

Delegating Resource Allocations in a Multidimensional World

Simone Galperti*

UCSD

November 23, 2015

Abstract

This paper studies multidimensional delegation problems with a resource constraint, using a consumption-savings model. An individual delegates his short-run, present-biased self to allocate his income across multiple consumption goods and savings. The short-run self observes information on the trade-offs across goods as well as between present and future utility. Unlike when consumption involves only one good or information is only on the intertemporal utility trade-off, the individual may want to inefficiently cap his spending on specific goods—in some cases without and in others on top of committing to a minimum-savings rule. The individual uses the caps when his bias is *weak*, but only a minimum-savings rule when his bias is *strong*. These results contribute to our understanding of phenomena such as “mental budgeting” and the demand of commitment devices. The paper also discusses other applications in the areas of corporate governance, public finance, workforce management, and organization design.

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

KEYWORDS: delegation, multidimensional decision, multidimensional information, resource constraint, cap, floor, budgeting, commitment device.

*I thank Nageeb S. Ali, Eddie Dekel, Alexander Frankel, Mark Machina, Alessandro Pavan, Carlo Prato, Ron Siegel, Joel Sobel, Bruno Strulovici, and Joel Watson for useful comments and suggestions. I also thank seminar participants at Columbia University, University of Pennsylvania, Penn State University, Boston College, and UC Riverside. John N. Rehbeck provided excellent research assistance. All remaining errors are mine. First version: February 2015.

1 Introduction

In many settings, a principal delegates a better informed agent to allocate finite resources across multiple categories. For example, the owner of a company delegates a manager to allocate its cash across promoting its products and investing in R&D; an employer delegates a worker to allocate his time across job tasks and resting; an individual delegates his myopic, short-run self to allocate income across consumption goods and savings.¹ In these settings, the richness of the allocation problem is often coupled with richness of the agent's information, which affects the trade-offs across all categories under his control: The manager knows which products need advertising, the worker which tasks deserve attention, and the short-run self which goods are on sale or best suit his tastes. Finally, the principal's and the agent's interests usually disagree, but this conflict may arise only from some categories: The manager may underweight R&D, the worker may overvalue taking breaks, and the short-run self may undervalue saving for the future.

The paper studies this class of delegation problems in which there is a finite resource constraint, the agent's allocation as well as information have multiple dimensions, and we can identify one clear source of the conflict between the agent's and the principal's interests. It examines how the principal designs delegation policies to solve the fundamental trade-off between rules and discretion in these problems: By restricting the available allocations, rules can limit the consequences of the agent's bias, but at the same time can reduce the benefits of letting him act on his valuable information. The paper reveals that, in the solution to this trade-off, the resource constraint plays an important role. It investigates when and how the principal can limit the effects of the agent's bias on the final allocation by regulating dimensions which cause *no* conflict of interest; in particular, it explains how this hinges on whether the trade-offs across those dimensions are affected by the agent's information and on the strength of his bias.

For the sake of concreteness, I will present the analysis in the context of a consumption-savings model with imperfect self-control. In each period, an individual has to allocate a finite amount of resources (his income) between buying a bundle of consumption goods for the present and saving for the future. The optimal allocation depends on his varying tastes and aspects of the environment, such as prices and qualities of the goods. These tastes and aspects constitute the individual's information, which we shall call the state of the world. Importantly, part of this information affects the *intra*temporal trade-offs across goods and part of it affects the *inter*temporal trade-off between present and future utility. The individual can be interpreted as a single person, a household, the government, or the head of a department in some organization. The model is similar to that in Amador et al. (2006) and Halac and Yared (2014), but differs from their work by considering settings in which consumption involves multiple goods and information also affects the intratemporal trade-offs; also, in contrast to the second paper, here information is not persistent over time. Allowing for persistency seems an interesting but challenging extension, which is left for future research.

¹Thaler and Shefrin (1981) were the first to recognize the similarities between self-control problems and delegation problems within organizations.

