A Theory of Personal Budgeting

Simone Galperti*

UCSD

October 7, 2016

Abstract

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 which renders budgets useful commitments, their interaction with minimum-savings rules (another widely-studied commitment technique), and how budgeting depends on the intensity of self-control problems. This theory matches a number of empirical findings and can guide marketing personal budgeting devices.

JEL CLASSIFICATION: D23, D82, D86, D91, E62, G31

KEYWORDS: budget, minimum-savings rule, commitment, flexibility, intratemporal trade-off, uncertainty, present bias.

*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, and Joel Watson for useful comments and suggestions. I also thank seminar participants at Boston College, Columbia University, LSE, Oxford University, Penn State University, University of Pennsylvania, University College London, UC 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. First version: February 2015.

1 Introduction

Many studies argue that personal budgeting is a pervasive part of consumer behavior.[1] Personal budgeting denotes the practice of grouping expenditures into categories and constraining each with an implicit or explicit spending limit that applies to a specified time period (a week, a month, etc.).[2] This practice cannot be explained by the classic life-cycle theory of the consumer. Nonetheless, it has important consequences. It can account for "mysterious" large differences in wealth accumulation between consumers, which cannot be explained by time or risk preferences (Ameriks et al. (2003)). By violating the principle of fungibility of money, it shapes consumer demand in ways which cannot be explained by satiation and income effects (Heath and Soll (1996)). It affects how firms promote their products so as to avoid falling into the same category with other firms and thus compete for the same budget (Wertenbroch (2002)). It is at the foundation of the economics of commitment devices (Bryan et al. (2010)) by creating a demand for personal budgeting services.

Almost all existing studies informally suggest that consumers use personal budgets to manage self-control problems, often caused by present bias, which interfere with their saving goals. Thaler (1999) argues that households group expenditures into category-specific budgets (housing, food, etc.) "to keep spending under control." According to Ameriks et al. (2003), "many households that set up regular budgets regard this activity as contributing to a reduction in their spending. These results support a theory in which the channel connecting wealth accumulation and the propensity to plan operates through a form of effortful self-control." Antonides et al. (2011) find a positive correlation between budgeting and having savings goals.

Despite this consensus on the existence of a link between budgeting and self-control problems, a formal investigation of such a link seems to be missing. The paper fills this gap and offers a solid foundation for personal budgeting in a precise aspect of time preferences: present bias. It shows, however, that present bias alone is not enough to explain personal budgeting. Present bias induces consumers to value commitment in the form of constraints on future choices. But for personal budgeting to emerge, this preference for commitment has to be combined with a preference for flexibility of a specific but plausible kind, namely, that resulting from uncertainty about *intra*temporal trade-offs—for instance, due to shocks in the taste for or the price of some goods. Moreover, the paper uncovers potential tensions between good-specific budgets and minimum-savings rules, another commitment technique often studied in the literature. In turn, this leads to a *negative* relationship between the intensity of present bias and the use of good-specific budgets. These novel predictions help organize the existing evidence on personal

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

[2] This paper uses the term "personal budgeting" rather than "mental accounting" because the latter has a much broader meaning, indicating a general process by which people frame and label events, outcomes, and decisions. As such, mental accounting includes phenomena like choice bracketing, narrow framing, and gain-loss utility, which differ from budgeting.

budgeting and can guide future empirical studies.

The paper obtains these results using a simple planner-doer model of consumption and savings.[3] As usual, the planner and the doer have time-additive utility functions with common per-period consumption utility; while the planner has time consistent preferences, the doer is present biased (though not fully myopic). In each period, the doer chooses how much to consume and save subject to the usual income constraint. The paper modifies this standard setup in two minimal but crucial ways. First, consumption involves a bundle of multiple goods, rather than a single uniform commodity. Second, the optimal income allocation depends on a state of the world (capturing taste or price shocks) which affects not only the rate of substitution between present and future utility, but also the rates of substitution between goods within periods. This is key to introduce the uncertainty about intratemporal trade-offs driving the results. The distribution of the state components is allowed to be general within periods, but satisfies independence across periods. At the beginning of each period, the planner can adopt a commitment plan that dictates which income allocations the doer is allowed to choose. In line with its motivation, the paper focuses on plans that can freely combine good-specific budgets as well as an overall limit on consumption expenditures via a savings floor. A trade-off between commitment and flexibility is introduced by assuming that only the doer observes the state.[4]

The core of this paper characterizes how the planner optimally uses good-specific budgets and savings floors.[5] Since the doer tends to overspend but always agrees with the planner on how to divide every dollar across goods—recall that they have the same per-period consumption utility—an intuitive conjecture is that the planner wants to set only a limit on total consumption expenditures, that is, a savings floor. However, this is not true. Under some conditions, the planner can benefit from setting binding budgets. It is tempting to say that this is simply because budgets offer additional tools to commit and they automatically increase savings. But things are significantly more subtle, as budgets work differently than does the savings floor. A binding floor distorts the division of income between spending and saving, but never distorts the chosen consumption bundle—in the usual sense that marginal rates of substitutions equal price ratios. By contrast, a binding budget also distorts such bundle. In particular, it can exacerbate the doer's overspending for some other goods, which may even reduce savings.

Nonetheless, in some cases the consumption distortions caused by budgets can result in higher savings, precisely because they curtail how much the doer can gain in terms of

[3]Thaler and Shefrin (1981) were the first to propose a dual-self approach to study self-control problems. Other papers on dual-self models include Benabou 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 Ali (2011).

[4]This paper takes the process of noticing an expense and reporting it to the corresponding budget as a defining aspect of budgeting itself. To focus on the issues of interest here, it also assumes that people stick to their plans. This is not a trivial assumption, of course, and the paper motivates it accordingly in Section 3.

[5]One may ask what properties characterize the optimal commitment plans if arbitrarily general plans are allowed. This is an important, but very challenging, question in the presence of multidimensional consumption and uncertainty, as explained in Section 6.2.

present utility by undersaving; and this key mechanism can dominate those distortions from the planner's viewpoint. This is true, as we will see, if the goods satisfy appropriate substitutability and normality conditions. In this case, optimal plans always involve good-specific budgets when present bias is *weak*, but only a savings floor when present bias is *strong*. Importantly, the optimality of such budgets crucially depends on the uncertainty about intratemporal trade-offs. If we remove it, then under weak conditions the optimal plans involve only a savings floor, so the extra tools offered by budgets are in fact useless.

The intuition for these results is as follows. Suppose that there are only two goods and that ex ante the individual is uncertain about his ex-post marginal utility of each good, which can be high or low. He then realizes that a savings floor will help him limit overconsumption if both marginal utilities will be high—which make him want to consume a lot of both goods—but may be ineffective if only one marginal utility turns out to be high. The latter is especially likely if he sets a relatively permissive floor and present bias is weak. In this case, however, the individual realizes that the good with high marginal utility will be the main determinant of his excessive spending; hence, capping that good with a budget can increase savings because overspending will have to occur inefficiently and on a good with low marginal utility. On the other hand, when present bias is stronger, undersaving becomes more severe, but also less responsive to the consumption distortions caused by budgets. As a result, ex ante the individual prefers to set only a savings floor, for it limits undersaving without distorting consumption.

Several aspects of this paper's results go against what one may expect at fist glance. One may think that if a present-biased individual adds good-specific spending limits to an aggregate limit, then *a fortiori* an individual with a stronger bias should do the same; but this turns out to be false. One may think that optimal plans always feature a savings floor, but we will see that in some settings they rely exclusively on good-specific budgets. One may think that the results are driven by an underlying substitutability between savings floor and good-specific budgets, but their interaction is more intricate. For instance, people who use budgets may also set tighter savings floors, as the budgets' distortions lower the value of leaving more income for consumption. This also means that, once people are allowed to use good-specific budgets, it is not possible to conclude that a stronger present bias always results in a stricter savings floor.

Other theories may give rise to personal budgeting, but struggle to account for some related evidence. One theory may argue that people set budgets for those goods which they find tempting ("vice goods"). This can be true in some cases, but cannot explain why people set budgets for "unobjectionable goods like sports tickets and blue jeans" (as found by Heath and Soll (1996)) or housing, food, and even charitable giving (as reported by Thaler (1985, 1999)). Moreover, this theory runs the risk of assuming the answer: It can always explain that an individual sets a budget for any good X by assuming that he finds X tempting, which is ultimately a subjective aspect of his preference. Another plausible theory is that some people use budgeting as a technique to simplify the complex matter of household finance (Simon (1965), Johnson (1984)). This theory is complementary to the one in this paper, but it again struggles to explain some evidence. For instance, it is not clear why computational complexity would lead people to systematically set budgets which seem too strict and to cause underconsumption, as found by Heath and Soll (1996).

By contrast, these findings are consistent with the theory proposed here. Present bias can lead people to set budgets on "unobjectionable goods," because doing so helps them manage that bias, and to optimally choose the budgets' levels so that they systematically bind, which makes them appear too strict and cause underconsumption. This also provides an argument against Heath and Soll's (1996) conjecture that individuals may benefit by allowing themselves to reallocate more freely. Another prediction which speaks to the evidence on personal budgeting is that budgets should be used only by individuals who have a *weak* present bias.[6] Antonides et al. (2011) find that people who exhibit a "short-term time orientation" (which according to their description corresponds to strong present bias) are *less* likely to use budgets than people who exhibit a "long-term time orientation" (a weak present bias). More generally, the paper suggests caution when interpreting evidence on personal budgeting; for instance, observing that a person uses rich commitment plans involving budgets need not suffice to conclude that he is more present biased than other people who do not use such plans. The result that present bias has to be weak to give rise to budgeting is also important for how we model individuals with self-control problems: It offers another, empirically supported, reason for modeling short-run selves as not completely myopic, as advocated by Fudenberg and Levine (2012).

This paper contributes to our understanding of consumption-savings behavior in the presence of self-control problems and of the resulting demand for commitment. According to Bryan et al. (2010) the latter deserves more work especially on "soft commitments" (like personal budgeting). Since Thaler and Shefrin's (1981) and Laibson's (1997) seminal work, the literature[7] has developed almost entirely assuming a per-period consumption utility of a single commodity, "money," which can be interpreted as an indirect utility function.[8] This paper not only shows that that assumption is not innocuous with present-biased consumers (in contrast to settings with time-consistent consumers), but also demonstrates that what renders the multiplicity of goods relevant is the uncertainty about intratemporal trade-offs between them. Even though present bias does *not* create a conflict between short- and long-run preferences regarding how to spend income within a time period, it can lead—through budgeting or other personal rules—to non-standard demand behavior. With regard to the demand for commitment, the literature has focused on the consumers' problem of curbing undersaving and often stressed the usefulness of devices like illiquid assets and savings accounts. As we will see, however, under the realistic assumption of multiple goods consumers can do strictly better by (also) adopting good-specific budgets. This opens the door to a demand for other commitment devices,

[6]This prediction continues to hold for partially naive individuals who incorrectly think ex ante, when choosing a plan, that their present bias is weak (O'Donoghue and Rabin (2001)).

[7]This literature has become so vast that it is impossible to list all relevant papers here.

[8]Brocas and Carrillo (2008) discuss a model where consumption involves two goods, only one good has ex-ante uncertain utility, and the doer is fully myopic. In this case, the optimal commitment strategy consists of a non-linear plan that punishes spending on one good by cutting spending on the other, which is not a budgeting plan. Even if one focuses on budgeting plans, the planner never sets good-specific budgets with a fully myopic doer (cf Proposition 2). Their model also does not allow for studying the roles of uncertainty about intratemporal vs. intertemporal trade-offs and of the intensity of present bias.

such as services that help people implement those budgets.[9] The existence and regulation of a market for personal budgeting services can have significant consequences on welfare, for instance via the large effects that budgeting can have on wealth accumulation as found by Ameriks et al. (2003). Moreover, the present paper offers predictions on which type of consumers will demand which type of devices, and this can be used by third-party providers to target their promotional efforts.[10]

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

Finally, the insights of this paper are relevant beyond the consumption-savings application. This application is an instance of situations where a principal delegates to a better informed agent the allocation of finite resources across multiple categories.[11] Section 7 outlines other instances in public finance, corporate governance, and workforce management, where the results can take a more normative connotation. One takeaway is that these multidimensional problems can introduce a novel economic rationale, absent in unidimensional settings, for reducing the agent's choice set: to change the trade-offs he faces between choice dimensions causing no conflict of interests so as to alleviate—through the resource constraint—the consequences of the conflict in other dimensions. For such restrictions to be beneficial, however, it is crucial that the agent's information also affects those trade-offs.

2 Related Literature

Existing explanations of personal budgeting are based on Thaler's (1985) seminal work. Combining the notions of "transaction utility" and gain-loss utility, he argues that individuals treat the consequences of each transaction in isolation. In this case, he shows that they can solve their consumption-savings problems by means of transaction-specific budgets, a result which echoes Strotz's (1957) explanation of budgeting based on separable consumption utilities. In reality, however, people set budgets for sufficiently long periods (a week or a month) so that each budget covers many transactions. Also, in his (and Strotz's) deterministic model, people can achieve the same utility with and without budgets. But if people faced uncertainty, they would never impose ex ante budgets which bind ex post; that is, they do not exhibit a strict demand for budgets as commitment devices. Finally, transaction and gain-loss utility seem to have no direct link with self-

[9] This kind of services are currently offered by firms like Mint, Quicken, and StickK.

[10] Of course, in reality it may be hard to observe each consumer's degree of present bias and offer commitment devices accordingly. For an analysis of some of the issues that arise when that degree is consumers' private information, see Galperti (2015).

[11] Thaler and Shefrin (1981) were the first to draw a connection between self-control problems and delegation problems.

control problem, which the literature usually views as the underlying cause of personal budgeting. Other papers have shown that gain-loss utility can explain other phenomena commonly classified as mental accounting, such as choice bracketing (Koch and Nafziger (2016)), which are however different from budgeting.

This paper relates to the mechanism-design literature on the trade-offs between commitment and flexibility—in particular to Amador et al. (2006) and Halac and Yared (2014).[12] It departs from both papers by introducing multiple consumption goods and uncertainty about intratemporal trade-offs, thus uncovering how this type of uncertainty affects qualitatively the commitment-flexibility trade-off and its solutions. Amador et al. (2006) showed that in a world with unidimensional consumption savings floors often coincide with the fully optimal commitment plans (within a very general class of plans). Halac and Yared (2014) also differ from the present paper by focusing on the role of information persistency. In their setting, optimal commitment plans can distort future consumption choices, even though those choices cause no conflict of interests between the individuals' selves once today's choice is fixed. The reason is that information persistency creates a link between the doer's expected utility from tomorrow's choices and his information today; hence those choices can be exploited to relax today's incentive constraints, as in other dynamic mechanism-design problems.[13] The results of the present paper do not depend on the correlation between the doer's pieces of information.

This paper is also related to the literature that studies how rationing affects consumer behavior (Howard (1977), Ellis and Naughton (1990), Madden (1991)). By imposing a savings floor or a good-specific budget, an individual is essentially rationing his future selves as the government of a centralized economy may ration its citizens. Unlike this literature, however, here rationing assumes the function of a commitment device. The rationing literature has shown that predicting the effects of good-specific budgets is far from trivial. Its insights will be useful to identify conditions under which budgets can help the individual.

Finally, this paper contributes to the rich literature on delegation following Holmström (1977, 1984). One key difference from unidimensional delegation problems is that, in those problems, the principal reduces the agent's choice set in only two ways: She removes extreme options that can badly hurt her or intermediate ones so as to render the agent's choice more sensitive to his information (Alonso and Matouschek (2008)). Regarding the case of multidimensional delegation, few papers examine it and none in the setting studied here.[14] Koessler and Martimort (2012) consider specific settings which render the delegation problem similar to a unidimensional screening exercise. In Frankel (2014), the agent has the same bias for all dimensions of his choice, but the principal does not know its properties (strength, direction, etc.). In this case, the best delegation policies against the worst-case bias may require the agent's average decision to meet some preset target. Frankel (2016) focuses on policies that cap the gap between the agent's and the

[12]This literature also includes Athey et al. (2005), Ambrus and Egorov (2013), and Amador and Bagwell (2013b).

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

[14]In Alonso et al. (2014), finite resources have to be allocated across multiple categories, but each category is controlled by a distinct agent whose information is unidimensional.

principal's realized *payoffs*, not his specific decisions. For the settings in the present paper, one can show that such policies can be implemented by imposing a limit only on how much the agent can allocate to the dimension of his choice causing the conflict of interest with the principal (for instance, only a savings floor but no budgets).

3 Baseline Consumption-Savings Model

This section introduces the baseline model of consumption and savings with imperfect self-control. We will keep the model simple, with some admittedly strong assumptions, to avoid distractions from issues that are of second-order importance for this paper and do not distinguish it from the literature. Section 6 discusses how to relax these assumptions.

Allocations. Consider an individual who lives for two periods. In the first period, he consumes a bundle of two goods $\mathbf{c} = (c_1, c_2) \in \mathbb{R}^2_+$. He receives his entire income—normalized to 1—at the beginning of the first period. To finance consumption in the second period, he then has to save some amount $s \in \mathbb{R}_+$. Period-2 consumption involves a single good and hence equals s. In the first period, the feasibility constraint is then $c_1 + c_2 + s \leq 1$. Think of each c_i and s as the share of income allocated to good i and to savings.[15]

Preferences and Information. As in existing dual-self models (cf Footnote 3), the individual consists of a long-run self, called "planner" (she), and a short-run self, called "doer" (he). Their preferences depend on some information represented by a 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, for each state the planner and the doer evaluate streams $(\mathbf{c}, s)$ using respectively the utility functions

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

To add clarity and tractability to the model, for now assume that

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

Additive separability will be relaxed in Section 6.1. The parameter $\beta \in (0, 1)$ captures the doer's present bias and hence the conflict with the planner, who knows the level of β.[16] This formulation is consistent with the single-agent, quasi-hyperbolic discounting model of Laibson (1997). It is also consistent with viewing the individual as a household, whose members have time-consistent but heterogeneous time preferences. In this case, under weak conditions the household's aggregate preference exhibits present bias (Jackson and

[15]Nothing significant changes if the individual receives some income in period 2 and can borrow in period 1, or if the consumption bundle $\mathbf{c}$ involve more than two goods. In fact, the proofs are carried out for the general case with $n \geq 2$ goods. Section 6.3 extends the analysis to the case in which the individual consumes multiple goods also in future periods. Section 6.5 discusses the extension to more than two periods.

[16]Allowing for partial naiveté (for example, as in O'Donoghue and Rabin (2001)) would not change the message of the paper, as discussed in Section 4.

Yariv (2015)). Finally, note that the component $\mathbf{r}$ of the state can be interpreted as shocks in tastes, but also as shocks in prices, which determine how money spent on good i translates into its physical units.[17]

A key, novel aspect of this model is that information affects *inter*temporal as well as *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. This structure involves 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 highlights the difference between uncertainty about intratemporal and intertemporal trade-offs; it will allow us to easily shut down the former kind of uncertainty and analyze its consequences; it simplifies the comparison with the literature. Regarding the state distribution, denoted by G, rich forms of dependence as well as full independence will be allowed across its components θ, r_1, r_2. Hereafter, let $\omega = (\theta, \mathbf{r})$ and let Ω denote the state space.

Commitment Plans. The planner delegates to the doer the choice of an income allocation. Knowing his bias, she would like to design a commitment plan dictating which allocations he is allowed to implement. In the case of budgeting, such a plan takes the form of spending limits on specific consumption categories, denoted by b_i (for budget), or of an overall limit on consumption expenditures, which can be implemented through a minimum-savings rule denoted by f (for savings floor). For instance, a plan can require that consumption never exceed 80% of income and "going out" and "house expenses" never exceed 10% and 15% each. Individuals and households typically commit to such plans for a specified period of time, like a week or a month.[18]

To formalize this, let the set of feasible allocations in period 1 be

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

A budgeting plan, B, can then be expressed as follows:

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

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. Note that the savings floor f or some budget b_i may never constrain the doer, or they may do so only in some states. Therefore, from the ex-ante viewpoint, we will call f and b_i *binding* if they constrain the doer with strictly positive probability under G.

To focus on the issues of interest for this paper, it is assumed that people stick to their plans. This is not a minor assumption, of course, but the literature has proposed several mechanisms which can justify it. These mechanisms include people's 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

[17] For instance, for $i = 1, 2$, let $\overline{u}^i(z_i) = \frac{z_i^{1-\gamma_i}}{1-\gamma_i}$ with $\gamma_i > 0$ be the utility from z_i units of good i, and let $\vartheta\rho_i > 0$ be its price, which has a common component (ϑ) and an idiosyncratic component (ρ_i). If we define $c_i = \vartheta\rho_i z_i$ for all i, we can write the usual resource constraint as $c_1 + c_2 + s \leq 1$. Letting $\theta r_i = [\vartheta\rho_i]^{\gamma_i - 1}$, we get $\theta u^i(c_i; r_i) = \theta r_i \overline{u}^i(c_i)$.

