline{\theta}'] \times [\underline{r}, \overline{r}]^2$, where $\overline{\theta}'$ is as in Lemma 10. For $\alpha \in [0, 1]$, let $G_\alpha^{\text{fb}} = \alpha G^{\text{fb}} + (1 - \alpha)\overline{G}$ and $G_\alpha^{\text{b}} = \alpha G^{\text{b}} + (1 - \alpha)\overline{G}'$. Proposition 4 follows from the next result.

Corollary 2. *(i) There exists $\alpha \in (0, 1)$ such that, given G_α^{fb}, f, b_1, and b_2 are all binding for every optimal B. (ii) There exists $\alpha' \in (0, 1)$ such that, given $G_{\alpha'}^{\text{b}}$, for every optimal B, both b_1 and b_2 bind, but f never binds.*

Proof. Let $\mathcal{B}_f \subset \mathcal{B}$ contain all Bs that can use only f, let $\mathcal{B}_{\mathbf{b}}$ contain all Bs that can use both b_1 and b_2, and let $\mathcal{B}_{b_i}$ contain all Bs that can use only b_i for $i = 1, 2$. To indicate that self-0's expected payoff is computed using some distribution $\hat{G}$, we will use the notation $\mathcal{U}(B; \hat{G})$.

Part (i). For every B and $\alpha \in [0, 1]$, $\mathcal{U}(B; G_\alpha^{\text{fb}}) = \alpha\mathcal{U}(B; G^{\text{fb}}) + (1 - \alpha)\mathcal{U}(B; \overline{G})$. Now define

$$\mathcal{W}_f^{\text{fb}}(\alpha) = \max_{B \in \mathcal{B}_f} \mathcal{U}\big(B; G_\alpha^{\text{fb}}\big) \quad \text{and} \quad \mathcal{W}_{\mathbf{b}}^{\text{fb}}(\alpha) = \max_{B \in \mathcal{B}_{\mathbf{b}}} \mathcal{U}\big(B; G_\alpha^{\text{fb}}\big), \quad \alpha \in [0, 1].$$

Both $\mathcal{W}_f^{\text{fb}}$ and $\mathcal{W}_{\mathbf{b}}^{\text{fb}}$ are well defined by the same argument of Lemma 2 and continuous functions of α by the maximum theorem.[26] Let B^{fb} denote the optimal plan in Lemma 9. Note that $\mathcal{U}(B^{\text{fb}}; \overline{G})$ is finite since self-1's resulting choices are bounded away from 0 in all dimensions. We have that $\lim_{\alpha\uparrow 1} \mathcal{U}(B^{\text{fb}}; G_\alpha^{\text{fb}}) - \mathcal{W}_j^{\text{fb}}(\alpha) > 0$ for both $j = f$ and $j = \mathbf{b}$. Therefore, there exists $\hat{\alpha} \in (0, 1)$ such that B^{fb} strictly dominates every $B \in \mathcal{B}_f \cup \mathcal{B}_{\mathbf{b}}$ given $G_{\hat{\alpha}}^{\text{fb}}$.

Part (ii). For every B and $\alpha \in [0, 1]$, $\mathcal{U}(B; G_\alpha^{\text{b}}) = \alpha\mathcal{U}(B; G^{\text{b}}) + (1 - \alpha)\mathcal{U}(B; \overline{G}')$. Let B^{b} denote the optimal plan in Lemma 10. By the same logic in the proof of part (i), there exists $\alpha'' \in (0, 1)$ such that, for $\alpha \in (\alpha'', 1)$, B^{b} strictly dominates every $B \in \mathcal{B}_f \cup \mathcal{B}_{\mathbf{b}} \cup \mathcal{B}_{b_1} \cup \mathcal{B}_{b_2}$ given G_α^{b}. It remains to show that there exists $\alpha' \in (\alpha'', 1)$ such that B^{b} strictly dominates every $B \in \mathcal{B}$ given $G_{\alpha'}^{\text{b}}$. To this end, define $\mathcal{B}(\alpha) = \arg\max_{B \in \mathcal{B}} \mathcal{U}(B; G_\alpha^{\text{b}})$. Set $\mathcal{B}(\cdot)$ is upper hemicontinuous by the maximum theorem. Note that $\mathcal{B}(1)$ is characterized by (f^*, b_1^*, b_2^*) such that b_1^* and b_2^* are unique and satisfy Lemma 10, and $f^* \in [0, \overline{f}]$, where $\overline{f} = 1 - b_1^* - b_2^*$. For every $\eta > 0$, there exists $\varepsilon > 0$ such that if $\alpha \in (1 - \varepsilon, 1]$, then $f \in [0, \overline{f} + \eta]$, $b_1 \in (b_1^* - \eta, b_1^* + \eta)$, and $b_2 \in (b_2^* - \eta, b_2^* + \eta)$ for every (f, b_1, b_2) corresponding to some $B \in \mathcal{B}(\alpha)$. Choosing η small enough ensures that for all $B \in \mathcal{B}(\alpha)$, we have that (a) $1 - b_1 - b_2 > s^p(\omega^0)$, (b) removing f leads to a plan such that b_1 and b_2