Following the literature on self-control, I model the individual as having two selves with conflicting interests: a long-run self or “planner” (the principal) and a short-run self or “doer” (the agent).² The planner (she) delegates the income allocation to the doer (he), who is present biased. This bias may arise from individual preferences in the case of a single person (as in Strotz (1955) and Laibson (1997)), from aggregating heterogeneous time-consistent preferences in the case of a household or department head (as in Jackson and Yariv (2015)), or from political turnover in the case of a government (as in Aguiar and Amador (2011), for example). Anticipating the doer’s bias, the planner wants to impose some rules which restrict his choices to a subset of all feasible allocations. However, such rules often have to be set up *ex ante*, prior to learning the state. To capture this, in the model only the doer observes the state, which is non-contractible. Consequently, the planner also wants to grant the doer some discretion to act on his information. This leads to the non-trivial delegation problem of finding commitment policies that optimally balance rules and discretion. Formally, a commitment policy defines a subset of the resource constraint containing all allocations that the doer is allowed to choose. One aspect is worth noting here: Due to his present bias, the doer disagrees with the planner on the value of saving for the future; however, given any level of savings, the planner and the doer always agree on how to divide the remaining resources across goods in the present. For this reason, we will refer to savings as the conflict dimension and to the consumption goods as the agreement dimensions of the allocation problem.

Similar delegation problems exist in many other settings, in which the principal-agent conflict need not arise from their time preferences. For instance, due to compensation schemes as well as career goals, a manager of a company may be more concerned, relative to its owner, about cash flows than about developing new products. As a result, the manager may assign more importance to promoting the company’s existing products than to investing in R&D. A worker may value his break time more than does his employer, which may induce the worker to shirk on his tasks. Universities delegate professors to allocate their time between research projects and teaching, and a professor may care less about teaching than does the university.

The planner’s problem of finding an optimal commitment policy is a mechanism-design problem with multidimensional decisions and information. Such problems are notoriously hard, even without a resource constraint; in particular, here standard methods based on direct mechanisms are not helpful (see Section 2). To make any progress, some structure needs to be added to the environment. This paper allows information to have general distributions, but assumes that in the preferences no interaction occurs across dimensions of the allocation. Besides for tractability, this assumption serves to highlight—as we will see—the key role that the resource constraint plays in the problems studied here. Also, the paper mostly focuses on policies that belong to a plausible, tractable class: The planner can impose either a *cap* or a *floor* (or both) on how much the doer is allowed to allocate to savings and to each consumption good. Such policies are

²See Thaler and Shefrin (1981) for a first dual-self approach to imperfect self-control. Other papers include Benabou and Pycia (2002), Bernheim and Rangel (2004), Benhabib and Bisin (2005), Fudenberg and Levine (2006a, 2012), Loewenstein and O’Donoghue (2007), Brocas and Carrillo (2008), and Chatterjee and Krishna (2009).

the multidimensional equivalent of Holmström’s (1977) “interval controls.” As he noted, interval controls “are simple to use with minimal amount of information and monitoring needed to enforce them” and “are widely used in practice” (Holmström (1977), p. 68).³ Simple policies may be attractive in organizational contexts as well as for self-control problems: According to Thaler and Shefrin (1981), commitment “rules by nature must be simple.”⁴

The main result of the paper characterizes how the planner uses caps and floors to design optimal commitment policies. Since the doer tends to overspend on consumption but always agrees with the planner on how to divide spendable income across goods, it is intuitive to expect that optimal policies set only an aggregate limit on total consumption expenditures, which can be implemented with a minimum-savings rule (that is, a savings floor). However, this is not true. Leveraging the resource constraint, the planner can benefit from imposing specific caps that restrict spending on some goods, even though such caps will end up distorting the bundles that the doer buys. This result crucially depends on the fact that the doer has information on the intratemporal trade-offs across goods, but the correlation across his pieces of information is irrelevant—for that matter, they can be fully independent. Perhaps counterintuitively, optimal policies must always involve good-specific caps when present bias is sufficiently *weak*; by contrast, they must involve only a savings floor when present bias is sufficiently *strong*. Finally, although one may expect that a savings floor is always part of an optimal policy, the paper shows that in some settings optimal policies rely exclusively on the caps.