[18] It is implicitly assumed here that people correctly notice and report every expense to the corresponding budget, a process which defines budgeting itself.

(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 on their own provided that such plans are not too stringent or too costly ex post. Thus, they may set good-specific budgets or savings floors not as strict as they would want to, yet still use them. Even in this case, it is worth understanding which forces lead people to find budgets and floors useful despite their ex-post inefficiency. This understanding can also be valuable for third parties which design commitment devices to help people stick to their plans (such as firms like Mint, Quicken, and StickK). For instance, some present-biased individuals may not use budgets not because they cannot stick to them, but simply because they do not find them useful; it would then be pointless to try to sell devices for implementing those budgets to such individuals.

Timing. In reality, individuals set their commitment plans prior to observing all the necessary information for making a decision. This creates a non-trivial trade-off between commitment and flexibility. To model this, we will assume that only the doer observes the state realizations. This corresponds to the following timing: First the planner commits to a plan B, then the doer observes ω and implements some allocation from B. The planner designs her plan to maximize her expected payoff taking into account the doer's future decisions.[19]

The goal of the paper is to understand whether and how the planner sets minimum-savings rules and good-specific budgets. In other words, it analyzes the problem of choosing $B \in \mathcal{B}$ so as to maximize

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

subject to

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

We will refer to a solution to this problem as an optimal plan, where "optimal" is of course from the planner's viewpoint.

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 all $i = 1, 2$. We will only assume that the joint probability distribution G of (θ, r_1, r_2) has full support on Ω (that is, $G(O) > 0$ for every open $O \subset \Omega$). The conditions on $\underline{\theta}$ and $\overline{\theta}$ are meant to rule out the implausible situation in which the individual does not care at all about the present or the future.[20] The conditions on $\underline{r}_i$ and $\overline{r}_i$ have bite only when combined with the properties of u listed next.

Differentiability, monotonicity, concavity: The function v is assumed to be twice continuously differentiable with $v' > 0$ and $v'' < 0$. For $i = 1, 2$ and all $r_i \in [\underline{r}_i, \overline{r}_i]$, the function $u^i(\cdot; r_i) : \mathbb{R}_+ \to \mathbb{R}$ is twice differentiable with $u^i_c(\cdot; r_i) > 0$ and $u^i_{cc}(\cdot; r_i) < 0$;

[19]When period-2 consumption involves multiple goods or the future involves multiple periods, the question arises of whether the planner may benefit by committing to plans that cover more than one period. Section 6.4 discusses this issue.

[20]A similar assumption appears in Amador et al. (2006), who point out that with unbounded support it may always be optimal to grant the doer full flexibility.

also, u_c^i and u_{cc}^i are continuous on $(0,+\infty) \times [\underline{r}_i, \overline{r}_i]$. In particular, this implies that u_c^i is bounded above and away from zero; this seems plausible to the extent that i refers to "food," "housing," or "entertainment" and a period corresponds to a week or a month.

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

4 Optimal Budgeting Plans

4.1 Preliminaries

First of all, the planner's problem has a solution.[21]

Lemma 1. *There exists B that maximizes $\mathcal{U}(B)$ over $\mathcal{B}$.*

Every plan B induces the doer to choose an income allocation state by state. Two allocation rules are focal benchmarks. First, for each ω, let $(\mathbf{c}^d(\omega), s^d(\omega))$ be the allocation that the doer would choose if granted full discretion, namely the solution to $\max_{(\mathbf{c},s)\in F}\{\theta u(\mathbf{c};\mathbf{r}) + \beta v(s)\}$. The second allocation, denoted by $(\mathbf{c}^p(\omega), s^p(\omega))$, represents what the planner would like the doer to choose in ω, which is the solution to $\max_{(\mathbf{c},s)\in F}\{\theta u(\mathbf{c};\mathbf{r}) + v(s)\}$. We will call $(\mathbf{c}^d, s^d)$ the *full-discretion allocation* and $(\mathbf{c}^p, s^p)$ the *first-best allocation*. They satisfy some useful properties, summarized in the following remark.

Remark 1.
1. $(\mathbf{c}^p, s^p)$ and $(\mathbf{c}^d, s^d)$ are continuous in ω;
2. Each component of $(\mathbf{c}^p, s^p)$ and of $(\mathbf{c}^d, s^d)$ takes values in a closed interval and is bounded away from zero;
3. For all $i = 1, 2$, c_i^p and c_i^d are strictly increasing in r_i and θ and strictly decreasing in r_j for $j \neq i$;
4. s^p and s^d are strictly decreasing in θ, r_1, and r_2;
5. For all $\omega \in \Omega$, $s^d(\omega) < s^p(\omega)$ and $s^d(\omega)$ is continuous and strictly increasing in β;
6. For all $\omega \in \Omega$ and $i = 1, 2$, $c_i^d(\omega)$ is continuous and strictly decreasing in β.

One last property of the model is worth noting. All consumption goods are normal: For both selves higher spendable income, $1-s$, always leads to a higher optimal allocation to each good.[22]

As an illustration, Figure 1 represents the first-best and full-discretion allocations for a model which is fully symmetric with respect to good 1 and 2. Since the doer will always save whatever he does not consume ($v' > 0$), we can focus on his choices of $\mathbf{c}$. Note that bundles which lie on negative-45° lines closer to the origin correspond to higher levels of savings. Figure 1 describes $\mathbf{c}^p$ and $\mathbf{c}^d$ as the regions inside the dashed and solid lines, whose shape follows from property 3 and 4 in Remark 1. To see this,

[21] The proofs of the main results are in the Appendix. All other proofs appear in the Online Appendix.
[22] This property follows, for instance, from Proposition 1 in Quah (2007).

consider $\mathbf{c}^p$ and suppose for the moment that θ takes only one value. Start from state $(\theta, \underline{r}_1, \underline{r}_2)$, which corresponds to the highest savings level. If we raise, say, r_1 continuously up to $\overline{r}_1$, c_1^p increases while c_2^p as well as s^p decrease, which means that we move along the south portion 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 as well as s^p decrease; that is, we move along the east portion of the dashed line. Proceeding in this way, we can describe the entire dashed line; continuity of $\mathbf{c}^p$ implies that its range has to be connected and hence equal to the entire region inside this line. Finally, the doer's systematic undersaving corresponds to a shift of the $\mathbf{c}^d$ region away from the origin; the stronger his bias, the bigger the shift.

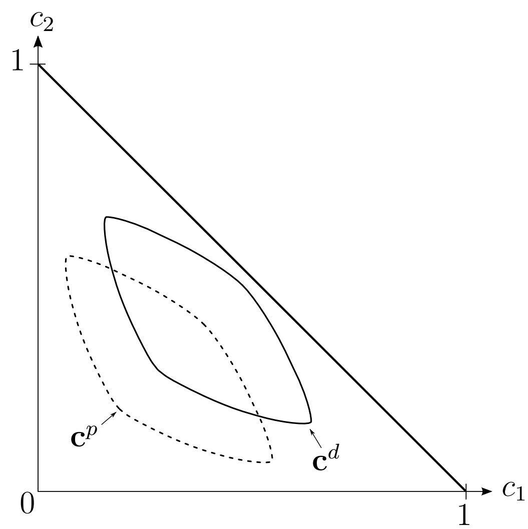

Figure 1: First-best and Full-discretion Allocations

4.2 Main Results

This section presents the main results of the paper, which relate the use of minimum-savings rules and good-specific budgets to the intensity of present bias. It also shows that reducing the uncertainty on intratemporal trade-offs renders the use of specific budgets less likely to be optimal, thereby highlighting the role of this kind of uncertainty.

Due to his present bias, the doer tends to systematically overspend on consumption and undersave from the planner's viewpoint. The planner would like to limit these departures from the first-best allocation. Therefore, it may seem straightforward that she always sets a savings floor as well as good-specific budgets. After all, when consumption involves multiple goods, more commitment tools should always help: The floor can limit *total* consumption expenditures, while specific budgets can limit splurging good by good.

The following observations, however, suggest that things are not so simple. If we fix savings, the planner and the doer always agree on how to divide the remaining income between goods. Thus, without uncertainty, the planner can achieve the first best both by relying exclusively on a savings floor and by using only good-specific budgets: Given ω, both the plan which sets $f = s^p(\omega)$ and $b_1 = b_2 = 1$ and the plan which sets $f = 0$,

$b_1 = c_1^p(\omega)$, and $b_2 = c_2^p(\omega)$ induce the doer to implement $(\mathbf{c}^p(\omega), s^p(\omega))$. On the other hand, with uncertainty, imposing only budgets cannot implement the same allocations as imposing only a savings floor.[23] Intuitively, when f binds, the doer continues to react to $\mathbf{r}$ by changing how he spends $1 - f$ across goods. Thus, good-specific budgets which are sufficiently strict to ensure that total spending never exceeds $1 - f$ will have to constrain either c_1 or c_2 strictly below the largest amount spent on that good subject to only f. Finally, a binding savings floor distorts the division of income between spending and saving, but never distorts the implemented consumption bundle; by contrast, binding budgets also distort consumption. When f binds, the doer allocates $1 - f$ so as to maximize $u(\cdot; \mathbf{r})$ and hence chooses a bundle $\mathbf{c}$ which equalizes marginal utilities between goods. By contrast, a binding budget on good i mitigates the doer's *aggregate* overconsumption, but without other rules it exacerbates overconsumption in good j.[24] Intuitively, overspending on good j comes at the cost of subtracting money from savings, which the doer undervalues, or from good i, which he values on par with j. Capping good i removes the second cost, so the doer overspends on j even more.

Despite these drawbacks of good-specific budgets, Proposition 1 shows that there always exists a sufficiently *weak* present bias such that every optimal plan must include them.

Proposition 1. *There exists $\beta^* \in (0, 1)$ such that, if $\beta > \beta^*$, then every optimal $B \in \mathcal{B}$ must include binding good-specific budgets.*

The Appendix describes an algorithm to derive β^*. Note that the conclusion of Proposition 1 does not hold only in the limit for $\beta \simeq 1$, and β^* can be significantly smaller than 1—for that matter, smaller than the levels of present bias usually considered plausible based on the empirical evidence. The next section explains the logic of Proposition 1, which should also clarify that the economics behind it does not require $\beta \simeq 1$. Thus, allowing (or enabling) the planner to use good-specific budgets on top of a savings floor can have significant first-order effects on her payoff. Concretely, for example, recall that Ameriks et al. (2003) find that detailed budgeting plans can contribute to increasing significantly households' wealth accumulation.

Although Proposition 1 says that good-specific budgets must be part of an optimal plan, it does not say that they are always combined with a savings floor. Indeed, it is possible to have situations in which they are as well as situations in which they are not. Section 4.4 illustrates this.

Do optimal commitment plans always require good-specific budgets? The answer is no. There always exists a sufficiently *strong* present bias such that, to be optimal, a plan should impose only a savings floor.

Proposition 2. *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_*. Note that the conclusion of Proposition 2 does not hold only in the limit for $\beta \simeq 0$ and β_* can be significantly larger than 0. The logic

[23] See Lemma 8 in the Online Appendix.

[24] See Lemma 9 in the Online Appendix.

behind Proposition 2, explained in the next section, should clarify this point. Also, β_* depends on the properties of the utility functions and on the distribution G only through its support. Given this, we can show that weaker biases suffice for the conclusion of Proposition 2 to hold, if we reduce the uncertainty on intratemporal trade-offs in the following sense.

Corollary 1. *Consider two individuals 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 thresholds in Proposition 2 corresponding to the two individuals. 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 $\beta'_* > \beta_*$.*

4.3 Intuition

Consider Proposition 1. The easiest way to understand it is to start by analyzing how the planner would use a savings floor f and each budget b_i *in isolation*. We can then build on the insights of this thought experiment to understand how she combines f, b_1, and b_2. It is important to keep in mind that f and some b_i can simultaneously bind, thereby affecting the doer's choices in possibly complicated ways.

Suppose first that the planner could use only f. In this case, she would always set f strictly between the largest and smallest first-best levels of savings, $\overline{s}^p$ and $\underline{s}^p$.[25] Moreover, she would raise f in response to a stronger present bias.[26] The intuition is simple. Under full discretion the doer saves less than $\underline{s}^p$ for some states, which is never justifiable for the planner. And f never distorts the chosen bundle $\mathbf{c}$, because the two selves have the same consumption utility u. Consequently, the planner always sets $f \geq \underline{s}^p$. She cannot optimally set $f = \underline{s}^p$, because by marginally increasing f she suffers only a second-order loss in states where $s^p(\omega) = \underline{s}^p$, but a first-order gain in states where $s^p(\omega) > \underline{s}^p$ and f binds. A similar logic explains why $f < \overline{s}^p$. Thus, the best f has to balance the benefits of limiting undersaving in some states with the cost of causing inefficient oversaving in others. Finally, regarding the dependence of f on β, a stronger tendency of the doer to undersave (lower β) strengthens the benefits of raising f, but does *not* change its cost: In states where $s^p(\omega) < f$, f binds for *any* bias because $s^d(\omega) < s^p(\omega)$.

Now suppose that the planner can use only a single budget. It turns out that capping spending on even only one good dominates granting the doer full discretion.[27] To see why, start from $b_i = \overline{c}_i^d = \max_\omega c_i^d(\omega)$. Lowering b_i a bit has two effects. First, when binding, b_i distorts the chosen bundle $\mathbf{c}$. This negative effect, however, is initially of second-order importance for the planner: Under full discretion the doer's choice of $\mathbf{c}$ is always efficient, in the sense that it equalizes marginal utilities between goods. To quantify this effect, the proof relies on the Lagrange multiplier for the constraint that b_i

[25]See Lemma 4 in the Appendix. Though reminiscent of Amador et al.'s (2006) main result, Lemma 4 differs in several respects. It assumes right away that plans can use only f, and its proof uses different techniques from the clever mechanism-design approach in Amador et al. (2006). That approach does not work in the present setting because consumption and information have multiple dimensions, as discussed in Section 6.2.

[26]See Lemma 5 in the Appendix, which also deals with the case of multiple optimal floors.

[27]See Lemma 6 in the Appendix.

adds to the doer's optimization problem. The second effect of b_i is to curtail the doer's undersaving with positive probability,[28] which causes a first-order gain for the planner. Overall b_i should then benefit her, but there is a subtlety here: The doer should not reallocate money to the unrestricted goods at a much faster rate than to savings, which is not obvious and need not be true. The proof leverages the additive structure of u to derive this key property, which however holds more generally (cf Section 6.1).

These observations bring to light the mechanism whereby the multidimensionality of consumption can help to curb the consequences of present bias. A budget b_i incentivizes the doer to save more because it forces him to choose inefficient bundles—not just to spend less money on good i, which he could fully shift to other goods—and these inefficiencies limit how much he can gain in terms of present utility by undersaving. Not all distorting rules work, however. For example, one can show that imposing a binding minimum-spending rule on any good, though distorting $\mathbf{c}$, never improves savings and hence is never part of an optimal plan.[29]

Combining these insights leads to Proposition 1. To see this, consider Figure 2, which reproduces Figure 1 for the cases of strong and weak biases (that is, low and high β). Recall that the shape of the regions $\mathbf{c}^p$ and $\mathbf{c}^d$ follows from property 3 and 4 in Remark 1. In particular, for $i = 1, 2$ and $k = p, d$, we have that

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

These properties imply that the states in which both selves want to spend the most on c_1 or c_2 are *not* the states in which they want to save the least; the latter states map to the dark-shaded areas in Figure 2, the former to the light-shaded areas. Finally, in Figure 2 budgets correspond to vertical and horizontal lines, while f to a line with slope -1.

Optimal plans must impose good-specific budgets, when the bias is sufficiently weak, because they help the planner improve savings when doing so via f would require an excessively tight floor. Recall that if the planner can use only f, she relaxes it as the bias weakens. In Figure 2, this corresponds to the f line moving farther away from the origin. Consequently, f primarily targets the doer's decisions in the dark-shaded area, but becomes less likely to affect them in the light-shaded areas where only one good is very valuable—compare panels (a) and (b). Undersaving also occurs in these states, however. To curb it, the planner prefers not to use f for weak biases, but she can add a budget b_i for each good that binds when f does not (as in panel (b)). This b_i will curb undersaving when the marginal rate of substitution between good i and j is very high, and despite distorting $\mathbf{c}$, it benefits the planner.

Consider now Proposition 2. Its proof hinges on showing that, no matter how the planner combines savings floor and budgets, she never lets the doer save less than $\underline{s}^p$.[30] Intuitively, if plan B allows this, then raising f up to $\underline{s}^p$ uniformly improves the planner's payoff with regard to savings. Now recall that all goods are normal. Therefore, the resulting lower spendable income renders the budgets of B (if any) less likely to bind and hence distort $\mathbf{c}$. Thus, the planner also gains on this front.

[28]See Lemma 9 in the Online Appendix.
[29]See Example 10.6 in the Online Appendix.
[30]See Lemma 7 in the Appendix.

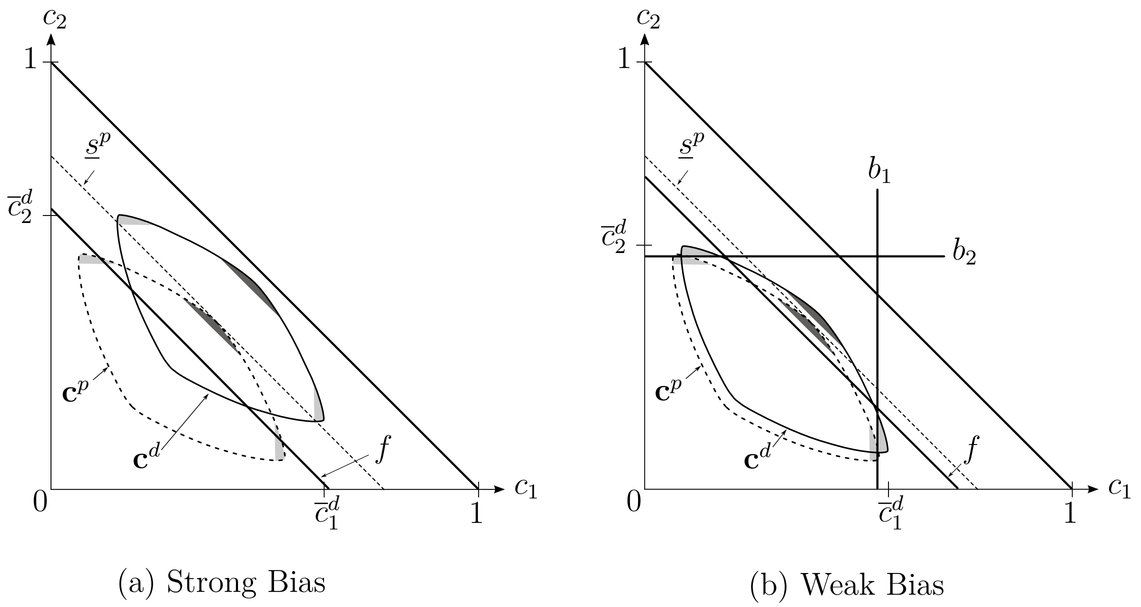

Figure 2: Optimal Delegation Policy – Intuition

We can now see the intuition behind Proposition 2. When β is sufficiently small, the doer wants to save less than $\underline{s}^p$, no matter what information he observes. Hence, when a budget forces him to consume less of good i, he reallocates all the unspent money across the other goods, but not to savings. Since binding budgets distort $\mathbf{c}$, they cannot benefit the planner if they do not increase s. This logic also leads to the following observation, which is useful to immediately identify certain plans as strictly suboptimal.

Remark 2. Suppose that $B \in \mathcal{B}$ involves binding budgets and always induces the same level of savings, say s'. Then B is *strictly* dominated by a plan which imposes only $f = s'$.

Concretely, suppose we observe a household that uses budgets which always bind, thereby resulting in the same level of savings. According to this model, we can immediately conclude that the household's plan is suboptimal.

Corollary 1 offers insights into how reducing uncertainty on intratemporal trade-offs affects the optimal plans. Shrinking the range of the doer's information on those trade-offs—without changing that on the intertemporal trade-off (θ)—expands the set of strong biases for which optimal plans use only f. This is because budgets are useful to curb undersaving in states with large asymmetry in the goods' marginal utilities (recall Figure 2). Thus, shrinking the range of this asymmetry reduces the scope for budgets to be used. One may wonder what happens in the limit when uncertainty is only about θ, but consumption continues to involves multiple goods. Do optimal plans always use only f? If not, under which conditions? Section 5 will provide the answer.