[26] Recall footnote 25.

bind in ω^0, and (c) f cannot bind in ω^1 and ω^2, as self-1's choice of s strictly exceeds $\overline{f}$ in ω^1 and ω^2 under all $B \in \mathcal{B}(1)$.

Take any $B \in \mathcal{B}(\alpha)$ and fix its b_1 and b_2. The f completing B must be optimally chosen given b_1 and b_2. We claim that it must satisfy $f \leq 1 - b_1 - b_2 = \overline{k}$ for α sufficiently close to 1. Suppose not. Consider self-0's expected gain from imposing $f > \overline{k}$. The gain in ω^0 is

$$(1-\beta)\big[v(f) - v(\overline{k})\big] + \overline{V}\big(f; \omega^0\big) - \overline{V}\big(\overline{k}; \omega^0\big),$$

and the expected gain under the distribution $\overline{G}'$ is

$$\int_{\Omega(f)} \big\{(1-\beta)\big[v(f) - v\big(\hat{s}(\omega)\big)\big] + \overline{V}(f; \omega) - \overline{V}\big(\hat{s}(\omega); \boldsymbol{\omega}\big)\big\}\, d\overline{G}', \tag{13}$$

where $\Omega(f) \subset \Omega' = [\underline{\theta}, \overline{\theta}'] \times [\underline{r}, \overline{r}]^2$ is the set of states in which f affects self-1's choices, $(\hat{\mathbf{c}}, \hat{s})$ is self-1's allocation function under the policy using only b_1 and b_2, and

$$\overline{V}(k; \omega) = \max_{\{(\mathbf{c},s)\in B: c_1 \leq b_1, c_2 \leq b_2, s \geq k\}} \big\{\theta u(\mathbf{c}; \mathbf{r}) + \beta v(s)\big\}, \quad k \in [\overline{k}, 1], \omega \in \Omega'.$$

Since $\overline{V}(f; \omega) \leq \overline{V}(\hat{s}(\omega); \omega)$ and $\hat{s}(\omega) \geq \overline{k}$ for all $\omega \in \Omega'$, for all $f \geq \overline{k}$, (13) is bounded above by

$$\int_{\Omega(f)} (1-\beta)\big[v(f) - v\big(\hat{s}(\omega)\big)\big]\, dG' \leq (1-\beta)\big[v(f) - v(\overline{k})\big].$$

Note that the right-hand side of this expression depends on α only via $\overline{k}$.