The intuition for the main result can be seen as follows. The information on intratemporal trade-offs implies that both the planner and the doer want to spend the most on one good in the states where its marginal utility is much higher than that of all other goods; by contrast, they want to spend the most in aggregate in states where *all* goods have high marginal utilities. An aggregate limit on total consumption deals with the doer’s bias in the latter states. However, it may not curb overspending in the former states if it is set relatively high—which is indeed what the planner wants to do when the bias is weak. To overcome this issue, she can add good-specific caps that bind only when the aggregate limit does not. Such caps force the doer to buy inefficient bundles (that is, marginal rates of substitutions do not equal price ratios), but in so doing they lower the consumption utility that he can achieve in the present and hence curb his incentive to undersave. This also highlights why the multidimensionality of consumption can help the planner design superior commitment policies, even though her consumption utility coincides with the doer’s. Nonetheless, distorting the doer’s consumption choices curbs his undersaving less and less as his bias becomes stronger. In this case, it is better to

³Alonso and Matouschek (2008) and Amador and Bagwell (2013b) provide conditions for interval controls to be fully optimal in settings with unidimensional decisions and information. No such condition is currently known for the multidimensional case.

⁴Discussing multidimensional delegation, Armstrong (1995) writes, “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, and that [...] simple results connecting the dispersion of tastes and the degree of discretion could be difficult to obtain. Moreover, in a multidimensional setting it will often be precisely the *shape* of the choice set that is of interest.” (p. 20, emphasis in the original). The results of the present paper will address the role of the dispersion of tastes as well as the shape of the choice set.

rely only on a savings floor, for it limits undersaving without distorting consumption.

The theory developed in this paper has several implications, whose details are discussed in Section 5.2. First, we may observe that an individual uses richer commitment policies with more rules than other individuals do and conclude that this is because he is more present biased. The paper shows, however, that the correlation between complexity of policies and intensity of the planner-doer conflict may actually be negative, as richer policies can work better for weaker conflicts and simpler ones for stronger conflicts.

Second, the theory provides an explicit foundation for the idea that some people use “mental budgeting” as a method to deal with their self-control problems involving savings. This idea has often been informally suggested or assumed in the behavioral literature (among others by Thaler (1985, 1999), Heath and Soll (1996), Prelec and Loewenstein (1998), and Antonides et al. (2011)).⁵ The theory also suggests some qualifications of existing views on mental budgeting. Budgets can involve goods that one would not classify as temptations, a property that Heath and Soll (1996) found puzzling in their data. Only individuals who (think to) have a *weak* present bias may rely on good-specific budgets, a prediction that finds consistent evidence in Antonides et al. (2011). This last point is also important for how we view and model self-control. To explain mental budgeting, the doer in a dual-self model cannot be fully myopic—that is, care only about the present (as in Thaler and Shefrin (1981) and Fudenberg and Levine (2006b), for example). This provides an argument in favor of modeling the doer as not completely myopic, as advocated by Fudenberg and Levine (2012).

Third, the theory expands our understanding of the demand for commitment devices (see Bryan et al. (2010) for an extensive literature review). Since the seminal work of Thaler and Shefrin (1981) and Laibson (1997), the literature has often emphasized the key role of devices like illiquid assets. In a model with one consumption good, Amador et al. (2006) showed that, under weak conditions, optimal commitment policies are fully characterized by a minimum-savings rule, which can be implemented via illiquid assets. However, when consumption involves multiple goods, minimum-savings rules can be strictly dominated by policies that (also) rely on good-specific budgets. This may explain why some individuals demand services that allow them to budget expenses by categories (such as those offered by companies like Mint.com, Quicken.com, and StickK.com). Markets of such services cannot be explained by existing theories.

The present theory also offers insights for the aforementioned, similar delegation problems in the areas of public finance, corporate governance, and organization design. Besides imposing budget-deficit ceilings, governments may inefficiently cap their spending on some public services so as to curb their tendency to overborrow. The owner of a company may inefficiently cap how much its manager is allowed to spend on promoting specific products so as to induce him to invest more in R&D—possibly in addition to requiring minimum levels of R&D. Besides restricting the number of allowed breaks per day, employers may require workers to spend a minimum amount of time on certain tasks

⁵Bénabou and Tirole (2004) provide an explanation based on self-reputation of why personal rules (such as mental budgets) may be effective commitment devices. Benhabib and Bisin (2005) model how ex-ante plans can trigger internal control processes that prevent impulsive processes from deviating from such plans. Bernheim et al. (2015) offer an alternative analysis of self-enforcing commitment rules.

so as to lead them to take few breaks. Section 7 discusses these applications in detail.