As should be expected, the weakest bias for which optimal plans include good-specific budgets depends on the details of the setting at hand. As β falls below β^*, for any B it increases the probability that the doer ends up in a state where he is constrained by B's actual lower bound on savings, which prevents them from falling below some level

$\sigma \geq \underline{s}^p$. Since in these states binding budgets only create inefficiencies, their appeal for the planner falls accordingly. How she balances the inefficiencies in those states with the budgets' benefits in other states ultimately depends on their distribution G. Nonetheless, since she can always set $f = \sigma$, for biases below some level $\hat{\beta} \geq \beta_*$ every optimal plan will again involve only f.

One may think that these results are driven by the fact that good-specific budgets and the savings floor are substitute tools—in the sense that, everything else equal, optimal plans set a slacker floor when they can also use budgets. The interaction between these tools is more subtle. By curbing undersaving in some states, binding budgets lower the return of tightening f to affect the doer's behavior in those states, which can result in a slacker f. At the same time, however, the budgets also lower the return of loosening f because they prevent the doer from consuming efficiently the extra spendable income, which can result in a tighter f. Therefore, when individuals can freely combine specific budgets and a savings floor, it is not possible to conclude that the floor's level varies monotonically with the intensity of present bias.

The prediction that good-specific budgets are part of an optimal commitment plan only for weak present biases is consistent with the qualitative findings in Antonides et al. (2011) on which kind of people adopt budgeting, discussed in the introduction.[31] One might wonder whether the finding that strongly present-biased individuals seem to not use budgets simply indicates that such individuals are less sophisticated or less able to commit than others. First, severe naiveté seems an implausible assumption in the settings considered here: For similar settings (of course with many periods), Ali (2011) concludes that an individual should learn his true bias through experimentation. Second, we saw that once a strongly biased individual can rely on a savings floor, the reason why good-specific budgets do not work for him is not that he cannot honor them: Even if he could, a simple economic logic shows that using them would strictly lower his utility. Finally, as long as an individual can commit to some degree of budgeting—however small—the theory in this paper predicts that partial naiveté can actually result in him setting up good-specific budgets.[32] Intuitively, by underestimating his bias the individual may incorrectly think ex ante that he faces the situation represented in Figure 2(b) (as opposed to (a)). In this case, he concludes that he strictly benefits by combining a savings floor with whatever budgets he can stick to.

More generally, allowing for partial naiveté would not change the message of the paper, since the entire analysis is from the ex-ante viewpoint of the planner. As noted, a naive individual may set up good-specific budgets ex ante only to realize, ex-post, that he should have used only a savings floor. Therefore, in the present model naiveté can be detrimental to the extent that it also leads individuals to adopt forms of commitment that rely on consumption distortions where there should be none. This does not mean, however, that naive individuals should be prevented from using budgets: Although ex ante they choose

[31] Note that Antonides et al. (2011) do not report evidence regarding how stringent budgets or floors are in relation to the intensity of present bias. For that matter, the present theory suggests that the relationship need not be monotonic. It would be interesting to run additional experiments designed specifically to test the theory in the present paper.

[32] Partial naiveté can be modeled as in O'Donoghue and Rabin (2001), for example.

a commitment plan that is strictly dominated (from the analyst's viewpoint), ex post the budgets may still provide some benefit in curbing the consequences of present bias. In this case—and especially in the leading case of sophistication—helping individuals stick to their voluntarily chosen budgets as well as savings floors can result in long-run welfare gains.

4.4 Budgets with Savings Floor or Only Budgets?

This section examines whether optimal commitment plans always involve a binding savings floor. While this is the case in settings with a single consumption good,[33] with multiple goods there exist both settings in which the planner combines good-specific budgets with a savings floor and settings in which she uses only the budgets. To show this, throughout this section we will focus on the following symmetric model: Let $u^1(c;r) = u^2(c;r) = r\ln(c)$, $\underline{r}_1 = \underline{r}_2 = \underline{r} > 0$, $\overline{r}_1 = \overline{r}_2 = \overline{r} > \underline{r}$, and $v(s) = \ln(s)$.[34]

To develop the intuitions, we will first consider a setting with three states and later extend the results to the general setting. Let $\omega^0 = (\overline{\theta}, \overline{r}_1, \overline{r}_2)$, $\omega^1 = (\underline{\theta}, \overline{r}_1, \underline{r}_2)$, and $\omega^2 = (\underline{\theta}, \underline{r}_1, \overline{r}_2)$, whose respective probabilities are g, $\frac{1}{2}(1-g)$, and $\frac{1}{2}(1-g)$. Remark 1 and symmetry imply that

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

similar properties hold for $(\mathbf{c}^p, s^p)$. By continuity, there exist $\beta < 1$ sufficiently high so that $s^d(\omega^1) = s^d(\omega^2) > s^p(\omega^0)$; hereafter, fix β at such a value. There also exists $\underline{\theta} < \overline{\theta}$ sufficiently close to $\overline{\theta}$ so 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 such a situation, focussing on consumption as in Figure 2. Concretely, we can think about this situation in the following terms. Imagine that Bob enjoys dining out (c_1) and live music (c_2). In a given period, his best friend Ann may visit him ($\overline{\theta}$) or not ($\underline{\theta}$). If on his own, depending on the circumstances Bob prefers either to dine at a fancy restaurant with a piano bar (in ω^1) or to grab a quick sandwich and attend a great concert (in ω^2). By contrast, when Ann is in town, Bob prefers to combine a good restaurant with a good concert (in ω^0), caring more about her company.

Letting g be the only free parameter, we obtain the following.

Proposition 3. *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 $b_2 = c_2^p(\omega^2)$.*

The intuition is as follows. If Bob knew Ann was not visiting, he could set f so as to eliminate splurging in both ω^1 and ω^2—this f corresponds to the dotted line in Figure 3(a). However, such an f is too stringent if Ann visits. Therefore, if Bob assigns high probability to her visit, ex ante he views setting f high enough to affect his choices in ω^1 and ω^2 as too costly, and hence prefers $f = s^p(\omega^0)$. This f grants full discretion in ω^1 and ω^2. In these states, however, by the same logic of Lemma 6 Bob can curb his undersaving by

[33] See Proposition 2 and 10 in Amador et al. (2006).

[34] The function $\ln(\cdot)$ violates the continuity and differentiability assumptions of Section 3 at 0, but this is irrelevant for the analysis.

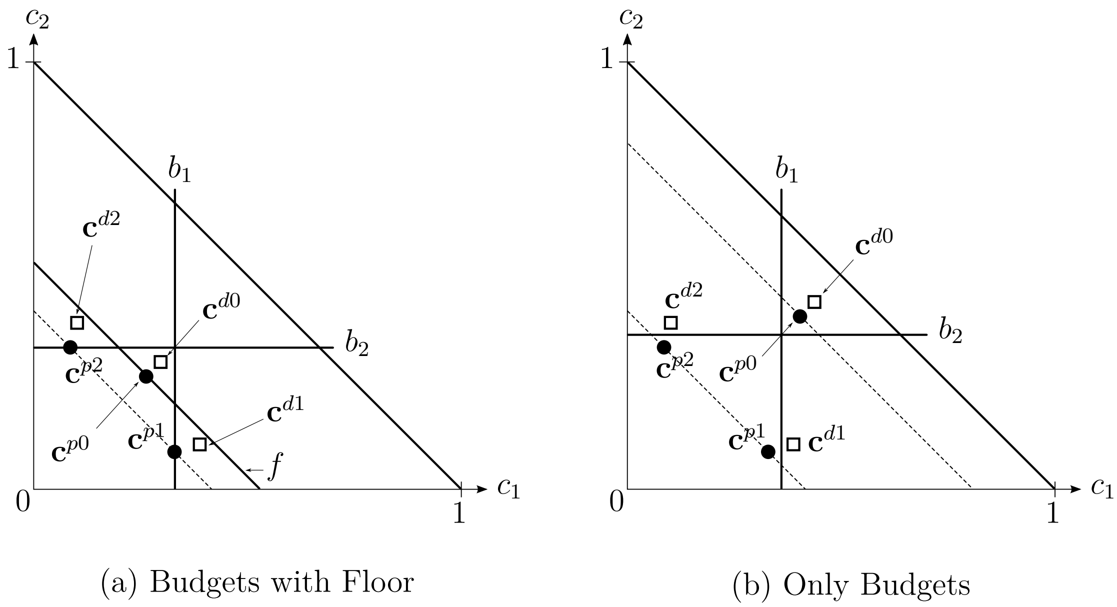

Figure 3: Three-State Example ($\mathbf{c}^{di} = \mathbf{c}^d(\omega^i)$ and $\mathbf{c}^{pi} = \mathbf{c}^p(\omega^i)$)

setting good-specific budgets; moreover, here he can do so without affecting his choice in ω^0. The specific levels of b_1 and b_2 are just a byproduct of logarithmic payoffs.

A simple change of the previous setting suffices to show that 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 of each good 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 $c_2^p(\omega^0) > c_2^p(\omega^2)$. Continuing our story, suppose that now, when she visits, Ann always insists on going to the very best restaurants and concerts and Bob caves in.

Proposition 4. *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 every $i = 1, 2$.*

Figure 3(b) helps us see the intuition. Now Bob is willing to spend more on both goods when Ann is in town than when she is not. Therefore, the optimal budgets that he would set to curb splurging in ω^1 and ω^2 already create an aggregate limit on consumption which binds in ω^0. As a result, Bob wants to relax such budgets; and if his first-best consumption in ω^0 is not too high, he may be able to keep them sufficiently low so as to continue to curb splurging in ω^1 and ω^2. Since these budgets already push savings above the first-best level in ω^0, Bob cannot benefit by adding a binding f. Note that, although Proposition 3 and 4 are intuitive, it takes some work to rule out the possibility of multiple, perhaps asymmetric, optimal plans featuring different properties from those stated in the propositions.[35]

[35] Some readers may wonder whether, with finitely many states, the planner could do better by defining the doer's choice set directly as a list of fully specified allocations, one for each state (for example, as the

It remains to argue that the qualitative properties of the plans in Propositions 3 and 4 can arise in settings with a continuum of states. This should be the case if the planner assigns sufficiently high probability to states that induce trade-offs *similar* to those in ω^0, ω^1, and ω^2. Corollary 2 in the Online Appendix formalizes and confirms this idea. One comment is in order for Proposition 4. In contrast to the three-state setting, now there will always be a set of states with positive probability where the optimal b_1 and b_2 do not affect the doer's choices. On the one hand, marginally increasing f above $1 - b_1 - b_2$ would curb undersaving in those states. On the other hand, doing so would force the doer to save a bit more in states similar to ω^0 than the level induced by b_1 and b_2, which already exceeds the first best. Both effects cause first-order changes in the planner's payoff, but one can show that the latter dominates if she cares enough about the states similar to ω^0.

To sum up, a weakly present-biased individual may adopt plans that involve only good-specific budgets for the following reasons. First of all, to limit undersaving in states with large asymmetry in consumption marginal utilities she prefers to use the budgets rather than a savings floor, since the latter would have to be too stringent. Together the budgets then impose a cap on total expenditures. If this cap already ensures that her savings will be sufficiently high in states where present consumption is very valuable overall, then any binding floor will have to cause additional oversaving (recall that for every optimal plan s exceeds $\underline{s}^p$) and this inefficiency can dominate the floor's commitment benefits.

5 The Role of Uncertainty about Intratemporal Trade-offs

The previous results illustrate how the multiplicity of consumption goods and the uncertainty about intratemporal trade-offs can render good-specific budgets useful commitment tools. This section further demonstrates that such uncertainty plays a key role in explaining budgeting and disentangles this role from the multidimensionality of consumption. To this end, it shuts down that uncertainty, while keeping multiple goods and uncertainty about the intertemporal utility trade-off. In this case, as we will see, most of the time the optimal plan will involve a savings floor but *no* good-specific budgets.

For the sake of the argument, imagine that now the planner observes the component $\mathbf{r}$ of the state (but not θ) before designing her plan, while only the doer continues to observe the whole state $(\theta, \mathbf{r})$. We can then examine the planner's problem defined by (1) and (2) treating $\mathbf{r}$ as fixed. Thus, hereafter we will omit $\mathbf{r}$ and let G denote the distribution of $\theta \in [\underline{\theta}, \overline{\theta}]$. For this section, we will assume that G has a density function g which is strictly positive and continuous on $[\underline{\theta}, \overline{\theta}]$. We will proceed in two steps. The first shows that, in this setting, it is possible to focus on commitment plans which regulate only savings and

range of $(\mathbf{c}^p, s^p)$ itself). First, such a plan would induce the doer to implement the first best only if his present bias is sufficiently weak; the logic is the same as that of Proposition 1 in Amador et al. (2006). Second, in real settings, which probably involve many goods and states, commitment plans of this form can become very complicated and (perhaps for this reason) people do not seem to adopt them.

total consumption expenditures, but not how expenditures are divided between goods. Thus, the multidimensionality of consumption becomes irrelevant. Given this, the second step argues that plans which use only a savings floor are optimal.

The logic behind the first step is as follows. Recall that budgets curb the doer's tendency to undersave by forcing him to choose inefficient consumption bundles, thereby lowering the utility he can get from the income he does *not* save, $1-s$. However, another method to lower this utility is simply to not let the doer spend all of $1-s$. In the literature, this is called "money burning."[36] Burning part of $1-s$ and spending the rest efficiently can achieve any utility level obtained by spending $1-s$ inefficiently: If we let $u^*(y)$ be the indirect utility of spending $y \in [0,1]$, then for every $\mathbf{c} \in \mathbb{R}^2_+$ there exists $y \leq c_1 + c_2$ that yields $u^*(y) = u(\mathbf{c})$. When intratemporal trade-offs are uncertain, the way in which the planner wants to "punish" the doer for undersaving depends on the configuration of those trade-offs, holding fixed the actual level of savings. Instead without that uncertainty, there is only one optimal punishment and this punishment can always be achieved with money burning, provided that it can flexibly depend on the level of savings. This requires allowing for more general commitment plans than the simple budgeting plans $\mathcal{B}$. Formally, let F^{tc} be the set of feasible allocations defined in terms of total consumption y and savings s:

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

Given an arbitrary subset $D^{\text{tc}} \subset F^{\text{tc}}$, for each θ the doer maximizes $\theta u(\mathbf{c}) + \beta v(s)$ subject to $c_1 + c_2 \leq y$ and $(y,s) \in D^{\text{tc}}$.

Lemma 2. *Suppose information 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 uncertainty is only about the intertemporal trade-off, whether consumption involves one or multiple goods is irrelevant, as long as we allow for general commitment plans.

What is more remarkable is that, under some simple conditions, the multidimensionality of consumption is irrelevant even when the planner can use only budgeting plans—in fact, only minimum-savings rules. To derive this second step, note that since the constraint $c_1 + c_2 \leq y$ will always bind for the doer, by Lemma 2 the planner's problem becomes to choose $D^{\text{tc}} \subset F^{\text{tc}}$ so as to maximize

$$\int_{\underline{\theta}}^{\overline{\theta}} [\theta u^*(y(\theta)) + v(s(\theta))] g(\theta) d\theta$$

subject to

$$(y(\theta), s(\theta)) \in \underset{(y,s)\in D^{\text{tc}}}{\arg\max} \{\theta u^*(y) + \beta v(s)\}, \quad \theta \in [\underline{\theta}, \overline{\theta}]. \tag{3}$$

[36] One way to do this is to commit to some charitable donations and link their amount to the level of savings. For such donations to qualify as "money burning," they should not be an argument of the individual's utility function (for example, through altruism). Besides Amador et al. (2006), papers that study money burning in delegation problems include Ambrus and Egorov (2009), Ambrus and Egorov (2013), Amador and Bagwell (2013a), Amador and Bagwell (2013b).

This problem coincides with that studied by Amador et al. (2006). To better understand its solution, it is helpful to follow their analysis and rewrite the problem into an equivalent formulation which (up to an additive scalar that is not essential at the moment) states the planner's objective as

$$\int_{\underline{\theta}}^{\overline{\theta}} H(\theta)u^*(y(\theta))d\theta,$$

where

$$H(\theta) = 1 - G(\theta) - (1-\beta)\theta g(\theta), \quad \theta \in [\underline{\theta}, \overline{\theta}].$$

To see what $H(\theta)$ captures, ignore feasibility and suppose that the planner lets the doer spend a bit more in state θ. To keep his payoff unchanged in θ—so that he does not pick another allocation—she also has to make him save a bit less. Overall this change harms the planner when θ occurs, because she cares more about savings. This explains the term $-(1-\beta)\theta g(\theta)$. The new allocation in θ becomes more attractive for the doer when he values present utility more, that is, in all states $\theta' > \theta$ whose mass is $1 - G(\theta)$. To keep his behavior unchanged in those states, the planner can induce him to save more, which is exactly what she wants. This explains the term $1 - G(\theta)$.

The solution to the planner's problem hinges on the following condition, where θ^* is defined as

$$\theta^* = \min\left\{\theta \in [\underline{\theta}, \overline{\theta}] : \textstyle\int_{\theta'}^{\overline{\theta}} H(\hat{\theta})d\hat{\theta} \leq 0 \text{ for all } \theta' \geq \theta\right\}. \tag{4}$$

Condition 1. The function H is non-increasing over $[\underline{\theta}, \theta^*]$.

Also, let

$$D^{\mathrm{tc}}(\theta^*) = \{(y,s) \in F^{\mathrm{tc}} : s \geq s^d(\theta^*)\},$$

which is essentially a plan that involves only the savings floor $f = s^d(\theta^*)$.

Proposition 5 (Amador et al. (2006))**.** *The plan $D^{\mathrm{tc}}(\theta^*)$ is optimal among* all *subsets of F^{tc} if and only if Condition 1 holds.*

As Amador et al. observed, many distributions—especially those commonly used in applications—satisfy Condition 1 for all $\beta \in [0,1]$. More generally, if the density g is uniformly bounded away from 0 and changes at a bounded rate, then Condition 1 always holds when the doer's bias is sufficiently weak (high β).[37] It is worth noting that, by Proposition 1, weak biases characterize those environments with uncertain intratemporal trade-offs where good-specific budgets improve on plans using only a savings floor.

For the sake of completeness, consider briefly the case in which Condition 1 fails. Intuitively this happens if, for instance, the density g suddenly drops over some intermediate interval (θ_1, θ_2) of states. Amador et al. (2006) argue that in this case the planner may have to rely on money burning: Roughly speaking, this deters the doer from undersaving when $\theta < \theta_1$ by pretending that θ is in (θ_1, θ_2). Proposition 7 in the Online Appendix shows that the planner can exploit the multidimensionality of consumption to replace money burning (sometimes entirely) with rules that force the doer to choose inefficient bundles. This may contribute to explaining why in reality we see less money

[37]Condition 1 is not necessary for the optimal $B \in \mathcal{B}$ to involve only a floor f because plans that improve on $D^{\mathrm{tc}}(\theta^*)$ may lie outside $\mathcal{B}$ (cf Amador et al. (2006) for an example).

burning than we might expect. Distortionary restrictions on consumption can entirely replace money burning if, for instance, zero consumption of some good is very inefficient and leads to a sufficiently low utility. Examples of such goods may include one's favorite drink or food, or going out with friends. One way to implement the necessary distortions in consumption is again to use good-specific budgets, which however may now have to vary based on how much the doer saves.

In summary, this section shows that uncertainty about intratemporal trade-offs plays a crucial part in the explanation of why people may find good-specific budgets useful commitment tools. This together with Corollary 1 suggests that the phenomenon of budgeting should be more prominent in settings where that type of uncertainty is particularly strong. Perhaps, this is the case for younger households with children in the early stages of their life, than for older households who settled down and whose children left home. Another relevant dimension may be the time horizon of reference: A month may entail more uncertainty than a week and this may lead people to set monthly, but not weekly, budgets.

6 Extensions

6.1 Non-Additively-Separable Utilities

To stress the role of the resource constraint in linking dimensions of the doer's choice and to focus on the core message of the paper, interactions between goods at the preference level were ruled out by assuming an additive consumption utility. This assumption can be relaxed without changing that message. Continue to assume that $u(\mathbf{c};\mathbf{r})$ is strictly concave in $\mathbf{c}$ for all $\mathbf{r}$ and twice differentiable with continuous $u_{c_i}(\mathbf{c};\mathbf{r}) > 0$ and $u_{c_i c_j}(\mathbf{c};\mathbf{r})$ in both arguments for all i and j. We saw that good-specific budgets help the planner curb the doer's undersaving if (a) they increase savings and (b) there exist states which 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 may exacerbate overspending on other goods, for it to benefit the planner. Given space constraints, we state these properties directly in terms of allocations.

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

$$s^*(\omega) - s^d(\omega) \geq \varepsilon[c_j^d(\omega) - c_j^*(\omega)], \quad \text{for all } \omega \in \Omega.$$

The Online Appendix presents an example which satisfies Condition 2.