Now focus on $\overline{V}(k; \omega^0)$. For every $f > \overline{k}$, (a) f always binds, because $\overline{k} > s^p(\omega^0)$ and, hence, self-1 wants to choose $s < f$, (b) only one b_i can bind, because if both bind, then $s(\omega^0) = \overline{k} < f$, which is impossible, and (c) one b_i never binds, because goods are normal, so for all $f > \overline{k}$, self-1's chooses $c_i(\omega^0) < b_i$ for at least one $i = 1, 2$. Without loss, suppose that b_2 never binds. If we remove b_2, $\overline{V}(k; \omega^0)$ equals self-1's indirect utility under the plan defined by $k \in [\overline{k}, 1]$ and b_1 only, denoted by $\overline{V}(k; \boldsymbol{\omega}^0, b_1)$. By the same argument in Lemma 4, $\overline{V}(k; \omega^0, b_1)$ is continuously differentiable in k for $k \in (0, 1]$ and $\overline{V}'(k; \boldsymbol{\omega}^0, b_1) = -\lambda(\omega^0; k)$, where $\lambda(\omega^0; k)$ is the Lagrange multiplier associated to $s \geq k$. Using the Lagrangian that defines $\overline{V}(k; \omega^0, b_1)$, we have $\lambda(\omega^0; k) = \overline{\theta}' u_c^2(c_2(\omega^0; k); \overline{\mathbf{r}}) - \beta v'(k)$. Note that $\lambda(\omega^0; k) > 0$ for all $k \in [\overline{k}, 1]$, because such floor levels must always bind for self-1. Moreover, $\lambda(\omega^0; k)$ is strictly increasing in $k \in [\overline{k}, 1]$ because v is strictly concave, $u_{cc}^i < 0$, and $c_2(\omega^0; k)$ is nonincreasing in k by normality of goods. Therefore, $\overline{V}'(k; \omega^0) = -\lambda(\omega^0; k)$ for all $k \in (\overline{k}, 1]$ and $\overline{V}'(\overline{k}+; \omega^0) = -\lambda(\omega^0; \overline{k})$, where the plus denotes the right derivative.[27] Moreover, $\overline{V}'(k; \boldsymbol{\omega}^0)$ is strictly decreasing in k.

Observe that

$$(1-\beta)v'(\overline{k}) + \overline{V}'\big(\overline{k}; \omega^0\big) = v'(\overline{k}) - \overline{\theta}' u_c^2\big(c_2\big(\omega^0; \overline{k}\big); \overline{\mathbf{r}}\big), \tag{14}$$

[27]The function $\overline{V}(k; \omega^0)$ is not differentiable at $k = \overline{k}$, as $\overline{V}(k; \omega^0)$ is constant for $k < \overline{k}$ and $\overline{V}'(\overline{k}-; \omega^0) = 0$.

which is strictly negative. This is because $b_1 < c_1^p(\omega^0)$ and $b_2 < c_2^p(\omega^0)$ by Lemma 10 since α is close to 1, which implies that b_1 and b_2 must bind for self-0; consequently, $f = \overline{k}$ and b_1 must also bind for self-0. The right-hand side of (14) coincides with the negative of the Lagrange multiplier associated with $s \geq \overline{k}$ in self-0's problem that also includes $c_1 \leq b_1$.

Recall that $\overline{k}$ depends on α—hence, denote it by $\overline{k}_\alpha$—and consider

$$g\overline{V}'\big(\overline{k}_\alpha; \omega^0\big) + \big[\alpha g + (1-\alpha)\big](1-\beta)v'(\overline{k}_\alpha). \tag{15}$$

This is strictly negative for $\alpha = 1$, in which case $\overline{k}_1 = 1 - b_1^* - b_2^*$. By continuity of (15) in (α, k) and upper hemicontinuity of $\mathcal{B}(\alpha)$, there exists $\varepsilon > 0$ such that (15) is strictly negative for $\alpha \in (1-\varepsilon, 1]$. By the monotonicity of v' and $\overline{V}'(\cdot; \omega^0)$, (15) is strictly decreasing for $k \geq \overline{k}_\alpha$.

Finally, for every $\alpha \in (1-\varepsilon, 1]$ and $f > \overline{k}_\alpha$,

$$\begin{aligned}
&\big[\alpha g + (1-\alpha)\big](1-\beta)\big[v(f) - v(\overline{k}_\alpha)\big] + g\big[\overline{V}\big(f; \omega^0\big) - \overline{V}\big(\overline{k}_\alpha; \omega^0\big)\big] \\
&\quad = \int_{\overline{k}_\alpha}^{f} \big\{\big[\alpha g + (1-\alpha)\big](1-\beta)v'(k) + g\overline{V}'\big(k; \boldsymbol{\omega}^0\big)\big\}\, dk \\
&\quad < \big\{\big[\alpha g + (1-\alpha)\big](1-\beta)v'(\overline{k}_\alpha) + g\overline{V}'\big(\overline{k}_\alpha; \omega^0\big)\big\}(f - \overline{k}_\alpha) < 0.
\end{aligned}$$

We conclude that self-0 is strictly worse off by imposing a binding f in addition to b_1 and b_2, and, hence, every optimal B must use binding budgets for both goods, but no binding f. □