Since standard direct mechanisms prove unsuited for solving the planner’s problem, to derive the above results the paper follows a different approach. First, it divides the problem into simpler parts, each examining how restricting one dimension of the doer’s decision at a time affects the planner’s payoff. Intuitively, to infer the effects of tightening a cap or a floor, the paper exploits the information contained in the Lagrange multipliers for the constraints that caps and floors add to the doer’s optimization problem. Having understood how the planner would use caps and floors for each dimension in isolation, the paper then shows how optimal policies combine them depending on the strength of the doer’s bias. The analysis uses the assumption that information has a distribution with continuous support, but otherwise allows for mass points and general forms of dependence across its components.

To highlight the role that information on intratemporal trade-offs plays in the theory, Section 6 studies the case in which information is only about the intertemporal utility trade-off (as in Amador et al. (2006) and Halac and Yared (2014)), but consumption is still multidimensional. In this case, most of the time the optimal commitment policy involves only a savings floor—especially when the doer’s bias is weak! Thus, the multiplicity of consumption goods no longer matters for designing commitment policies. The key step here is to show that when information is only on the intertemporal trade-off, it is without loss of generality to focus on policies that take the form of *any* subset of feasible allocations defined only in terms of savings and *total* consumption expenditures. Given this, we can invoke Amador et al.’s (2006) main result which says that, under a weak condition on the information distribution, optimal policies coincide with a minimum-savings rule.

Finally, one might wonder what happens when Amador et al.’s (2006) condition fails. In this case, they argue that optimal policies may have to involve “money burning,” in the sense of inducing allocations strictly inside the resource constraint.⁶ The present paper shows that, by exploiting the multidimensionality of consumption, the planner can design optimal policies which require less money burning. It also provides a sufficient condition for money burning to be superfluous.

2 Related Literature

This paper contributes to the mechanism-design literature on the trade-off between rules (commitment) and discretion (flexibility) and its numerous applications. The closest papers are Amador et al. (2006) and Halac and Yared (2014).⁷ Relative to both papers, the present theory shows how dimensions of the agent’s decisions causing no conflict with the principal and idiosyncratic information on them can play an important role in the solution to the rules-discretion dilemma, opening the door to superior (yet simple) commitment policies. Relative to Halac and Yared (2014), another difference should be

⁶Other papers that study money burning as a tool to shape incentives in delegation problems include Ambrus and Egorov (2009), Ambrus and Egorov (2013), Amador and Bagwell (2013a), Amador and Bagwell (2013b).

⁷See also Athey et al. (2005), Ambrus and Egorov (2013), Amador and Bagwell (2013b).

noted. They also consider consumption-savings problems with a present-biased doer in which, however, consumption involves one good and information is only on the intertemporal trade-off and is persistent over time. In such environments, optimal commitment policies can distort future consumption levels, even though they cause no conflict between planner and doer once present consumption (the conflict dimension) is fixed. This may resemble the results of the present paper, but the economics is totally different. With information persistency, the doer’s expected utility from future choices is linked to his present information, and hence the planner can leverage those choices to relax his incentive constraints, as in other dynamic mechanism-design problems.⁸ By contrast, in the present paper agreement dimensions may be distorted to exploit the link with other dimensions created by the resource constraint, and this is true even if the doer’s pieces of information are fully independent. Finally, Brocas and Carrillo (2008) briefly discuss a consumption-savings model in which consumption involves two goods, one of which has ex-ante uncertain utility, and the doer is fully myopic. This model does not allow for an analysis that examines the roles of information on intratemporal and intertemporal trade-offs and of the intensity of the doer’s bias.

This paper also relates to the rich literature on principal-agent delegation problems following Holmström (1977, 1984). The work on problems with multidimensional decisions and information, however, is very scarce and does not examine the kind of settings studied here.⁹ In Koessler and Martimort (2012), information has one dimension, decisions have two unconstrained dimensions, and payoffs are quadratic, which implies that they ultimately depend on the decision’s mean and spread. This turns the problem into a screening exercise where the spread works as a pseudotransfer. In Frankel (2014, 2015), both information and decisions are multidimensional. In Frankel (2014), the agent has the *same* bias for all dimensions, but the principal is uncertain about its properties (strength, direction, etc.). In this case, the best policies against the worst-case bias (max-min policies) may require the agent’s *average* decision to satisfy a preset value (called “budget”). Such policies should not be confused with the caps and floors in the present paper. Frankel (2015) considers policies that set a cap not directly on decisions, but on the gap between the agent’s and the principal’s final payoffs. Under the assumption that information is i.i.d. across dimensions, such policies are optimal in some settings and, more generally, ensure that with a large number of dimensions the *per-dimension* loss from the first best is small. By contrast, the present paper allows for general distributions and considers a different class of delegation policies, which can dominate those in Frankel (2015). It also provides useful results for settings with a small number of dimensions.