To state the next result, consider a more general class of budgeting plans, denoted by $\overline{\mathcal{B}}$, which allow the planner to also set good-specific floors and a savings cap: Such a plan is defined as

$$\overline{B} = \{(\mathbf{c}, s) \in F : f_0 \le s \le b_0, f_1 \le c_1 \le b_1, f_2 \le c_2 \le b_2\}, \tag{5}$$

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

Proposition 6. *Under Condition 2, there exists $\beta^* \in (0,1)$ such that, if $\beta > \beta^*$, then every optimal $\overline{B} \in \overline{\mathcal{B}}$ must use distorting good-specific restrictions.*[38]

The proof is omitted, because using Condition 2, one can adapt the proof of Lemma 6 and Proposition 1 to show that plans which use only f_0 are strictly dominated for sufficiently weak biases. Since imposing a binding b_0 is never optimal, the result follows.

Do optimal good-specific restrictions always take the form of budgets? In general, the answer depends on the substitutability and complementarity between goods and between each good and savings, which can be affected by the plan restrictions themselves. A sufficient condition for optimal plans to never use good-specific floors is that all goods are Hicks substitutes and collectively sufficiently normal (see Ellis and Naughton (1990) for a formal statement of this property). Under this condition, 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 set of binding good-specific floors, relaxing them increases savings. Hence, since those floors distort consumption, they strictly harm the planner. The property that optimal plans can involve only good-specific budgets holds for the example presented in the Online Appendix.

6.2 General Commitment Plans

Motivated by the phenomenon of personal budgeting, the paper has focused on the class of commitment plans $\mathcal{B}$. From a normative perspective, one may wonder how the best among all conceivable plans looks like and whether, perhaps in some cases, it belongs to $\mathcal{B}$. Answering these questions is obviously important; unfortunately, it is also hard in the presence of multidimensional consumption and information. In general, a commitment plan can be any $D \subset F$ from which the doer is allowed to choose an allocation. The problem is then to find a $D \subset F$ that maximizes the planner's expected utility from the doer's resulting choices. The usual mechanism-design approach is to turn this problem into an equivalent problem of finding an optimal direct mechanism, which consists of an incentive-compatible and resource-feasible allocation function of the state. The latter problem is usually easier, but in the present setting, it remains intractable.

The main challenges come from the combination of the income constraint with the complexity of the incentive constraints. It is well-known that, with multidimensional types, one cannot focus on local incentive constraints—even if the doer cannot benefit from misreporting his information locally, he may benefit from global misreports to the

[38] Existence of an optimal $\overline{B}$ can be established along the lines of the proof of Lemma 1.

planner's mechanism. One can try to apply the insights of the literature on multidimensional screening (cf Rochet and Stole (2003) for a survey) to the present problem, but substantive differences preclude this. First, in screening problems the mechanism designer can use transfers. Here, one can view the expected continuation utility from savings as a transfer and use Rochet and Choné's (1998) dual approach to simplify the incentive constraints and the planner's objective. But the state-wise income constraint, the second key difference from screening problems, cannot be simplified. General techniques exist for handling such constraints (for example, Luenberger (1969)). Unlike in the case of unidimensional consumption (Amador et al. (2006)), however, here those techniques do not help to characterize the optimal mechanisms, which requires to jointly determine the mechanism and the state-wise Lagrange multiplier for the income constraint. The solution need not follow the logic of the unidimensional case or of multidimensional screening, for which we know that solutions rarely exist in closed form and tend to have quite intricate structures even for simple settings.

Despite the loss of generality, the class of "interval plans" $\mathcal{B}$ (or $\overline{\mathcal{B}}$) remains of interest for several reasons. First, $\mathcal{B}$ is not fully general even in unidimensional settings, but in these settings, under weak conditions there exist optimal $D \subset F$ which belong to $\mathcal{B}$ (Amador et al. (2006)).[39] Hence, $\mathcal{B}$ represents a natural starting point to analyze multidimensional settings. Second, in his seminal work on delegation problems—which formally include the planner's problem—Holmström (1977) noted that "one might want to restrict D to [...] only certain simple forms of [policies], due to costs of using other and more complicated forms or due to the fact that the delegation problem is too hard to solve in general." In his unidimensional settings, Holmström (1977) focused on interval policies, because they "are simple to use with minimal amount of information and monitoring needed to enforce them" and "are widely used in practice." In a similar spirit, discussing multidimensional delegation, Armstrong (1995) acknowledged that "in order to gain tractable results it may be that *ad hoc* families of sets such as rectangles or circles would need to be considered" (p. 20, emphasis in the original). Finally, for consumption-savings problems, Thaler and Shefrin (1981) argued that commitment "rules by nature must be simple." Simplicity is a property of budgeting plans, which also makes it easier for people to stick to them[40] and for third-party providers to market devices that implement them.

It may be worth considering other subclasses of commitment plans. The task, however, will be to identify classes which not only are sensible, but also have enough structure to render the search of an optimal plan tractable. For instance, one may consider commitment plans defined by a binding boundary that smooths the kinks created by the budgets and the savings floor in Figure 2(b). However, it is not obvious how to mathematically describe such a class and optimize over it. It is also not immediate that this class contains the overall optimal plan, because such a plan may involve interior bunching regions as it often happens in multidimensional screening problems (Rochet and Choné

[39] Alonso and Matouschek (2008) and Amador and Bagwell (2013b) provide conditions for interval policies to be globally optimal in other delegation settings with unidimensional decisions and information.

[40] Benhabib and Bisin (2005) provide a rationale for why people may prefer simple commitment rules, based on higher psychological costs of complying with complex rules.

(1998)). Finally, it may be hard for individuals to implement such complex plans, even with the help of third-party commitment devices.

6.3 Multiple Goods in Each Period

It is straightforward to generalize the model of Section 3 so that the individual consumes a bundle of multiple goods in each period. Suppose that the function $u(\mathbf{c};\mathbf{r})$ now represents the consumption utility of both the planner and the doer in each period, where $\mathbf{r}$ continues to represent the information affecting the intratemporal trade-offs. A state now includes the realization of $\mathbf{r}$ in period 2 and so is given by $(\theta, \mathbf{r}^1, \mathbf{r}^2)$, where the superscripts refer to the time period. In period 1, for each state the planner and the doer evaluate streams $(\mathbf{c}^1, \mathbf{c}^2)$ using respectively the utility functions

$$\theta u(\mathbf{c}^1;\mathbf{r}^1) + u(\mathbf{c}^2;\mathbf{r}^2) \quad \text{and} \quad \theta u(\mathbf{c}^1;\mathbf{r}^1) + \beta u(\mathbf{c}^2;\mathbf{r}^2).$$

The set of feasible allocations in period 1 and in period 2 are

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

and

$$F(s) = \{\mathbf{c}^2 \in \mathbb{R}^2_+ : c_1 + c_2 \leq s\}, \quad s \in [0,1].$$

The family of budgeting plans for period 1, $\mathcal{B}^1$, is the same as $\mathcal{B}$ defined in Section 3. For period 2, a budgeting plan is defined as

$$B^2 = \{\mathbf{c} \in F(s) : c_1 \leq b_1, c_2 \leq b_2\},$$

where $b_1, b_2 \in [0,1]$.

As before, only the doer observes the state in each period. The timing changes as follows. In period 1, first the planner commits to a plan B^1, then the doer observes $(\theta, \mathbf{r}^1)$ and implements some allocation from B^1. Given the chosen s, in period 2, first the planner commits to a plan B^2, then the doer observes $\mathbf{r}^2$ and implements some allocation from B^2. As Halac and Yared (2014) showed, information persistency over time plays an important role in consumption-savings problems with imperfect self-control. However, to keep things simple, assume that here information is not persistent: $\mathbf{r}^2$ is independently distributed from $(\theta, \mathbf{r}^1)$ according to the full-support probability measure G^2 over the space $\Omega^2 = [\underline{r}_1, \overline{r}_1] \times [\underline{r}_2, \overline{r}_2]$. Note that the timing involves the assumption that the individual has the power to commit within a period, but not across periods. This assumption is not innocuous, even without information persistency: By committing in period 1 to how B^2 will depend on the chosen $(\mathbf{c}^1, s)$, the planner may be able to manage the doer's bias in more effective ways. We will come back to this in Section 6.4.

Under these assumptions, it is easy to see that the planner will grant the doer full flexibility in the last period. The reason is twofold: First, in period 2 both selves have the same preferences; second, in period 1 the planner cannot commit to curtailing the doer's flexibility in period 2, so as to reward or punish his decisions in period 1. Formally, for every s inherited from period 1, the plan $B^2 = F(s)$ is optimal in period 2. Because of this, we can define the period-1 utility from savings $v : [0,1] \to \mathbb{R}$ by

$$v(s) = \int_{\Omega^2} \left[\max_{\mathbf{c} \in F(s)} u(\mathbf{c};\mathbf{r}) \right] dG^2(\mathbf{r}). \tag{6}$$

This function satisfies some useful properties—summarized in Lemma 10 in the Online Appendix—which lead to the same analysis as in the simpler model used throughout the paper.

6.4 Multi-Period Commitments

Consider now a model in which consumption involves multiple goods in each period (as in Section 6.3) and information is not persistent, but the planner can commit in period 1 to the plans she will impose in period 2. In principle, by linking her period-2 plan to the doer's entire choice in period 1, $(\mathbf{c}^1, s)$, the planner could use future punishments and rewards to better discipline him in the present; moreover, she could do so in ways that are suboptimal ex post.

For such a model, one-period commitments entail no loss of generality if in each period plans can be any set D of feasible allocations. This depends on two things (cf also Amador et al. (2003)): (a) such plans can flexibly regulate how much the doer is allowed to save, possibly via money burning; (b) information is not persistent over time. To see this, imagine that for some $(\mathbf{c}^1, s)$ the planner sets $D(\mathbf{c}^1, s) \subsetneq F(s)$ to punish the doer in period 1. By (b), both selves assign the same expected payoff to $D(\mathbf{c}^1, s)$, which lies in the interval $[u(\mathbf{0}), v(s)]$. Therefore, the planner can replace every $(\mathbf{c}^1, s)$ in her period-1 plan with $(\mathbf{c}^1, s')$ and $D(\mathbf{c}^1, s)$ with $F(s')$ for some $s' \leq s$—by (a)—so that she and the doer get the same payoff from $(\mathbf{c}^1, s)$ and $(\mathbf{c}^1, s')$ for every state; it is then optimal for the doer to pick $(\mathbf{c}^1, s')$ whenever he was choosing the corresponding $(\mathbf{c}^1, s)$, thereby producing the same overall payoff for the planner. This is because the doer's expected payoff from $D(\mathbf{c}^1, s)$ does not depend on his period-1 information—by (b)—and hence cannot be used, *in addition to* $\mathbf{c}^1$ itself, to incentivize him to pick $(\mathbf{c}^1, s)$ in some states but not in others. This point about one-period commitments is related to the finding in models with unidimensional consumption that the ex-ante and sequentially optimal plans coincide under non-persistent information (Amador et al. (2003, 2006)), but differ under persistent information (Halac and Yared (2014)). Lack of information persistency, however, is only part of the story here; property (a) also plays a key role, as it gives great flexibility in managing continuation payoffs using one-period commitments. For instance, Breig (2015) studies another class of dynamic delegation problems with non-persistent information where ex-ante optimal mechanisms differ from the sequentially optimal ones.

If the planner can use only budgeting plans in $\mathcal{B}$, the inability to commit for future periods may prevent her from achieving higher payoffs. This is because $\mathcal{B}$ violates property (a). Note, however, that in settings where information is not persistent and affects only the intertemporal trade-off, weak conditions ensure that the sequentially optimal, unrestricted, plans D take the form of a minimum-savings rule in each period.[41] Therefore, the main messages of the paper regarding good-specific budgets is unaffected. Furthermore, if in the present model the ex-ante optimal budgeting plans differed from the sequentially optimal ones, the substantive cause would again be the uncertainty about

[41] Amador et al. (2003) show this for the case of unidimensional consumption and the result can be extended to the case of multidimensional consumption using the ideas in Section 5.

intratemporal trade-offs.

6.5 Multiple Future Periods

Not surprisingly, extending the model beyond two periods adds intricacies to the analysis, which however do not distinguish this paper from the literature. The following are the most substantive.

When the future has multiple periods, the expected continuation utility from saving, v, depends on the future commitment plans as well as the intensity of present bias, β. In each period (except possibly the last one), the planner commits to plans that curtail the doer's discretion and hence influence the value of entering that period with savings s. Such a value also depends on β, as β affects the chosen plan as well as the doer's behavior. The dependence on β of the marginal value of having one extra dollar next period ultimately affects the planner's desire to curb undersaving in this period: A stronger bias induces the doer to undersave more in the present, thereby strengthening that desire; at the same time, it may also reduce the marginal value of savings—due to worse future allocations—thereby weakening that desire.

Nonetheless, optimal budgeting plans should continue to involve good-specific budgets for sufficiently weak biases, but only a savings floor f for sufficiently strong biases. Consider any intermediate period. Intuitively, for β close to 1 the planner should again want to set a very slack f when this is the only tool available. Such an f would again grant the doer full discretion in states that call for high consumption of only some good. In this case, setting a good-specific budget strictly improves on using only f for the same reasons highlighted above. On the other hand, for β close to zero budgets cannot increase savings but distort consumption; hence, plans which use budgets are strictly dominated by those which use only f.

A second issue, discussed also by Amador et al. (2003), is that an infinite-horizon model may have multiple equilibria of the game between the individual's selves. This issue is well-known and orthogonal to the message of the present paper. The point remains valid that distorting restrictions on consumption (on top of savings restrictions) can be useful tools to discipline the doer. Such distortions should therefore appear as a novel feature in at least some equilibria of the multi-self game.

7 Other Applications

The insights offered by the planner-doer problem can be applied to other settings. For these applications the theory can assume a more normative flavor, while keeping in mind the limitations of the simple class of budgeting policies considered in this paper.

Public Finance. Consider first a public-finance application along the lines of Halac and Yared (2014). In each period, a government chooses how much to spend on a list of public goods and services, $\mathbf{c}$, and how much to save or borrow, 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 due to uncertainty in the political turnover (Aguiar and Amador (2011)). For this setting, the theory developed in Section 4 can explain why governments often set limits on borrowing (for example, via budget-deficit ceilings), but also specific limits on some good or service expenses (for example, via fiscal budgets). The theory shows when and why combining deficit ceilings with fiscal budgets can dominate policies that rely only on the ceilings. At the same time, it shows that fiscal budgets are no free lunch: They can cause inefficiencies in public spending, which however play a key role in curtailing overborrowing. Of course, fiscal budgets can just be part of accounting techniques. This theory, however, adds a different perspective on the functions that such budgets can have.

More generally, the problem studied in this paper can be viewed as an instance of a multidimensional delegation problem where a principal delegates to a better-informed agent the allocation of a finite amount of resources (money or time) to multiple categories of expenditures. Also, it is not essential that the consequences of the agent's allocation happen at different dates: They can all happen within the same period, and the conflict of interest between the principal and the agent need not stem from time preferences. The following situations illustrate this.

Corporate Governance. The owner of a company appoints a manager, who each period decides how to allocate some total resources between spending on sales activities, y, and investment in R&D, x_0. The company sells two products and the manager also chooses which share of y goes to promoting which product (x_1 and x_2). Let $F = \{\mathbf{x} \in \mathbb{R}^3_+ : x_0+x_1+x_2 \leq 1\}$ be the set of all feasible allocations. The manager privately observes information on the returns to marketing each product and from investing in R&D. Let this information be represented by state $\omega = (\omega_0, \omega_1, \omega_2)$, where $0 < \underline{\omega}_i \leq \omega_i \leq \overline{\omega}_i < +\infty$ for all i. Finally, as a result of compensation schemes or career goals, the manager may care more or less about cash flows—and hence about sales vs. R&D—than does the owner.

We can express the owner's problem of delegating to the manager the allocation of resources as follows: Choose $D \subset F$ to maximize

$$\mathcal{U}(D) = \int_{\Omega} [\omega_0 u^0(x_0(\omega)) + \omega_1 u^1(x_1(\omega)) + \omega_2 u^2(x_2(\omega))] dG(\omega) \tag{7}$$

subject to

$$\mathbf{x}(\omega) \in \arg\max_{\mathbf{x}\in D} \left[\beta\omega_0 u^0(x_0) + \omega_1 u^1(x_1) + \omega_2 u^2(x_2)\right], \quad \omega \in \Omega. \tag{8}$$

The parameter $\beta > 0$ represents the conflict of interest between parties. Note that, as in the main model the planner and the doer disagreed on the importance of consuming vs. savings within each period, here the owner and the manager disagree regarding the importance of sales vs. R&D activities; but given R&D investment x_0, they agree on how to divide the remaining resources to promote each product.

To illustrate, suppose that the manager undervalues R&D (that is, $\beta < 1$) and that delegation policies are restricted to $\mathcal{B}$. This paper suggests that to best incentivize the manager, the owner may have to impose limits on how much can be spent each period on promoting specific products, possibly in addition to requiring a minimum investment in

R&D. Due to these limits, the manager may end up promoting some product too little and the other too much. This, however, is a risk the owner should take, as it is more than compensated in expectation by better allocations to R&D. A detailed budget plan with rules applying to specific products is more likely to benefit the owner when she agrees sufficiently with the manager on how important R&D is. This may be true, for instance, if the manager himself has significant stakes in the company. Otherwise, the owners should simply demand only a minimum investment in R&D.

Fiscal-constitution Design. Society delegates a government to divide the economy resources between private consumption, x_0, and public spending, y. The government incorporates the preference of a representative agent in society, but is biased in favor of public spending ($\beta < 1$).[42] This bias may depend on incentives created by political institutions, for instance. The government spends y to fund two services (x_1 and x_2). Despite its bias, at a first approximation it may not favor any service more than others, and thus agrees with society on how to allocate any level of public spending between services.[43] The government acts on non-contractible information, ω, which affects the social value of each service (for example, the gravity of national-security threats or natural disasters), as well as the overall trade-off between private consumption and public spending (for example, the state of the business cycle). In this case, behind a veil of ignorance society may want to design a fiscal constitution—a delegation policy D—that restricts which allocations the government can choose. This problem can again be expressed by (7) and (8).

Suppose society can choose only policies in $\mathcal{B}$. The analysis shows that if the government's bias is weak, then an optimal fiscal constitution must involve service-specific spending caps. Such caps work because they cause inefficiencies in the composition of public spending, thereby weakening the government's desire to spend. But from society's viewpoint, these inefficiencies are more than compensated by the resulting higher level of private consumption. On the other hand, if the government is strongly biased, the constitution should involve only an aggregate spending cap (or a private-consumption floor). This is because any binding specific cap distorts public spending without sufficiently improving private consumption.

Workforce Management. An employer (she) hires a worker (he) under a contract specifying a fixed wage and number of hours per workday. The worker performs two tasks and chooses how to allocate time between them (x_1 and x_2). He can also take breaks, represented by x_0. The worker privately observes information on which task demands more attention. Given this, the employer delegates to him his time allocation. However, the worker weighs his break time more than does the employer ($\beta > 1$). Thus, she may want to set up some rules to restrict his choices, modeled as a subset D of all feasible time allocations. Expressions (7) and (8) again capture the employer's problem.

Suppose that she can use only budgeting policies. An analysis similar to that above

[42] This hypothesis is supported by theoretical as well as empirical work in the political-economy literature (Niskanen (1975), Romer and Rosenthal (1979), Peltzman (1992), Funk and Gathmann (2011)).

[43] This property is arguably strong, but is consistent with some empirical evidence. For example, Peltzman (1992) finds that U. S. voters penalize federal spending growth, but its composition seems irrelevant, which suggests that "every extra dollar is equally bad."

leads to results in which caps and floors swap roles. A floor on task i induces the worker to allocate less time to x_0, but also less time to the other task. Nonetheless, if the worker's tendency to indulge in breaks is weak, an optimal policy must involve binding task-specific floors. By contrast, if the worker's bias is strong, the employer should impose only a cap b_0 on x_0. In practice, it may be impossible to monitor the worker's breaks; however, b_0 can be implemented via an overall minimum time that the worker has to allocate to his tasks, which may be easier to monitor. The theory shows that the possibility of monitoring each task *individually* may allow the employer to design strictly superior policies by adding individual floors.