A.6 *Example of non-additive utility*

Suppose that

$$u(\mathbf{c}; \mathbf{r}) = \frac{1}{1-\gamma}\big(r_1 c_1^{\frac{e-1}{e}} + r_2 c_2^{\frac{e-1}{e}}\big)^{\frac{e}{e-1}(1-\gamma)} \quad \text{and} \quad v(s) = \frac{s^{1-\gamma}}{1-\gamma}.$$

Assume that $e > 1$, $0 < \gamma < 1$, and $e \leq 1/\gamma$. By standard calculations, given total expenditures $y \in [0, 1]$ in a period, the optimal allocation to good i is

$$c_i(\mathbf{r}; y) = y\frac{r_i^e}{r_1^e + r_2^e}. \tag{16}$$

We now show that $(\mathbf{c}^d, s^d)$ satisfies Condition 1; similar steps establish the desired properties of $(\mathbf{c}^p, s^p)$. For every $\omega \in \Omega$, maximizing

$$\frac{\theta\tau(\mathbf{r})}{1-\gamma}(1-s)^{1-\gamma} + \frac{\beta}{1-\gamma}s^{1-\gamma}$$

yields

$$s^d(\omega) = \frac{\beta^{\frac{1}{\gamma}}}{\big[\theta\tau(\mathbf{r})\big]^{\frac{1}{\gamma}} + \beta^{\frac{1}{\gamma}}}.$$

Clearly, $s^d(\mathbf{s})$ is always interior and strictly decreasing in θ, r_1, and r_2. Replacing y with $1 - s^d(\omega)$ into (16), we get that $c_i^d(\omega)$ is strictly increasing in θ and r_i, and satisfies

$$\frac{\partial}{\partial r_j} c_i^d(\omega) \propto s^d(\omega) \frac{1-\gamma}{\gamma(e-1)} - 1.$$

Thus, for $\partial c_i^d(\omega)/\partial r_j$ (and similarly $\partial c_i^p(\omega)/\partial r_j$) to be strictly negative for all ω, a sufficient condition is that

$$\frac{1-\gamma}{\gamma(e-1)} < \frac{1}{\bar{s}^p},$$

because $s^d(\omega) < \bar{s}^d < \bar{s}^p < 1$.[28]

Now consider setting only a budget on good 1 (or, equivalently, on good 2). We will show that whenever b_1 binds, increasing it leads to lower savings and that part (ii) of Condition 1 holds. Suppose b_1 binds in state ω. Then self-1's choice satisfies $c_2^*(\omega) = 1 - s^*(\omega) - b_1$ and the optimal $s^*(\omega)$ solves the first-order condition

$$\theta u_2\big(b_1, 1 - s^*(\omega) - b_1; \mathbf{r}\big) = \beta\big[s^*(\omega)\big]^{-\gamma}.$$

Therefore,

$$\frac{\partial}{\partial b_1} s^*(\omega) = \left.\frac{\theta\big[u_{21}(\mathbf{c}; \mathbf{r}) - u_{22}(\mathbf{c}; \mathbf{r})\big]}{\theta u_{22}(\mathbf{c}; \mathbf{r}) - \dfrac{\gamma\beta}{s^{1+\gamma}}}\right|_{(\mathbf{c},s)=(\mathbf{c}^*(\omega),s^*(\omega))}.$$

A sufficient condition for this to be strictly negative is that $u_{21}(\mathbf{c}; \mathbf{r}) \geq 0$, which holds under our assumptions since $u_{21}(\mathbf{c}; \mathbf{r}) \propto 1 - \gamma e$. Finally, since self-1's allocation to s and c_2 is bounded away from zero for every $b_1 \leq \bar{c}_1^d$, and since $u_{22}(\mathbf{c}; \mathbf{r})$ and $u_{21}(\mathbf{c}; \mathbf{r})$ are continuous in both arguments, it follows that $\partial s^*/\partial b_1$ is uniformly bounded away from zero. Therefore, part (ii) of Condition 1 holds. Note that in the previous argument we can replace b_1 with f_1. Therefore, in this example, binding good-specific floors lead to lower savings, and, hence, they are never part of optimal plans.