Finally, this paper relates to the literature on multidimensional screening (see Stole and Rochet (2003) for a detailed survey). Screening and delegation problems differ because in the latter the principal cannot use transfers. This prevents us from applying the insights of that literature here and forces us to restrict the class of delegation policies. In the present paper, we could use the dual approach and other methods in Rochet and Choné (1998) to simplify the agent’s incentive constraints and the principal’s objective.

⁸See, for example, Courty and Li (2000), Battaglini (2005), Pavan et al. (2014).

⁹In Alonso et al. (2013), a fixed amount of resources has to be allocated across multiple categories, but each category is controlled by a different agent with a unidimensional piece of information.

However, the fixed resources add a state-wise constraint to the problem. General techniques exist for such problems (for example, Luenberger (1969)), but are not helpful in our case to characterize the optimal unrestricted mechanisms, which need not follow the same logic of the unidimensional case—as suggested by Rochet and Choné (1998)—or, for that matter, of monopolistic screening. One benefit of the approach of this paper is that its results are insensitive to details of the information structure. By contrast, such details can matter significantly in the theory of multidimensional screening (see Manelli and Vincent (2007)).

3 The Model

This section introduces the model as a consumption-savings problem with imperfect self-control. Section 3.1 explains how this model can capture many other settings that lead to similar delegation problems.

Consider an individual who lives for multiple periods and in each period, given his income, chooses a bundle of consumption goods for the present and a level of savings for the future. Hereafter, we focus on one of these periods. The individual’s per-period income is known and is normalized to 1. His consumption bundle involves $n > 1$ goods and is represented by $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}_+^n$; his level of savings is represented by $x_0 \in \mathbb{R}_+$.¹⁰ The set of feasible allocations is then

$$B = \{(\mathbf{x}, x_0) \in \mathbb{R}_+^{n+1} : \sum_{i=0}^n x_i \leq 1\}.$$

How the individual wants to allocate his income depends on his tastes as well as other aspects of the environment. These tastes and aspects define the individual’s information, which affects in part how he trades off consumption goods in the present and in part how he trades off present vs. future utility. Let the first part on the *intra*temporal trade-offs be represented by $\mathbf{r} \in \mathbb{R}^n$, where r_i is the idiosyncratic component on good i ; let the second part on the *inter*temporal trade-off be represented by $\theta \in \mathbb{R}$. This setting is similar to that in Amador et al. (2006) (hereafter, AWA), but departs from it in two ways: First, consumption involves multiple goods, whereas in AWA it involves only one good (that is, $n = 1$); second, information also affects the trade-offs across consumption goods, whereas in AWA it consists only of θ .

As in existing dual-self models of self-control (see Footnote 2), the individual consists of a long-run self or “planner” (she) and a short-run self or “doer” (he) with conflicting preferences. Given (θ, \mathbf{r}) , the planner’s and the doer’s utility functions are respectively

$$\theta u(\mathbf{x}; \mathbf{r}) + v(x_0) \quad \text{and} \quad \theta u(\mathbf{x}; \mathbf{r}) + bv(x_0),$$

where the function $v : \mathbb{R}_+ \rightarrow \mathbb{R}$ represents the expected utility from saving x_0 for the future. The assumption here is that information is independently distributed across periods (in contrast to Halac and Yared (2014)) and that the individual anticipates how he will use his savings in the future.¹¹ The parameter $b \in (0, 1)$ captures in a tractable

¹⁰The model can be extended to allow for negative savings (or borrowing) along the lines of Halac and Yared (2014) (see Section 3.1).