8 Concluding Remarks

This paper provides a theoretical analysis of the relationship between self-control problems and personal budgeting using a simple consumption-savings model which introduces no "behavioral" or ad-hoc feature besides a standard form of present bias. Unlike minimum-savings rules, good-specific spending caps help to curtail overspending because they cause inefficiencies in consumption which lower the return from undersaving, thereby counteracting present bias. Consequently, good-specific budgets are no free lunch and are used only by consumers who are weakly biased and ex-ante 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 solid insights into the subtle forces underlying a widely observed phenomenon, which has far-reaching consequences for consumer behavior and welfare by affecting demand in different ways from satiation and income effects and by significantly contributing to households' wealth accumulation. The theory matches existing empirical findings, such as that many people set budgets for goods normally not viewed as temptations and that only people who exhibit weak present bias seem to use budgets. The theory also suggests new directions for enriching this limited evidence on personal budgeting by demonstrating its key dependence on uncertain intratemporal trade-offs, and for designing commitment devices whose functions are targeted to the right type of present-biased individuals.

From a normative perspective, the paper does not show that optimal commitment plans within the universe of all possible plans turn out to take the form of budgeting. This result seems very unlikely to hold and was not the goal of the paper to begin with. Technical difficulties aside, it seems more plausible that fully optimal plans are quite complicated and sensitive to details of the environment. This complexity may undermine an individual's ability to stick to a plan and hence should be traded off with its capacity of solving self-control problems. Analyzing this trade-off is beyond the scope of this paper. In practice, individuals may have settled for personal budgeting as a viable compromise based on robustness, learning, or some heuristics. Of course, this explanation is less satisfactory for the application of the theory to other, formally equivalent, delegation problems within organizations—which may, however, still value simplicity. Future research may benefit from the insights of this paper to build a general normative theory of how to delegate multidimensional allocations of finite resources.

9 Appendix: Proofs

9.1 Proof of Proposition 1

The proof relies on the following four lemmas. Hereafter, for $k = p, d$, let $\underline{s}^k = \min_\omega s^k(\omega)$ and $\overline{s}^k(\omega) = \max_\omega s^k(\omega)$.

For any floor $f \in [\underline{s}^d, \overline{s}^p]$,[44] denote by B_f the corresponding policy in $\mathcal{B}$.

Lemma 3. *Define* $\overline{\Omega}(f) = \{\omega \in \Omega : s^d(\omega) \leq f\}$ *and*

$$\mathbf{c}^f(\omega) = \underset{\{\mathbf{c}\in\mathbb{R}^n_+:\sum_{i=1}^n c_i \leq 1-f\}}{\arg\max} u(\mathbf{c}; \mathbf{r}), \quad \omega \in \Omega.$$

The payoff $\mathcal{U}(B_f)$ *is differentiable in* f *over the domain* $[\underline{s}^d, \overline{s}^p]$ *with*

$$\frac{d}{df}\mathcal{U}(B_f) = \int_{\overline{\Omega}(f)} \left[v'(f) - \theta \frac{\partial}{\partial c_i} u(\mathbf{c}^f(\omega); \mathbf{r})\right] dG,$$

for any $i = 1, \ldots, n$.

Proof. For simplicity, let $\Psi(f) = \mathcal{U}(B_f)$. Also, we will consider only $f \in [\underline{s}^d, \overline{s}^p]$ without repeating this every time. Given f and any ω, define

$$\tilde{u}(f; \omega) \equiv u(\mathbf{c}^f(\omega); \mathbf{r}) = \max_{\{\mathbf{c}\in\mathbb{R}^n_+:\sum_{i=1}^n c_i \leq 1-f\}} u(\mathbf{c}; \mathbf{r}). \tag{9}$$

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. This implies that $\tilde{U}(\cdot; \omega)$ is also strictly concave in f.

Now consider the derivative of $\tilde{U}(f; \omega)$ with respect to f. Whenever it is defined, we have

$$\tilde{U}'(f; \omega) = \theta\tilde{u}'(f; \omega) + v'(f).$$

The first-order conditions of the Lagrangian defining $\tilde{u}(f; \mathbf{r})$ say that $\frac{\partial}{\partial c_i} u(\mathbf{c}^f(\omega); \mathbf{r}) = \lambda(\omega; f)$ for $i = 1, \ldots, n$, where $\lambda(\omega; f)$ is the Lagrange multiplier for the constraint $\sum_{i=1}^n c_i \leq 1 - f$. Since $\mathbf{c}^f(\omega)$ is continuous in f for every ω, so is $\lambda(\omega; f)$ given our assumptions on u. By Theorem 1, p. 222, of Luenberger (1969), for every $f', f'' \in (0, 1)$ we have

$$\lambda(\omega; f')(f'' - f') \leq \tilde{u}(f'; \mathbf{r}) - \tilde{u}(f''; \mathbf{r}) \leq \lambda(\omega; f'')(f'' - f').$$

Continuity of $\lambda(\omega; \cdot)$ then implies that $\frac{\partial}{\partial f}\tilde{u}(f; \mathbf{r})$ exists for every $f \in (0, 1)$ and satisfies

$$\frac{\partial}{\partial f}\tilde{u}(f; \mathbf{r}) = -\lambda(\omega; f) = -\frac{\partial}{\partial c_i} u(\mathbf{c}^f(\omega); \mathbf{r}).$$

Therefore,

$$\tilde{U}'(f; \omega) = v'(f) - \theta\frac{\partial}{\partial c_i} u(\mathbf{c}^f(\omega); \mathbf{r}), \quad \omega \in \Omega. \tag{10}$$

[44] Any other floor is dominated by one in this range.

For any f, denote by $(\mathbf{c}^f, s^f)$ the doer's behavior as a function of ω under policy B_f. Note that $(\mathbf{c}^f(\omega), s^f(\omega))$ is continuous in both f and ω by the Maximum Theorem. Since, given any choice of s, the planner and the doer would choose the same $\mathbf{c}$ in every ω, by definition we have

$$\Psi(f) = \int_\Omega \tilde{U}(s^f(\omega); \omega) dG.$$

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

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

where the second equality follows because $s^f(\omega) = s^{\hat{f}}(\omega)$ for $\omega \notin \overline{\Omega}(f)$ and $s^f(\omega) = f$ for $\omega \in \overline{\Omega}(f)$. Dividing both sides by $f - \hat{f}$, we get

$$\begin{aligned}
\lim_{f \downarrow \hat{f}} \frac{\Psi(f) - \Psi(\hat{f})}{f - \hat{f}} &= \lim_{f \downarrow \hat{f}} \int_{\overline{\Omega}(\hat{f})} \frac{\tilde{U}(f; \omega) - \tilde{U}(\hat{f}; \omega)}{f - \hat{f}} dG \\
&+ \lim_{f \downarrow \hat{f}} \int_{\overline{\Omega}(f) \cap \left(\overline{\Omega}(\hat{f})\right)^c} \frac{\tilde{U}(f; \omega) - \tilde{U}(s^{\hat{f}}(\omega); \omega)}{f - \hat{f}} dG.
\end{aligned} \tag{11}$$

Consider the first limit. For all ω, we have

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

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

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

Since $\tilde{U}'(f; \omega)$ is continuous in ω and f as illustrated by (10),

$$\max_{\{(f,\omega) \in [\underline{s}^d, \overline{s}^p] \times \Omega\}} \left|\tilde{U}'(f; \omega)\right| = M < +\infty.$$

Therefore, by Lebesgue's Bounded Convergence Theorem, we have

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

Consider now the second limit in (11). 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 that

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

Therefore,

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

Now, observe that $\overline{\Omega}(f) \cap \left(\overline{\Omega}(\hat{f})\right)^c = \{\omega : \hat{f} < s^{\hat{f}}(\omega) \leq f\}$ which converges to an empty set as $f \downarrow \hat{f}$. It follows that the second limit in (11) converges to zero as $f \downarrow \hat{f}$. We conclude that for every $\hat{f} \in [\underline{s}^d, \overline{s}^p)$, we have

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

Now consider any $f < \hat{f}$. Then,

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

where the second equality follows because $s^f(\omega) = s^{\hat{f}}(\omega)$ for $\omega \notin \overline{\Omega}(\hat{f})$ and $s^{\hat{f}}(\omega) = \hat{f}$ for $\omega \in \overline{\Omega}(\hat{f})$, and the last equality follows because $s^f(\omega) = f$ for $\omega \in \overline{\Omega}(f)$. By the same argument as before,

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

$$\lim_{f \uparrow \hat{f}} \int_{\overline{\Omega}(\hat{f}) \cap \left(\overline{\Omega}(f)\right)^c} \frac{\tilde{U}(s^f(\omega);\omega) - \tilde{U}(f;\omega)}{f - \hat{f}} dG = 0.$$

We conclude that for every $\hat{f} \in (\underline{s}^d, \overline{s}^p]$, we have

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

Hence, $\Psi(f)$ is differentiable over the restricted domain $[\underline{s}^d, \overline{s}^p]$.

□

We now consider how the planner would use only the floor f.

Lemma 4. *When the planner can set only a floor on s, every optimal f lies strictly between $\underline{s}^p$ and $\overline{s}^p$.*

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 $\Psi(f)$ is continuous in f. These observations will imply that every optimal f^* is in $(\underline{s}^p, \overline{s}^p)$.

For any $f \in (\underline{s}^d, \overline{s}^p]$, define

$$\Omega^+(f) = \{\omega : s^p(\omega) > f\} \quad \text{and} \quad \Omega^-(f) = \{\omega : s^p(\omega) \leq f\}.$$

For $\omega \in \Omega^+(f)$, consider the the following fictitious problem in which the planner is forced to save less than f:

$$\max_{(\mathbf{c},s)\in\mathbb{R}^{n+1}_+} \{\theta u(\mathbf{c};\mathbf{r}) + v(s)\}$$

subject to $s + \sum_i c_i \leq 1$ and $s \leq f$. The associated Lagrangian is

$$\theta u(\mathbf{c};\mathbf{r}) + v(s) + \mu(\omega)\left[1 - s - \sum_{i=1}^{n} c_i\right] + \phi^+(\omega)[f - s].$$

Hence, the first-order conditions are[45]

$$v'(s) = \mu(\omega) + \phi^+(\omega) \quad \text{and} \quad \theta\frac{\partial}{\partial c_i}u(\mathbf{c};\mathbf{r}) = \mu(\omega) \quad \text{for all } i.$$

Clearly, the constraint $s \leq f$ must bind for $\omega \in \Omega^+(f)$, which implies that $s = f$ and $\phi^+(\omega) > 0$. Also, conditional on choosing $s = f$, the planner and the doer would choose the same $\mathbf{c}$ in state ω, which therefore equals $\mathbf{c}^f(\omega)$. Using (10), it follows that, for every i,

$$\phi^+(\omega) = v'(f) - \theta\frac{\partial}{\partial c_i}u(\mathbf{c}^f(\omega);\omega) = \tilde{U}'(f;\omega) \tag{12}$$

when $\omega \in \Omega^+(f)$.

For $\omega \in \Omega^-(f)$, consider the following fictitious problem in which the planner is forced to save more than f:

$$\max_{(\mathbf{c},s)\in\mathbb{R}^{n+1}_+} \{\theta u(\mathbf{c};\mathbf{r}) + v(s)\}$$

subject to $s + \sum_i c_i \leq 1$ and $s \geq f$. Again, the associated Lagrangian is

$$\theta u(\mathbf{c};\mathbf{r}) + v(s) + \mu(\omega)\left[1 - s - \sum_{i=1}^{n} c_i\right] + \phi^-(\omega)[s - f].$$

Hence, the first-order conditions are

$$v'(s) = \mu(\omega) - \phi^-(\omega) \quad \text{and} \quad \theta\frac{\partial}{\partial c_i}u(\mathbf{c};\mathbf{r}) = \mu(\omega) \quad \text{for all } i,$$

Clearly, the constraint $s \geq f$ must bind for $\omega \in \Omega^-(f)$ except when $s^p(\omega) = f$, which implies that $s = f$ and $\phi^-(\omega) \geq 0$. Also, conditional on choosing $s = f$, the planner and the doer would choose the same $\mathbf{c}$ in state ω, which therefore equals $\mathbf{c}^f(\omega)$. Using (10), it follows that, for every i,

$$\phi^-(\omega) = \theta\frac{\partial}{\partial c_i}u(\mathbf{c}^f(\omega);\mathbf{r}) - v'(f) = -\tilde{U}'(f;\omega)$$

[45]Here, as well as in the other proofs, the complementary slackness conditions are omitted for simplicity.

when $\omega \in \Omega^-(f)$.

Consider any $f \in (\underline{s}^d, \overline{s}^p]$. Recall that $\overline{\Omega}(f) = \{\omega : s^d(\omega) \leq f\}$. Using Lemma 3, we have

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

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 this set has strictly positive measure, which implies $\Psi'(f) > 0$. This is immediate if $f \in (\underline{s}^d, \underline{s}^p)$, because in this case $\Omega^+(f) = \Omega$. 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 we can show that this set is non-empty, we are done because G assigns strictly positive probability to it. Both $\overline{\Omega}^\circ(\underline{s}^p)$ and $\Omega^+(\underline{s}^p)$ are nonempty. Suppose that there is no $\omega \in \Omega^+(\underline{s}^p)$ such that we also have $\omega \in \overline{\Omega}^\circ(\underline{s}^p)$. Then, it means that 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

$$\lim_{\omega_n \to \hat{\omega}} \inf s^d(\omega_n) \geq \underline{s}^p > s^d(\hat{\omega}).$$

But this violates the continuity of s^d and hence leads to a contradiction.

Now consider $f = \overline{s}^p$. Using again Lemma 3, we have

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

where $\phi^-(\omega) > 0$ for all ω such that $s^p(\omega) < \overline{s}^p$. Therefore, $\Psi'(\overline{s}^p-) < 0$.[46]

□

We now consider how an optimal f varies with β. Recall that $u^*(y;\mathbf{r})$ is the indirect utility of spending $y \in [0,1]$ on consumption.

Lemma 5. *The set of optimal floors, denoted by $E(\beta)$, is decreasing in β in the strong set order.*[47] *In particular, the largest optimal floor converges monotonically to $\underline{s}^p$ as $\beta \uparrow 1$. Moreover, there exists $\underline{\beta} > 0$ such that $E(\beta) = \{\overline{f}\}$ for all $\beta \leq \underline{\beta}$, where $\overline{f}$ satisfies $\overline{f} < \overline{s}^p$ and*

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

[46] 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, the doer 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 the planner's payoff is strictly larger under f' than under f.

[47] Given two sets E and E' in $\mathbb{R}$, $E \geq E'$ 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)).

Proof. Fix $f \in [\underline{s}^d, \overline{s}^p]$. The set $\overline{\Omega}(f)$ in Lemma 3 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 $\overline{\Omega}(f; \beta') \subset \overline{\Omega}(f; \beta)$. On the other hand, for every $\beta < 1$, we have $\Omega^-(f) \subset \overline{\Omega}(f; \beta)$ because $s^d(\omega; \beta) < s^p(\omega)$ for every ω. So, if $\beta < \beta' < 1$, we have

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

where the inequality follows from (12). Standard monotone-comparative-static results then imply that $E(b)$ is increasing in the strong set order.

Define $\overline{f}(\beta) = \max\{f : f \in E(\beta)\}$. Since $\overline{f}(\beta) \geq \underline{s}^p$ for all β and $\overline{f}(\cdot)$ is decreasing, $\lim_{\beta\uparrow 1} \overline{f}(\beta)$ exists; denote it by $\overline{f}(1-) \geq \underline{s}^p$. 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 satisfies

$$\lim_{\beta\uparrow 1} \Psi'(f; \beta) = -\int_{\Omega^-(f)} \phi(\omega) dG < 0.$$

This implies that for β close enough to 1, $\overline{f}(\beta) \geq \overline{f}(1-)$ cannot be optimal, a contradiction which implies that $\overline{f}(1-) = \overline{f}(1)$.

It is easy to see that, for all $\omega \in \Omega$, $s^d(\omega; \beta) \to 0$ as $\beta \downarrow 0$. Therefore, $\overline{s}^d(\beta) = \max_\Omega s^d(\mathbf{s}; \beta)$ also decreases monotonically to 0 as $\beta \downarrow 0$. Let $\underline{\beta} = \max\{\beta \in [0, 1] : \overline{s}^d(\beta) \leq \underline{s}^p\}$ which is strictly positive because $\underline{s}^p > 0$. Then, $\overline{\Omega}(f) = \Omega$ for all $\beta \leq \underline{\beta}$ and $f \in [\underline{s}^p, \overline{s}^p]$, which implies that

$$\Psi(f; \beta) = \int_\Omega [\theta u(\mathbf{c}^f(\omega); \mathbf{r}) + v(f)] dG. \tag{13}$$

From the proof of Lemma 3, we have that $u(\mathbf{c}^f(\omega); \mathbf{r}) = \tilde{u}(f; \omega)$ is strictly concave in f for all $\omega \in \Omega$. This implies that the maximizer of (13) is unique. From the proof of Lemma 4, we know that the derivative of (13) is negative at $\overline{s}^p$ and hence $\overline{f} < \overline{s}^p$.

□

We now show that the planner benefits from imposing only one budget on any good.

Lemma 6. *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 the planner strictly benefits from it, that is, $\mathcal{U}(B_{b_i}) > \mathcal{U}(F)$.*

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

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

Let $\Omega(b_1) = \{\omega : c_1^d(\omega) > b_1\}$. Note that for any $b_1 < \overline{c}_1^d$, since c_1^d is continuous, $\Omega(c_1)$ is non-empty and open, and hence has strictly positive probability under G. We have

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

where

$$\tilde{V}(\hat{b}_1;\omega) = \max_{\{(\mathbf{c},s)\in\mathbb{R}_+^{n+1}:\sum_{j=1}^n c_j\le 1, c_1\le \hat{b}_1\}} \{\theta u(\mathbf{c};\mathbf{r}) + \beta v(s)\}.$$

Clearly, $\tilde{V}(c_1^d(\omega);\omega) \ge \tilde{V}(b_1;\omega)$ for every 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 on the constraint $c_1 \le \hat{b}_1$. Since $(\mathbf{c}^{\hat{b}_1}(\omega), s^{\hat{b}_1}(\omega))$ is continuous in $\hat{b}_1$ as well as ω, so is $\lambda_1(\omega;\hat{b}_1)$. Relying again on Theorem 1, p. 222, of Luenberger (1969), we conclude that $\tilde{V}'(\hat{b}_1;\omega)$ exists for every $\hat{b}_1$ and equals $\lambda_1(\omega;\hat{b}_1)$. It follows that $\tilde{V}'(c_1^d(\omega);\omega) = 0$ for every ω by the definition of $(\mathbf{c}^d, s^d)$. Therefore, by the Mean Value Theorem (MVT),

$$\tilde{V}(c_1^{b_1}(\omega);\omega) - \tilde{V}(c_1^d(\omega);\omega) = \tilde{V}'(\chi(\omega);\omega)(c_1^{b_1}(\omega) - c_1^d(\omega)),$$

$$v(s^{b_1}(\omega)) - v(s^d(\omega)) \;=\; v'(\xi(\omega))(s^{b_1}(\omega) - s^d(\omega)),$$

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 small $\varepsilon > 0$. Fix any $\omega \in \Omega(b_1^\varepsilon)$ and, for now, suppress the dependence on ω for simplicity. 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 any $\varepsilon > 0$ (cf Lemma 9), we can write

$$-\frac{c_1^{b_1^\varepsilon} - c_1^d}{s^{b_1^\varepsilon} - s^d} = 1 + \sum_{j\ne 1} \frac{c_j^{b_1^\varepsilon} - c_j^d}{s^{b_1^\varepsilon} - s^d}. \tag{14}$$

Now, for any b_1^ε, the following first order condition must hold for every $j \ne 1$:

$$\beta v'(s) - \theta u_c^j(c_j;r_j) = 0.$$

Therefore, using again the MVT, for all $j \ne 1$ we have

$$c_j^{b_1^\varepsilon} - c_j^d = \frac{\beta[v'(s^{b_1^\varepsilon}) - v'(s^d)]}{\theta u_{cc}^j(\zeta_j;r_j)} \tag{15}$$

for some $\zeta \in [c_j^d, c_j^{b_1^\varepsilon}]$. Now, since v'' is continuous, we have that $v'(y) - v'(\hat{y}) \ge \underline{v}''[y - \hat{y}]$ for every $y > \hat{y} \ge \underline{s}^d$, where $\underline{v}'' = \min_{\xi\in[\underline{s}^d,1]} v''(\xi) < 0$. Therefore, using (14) and (15), we have that

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

where the first inequality uses the fact that $u_{cc}^j < 0$ and $\overline{u}_{cc}^j = \max_{\xi\in[\underline{s}^d,1], r_j\in[\underline{r}_j,\overline{r}_j]} u_{cc}^j(\xi;r_j) < 0$ (which is well defined by continuity of u_{cc}^j). Thus, letting $K = \left[1 + \frac{\beta\underline{v}''}{\underline{\theta}} \sum_{j\ne 1} \frac{1}{\overline{u}_{cc}^j}\right]^{-1} > 0$, it follows that for every $\omega \in \Omega(b_1^\varepsilon)$,

$$s^{b_1^\varepsilon}(\omega) - s^d(\omega) \ge K\left[c_1^d(\omega) - c_1^{b_1^\varepsilon}(\omega)\right].$$

Using these observations, we have that $\Phi(b_1^\varepsilon) - \Phi(\overline{c}_1^d)$ is bounded below by

$$\int_{\Omega(b_1^\varepsilon)} \left[K(1-\beta)v'(\xi(\omega)) - \tilde{V}'(\chi(\omega);\omega)\right](c_1^d(\omega) - b_1^\varepsilon)dG. \tag{16}$$

Since v' is continuous and strictly positive everywhere and $\xi(\omega) \in [\underline{s}^d, 1]$ with $\underline{s}^d > 0$ for all $\omega \in \Omega(b_1^\varepsilon)$, there exists a finite $\kappa > 0$ such that $v'(\xi(\omega)) \geq \kappa$ for all $\omega \in \Omega(b_1^\varepsilon)$.