References

Aguiar, Mark and Manuel Amador (2011), "Growth in the shadow of expropriation." *Quarterly Journal of Economics*, 126, 651–697. [179]

Amador, Manuel and Kyle Bagwell (2013), "The theory of optimal delegation with an application to tariff caps." *Econometrica*, 81, 1541–1599. [176, 182]

Amador, Manuel and Kyle Bagwell (2016), "Money burning in the theory of delegation." Unpublished paper, University of Minnesota. [182]

Amador, Manuel, Ivan Werning, and George-Marios Angeletos (2006), "Commitment vs. flexibility." *Econometrica*, 74, 365–396. [174, 175, 176, 178, 181, 182, 183, 185]

[28] Recall that $\bar{s}^k = \max_\omega s^k(\omega)$ for $k = d, p$.

Ambrus, Attila and Georgy Egorov (2013), "Comment on "Commitment vs. flexibility"." *Econometrica*, 81, 2113–2124. [176, 182]

Ambrus, Attila and Georgy Egorov (2017), "Delegation and nonmonetary incentives." *Journal of Economic Theory*, 171, 101–135. [182]

Ameriks, John, Andrew Caplin, and John Leahy (2003), "Wealth accumulation and the propensity to plan." *Quarterly Journal of Economics*, 118, 1007–1047. [173, 174]

Antonides, Gerrit, I. Manon de Groot, and W. Fred van Raaij (2011), "Mental budgeting and the management of household finance." *Journal of Economic Psychology*, 32, 546–555. [173, 174, 188]

Athey, Susan, Andrew Atkeson, and Patrick J. Kehoe (2005), "The optimal degree of discretion in monetary policy." *Econometrica*, 73, 1431–1475. [176]

Bakke, Edward Wight (1940), *The Unemployed Worker; Citizens Without Work*. Yale University Press. [173]

Battaglini, Marco (2005), "Long-term contracting with Markovian consumers." *American Economic Review*, 95, 637–658. [177]

Baumeister, Roy F., Todd F. Heatherton, and Dianne M. Tice (1994), *Losing Control: How and Why People Fail at Self-Regulation*. Academic Press, San Diego, California. [173]

Bénabou, Roland and Marek Pycia (2002), "Dynamic inconsistency and self-control: A planner–doer interpretation." *Economics Letters*, 77, 419–424. [174]

Bénabou, Roland and Jean Tirole (2004), "Willpower and personal rules." *Journal of Political Economy*, 112, 848–886. [173, 179]

Benhabib, Jess and Alberto Bisin (2005), "Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption-saving decisions." *Games and Economic Behavior*, 52, 460–492. [174, 179]

Bernheim, B. Douglas and Antonio Rangel (2004), "Addiction and cue-triggered decision processes." *American Economic Review*, 94, 1558–1590. [174]

Bernheim, B. Douglas, Debraj Ray, and Şevin Yeltekin (2015), "Poverty and self-control." *Econometrica*, 83, 1877–1911. [179]

Beshears, John, Katherine L. Milkman, and Joshua Schwartzstein (2016), "Beyond beta-delta: The emerging economics of personal plans." *American Economic Review*, 106, 430–434. [173]

Brocas, Isabelle and Juan D. Carrillo (2008), "The brain as a hierarchical organization." *American Economic Review*, 98, 1312–1346. [174, 176]

Bryan, Gharad, Dean Karlan, and Scott Nelson (2010), "Commitment devices." *Annual Review of Economics*, 2, 671–698. [174]

Chatterjee, Kalyan and R. Vijay Krishna (2009), "A "dual self" representation for stochastic temptation." *American Economic Journal: Microeconomics*, 1, 148–167. [174]

Courty, Pascal and Li Hao (2000), "Sequential screening." *Review of Economic Studies*, 67, 697–717. [177]

Ellis, Christopher J. and Barry J. Naughton (1990), "On the theory of household saving in the presence of rationing." *Journal of Comparative Economics*, 14, 269–285. [177, 189]

Festinger, Leon (1962), *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford, California. [179]

Fudenberg, Drew and David K. Levine (2006), "A dual-self model of impulse control." *American Economic Review*, 96, 1449–1476. [174]

Fudenberg, Drew and David K. Levine (2012), "Timing and self-control." *Econometrica*, 80, 1–42. [174]

Galperti, Simone (2015), "Commitment, flexibility, and optimal screening of time inconsistency." *Econometrica*, 83, 1425–1465. [176]

Halac, Marina and Pierre Yared (2014), "Fiscal rules and discretion under persistent shocks." *Econometrica*, 82, 1557–1614. [176, 177, 179]