¹¹Under this independence assumption, this reduced-form one-period formulation can be derived from a full-fledged multi-period model along the lines of Amador et al. (2003).

way the doer’s bias towards present consumption, and hence the conflict with the planner. The level of b is known to the planner.¹² This formulation of present bias is consistent with the single-agent, quasi-hyperbolic discounting model (Laibson (1997)), but also with interpreting the individual as the representative of a group—like a household, a political committee, or the board of a firm—who aggregates the preferences of its heterogeneous, time-consistent members. In this case, Jackson and Yariv (2015) show that under weak conditions the resulting aggregate preference exhibits present bias. To add clarity and tractability to the model, assume that

$$u(\mathbf{x}; \mathbf{r}) = \sum_{i=1}^n u^i(x_i; r_i) \quad \text{with} \quad u_{12}^i > 0 \text{ for all } i.$$

Additive separability may rule out realistic interactions across goods, but it will help to isolate the mechanisms of interest for this paper. This assumption is superfluous for some of the results below, which I will point out; I expect that the other results are robust at least to moderate interactions across goods (see below).

The information structure involves some degree of redundancy, as both an increase in θ and an increase in all components of \mathbf{r} render consumption more valuable. This approach, however, has several advantages: It clarifies the conceptual distinction between information on the intratemporal and intertemporal trade-offs; it will allow us to keep the same model for the entire analysis, thereby focusing on the core messages of the paper; it simplifies the comparison with the literature. Note that the model allows us to interpret information as taste shocks or as the observation of prices, which determine how dollars spent on good i , x_i , translate into its physical units.¹³

The planner (principal) delegates the doer (agent) to choose an income allocation. Knowing his bias, the planner would like to design a commitment policy dictating which allocations the doer is allowed to implement. In general, such a policy defines a nonempty subset D of the feasibility set B . For example, $D = B$ grants the doer full discretion, whereas a set D containing only one element grants him no discretion at all. If the planner can condition D on the realization of (θ, \mathbf{r}) —either because she observes it or because it is contractible—her problem would be trivial, as she can let D contain only her most preferred allocation given that information. However, in reality commitments are usually chosen prior to observing all the necessary information to make a decision. One way to capture this is to assume that information (θ, \mathbf{r}) is non-contractible and only the doer observes it. This creates a non-trivial delegation problem in which the planner faces a trade-off between rules and discretion: She wants to limit the doer’s freedom, but at the same time let him act on his information. In short, the timing is as follows:

¹²Ali (2011) concludes that sophistication is a plausible assumption in frameworks similar to the one considered here, as by experimenting the individual can learn his true bias. Allowing for partial naiveté (for example, as in O’Donoghue and Rabin (2001)) would not change the message of the paper, as the entire analysis is from the ex-ante viewpoint of the planner.

¹³For instance, for all $i = 1, \dots, n$, let $\bar{u}^i(z_i) = \frac{z_i^{1-\gamma_i}}{1-\gamma_i}$ with $\gamma_i > 0$ be the utility from z_i units of good i , and let $\pi_i > 0$ be its price. If we define $x_i = \pi_i z_i$ for all i , we can write the usual resource constraint as $\sum_{i=1}^n x_i + x_0 \leq 1$. Letting $r_i = \pi_i^{\gamma_i - 1}$, we can define $u^i(x_i; r_i) = r_i \bar{u}^i(x_i)$, which satisfies all our assumptions. This example can be generalized by allowing each \bar{u}^i to be a smooth, strictly increasing, and strictly concave function.

First, the planner commits to a policy D . Then, the doer observes (θ, \mathbf{r}) and implements some $(\mathbf{x}, x_0) \in D$, which determines the payoffs. The planner designs D to maximize her expected payoff from the doer's resulting decisions. Section 4 will be more precise about the class of policies examined in this paper.

The following assumptions complete the model and are mostly technical:

- *Information:* Let $S = [\underline{\theta}, \bar{\theta}] \times [\underline{r}_1, \bar{r}_1] \times \cdots \times [\underline{r}_n, \bar{r}_n]$, where $0 < \underline{\theta} < \bar{\theta} < +\infty$ and $-\infty < \underline{r}_i < \bar{r}_i < +\infty$ for all $i = 1, \dots, n$. The joint distribution of (θ, \mathbf{r}) is represented by the probability measure G which has full support over S ; that is, $G(S') > 0$ for every open $S' \subset S$.¹⁴ Note that this assumption allows for rich forms of dependence across information components, but of course also for full independence across all of them.
- *Differentiability, monotonicity, concavity:* The function $v : \mathbb{R}_+ \rightarrow \mathbb{R}$ is twice continuously differentiable with $v' > 0$ and $v'' < 0$. For all $i = 1, \dots, n$, the function $u^i : \mathbb{R}_+ \times [\underline{r}_i, \bar{r}_i] \rightarrow \mathbb{R}$ is twice differentiable with $u_1^i(\cdot; r_i) > 0$ and $u_{11}^i(\cdot; r_i) < 0$ for all $r_i \in [\underline{r}_i, \bar{r}_i]$; also, u_1^i and u_{11}^i are continuous in both arguments.
- *Boundary conditions:* $\lim_{x_0 \rightarrow 0} v'(x_0) = +\infty$ and $\lim_{x_i \rightarrow 0} u_1^i(x_i; r_i) = +\infty$ for all $r_i \in [\underline{r}_i, \bar{r}_i]$ and $i = 1, \dots, n$. This will allow us to focus on interior solutions.