Next let $\overline{\Omega}(b_1^\varepsilon) = \{\omega : c_1^d(\omega) \geq b_1^\varepsilon\}$ which is a closed and bounded set by continuity of c_1^d and hence is compact. As a function of b_1^ε, the correspondence $\overline{\Omega}(\cdot)$ is continuous by continuity of c_1^d. Note that, if $c_1^d(\omega) = b_1^\varepsilon$, then $\tilde{V}'(\chi(\omega);\omega) = \tilde{V}'(c_1^d(\omega);\omega) = 0$. We have

$$\sup_{\omega\in\Omega(b_1^\varepsilon)} \tilde{V}'(\chi(\omega);\omega) \quad = \quad \sup_{\omega\in\overline{\Omega}(b_1^\varepsilon)} \tilde{V}'(\chi(\omega);\omega) \leq \max_{b_1^\varepsilon \leq \zeta \leq \overline{c}_1^d, \omega\in\overline{\Omega}(b_1^\varepsilon)} \tilde{V}'(\zeta;\omega) \equiv \kappa(b_1^\varepsilon).$$

Clearly, $\kappa(b_1^\varepsilon) \geq 0$ for every $\varepsilon > 0$, $\varepsilon' > \varepsilon > 0$ implies that $\kappa(b_1^\varepsilon) \leq \kappa(b_1^{\varepsilon'})$, and $\lim_{\varepsilon\to 0}\kappa(b_1^\varepsilon) = 0$ because $\kappa(\cdot)$ is also continuous. Therefore, there exists $\varepsilon^* > 0$ such that

$$\kappa(b_1^{\varepsilon^*}) < \kappa(1-\beta)K.$$

It follows that for ε^*, expression (16) is strictly positive, and hence $\Phi(b_1^{\varepsilon^*}) > \Phi(\overline{c}_1^d)$. This also holds for all $\varepsilon \in (0, \varepsilon^*)$.

□

We can now complete the proof of Proposition 1. By Lemma 5, $\overline{f}(\beta)$ decreases monotonically to $\underline{s}^p$ when $\beta \uparrow 1$. Also, for every $i = 1, \ldots, n$, we have that $s^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}; \beta)$ increases 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

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

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 $b_i^\varepsilon = \overline{c}_i^d - \varepsilon$ as in Lemma 6 where $\overline{c}_i^d = c_i^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i})$ by Remark 1. Let $\Phi(b_i^\varepsilon, \overline{f}(\beta))$ be the planner's expected payoff from adding b_i^ε to the existing optimal floor $\overline{f}(\beta)$. We will show that there exists $\varepsilon > 0$ such that $\Phi(b_i^\varepsilon, \overline{f}(\beta)) > \Phi(b_i^0, \overline{f}(\beta))$ where $\Phi(b_i^0, \overline{f}(\beta)) = \mathcal{U}(B_{\overline{f}(\beta)})$. To do so, for any $\varepsilon \geq 0$, let $(\mathbf{c}^\varepsilon, s^\varepsilon)$ be the doer's allocation function under $(b_i^\varepsilon, \overline{f}(\beta))$ and $\Omega(b_i^\varepsilon) = \{\omega \in \Omega : c_i^0(\omega) > b_i^\varepsilon\}$. Then,

$$\Phi(b_i^\varepsilon, \overline{f}(\beta)) - \Phi(b_i^0, \overline{f}(\beta)) = \int_{\Omega(b_i^\varepsilon)} \{[\theta u(\mathbf{c}^\varepsilon(\omega); \mathbf{r}) + v(s^\varepsilon(\omega))] - [\theta u(\mathbf{c}^0(\omega); \mathbf{r}) + v(s^0(\omega))]\}dG.$$

Note that, if there exists $\overline{\varepsilon} > 0$ such that for all $0 < \varepsilon < \overline{\varepsilon}$ we have $(\mathbf{c}^\varepsilon(\omega), s^\varepsilon(\omega)) = (\mathbf{c}^d(\omega), s^d(\omega))$ for all $\omega \in \Omega(b_i^\varepsilon)$, then for such ε's the previous difference equals $\Phi(b_i^\varepsilon) - \Phi(\overline{c}_i^d)$ in the proof of Lemma 6. The conclusion of that proof then implies that there exists $\varepsilon^{**} \in (0, \overline{\varepsilon})$ such that $\Phi(b_i^{\varepsilon^{**}}, \overline{f}(\beta)) > \Phi(b_i^0, \overline{f}(\beta))$.

Thus we only need 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 $\tilde{c}_i = \max_{\overline{\Omega}(\overline{f}(\beta))} c_i^0(\omega)$ which is well defined by continuity of $\mathbf{c}^0$. Since $s^d(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}) > \overline{f}(\beta)$, it follows that $(\overline{\theta}, \overline{r}_i, \underline{r}_{-i}) \notin \overline{\Omega}(\overline{f}(\beta))$ 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 $\tilde{c}_i < \overline{c}_i^d$: for all $\omega \in \overline{\Omega}(\overline{f}(\beta))$, optimality requires

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

where $\lambda_0(\omega) \geq 0$ is the Lagrange multiplier for constraint $s \geq \overline{f}(\beta)$. If $\omega \in \Omega$ 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 therefore $\mathbf{c}^0(\omega) = \mathbf{c}^d(\omega)$. Now define $\overline{\varepsilon} = \overline{c}_i^d - \tilde{c}_i > 0$. By construction for every $\varepsilon \in (0, \overline{\varepsilon})$, $c_i^0(\omega) > b_i^\varepsilon$ implies that $\mathbf{c}^0(\omega) = \mathbf{c}^d(\omega)$, as desired.

9.2 Proof of Proposition 2 and Corollary 1

Lemma 7. *For every $\beta \in (0,1)$, if $B \in \mathcal{B}$ is optimal, then* $\max\{f, 1 - \sum_{i=1}^n b_i\} \geq \min_{\omega \in \Omega} s^p(\omega)$.

Proof. Define

$$\sigma = \max\{f, 1 - \sum_{i=1}^n b_i\}$$

Given this, we have that $s(\omega) + \sum_{i=1}^n c_i(\omega) = 1$, and hence $s(\omega) \geq \sigma$ for all $\omega \in \Omega$. Without loss of generality, we can let $\sigma = \min_\Omega s(\omega)$: If $\min_\Omega s(\omega) > \sigma$, we could simply 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 $B'' \in \mathcal{B}$ identical to B', except that $f'' = \underline{s}^p$. Since B' is convex and compact, the ensuing allocation $(\mathbf{c}', s')$ is a continuous function of ω. Hence, the set $\Omega(\underline{s}^p) = \{\omega \in \Omega : s'(\omega) < \underline{s}^p\}$ contains an open subset and hence has strictly positive probability under G.

Consider any $\omega \in \Omega(\underline{s}^p)$. Suppose the planner faces the following problem:

$$\max_{(\mathbf{c},s) \in \mathbb{R}_+^{n+1}} \{\theta u(\mathbf{c}; \mathbf{r}) + v(s)\}$$

subject to $c_i \leq b_i'$, and $s \leq f$. For any $f < \underline{s}^p$, the latter constraint must bind because, by the same logic of Lemma 9, the planner would want to save at least $s^p(\omega) \geq \underline{s}^p$ if facing only the constraints $c_i \leq b_i'$ for $i = 1, \ldots, n$. Therefore, the planner's payoff from this fictitious problem is strictly increasing in f for $f \leq \underline{s}^p$. When the doer faces B'', the constraint $s \geq \underline{s}^p$ must bind. Hence, 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 the planner's allocation under the fictitious problem with $f = \underline{s}^p$. Hence, in ω, $(\mathbf{c}''(\omega), \underline{s}^p)$ is strictly better for the planner than $(\mathbf{c}''(\omega), s'(\omega))$.

We conclude that, for all $\omega \in \Omega(\underline{s}^p)$, the planner's payoff is strictly larger under B'' than under B'. Since for $\omega \notin \Omega(\underline{s}^p)$ the doer's allocation is unchanged, we must have $\mathcal{U}(B'') > \mathcal{U}(B')$, which contradicts the optimality of B'.

□

Given Lemma 7, we can complete the proof of Proposition 2. 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 the doer's problem in $\underline{\omega} = (\underline{\omega}, \underline{\mathbf{r}})$, we have that $(\mathbf{c}(\underline{\omega}), s(\underline{\omega}))$ must satisfy

$$\beta v'(\overline{s}) + \phi_0(\underline{\omega}) + \gamma_i(\underline{\omega}) = \underline{\theta} u_c^i(c_i(\underline{\omega}); \underline{r}_i) \quad \text{for all } i = 1, \ldots, n,$$

where $\phi_0(\underline{\omega}) \geq 0$ and $\gamma_i(\underline{\omega}) \geq 0$ are the Lagrange multipliers for constraints $s \geq f$ 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, $u^i_c(c_i(\underline{\omega}); \underline{r}_i) \geq u^i_c(1; \underline{r}_i) > 0$. Now let

$$\beta_{**} = \min_i \frac{\underline{\theta} u^i_c(1; \underline{r}_i)}{v'(\underline{s}^p)} > 0. \tag{17}$$

Then, for every $\beta < \beta_{**}$, we have $\beta v'(s(\underline{\omega})) < \underline{\theta} u^i_c(c_i(\underline{\omega}); \underline{r}_i)$ for all $i = 1, \dots, n$. Therefore, $\phi_0(\underline{\omega}) + \gamma_i(\underline{\omega}) > 0$ for all $i = 1, \dots, n$. Hence, either $\phi_0(\underline{\omega}) > 0$, in which case $\overline{s} = f = \sigma$; or $\gamma_i(\underline{\omega}) > 0$ for all $i = 1, \dots, n$, 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 5. Let $B^\beta \in \mathcal{B}$ be an optimal plan for β. By Lemma 7, $\sigma^\beta \geq \underline{s}^p$. The previous result then implies that

$$\mathcal{U}(B^\beta) = \int_\Omega [\theta u(\mathbf{c}(\omega); \mathbf{r}) + v(\sigma^\beta)] dG.$$

Hence,

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

where the first inequality follows 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 5. It is immediate to see that if B^β involves budgets that bind for a set Ω' whose probability is strictly positive, then $u(\mathbf{c}(\omega); \mathbf{r}) < u(\mathbf{c}^{\sigma^\beta}(\omega); \mathbf{r})$ for all $\omega \in \Omega'$, and hence $\mathcal{U}(B^\beta) < \mathcal{U}(B_{\overline{f}})$. Therefore, optimal plans can only involve a savings floor.

Finally, let $\underline{\mathbf{r}}'$, $\underline{\mathbf{r}}$, $\overline{\mathbf{r}}'$, and $\overline{\mathbf{r}}$ satisfy the properties in Proposition 2. It follows that $\underline{s}^{p\prime} = s^p(\overline{\theta}, \overline{\mathbf{r}}') \geq s^p(\overline{\theta}, \overline{\mathbf{r}}) = \underline{s}^p$ with strict inequality if $\overline{\mathbf{r}} \neq \overline{\mathbf{r}}'$ (Remark 1). Similarly, for each $\beta \in (0,1)$, $\overline{s}^{d\prime}(\beta) = s^d(\underline{\theta}, \underline{\mathbf{r}}'; \beta) \leq s^d(\underline{\theta}, \underline{\mathbf{r}}; \beta) = \overline{s}^d(\beta)$ again with strict inequality if $\underline{\mathbf{r}}' \neq \underline{\mathbf{r}}$. Using the definition of β_{**} in (17), the strict concavity of the function v, and that $\underline{r}'_i \geq \underline{r}_i$, we have that $\beta'_{**} > \beta_{**}$. Using the definition of $\underline{\beta}$ in the proof of Lemma 5 and that $\overline{s}^d$ is strictly increasing in β, we have that $\underline{\beta}' > \underline{\beta}$. Therefore $\beta'_* > \beta_*$.

9.3 Proof of Proposition 3

By an argument similar to the proof of Lemma 1, we can conclude that an optimal plan $B^* \in \mathcal{B}$ exists in this three-state setting. The following claims characterize its properties.

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

Proof. If B can use only f, 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 the planner'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 $f \in [s^d(\omega^1), s^p(\omega^1)]$, her 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 3. Thus, $f = s^p(\omega^0)$ identifies the best policy that involves only f if

$$\frac{g}{1-g} > \max_{f\in[s^d(\omega^1),s^p(\omega^1)]} \frac{U(\mathbf{c}^f(\omega^1), f; \omega^1) - U(\mathbf{c}^d(\omega^1), s^d(\omega^1); \omega^1)}{U(\mathbf{c}^p(\omega^0), s^p(\omega^0); \omega^0) - U(\mathbf{c}^f(\omega^0), f; \omega^0)} \geq 0. \tag{18}$$

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 the planner knows that the state is ω^i and can only impose a budget b_i. Then, it is optimal to set $b_i = c_i^p(\omega^i)$.

Proof. Let $i = 1$—the other case is similar. Replicating the argument in the proof of Lemma 6, we can conclude that it is optimal to set $b_1 < c_1^d(\omega^1)$. To find the optimal $b_1 \in (0, c_1^d(\omega^1))$, consider first the doer's problem to maximize $\underline{\theta}[\overline{r}\ln(c_1)+\underline{r}\ln(c_2)]+\beta\ln(s)$ subject to $s+c_1+c_2 \leq 1$ and $c_1 \leq b_1$. Since both constraints must bind, this problem becomes

$$\max_{s\in[0,1]} \{\underline{\theta}\,\underline{r}\ln(1 - b_1 - s) + \beta\ln(s)\}.$$

The solution is characterized by the first-order condition, which leads to

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

Given this, we can compute the planner's payoff in ω^1 as a function of b_1, which equals (up to a constant)

$$\underline{\theta}[\overline{r}\ln(b_1) + \underline{r}\ln(1 - b_1)] + \ln(1 - b_1). \tag{19}$$

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

$$b_1 = \frac{\underline{\theta}\overline{r}}{1 + \underline{\theta}[\underline{r} + \overline{r}]}. \tag{20}$$

To complete the proof, we need to find $c_1^p(\omega^1)$, which results from maximizing $\underline{\theta}[\overline{r}\ln(c_1) + \underline{r}\ln(c_2)] + \ln(s)$ subject to $s + c_1 + c_2 \leq 1$. Substituting $s = 1 - c_1 - c_2$, taking first-order conditions, and combining them, we get

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

□

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

Proof. Let $i = 1$—the other case is similar. Mimicking the calculations in the proof of Claim 2, one can show that if the planner can impose only b_2, then she sets

$$b_2 = \frac{\underline{\theta}\,\underline{r}}{1 + \underline{\theta}[\underline{r} + \overline{r}]}. \tag{21}$$

We want to argue that her payoff in ω^1 is strictly larger if she imposes only b_1 as in (20) than if she imposes only b_2 as in (21). Substituting the allocations implied by b_1 and b_2 into the planner's utility function and simplifying, one can show that b_1 in (20) is strictly better than b_2 in (21) if and only if

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

To show that this condition holds, consider the function $\varphi(\beta,r)=(1+\underline{\theta}r)\ln(\beta+\underline{\theta}r)$, where $0<\beta<1$ and $r>0$. This function satisfies

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

Therefore, $\varphi(\beta,\overline{r})-\varphi(\beta,\underline{r})$ 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 states, then B is weakly dominated by a policy that involves only f and no budgets, as the budgets distort consumption without improving savings. Given $g>g^*$, by Claim 1 the latter plan is strictly dominated by one imposing only the floor $s^p(\omega^0)$. Clearly, if f binds in ω^1 and ω^2, then it must also bind in ω^0.

Now suppose that f binds only in ω^0 and another state, say, ω^1—the same argument applies for ω^2. There are two cases to consider:
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 the doer'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 the floor at $s^p(\omega^0)$.
Case 2: b_1 binds also in ω^2. This implies that f has to bind in all states. Indeed, since b_1 binds in both ω^1 and ω^2, the doer chooses $c_1=b_1$ in both states; moreover, since in ω^2 good 2 is more valuable than in ω^1, he wants to allocate more income to good 2 than to savings relative to ω^1, and therefore f also binds in ω^2. However, we have already argued that such a policy is strictly dominated by one that imposes only the floor $s^p(\omega^0)$.

□

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 bind in any state. Since $\overline{\theta}>\underline{\theta}$, we have $c_2(\omega^0)>c_2(\omega^2)$. The plan cannot be optimal because, given b_1, the planner would be strictly better off by adding a floor that binds only in ω^0: Even if b_1 were binding for her in ω^0, she would strictly prefer a level $c_2<c_2(\omega^0)$ 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 a floor equal to $1-b_1-b_2$—because budgets are distorting—which is in turn strictly dominated by the plan with only the floor $s^p(\omega^0)$ given $g>g^*$.
Case 3: b_2 binds in no state. Then, as in case 1, for B to be optimal f must bind at least in ω^0 and only in that state 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, the planner 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, the planner 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 that state by $\hat{b}_2$.
Case 4: b_2 binds only in ω^0. Since the doer's choices satisfy $c_2(\omega^0) > c_2(\omega^2)$ if the plan used only b_1, it follows that the planner can obtain in all states the same allocations induced by B if she imposes a floor that binds only in ω^0. Such a plan, however, is again strictly dominated for the same reasons as in case 3.
Case 5: b_2 binds in ω^0 and in ω^2. Since the doer's choices satisfy $c_2(\omega^0) > c_2(\omega^2)$ if the plan used only b_1, the planner could again obtain the same allocation in all states with a floor that binds only in ω^0 and ω^2. By Claim 4, however, such a plan cannot be optimal.

Now suppose that b_1 binds in only two states. If b_1 binds only in ω^1 and in ω^0, then by the same argument as in case 3 above the planner 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, then 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 consumption goods are more valuable. The case left is if b_1 binds only in ω^0 and ω^2, but this is impossible: It would have to bind also in ω^1, since in that state good 1 is more valuable than in ω^2.

Finally, suppose that b_1 binds in only one state. We have just argued that if b_1 binds in ω^2, then it must also bind in ω^1. Thus, we only have to rule out the case in which b_1 binds only in ω^0. This property is possible only if in ω^0 the budget b_2 also binds, inducing the doer to overconsume in good 1. However, such a b_2 must also bind in ω^2; hence, it cannot be part of an optimal B, because we just showed that a budget cannot bind in more than one state.

□

Combining Claims 1-5, we conclude that the optimal policy $B \in \mathcal{B}$ satisfies $f = s^p(\omega^0)$, $b_1 = c_1^p(\omega^1)$, and $b_2 = c_2^p(\omega^2)$.

9.4 Proof of Proposition 4

Start from the value of $\overline{\theta}$ which implies that $c_1^p(\omega^1) > c_1^p(\omega^0)$ and $c_2^p(\omega^2) > c_2^p(\omega^0)$ and hence leads to the optimal plan in Proposition 3. If we increase $\overline{\theta}$, both $c_1^p(\omega^0)$ and $c_2^p(\omega^0)$ increase continuously while always satisfying $c_1^p(\omega^0) = c_2^p(\omega^0)$. Therefore, there exists a unique $\overline{\theta}^\dagger$ such that, when $\overline{\theta} = \overline{\theta}^\dagger$, we have $c_1^p(\omega^1) = c_1^p(\omega^0)$ and $c_2^p(\omega^2) = c_2^p(\omega^0)$. For every $\overline{\theta} \leq \overline{\theta}^\dagger$, the optimal $B \in \mathcal{B}$ remains $b_1 = c_1^p(\omega^1)$, $b_2 = c_2^p(\omega^2)$, and $f = s^p(\omega^0)$, where the latter of course falls continuously as $\overline{\theta}$ rises towards $\overline{\theta}^\dagger$.