Heath, Chip, Richard P. Larrick, and George Wu (1999), "Goals as reference points." *Cognitive Psychology*, 38, 79–109. [179]

Heath, Chip and Jack B. Soll (1996), "Mental budgeting and consumer decisions." *Journal of Consumer Research*, 23, 40–52. [173, 174, 187]

Henderson, Pamela W. and Robert A. Peterson (1992), "Mental accounting and categorization." *Organizational Behavior and Human Decision Processes*, 51, 92–117. [173]

Howard, David H. (1977), "Rationing, quantity constraints, and consumption theory." *Econometrica*, 45, 399–412. [177, 188]

Hsiaw, Alice (2013), "Goal-setting and self-control." *Journal of Economic Theory*, 148, 601–626. [179]

Jackson, Matthew O. and Leeat Yariv (2015), "Collective dynamic choice: The necessity of time inconsistency." *American Economic Journal: Microeconomics*, 7, 150–178. [179]

Johnson, Michael D. (1984), "Consumer choice strategies for comparing noncomparable alternatives." *Journal of Consumer Research*, 11, 741–753. [187]

Koch, Alexander K. and Julia Nafziger (2016), "Goals and bracketing under mental accounting." *Journal of Economic Theory*, 162, 305–351. [176]

Laibson, David (1997), "Golden eggs and hyperbolic discounting." *Quarterly Journal of Economics*, 112, 443–478. [176, 179]

Lee, Rainwater, Richard Patrick Coleman, and Gerald Handel (1962), *Workingman's Wife: Her Personality, World and Life Style*. MacFadden-Bartell, New York, New York. [173]

Loewenstein, George and Ted O'Donoghue (2007), "The heat of the moment: Modeling interactions between affect and deliberation." Unpublished paper, Carnegie Mellon University. [174]

Luenberger, David G. (1969), *Optimization by Vector Space Methods*. Wiley, New York, New York. [176, 189, 191, 196]

Madden, Paul (1991), “A generalization of Hicksian q substitutes and complements with application to demand rationing.” *Econometrica*, 59, 1497–1508. [177, 188, 189]

Milgrom, Paul and Chris Shannon (1994), “Monotone comparative statics.” *Econometrica*, 62, 157–180. [195]

Nageeb, Ali S. (2011), “Learning self-control.” *Quarterly Journal of Economics*, 126, 857–893. [174]

Pavan, Alessandro, Ilya Segal, and Juuso Toikka (2014), “Dynamic mechanism design: A Myersonian approach.” *Econometrica*, 82, 601–653. [177]

Quah, John K-H. (2007), “The comparative statics of constrained optimization problems.” *Econometrica*, 75, 401–431. [180]

Rochet, Jean-Charles and Philippe Choné (1998), “Ironing, sweeping, and multidimensional screening.” *Econometrica*, 66, 783–826. [189]

Rochet, Jean-Charles and Lars A. Stole (2003), “The economics of multidimensional screening.” *Econometric Society Monographs*, 35, 150–197. [189]

Rockafellar, Tyrrell R. (1997), *Convex Analysis*. Princeton University Press. [190]

Simon, Herbert A. (1997), *Administrative Behavior*, fourth edition. The Free Press, New York, New York. [187]

Strotz, Robert H. (1957), “The empirical implications of a utility tree.” *Econometrica*, 25, 269–280. [176]

Thaler, Richard (1985), “Mental accounting and consumer choice.” *Marketing Science*, 4, 177–266. [173, 176, 187]

Thaler, Richard H. (1999), “Mental accounting matters.” *Journal of Behavioral Decision Making*, 12, 183–206. [173, 174, 187]

Thaler, Richard H. and Hersh M. Shefrin (1981), “An economic theory of self-control.” *Journal of Political Economy*, 89, 392–406. [173, 174, 175, 176]

Wertenbroch, Klaus (2002), “Self-rationing: Self-control in consumer choice.” Unpublished paper, INSEAD Working Paper No. 2001/63/MKT, SSRN 296954. [173, 174]

Zelizer, Viviana A. Rotman (1997), *The Social Meaning of Money*. Princeton University Press. [173]

Co-editor Thomas Mariotti handled this manuscript.

Manuscript received 16 May, 2017; final version accepted 24 April, 2018; available online 1 May, 2018.