Notation and Definitions

In the rest of the paper, we shall call every element of S state and denote it by $\mathbf{s} = (\theta, r_1, \dots, r_n)$. Given this, express the planner's and the doer's payoffs from allocation (\mathbf{x}, x_0) as

$$U(\mathbf{x}, x_0; \mathbf{s}) = \hat{u}(\mathbf{x}; \mathbf{s}) + v(x_0) \quad \text{and} \quad V(\mathbf{x}, x_0; \mathbf{s}) = \hat{u}(\mathbf{x}; \mathbf{s}) + bv(x_0), \quad \mathbf{s} \in S, \quad (1)$$

where $\hat{u}(\mathbf{x}; \mathbf{s}) = \theta \sum_{i=1}^n u^i(x_i; r_i)$. For each \mathbf{s} , let $\mathbf{p}^*(\mathbf{s})$ denote the allocation that the planner would like the doer to choose in that state and $\mathbf{a}^*(\mathbf{s})$ denote the allocation that the doer (the agent) actually chooses under full discretion:

$$\mathbf{p}^*(\mathbf{s}) = \arg \max_B U(\mathbf{x}, x_0; \mathbf{s}) \quad \text{and} \quad \mathbf{a}^*(\mathbf{s}) = \arg \max_B V(\mathbf{x}, x_0; \mathbf{s}), \quad \mathbf{s} \in S. \quad (2)$$

Hereafter, we refer to \mathbf{p}^* as the *first-best allocation* and to \mathbf{a}^* as the *full-discretion allocation*. These allocations satisfy some properties, summarized in the following lemma, which we will use later.

Lemma 1.

- Both \mathbf{p}^* and \mathbf{a}^* are continuous in \mathbf{s} .
- For all $i = 0, \dots, n$, the range of p_i^* (resp. a_i^*) equals an interval $[\underline{p}_i^*, \bar{p}_i^*]$ (resp. $[\underline{a}_i^*, \bar{a}_i^*]$), with $0 < \underline{p}_i^* < \bar{p}_i^* < 1$ (resp. $0 < \underline{a}_i^* < \bar{a}_i^* < 1$).
- For all $\mathbf{s} \in S$, $a_0^*(\mathbf{s})$ is continuous and strictly increasing in b and $a_0^*(\mathbf{s}) < p_0^*(\mathbf{s})$ if and only if $b < 1$.
- For all $\mathbf{s} \in S$, each $a_i^*(\mathbf{s})$ is continuous and strictly decreasing in b .

These properties follow immediately from (1) the assumptions on U and V , (2) compactness, connectedness, and convexity of S , and (3) standard comparative-statics arguments.

¹⁴This holds, for instance, if G has a strictly positive and continuous density function over S .

Two other observations will be useful later. First, all consumption goods are normal: For both the planner and the doer, higher spendable income $(1 - x_0)$ always leads to a higher optimal allocation to each good.¹⁵ Second, if we fix savings, the planner and the doer always agree on how to divide the remaining resources across goods: Given $\hat{x}_0 \in [0, 1]$,

$$\arg \max_{\{(\mathbf{x}, x_0) \in B: x_0 = \hat{x}_0\}} U(\mathbf{x}, x_0; \mathbf{s}) = \arg \max_{\{(\mathbf{x}, x_0) \in B: x_0 = \hat{x}_0\}} V(\mathbf{x}, x_0; \mathbf{s}), \quad \mathbf{s} \in S.$$

Hence, we will refer to x_0 as the conflict dimension and to \mathbf{x} as the agreement dimensions.