Now, let $\mathcal{B}(\overline{\theta}) \subset \mathcal{B}$ be the nonempty set of optimal plans as a function of $\overline{\theta}$. By Proposition 3 and the previous argument, $\mathcal{B}(\overline{\theta})$ is singleton for $\overline{\theta} \leq \overline{\theta}^\dagger$. Define the distance between any two plans B and B' as the Euclidean distance between the vector (f, b_1, b_2) describing B and the vector (f', b_1', b_2') describing B'. By the Maximum Theorem, $\mathcal{B}(\overline{\theta})$ is upper hemicontinuous in $\overline{\theta}$.[48] Hence, by choosing $\overline{\theta} > \overline{\theta}^\dagger$ sufficiently close to $\overline{\theta}^\dagger$, we can render the distance between $B(\overline{\theta}^\dagger)$

[48]Although the planner's and doer'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 $B \in \mathcal{B}$ that forces 0

and every $B \in \mathcal{B}(\overline{\theta})$ arbitrarily small. Thus, there exists $\varepsilon > 0$ such that, if $\overline{\theta} \in (\overline{\theta}^{\dagger}, \overline{\theta}^{\dagger} + \varepsilon)$, then for every $B \in \mathcal{B}(\overline{\theta})$ the following holds: (1) $b_i(\overline{\theta}) < c_i^d(\omega^i)$ for $i = 1, 2$; and (2) $f(\overline{\theta})$ can bind neither in ω^1 nor in ω^2. To see property (2), note 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 actually redundant. Thus, $\mathcal{B}(\overline{\theta})$ contains no plan with $f(\overline{\theta}) > 1 - b_1(\overline{\theta}) - b_2(\overline{\theta})$, because such plans are strictly dominated for the same argument that rules them out in the proof of Proposition 3. Since the largest value of $f(\overline{\theta})$ must be close to $f(\overline{\theta}^{\dagger})$ for $\overline{\theta} \in (\overline{\theta}^{\dagger}, \overline{\theta}^{\dagger} + \varepsilon)$, it follows that $f(\overline{\theta})$ cannot bind in ω^1 and ω^2 as well.

Hereafter, fix $\overline{\theta} \in (\overline{\theta}^{\dagger}, \overline{\theta}^{\dagger} + \varepsilon)$. The following claims characterize the properties of every $B \in \mathcal{B}(\overline{\theta})$.

Claim 6. For every $B \in \mathcal{B}(\overline{\theta})$, both $b_1(\overline{\theta})$ 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 the planner's objective in state ω^i as a function of b_i is strictly concave and decreasing for $b_i > c_i^p(\omega^i)$ (see equation (19)). Thus, if for example $b_1(\overline{\theta})$ is not binding for the doer in state ω^0—that is, $b_1(\overline{\theta}) > c_1(\omega^0)$—the planner can lower b_1 without affecting the doer's choice in ω^0 and ω^2 and strictly improve her payoff in ω^1. Hence, the initial plan would not be optimal.

□

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

Proof. Without loss, suppose that $b_1(\overline{\theta}) > b_2(\overline{\theta})$. Note that $b_2(\overline{\theta}) < c_2^d(\omega^0)$ because, otherwise, we would have $b_1(\overline{\theta}) > c_1^d(\omega^0) = c_2^d(\omega^0)$, which contradicts the previous point. Consider the alternative plan with $b_1^\varepsilon = b_1(\overline{\theta}) - \varepsilon$ and $b_2^\varepsilon = b_2(\overline{\theta}) + \varepsilon$, where $\varepsilon > 0$. For ε sufficiently small, both b_1^ε and b_2^ε continue to be binding in ω^0, and hence $1 - b_1^\varepsilon - b_2^\varepsilon = s(\omega^0)$. In ω^0, the planner's payoff is higher, because given $s(\omega^0)$ the consumption bundle is closer to being symmetric and hence to the best one according to the planner's preference. Due to symmetry and the strict concavity in the planner's payoff induced by b_i in ω^i for $i = 1, 2$ (see (19)), we have that the decrease in the her payoff in ω^2 resulting from the slacker b_2 is more than compensated by the increase in her payoff in ω^1 resulting from the tighter b_1. Hence, overall the planner's payoff is strictly larger with $(b_1^\varepsilon, b_2^\varepsilon)$ than with $(b_1(\overline{\theta}), b_2(\overline{\theta}))$, which contradicts the optimality of the latter plan.

□

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

Proof. If $1 - b_1(\overline{\theta}) - b_2(\overline{\theta}) < s^p(\omega^0)$, then the planner can set $f = s^p(\omega^0)$ and achieve a strictly higher payoff in ω^0 without affecting the doer'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 $i = 1, 2$, which means that $(\mathbf{c}(\omega^0), s(\omega^0)) = (\mathbf{c}^p(\omega^0), s^p(\omega^0))$. Therefore, it would be possible to lower both $b_1(\overline{\theta})$ and $b_2(\overline{\theta})$ by the same small amount ε, in order to induce a first-oder gain in the planner's payoff for both 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.

□

allocation to some dimension. Formally, there exists $\varepsilon > 0$ such that, if we required $f \leq 1 - \varepsilon$ and $b_i \geq \varepsilon$ for all $i = 1, 2$, we would never affect the planner's problem.

Claim 9. Every $B \in \mathcal{B}(\overline{\theta})$ is unique as far as b_1 and b_2 are concerned and satisfies the properties in Proposition 4.

Proof. Let $b_1 = b_2 = b$. The planner's payoff in ω^1 and ω^2 is given by (19) 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, the optimal b maximizes

$$(1-g)\left\{\underline{\theta}[\overline{r}\ln(b) + \underline{r}\ln(1-b)] + \ln(1-b)\right\} + g\left\{2\underline{\theta}\overline{r}\ln(b) + \ln(1-2b)\right\}.$$

Since this function is strictly concave, there is a unique optimal b. To see that $c_i^p(\omega^i) > b_i > c_i^p(\omega^0)$ for every $i = 1, 2$, consider the following observations. Note that $b_i > c_i^p(\omega^i)$ would be strictly dominated by $b_i = c_i^p(\omega^i)$ for every i, because this is the optimal level of the budget in the corresponding state. 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 hence reducing b_i below $c_i^p(\omega^i)$ by the same small amount for all $i = 1, 2$ causes a first-order gain in ω^0 and only a second-oder loss in ω^1 and ω^2. □

9.5 Proof of Lemma 2

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

$$\theta u(\chi(\theta)) + \beta v(t(\theta)) \geq \theta u(\chi(\theta')) + \beta v(t(\theta'))$$

and

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

(2) the pair (χ, t) maximizes

$$\int_{\underline{\theta}}^{\overline{\theta}} [\theta u(\chi(\theta)) + v(t(\theta))]\, \hat{g}(\theta) d\theta.$$

On the other hand, 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 two conditions:
(1') for all $\theta, \theta' \in [\underline{\theta}, \overline{\theta}]$

$$\theta u^*(\varphi(\theta)) + \beta v(\tau(\theta)) \geq \theta u^*(\varphi(\theta')) + \beta v(\tau(\theta')),$$

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

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

(2') the pair (φ, τ) maximizes

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[\theta u^*(\varphi(\theta)) + v(\tau(\theta))\right] \hat{g}(\theta) d\theta.$$

Suppose (χ, t) that satisfies condition **(1)** and **(2)**. Then, by our discussion on money burning before the statement of Lemma 2, there exists a function $\varphi : [\underline{\theta}, \overline{\theta}] \to \mathbb{R}_+$ such that $u^*(\varphi(\theta)) = u(\chi(\theta))$ 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 **(1')** and **(2')**.

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

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

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

References

Aguiar, M. and M. Amador (2011). Growth in the shadow of expropriation. *Quarterly Journal of Economics 126*(2), 651–697.

Ali, S. N. (2011). Learning self-control. *Quarterly Journal of Economics 126*(2), 857–893.

Alonso, R., I. Brocas, and J. D. Carrillo (2014). Resource allocation in the brain. *The Review of Economic Studies 81*(2), 501–534.

Alonso, R. and N. Matouschek (2008). Optimal delegation. *The Review of Economic Studies 75*(1), 259–293.

Amador, M. and K. Bagwell (2013a). Money burning in the theory of delegation. *Mimeo, Stanford University.*

Amador, M. and K. Bagwell (2013b). The theory of optimal delegation with an application to tariff caps. *Econometrica 81*(4), 1541–1599.

Amador, M., I. Werning, and G.-M. Angeletos (2003). Commitment vs. flexibility. Technical report, National Bureau of Economic Research.

Amador, M., I. Werning, and G.-M. Angeletos (2006). Commitment vs flexibility. *Econometrica 74*(2), 365–396.

Ambrus, A. and G. Egorov (2009). Delegation and nonmonetary incentives. *Unpublished working paper*.

Ambrus, A. and G. Egorov (2013). Comment on "commitment vs. flexibility". *Econometrica 81*(5), 2113–2124.

Ameriks, J., A. Caplin, and J. Leahy (2003). Wealth accumulation and the propensity to plan. *The Quarterly Journal of Economics 118*(3), 1007–1047.

Antonides, G., I. Manon de Groot, and W. Fred van Raaij (2011). Mental budgeting and the management of household finance. *Journal of Economic Psychology 32*(4), 546–555.

Armstrong, M. (1995). Delegation and discretion. *Mimeograph, University College London.*

Athey, S., A. Atkeson, and P. J. Kehoe (2005). The optimal degree of discretion in monetary policy. *Econometrica 73*(5), 1431–1475.

Bakke, E. W. (1940). *The unemployed worker.* Yale University Press.

Battaglini, M. (2005). Long-term contracting with markovian consumers. *American Economic Review 95*(3), 637–658.

Baumeister, R. F., T. F. Heatherton, and D. M. Tice (1994). *Losing control: How and why people fail at self-regulation.* Academic press.

Benabou, R. and M. Pycia (2002). Dynamic inconsistency and self-control: A planner–doer interpretation. *Economics Letters 77*(3), 419–424.

Bénabou, R. and J. Tirole (2004). Willpower and personal rules. *Journal of Political Economy 112*(4), 848–886.

Benhabib, J. and A. Bisin (2005). Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption–saving decisions. *Games and Economic Behavior 52*(2), 460–492.

Bernheim, B. D. and A. Rangel (2004). Addiction and cue-triggered decision processes. *The American Economic Review 94*(5), 1558–1590.

Bernheim, B. D., D. Ray, and S. Yeltekin (2015). Poverty and self-control. *Econometrica 83*(5), 1877–1911.

Beshears, J., K. L. Milkman, and J. Schwartzstein (2016). Beyond beta-delta: The emerging economics of personal plans. *The American Economic Review 106*(5), 430–434.

Breig, Z. (2015). Dynamic delegation. *Mimeo, UCSD.*

Brocas, I. and J. D. Carrillo (2008). The brain as a hierarchical organization. *The American Economic Review 98*(4), 1312–1346.

Bryan, G., D. Karlan, and S. Nelson (2010). Commitment devices. *Annu. Rev. Econ. 2*(1), 671–698.

Chatterjee, K. and R. V. Krishna (2009). A"dual self" representation for stochastic temptation. *American Economic Journal: Microeconomics 1*(2), 148–167.

Courty, P. and H. Li (2000). Sequential screening. *The Review of Economic Studies 67*(4), 697–717.

Ellis, C. J. and B. J. Naughton (1990). On the theory of household saving in the presence of rationing. *Journal of Comparative Economics 14*(2), 269–285.

Festinger, L. (1962). *A theory of cognitive dissonance*. Stanford university press.

Frankel, A. (2014). Aligned delegation. *American Economic Review 104*(1), 66–83.

Frankel, A. P. (2016). Delegating multiple decisions. *forthcoming AEJ: Microeconomics*.

Fudenberg, D. and D. K. Levine (2006). A dual-self model of impulse control. *American Economic Review 96*(5), 1449–1476.

Fudenberg, D. and D. K. Levine (2012). Timing and self-control. *Econometrica 80*(1), 1–42.

Funk, P. and C. Gathmann (2011). Does direct democracy reduce the size of government? new evidence from historical data, 1890–2000*. *Economic Journal 121*(557), 1252–1280.

Galperti, S. (2015). Commitment, flexibility, and optimal screening of time inconsistency. *Econometrica 83*(4), 1425–1465.

Halac, M. and P. Yared (2014). Fiscal rules and discretion under persistent shocks. *Econometrica 82*(5), 1557–1614.

Heath, C., R. P. Larrick, and G. Wu (1999). Goals as reference points. *Cognitive psychology 38*(1), 79–109.

Heath, C. and J. B. Soll (1996). Mental budgeting and consumer decisions. *Journal of Consumer Research 23*(1), 40–52.

Henderson, P. W. and R. A. Peterson (1992). Mental accounting and categorization. *Organizational Behavior and Human Decision Processes 51*(1), 92–117.

Holmström, B. (1984). On the theory of delegation. *In Boyer M. and Kiehlstrom, R.(eds) Bayesian Models in Economic Theory, New York, NorthMHolland*.

Holmström, B. R. (1977). *On incentives and control in organizations*. Stanford University.

Howard, D. H. (1977). Rationing, quantity constraints, and consumption theory. *Econometrica 45*(2), 399–412.

Hsiaw, A. (2013). Goal-setting and self-control. *Journal of Economic Theory 148*(2), 601–626.

Jackson, M. O. and L. Yariv (2015). Collective dynamic choice: The necessity of time inconsistency. *American Economic Journal: Microeconomics 7*(4), 150–78.

Johnson, M. D. (1984). Consumer choice strategies for comparing noncomparable alternatives. *Journal of Consumer Research 11*(3), 741–753.

Koch, A. K. and J. Nafziger (2016). Goals and bracketing under mental accounting. *Journal of Economic Theory 162*, 305–351.

Koessler, F. and D. Martimort (2012). Optimal delegation with multi-dimensional decisions. *Journal of Economic Theory 147*(5), 1850–1881.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics 112*(2), 443–478.

Loewenstein, G. and T. O'Donoghue (2007). The heat of the moment: Modeling interactions between affect and deliberation. *Mimeo*.

Luenberger, D. G. (1969). *Optimization by vector space methods*. John Wiley & Sons.

Madden, P. (1991). A generalization of Hicksian q substitutes and complements with application to demand rationing. *Econometrica 59*(5), 1497–1508.

Milgrom, P. and C. Shannon (1994). Monotone comparative statics. *Econometrica 62*(1), 157–180.

Niskanen, W. A. (1975). Bureaucrats and politicians. *Journal of Law and Economics 18*(3), 617–643.

O'Donoghue, T. and M. Rabin (2001). Choice and procrastination. *Quarterly Journal of Economics 116*(1), 121–160.

Pavan, A., I. Segal, and J. Toikka (2014). Dynamic mechanism design: A Myersonian approach. *Econometrica 82*(2), 601–653.

Peltzman, S. (1992). Voters as fiscal conservatives. *Quarterly Journal of Economics*, 327–361.

Quah, J. K.-H. (2007). The comparative statics of constrained optimization problems. *Econometrica 75*(2), 401–431.

Rainwater, L., R. P. Coleman, and G. Handel (1962). *Workingman's Wife: her personality, world and life style*. MacFadden-Bartell.

Rochet, J.-C. and P. Choné (1998). Ironing, sweeping, and multidimensional screening. *Econometrica 66*(4), 783–826.

Rochet, J.-C. and L. A. Stole (2003). The economics of multidimensional screening. *Econometric Society Monographs 35*, 150–197.

Rockafellar, R. T. (1997). *Convex Analysis.* Princeton University Press.

Romer, T. and H. Rosenthal (1979). Bureaucrats versus voters: On the political economy of resource allocation by direct democracy. *Quarterly Journal of Economics 93*(4), 563–587.

Simon, H. A. (1965). *Administrative behavior*, Volume 4. Cambridge Univ Press.

Strotz, R. H. (1957). The empirical implications of a utility tree. *Econometrica*, 269–280.

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science 4*(3), 199–214.

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making 12*(3), 183–206.

Thaler, R. H. and H. M. Shefrin (1981). An economic theory of self-control. *Journal of Political Economy 89*(2), 392–406.

Wertenbroch, K. (2002). Self-rationing: Self-control in consumer choice. *SSRN Working Paper Series.*

Zelizer, V. A. R. (1997). *The social meaning of money.* Princeton University Press.

10 Online Appendix: Supplementary Material (FOR ONLINE PUBLICATION ONLY)

10.1 Proof of Lemma 1

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 the planner 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 the doer'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 the doer's utility function. Clearly, the correspondence that for each $(f, \mathbf{b}) \in [0,1]^{n+1}$ maps to $B_{f,\mathbf{b}}$ is non-empty, 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 the planner's payoff is continuous in $(f, \mathbf{b})$. For each $(f, \mathbf{b}) \in [0,1]^{n+1}$, let

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

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.

10.2 Lemma 8

Lemma 8. *Fix $\hat{f} > \underline{s}^d = \min_\omega s^d(\omega)$. Let $B_{\hat{f}}$ be a plan that satisfies $f = \hat{f}$ and $b_i = 1$ for all i; let $B' \in \mathcal{B}$ be a plan that satisfies $\sum_{i=1}^n b_i = 1 - \hat{f}$. Every B' implements allocations that differ with positive probability from those implemented by $B_{\hat{f}}$.*

Proof. Fix $\hat{f} > \underline{s}^d$. Let $\overline{\mathbf{r}} = (\overline{r}_1, \ldots, \overline{r}_n)$. It is easy to see that $s^d(\overline{\theta}, \overline{\mathbf{r}}) = \underline{s}^d$. Therefore, $\hat{f}$ must bind in $\overline{\omega} = (\overline{\theta}, \overline{\mathbf{r}})$. Since $(\mathbf{c}^d, s^d)$ is continuous in ω, there exists $\varepsilon > 0$ such that, if $|\mathbf{r} - \overline{\mathbf{r}}| < \varepsilon$, then $s^d(\overline{\theta}, \mathbf{r}) > \hat{f}$, and hence $\hat{f}$ is still binding. When $\hat{f}$ binds, the doer's consumption allocation $\hat{\mathbf{c}}$ must maximize $\sum_{i=1}^n u^i(c_i; r_i)$ subject to $\sum_{i=1}^n c_i \leq 1 - \hat{f}$. So, for all $\mathbf{r}$ with $|\mathbf{r} - \overline{\mathbf{r}}| < \varepsilon$, we must have

$$u_c^i(\hat{c}_i(\overline{\theta}, \mathbf{r}); r_i) = u_c^j(\hat{c}_j(\overline{\theta}, \mathbf{r}); r_j) \quad \text{for all } i, j.$$

Hence, there exists $\mathbf{r}'$ with $|\mathbf{r}' - \overline{\mathbf{r}}| < \varepsilon$ such that $\hat{\mathbf{c}}(\overline{\theta}, \mathbf{r}') \neq \hat{\mathbf{c}}(\overline{\theta}, \overline{\mathbf{r}})$. Since $\sum_{i=1}^n \hat{c}_i(\overline{\theta}, \mathbf{r}') = \sum_{i=1}^n \hat{c}_i(\overline{\theta}, \overline{\mathbf{r}}) = 1 - \hat{f}$, there exists $i \neq j$ such that $\hat{c}_i(\overline{\theta}, \mathbf{r}') > \hat{c}_i(\overline{\theta}, \overline{\mathbf{r}})$ and $\hat{c}_j(\overline{\theta}, \mathbf{r}') < \hat{c}_j(\overline{\theta}, \overline{\mathbf{r}})$. Now let $\Omega(\hat{f})$ be the set of states for which $\hat{s}(\omega) = \hat{f}$. By the previous argument, $\hat{c}_i$ and $\hat{c}_j$ cannot be constant over $\Omega(\hat{f})$.

For each $k = 1, \ldots, n$, let $\hat{b}_k = \max_\Omega \hat{c}_k(\omega)$. When $\hat{c}_i(\omega) = \hat{b}_i$, we must have $\hat{c}_j(\omega) < \hat{b}_j$, and when $\hat{c}_j(\omega) = \hat{b}_j$, we must have $\hat{c}_i(\omega) < \hat{b}_i$. Therefore, $\sum_{i=1}^n \hat{b}_i > 1 - \hat{f}$. It follows that any collection of caps $\mathbf{b} = \{b_i\}_{i=1}^n$ satisfying $\sum_{i=1}^n b_i = 1 - \hat{f}$ must involve $b_i < \hat{b}_i$ for some $i = 1, \ldots, n$. So, when the doer faces the constraints $\mathbf{b}$, for some i and state ω, $c_i(\omega) \leq b_i$ for all states ω such that the doer chooses $c_i > b_i$ when facing only $\hat{f}$. Since $(\hat{\mathbf{c}}, \hat{s})$ is continuous in ω, the set $\Omega(\mathbf{b}) = \{\omega : \hat{c}_i(\omega) > b_i\}$ is open and hence it has strictly positive probability under G.

□

10.3 Lemma 9

Lemma 9. *Fix some $i \in \{1, \ldots, n\}$ and consider $B \in \mathcal{B}$ with $b_j = 1$ for all $j \neq i$. In any state ω, if $b_i < c_i^d(\omega)$, then the doer chooses $s > s^d(\omega)$, but also $c_j > c_j^d(\omega)$ for all $j \neq i$.*

Proof. Without loss of generality, let $i = 1$ and take any $b_1 \in (0, c_1^d(\omega))$. Consider the doer's problem in state ω subject to b_1:

$$\max_{\{(\mathbf{c},s) \in F : c_1 \leq b_1\}} \theta u(\mathbf{c}, \mathbf{r}) + \beta v(s).$$

The first-order conditions of the associated Lagrangian are

$$\begin{aligned} \beta v'(s(\omega)) &= \mu(\omega), \\ \theta u_c^1(c_1(\omega); r_1) &= \mu(\omega) + \lambda_1(\omega), \\ \theta u_c^i(c_i(\omega); r_i) &= \mu(\omega) \quad \text{for all } i \neq 1, \end{aligned}$$

where $\mu(\omega) \geq 0$ and $\lambda_1(\omega) \geq 0$ are the Lagrange multipliers for constraints $\sum_{i=1}^n c_i \leq 1$ 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 for $j \neq i$

$$\theta u_c^j(c_j(\omega); r_j) = \beta v'(s(\omega)) < \beta v'(s^d(\omega)) = \theta u_c^j(c_j^d(\omega); r_j).$$

By concavity, we have $c_j(\omega) > c_j^d(\omega)$ for $j \neq 0, 1$.

□

10.4 Budgets with Savings Floor or Only Budgets: General Setting

Let G^{fb} be a distribution over $(\omega^0, \omega^1, \omega^2)$ that leads to Proposition 3 and $\overline{G}$ the uniform distribution over $[\underline{\theta}, \overline{\theta}] \times [\underline{r}, \overline{r}]^2$. Similarly, let G^{b} be a distribution that leads to Proposition 4 and $\overline{G}'$ the uniform distribution over $[\underline{\theta}, \overline{\theta}'] \times [\underline{r}, \overline{r}]^2$, where $\overline{\theta}'$ is as in Proposition 4. Finally, let

$$G_\alpha^{\text{fb}} = \alpha G^{\text{fb}} + (1-\alpha)\overline{G} \qquad \text{and} \qquad G_\alpha^{\text{b}} = \alpha G^{\text{b}} + (1-\alpha)\overline{G}', \quad \alpha \in [0,1].$$

Corollary 2.
(1) There exists $\alpha \in (0,1)$ such that, given G_α^{fb}, every optimal $B \in \mathcal{B}$ involves a binding f as well as a binding b_i for both goods.
(2) There exists $\alpha' \in (0,1)$ such that, given $G_{\alpha'}^{\text{b}}$, for every optimal $B \in \mathcal{B}$ both b_1 and b_2 bind, but f never binds.

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

Part 1: Consider G_α^{fb}. For every $B \in \mathcal{B}$ and $\alpha \in [0,1]$, the planner's expected payoff is given by

$$\mathcal{U}(B; G_\alpha^{\text{fb}}) = \alpha \mathcal{U}(B; G^{\text{fb}}) + (1-\alpha)\mathcal{U}(B; \overline{G}).$$

Now define

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

Both $\mathcal{W}_f^{\text{fb}}$ and $\mathcal{W}_{\mathbf{b}}^{\text{fb}}$ are well defined by the same argument as in the proof of Lemma 1; moreover, by the Maximum Theorem, they are continuous functions of α.[49] Let B^{fb} denote the optimal plan in Proposition 3. Note that $\mathcal{U}(B^{\text{fb}}; \overline{G})$ is finite since the doer's resulting allocations 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 the distribution $G_{\hat{\alpha}}^{\text{fb}}$.

Part 2: Consider G_α^{b}. For every $B \in \mathcal{B}$ and $\alpha \in [0,1]$, the planner's expected payoff is given by

$$\mathcal{U}(B; G_\alpha^{\text{b}}) = \alpha \mathcal{U}(B; G^{\text{b}}) + (1-\alpha)\mathcal{U}(B; \overline{G}').$$

Let B^{b} denote the optimal plan in Proposition 4. By the same logic of the proof of Part 1, there exists $\alpha'' \in (0,1)$ such that, for every $\alpha \in (\alpha'', 1)$, the policy B^{b} strictly dominates

[49] Recall Footnote 48.

every $B \in \mathcal{B}_f \cup \mathcal{B}_{\mathbf{b}} \cup \mathcal{B}_{b_1} \cup \mathcal{B}_{b_2}$ given the distribution G^{b}_{α}. It remains to show that there exists $\alpha' \in (\alpha'', 1)$ such that B^{b} strictly dominates every $B \in \mathcal{B}$ given $G^{\mathrm{b}}_{\alpha'}$.

To this end, define

$$\mathcal{B}(\alpha) = \arg\max_{B \in \mathcal{B}} \mathcal{U}(B; G^{\mathrm{b}}_{\alpha}).$$

Another application of the Maximum Theorem implies that $\mathcal{B}(\cdot)$ is upper hemicontinuous. Note that $\mathcal{B}(1)$ is characterized by vectors (f^*, b_1^*, b_2^*) such that b_1^* and b_2^* are unique and satisfy the properties in Proposition 4, and $f^* \in [0, \overline{f}]$ where $\overline{f} = 1 - b_1^* - b_2^*$. Therefore, 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 $B \in \mathcal{B}(\alpha)$. This means that, by choosing η sufficiently small, we can ensure that for every $B \in \mathcal{B}(\alpha)$ the following holds: (1) $1 - b_1 - b_2 > s^p(\omega^0)$; (2) removing f leads to a plan such that both b_1 and b_2 bind in ω^0; and (3) f cannot bind in ω^1 and ω^2, since $\overline{f}$ is strictly smaller than the doer's choice of s in those states under every plan in $\mathcal{B}(1)$.

Take any $B \in \mathcal{B}(\alpha)$ and fix its b_1 and b_2. The f that completes B must be optimally chosen given b_1 and b_2. We claim that such an f must satisfy $f \le 1 - b_1 - b_2 = \overline{k}$ for α sufficiently close to 1. Suppose this is not true and consider the gain in the planner's expected payoff from imposing $f > \overline{k}$. Her gain in ω^0 would be

$$(1-\beta)[v(f) - v(\overline{k})] + \overline{V}(f; \omega^0) - \overline{V}(\overline{k}; \omega^0), \tag{22}$$

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

$$\int_{\Omega(f)} \left\{(1-\beta)[v(f) - v(\hat{s}(\omega))] + \overline{V}(f;\omega) - \overline{V}(\hat{s}(\omega); \omega)\right\} d\overline{G}', \tag{23}$$

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

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

Note that $\overline{V}(f;\omega) \le \overline{V}(\hat{s}(\omega);\omega)$ and $\hat{s}(\omega) \ge \overline{k}$ for all $\omega \in \Omega'$; therefore, for every $f \ge \overline{k}$, the quantity (23) is bounded above by

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

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

Now focus on $\overline{V}(k;\omega^0)$. For every $f > \overline{k}$, the following holds: (1) f always binds, because $\overline{k} > s^p(\omega^0)$ and hence the doer wants to save strictly less than f; (2) only one budget can bind, because if both bind, then $s(\omega^0) = \overline{k} < f$, which is impossible; (3) one budget never binds, because consumption goods are normal, so for every $f > \overline{k}$ the doer's chooses $c_i(\omega^0) < b_i$ for at least one $i = 1, 2$. Without loss, suppose that the budget that never binds is b_2. Therefore, if we remove b_2, $\overline{V}(k;\omega^0)$ coincides with the doer's indirect utility under the plan defined by $k \in [\overline{k}, 1]$ and b_1 only, which we denote by $\overline{V}(k;\omega^0, b_1)$. By the same argument as in the proof of Lemma 3, $\overline{V}(k;\omega^0, b_1)$ is continuously differentiable in k for $k \in (0,1]$ and $\overline{V}'(k;\omega^0, b_1) = -\lambda(\omega^0; k)$, where $\lambda(\omega^0;k)$ is the Lagrange multiplier associated to the constraint $s \ge k$. Using the Lagrangian defining $\overline{V}(k;\omega^0, b_1)$, we have that

$$\lambda(\omega^0; k) = \overline{\theta}' u_c^2(c_2(\omega^0; k); \overline{\mathbf{r}}) - \beta v'(k).$$

Note that $\lambda(\omega^0; k) > 0$ for all $k \in [\overline{k}, 1]$, because such levels of the floor must always bind for the doer. Moreover, $\lambda(\omega^0; k)$ is strictly increasing in $k \in [\overline{k}, 1]$ because v is strictly concave, $u^i_{cc} < 0$, and $c_2(\omega^0; k)$ is non-increasing in k by normality of goods. We conclude that $\overline{V}'(k; \omega^0) = -\lambda(\omega^0; k)$ for every $k \in (\overline{k}, 1]$ and $\overline{V}'(\overline{k}+; \omega^0) = -\lambda(\omega^0; \overline{k})$, where the plus denotes the right derivative.[50] Moreover, $\overline{V}'(k; \omega^0)$ is strictly decreasing in k.

Observe that

$$(1-\beta)v'(\overline{k}) + \overline{V}'(\overline{k}; \omega^0) = v'(\overline{k}) - \overline{\theta}' u^2_c(c_2(\omega^0; \overline{k}); \overline{\mathbf{r}}), \tag{24}$$

which is strictly negative. This is because $b_1 < c^p_1(\omega^0)$ and $b_2 < c^p_2(\omega^0)$ by Proposition 4 since α is close to 1, which implies that both budgets must bind for the planner; consequently, $f = \overline{k}$ and b_1 must also bind for the planner. The right-hand side of (24) coincides with the negative of the Lagrange multiplier associated with the constraint $s \geq \overline{k}$ in the planner's problem that also includes the constraint $c_1 \leq b_1$.

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

$$g\overline{V}'(\overline{k}_\alpha; \omega^0) + [\alpha g + (1-\alpha)](1-\beta)v'(\overline{k}_\alpha). \tag{25}$$

This quantity is strictly negative for $\alpha = 1$, which corresponds to $\overline{k}_1 = 1 - b^*_1 - b^*_2$. By continuity of (25) as a function of (α, k) and upper hemicontinuity of $\mathcal{B}(\alpha)$, there exists $\varepsilon > 0$ such that (25) remains strictly negative for all $\alpha \in (1-\varepsilon, 1]$. Given the monotonicity properties of v' and $\overline{V}'(\cdot; \omega^0)$, (25) is strictly decreasing for all $k \geq \overline{k}_\alpha$.

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

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

We conclude that the planner is strictly worse off by imposing a binding savings floor in addition to the budgets b_1 and b_2, and hence every optimal plan must involve binding budgets for both goods, but no binding floor on savings.

□

10.5 Proposition 7

To state the result, define

$$u_*(y) = \min_{\{\mathbf{c} \in \mathbb{R}^2_+ : c_1 + c_2 = y\}} u(\mathbf{c}), \qquad y \in [0, 1].$$

Proposition 7. *Suppose that the optimal $D^{\text{tc}} \subset F^{\text{tc}}$ induces money burning over some set $\Theta \subset [\underline{\theta}, \overline{\theta}]$ and positive consumption: $y(\theta) < 1 - s(\theta)$ for $\theta \in \Theta$ and $y(\theta) > 0$ for $\theta \in [\underline{\theta}, \overline{\theta}]$.*

(1) There exists an optimal $D' \subset F$ which uses less money burning: The induced allocation satisfies $s'(\theta) = s(\theta)$ and $c'_1(\theta) + c'_2(\theta) \geq y(\theta)$ for all θ, with strict inequality over all Θ.

(2) If $u_*(1 - s(\theta)) \leq u^*(y(\theta))$ for all $\theta \in \Theta$, then D' can be chosen so that money burning never occurs: $c'_1(\theta) + c'_2(\theta) = 1 - s(\theta)$ for $\theta \in [\underline{\theta}, \overline{\theta}]$.

[50] In fact, $\overline{V}(k; \omega^0)$ is not differentiable at $k = \overline{k}$ since $\overline{V}(k; \omega^0)$ is constant for $k < \overline{k}$ and hence $\overline{V}'(\overline{k}-; \omega^0) = 0$.

Proof. Let $D^{\text{tc}} \subset F^{\text{tc}}$ satisfy the premise of Proposition 7. Then, as noted in the proof of Lemma 2, we can describe the doer's allocation from D^{tc} with the functions (φ, τ) that satisfy condition **(1')** and such that $0 < \varphi(\theta) < 1 - \tau(\theta)$ for all $\theta \in \Theta$ and

$$\mathcal{U}(D^{\text{tc}}) = \int_{\underline{\theta}}^{\overline{\theta}} \left[\theta u^*(\varphi(\theta)) + v(\tau(\theta))\right] \hat{g}(\theta) d\theta.$$

Now, since u is continuous and $E_y = \{\mathbf{c} \in \mathbb{R}^n_+ : \sum_{i=1}^n c_i = y\}$ is connected, $u(E_y) = [u_*(y), u^*(y)]$. Since u is strictly concave, $u_*(y) < u^*(y)$ for $y > 0$. Since u is strictly increasing, so are u_* and u^*. Clearly, u_* is continuous.

These properties imply that, for every $\theta \in \Theta$, there exists $y(\theta) \in (\varphi(\theta), 1 - \tau(\theta)]$ and $\mathbf{c}(\theta) \in E_{y(\theta)}$ such that $u(\mathbf{c}(\theta)) = u^*(\varphi(\theta))$. So, for every $\theta \in [\underline{\theta}, \overline{\theta}]$, define $t(\theta) = \tau(\theta)$ and

$$\chi(\theta) = \begin{cases} \mathbf{c}(\theta) & \text{if } \theta \in \Theta \\ \arg\max_{\{\mathbf{c} \in \mathbb{R}^n_+ : \sum_{i=1}^n c_i \le \varphi(\theta)\}} u(\mathbf{c}) & \text{if } \theta \notin \Theta \end{cases}.$$

Then, by construction the pair (χ, t) satisfy conditions **(1)** and **(2)** in the proof of Lemma 2. Now, let $D' = \{(\mathbf{c}, s) \in \mathbb{R}^n_+ : (\mathbf{c}, s) = (\chi(\theta), t(\theta)),\ \text{for some}\ \theta \in [\underline{\theta}, \overline{\theta}]\}$. We have $D' \subset F$, $\mathcal{U}(D') = \mathcal{U}(D^{\text{tc}})$, and the doer's allocation satisfies $\mathbf{c}'(\theta) = \chi(\theta)$ and $s'(\theta) = \tau(\theta)$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$. By construction, $(\mathbf{c}', s')$ satisfies the stated relationship with $(\mathbf{c}, s)$.

The last part is immediate because we can choose $y(\theta) = 1 - \tau(\theta)$ for all $\theta \in \Theta$ in the previous construction.

□

10.6 Example of Non-Additive Utility

Suppose that

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

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

$$c_i(\mathbf{r}; y) = y \frac{r_i^e}{r_1^e + r_2^e}. \tag{26}$$

We now show that $(\mathbf{c}^d, s^d)$ satisfies Condition 2—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

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

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 (26), 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 $\frac{\partial}{\partial r_j} c_i^d(\omega)$ (and similarly $\frac{\partial}{\partial r_j} c_i^p(\omega)$) to be strictly negative for all ω, a sufficient condition is that $\frac{1-\gamma}{\gamma(e-1)} < \frac{1}{\overline{s}^p}$, because $s^d(\omega) < \overline{s}^d < \overline{s}^p < 1$.[51]

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 (2) of Condition 2 holds. Suppose b_1 binds in state ω. Then, the doer's choice satisfies $c_2^*(\omega) = 1 - s^*(\omega) - b_1$ and the optimal $s^*(\omega)$ solves the first-order condition

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

Therefore,

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

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 the doer's allocation to s and c_2 is bounded away from zero for every $b_1 \leq \overline{c}_1^d$ and $u_{22}(\mathbf{c};\mathbf{r})$ and $u_{21}(\mathbf{c};\mathbf{r})$ are continuous in both arguments, it follows that $\frac{\partial s^*}{\partial b_1}$ is uniformly bounded away from zero. Therefore, part (2) of Condition 2 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.

10.7 Lemma 10

Lemma 10. *The function $v : [0,1] \to \mathbb{R}$ defined in (6) is differentiable with $v' > 0$ and strictly concave. Moreover, v' is continuous on $(0,1]$. Finally, for every $\underline{s} > 0$, there exists $y < 0$ such that $v'(s) - v'(\hat{s}) \geq y[s - \hat{s}]$ whenever $s > \hat{s} \geq \underline{s}$.*

Proof. For every $s \in [0,1]$ and $\mathbf{r} \in \Omega^2$, let $\tilde{u}(s;\mathbf{r}) = \max_{\mathbf{c} \in F(s)} u(\mathbf{c};\mathbf{r})$ and $\mathbf{c}^s(\mathbf{r})$ the unique solution to this problem, which is continuous in s and $\mathbf{r}$ by the Maximum Theorem. Since $u(\cdot;\mathbf{r})$ is strictly concave in $\mathbf{c}$, so is $\tilde{u}(\cdot;\mathbf{r})$ in s by standard arguments, which implies strict concavity of v.

Consider now differentiability of v at $s > 0$. The first-order conditions of the Lagrangian defining $\tilde{u}(s;\mathbf{r})$ say that $\frac{\partial}{\partial c_i} u(\mathbf{c}^s(\mathbf{r});\mathbf{r}) = \lambda(\mathbf{r};s)$ for $i = 1,\ldots,n$, where $\lambda(\mathbf{r};s)$ is the Lagrange multiplier for the constraint $\sum_{i=1}^n c_i \leq s$. Since $\mathbf{c}^s(\mathbf{r})$ is continuous in s for every $\mathbf{r}$, so is $\lambda(\mathbf{r};s)$ given our assumptions on u. By Theorem 1, p. 222, of Luenberger (1969), for every $s', s'' \in (0,1]$ we have

$$\lambda(\mathbf{r};s')(s' - s'') \leq \tilde{u}(s';\mathbf{r}) - \tilde{u}(s'';\mathbf{r}) \leq \lambda(\mathbf{r};s'')(s' - s'').$$

Continuity of $\lambda(\mathbf{r};\cdot)$ then implies that $\frac{\partial}{\partial s}\tilde{u}(s;\mathbf{r})$ exists for every $s > 0$ and satisfies

$$\frac{\partial}{\partial s}\tilde{u}(s;\mathbf{r}) = \frac{\partial}{\partial c_i} u(\mathbf{c}^s(\mathbf{r});\mathbf{r}) > 0.$$

Continuity of $\frac{\partial u}{\partial c_i}$ in both arguments implies that $\frac{\partial \tilde{u}}{\partial s}$ is continuous in both arguments as well. Using continuity of $\tilde{u}$ and $\frac{\partial \tilde{u}}{\partial s}$, it is immediate to show that, for every $s > 0$,

$$v'(s) = \int_{\Omega^2} \frac{\partial}{\partial s}\tilde{u}(s;\mathbf{r}) dG^2 > 0. \tag{27}$$

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

Finally, since $\frac{\partial \tilde{u}}{\partial s}$ is continuous over $(0, +\infty)$ for every $\mathbf{r}$, it follows that v' is also continuous at every $s > 0$.

Finally, using (27), note that for every $s > \hat{s} \geq \underline{s} > 0$,

$$\begin{aligned} v'(s) - v'(\hat{s}) &= \int_{\Omega^2} \left[u_c^i(c_i^s(\mathbf{r}); r_i) - u_c^i(c_i^{\hat{s}}(\mathbf{r}); r_i) \right] dG^2 = \int_{\Omega^2} u_{cc}^i(\xi(\mathbf{r}); r_i) \left[c_i^s(\mathbf{r}) - c_i^{\hat{s}}(\mathbf{r}) \right] dG^2 \\ &\geq \underline{u}_{cc}^i \int_{\Omega^2} \left[c_i^s(\mathbf{r}) - c_i^{\hat{s}}(\mathbf{r}) \right] dG^2 \geq \underline{u}_{cc}^i [s - \hat{s}], \end{aligned}$$

where the first equality uses the MVT (with $\xi(\mathbf{r}) \in [c_i^s(\mathbf{r}), c_i^{\hat{s}}(\mathbf{r})]$), the first inequality uses $\underline{u}_{cc}^i = \min_{\xi \in [\underline{s}, 1], r_i \in [\underline{r}_i, \bar{r}_i]} u_{cc}^i(\xi; r_i) < 0$ (which is well defined and bounded by continuity of u_{cc}^i), and the second inequality uses the observation that increasing savings from $\hat{s}$ to s can increase the *optimal* consumption of good i in period 2 at most by $s - \hat{s}$.

□