3.1 Alternative Interpretations and Applications

This section outlines other settings that create delegation problems which can be modeled using the above framework. The reader interested in the results can skip this part.

Corporate governance. The owner of a company appoints a manager, who each year decides how to divide some total resources, I , between investment in R&D, x_0 , and overall spending on sales activities, y . The company sells multiple products and the manager also chooses which share of y goes to promoting which product (x_1, x_2 , etc.). As a result of compensation schemes and career goals, the manager may be more concerned with the company's cash flows than the owner is; hence, the manager may assign relatively more importance to current sales than to R&D ($b < 1$). Since he manages the business on a daily basis, he has better information on the returns from marketing each product (\mathbf{r}) as well as from funding R&D (θ), information that the owner would like to be incorporated in the allocation of I . However, due to the manager's bias, the owner may restrict his choices to some subset D of all feasible allocations B .

Workers' time management. An employer hires a worker under a contract that specifies a workday of I hours and a fixed wage. The worker is in charge of multiple tasks and chooses how to allocate his time across them (x_1, x_2 , etc.). Moreover, he can take breaks during the day, represented by x_0 . Being on the shop floor, the worker has firsthand information on which task demands more attention at each moment. Given this, the employer would like to let him choose how to allocate his time. However, the worker is likely to weigh his benefits from taking breaks more than does the employer (that is, $b > 1$ in the model).¹⁶ Thus, she may also want to set up some rules to avoid that the worker spends too much time on breaks. We can model such rules with a subset D of the feasible time allocations B .

Fiscal-constitution design. Society delegates a government to allocate the economy resources, I , between private consumption, x_0 , and total public spending, y . The government incorporates the preference of a representative agent in society, but is biased in favor of public spending ($b < 1$).¹⁷ The government spends y to fund multiple services under its control (x_1, x_2 , etc.). Although it is biased in favor of public spending, at a first approximation the government may not favor any specific service more than others rel-

¹⁵Given our assumptions, this property follows, for instance, from Proposition 1 in Quah (2007).

¹⁶The paper focuses on the case of $b < 1$, but Section 7 explains how the analysis works for $b > 1$.

¹⁷This 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)).

ative to society’s representative agent.¹⁸ Thus, given any level of y , the parties agree on how to allocate y across services. Before the government chooses an allocation, it observes non-contractible information which affects the social value of each service, captured by \mathbf{r} (for example, threats to national security or needs for natural-disaster relief), as well as the overall trade-off between private consumption and public spending, captured by θ (for example, the state of the business cycle). Due to these different goals and information, behind a veil of ignorance society may want to design a fiscal constitution—that is, a delegation policy D —that specifies which allocations the government can choose.

Research vs. teaching in academia. A university employs a professor to teach and conduct research. Each week, the professor has a total amount of hours I that he can allocate to research, y , or teaching, x_0 . Also, he works on several research projects and has to choose how much of y to spend on each of them (x_1, x_2 , etc.). The professor may care about teaching less than does the university ($b < 1$). Nonetheless, he has better information on which activity is more likely to advance his as well as the university’s interests on an ongoing basis. Thus, the university would like to let the professor choose how to allocate his time, but also establish some rules to limit the risk that he overlooks teaching. We can capture such rules with a subset D of the professor’s feasible time allocations B .

Public finance. As in Halac and Yared (2014), each year t the government chooses how much to borrow, z^t , and spend, y^t , subject to the constraint $y^t \leq \tau + z^t/\rho - z^{t-1}$, where z^{t-1} is the nominal debt inherited from period $t-1$, τ is a fixed tax revenue, and ρ is an exogenous (gross) interest rate. Differently from their setup, here the government divides y^t across multiple services, \mathbf{x}^t , and its information, (θ^t, \mathbf{r}^t) , is i.i.d. over time. At the beginning of each year, before observing (θ^t, \mathbf{r}^t) , the government evaluates allocations using the function $\theta^t u(\mathbf{x}^t; \mathbf{r}^t) + \hat{v}(z^t)$, where $\hat{v}(z^t)$ is the expected payoff from entering period $t+1$ with debt z^t . After observing (θ^t, \mathbf{r}^t) , however, it chooses allocations using the function $\theta^t u(\mathbf{x}^t; \mathbf{r}^t) + b\hat{v}(z^t)$ with $b \in (0, 1)$. The government’s present bias can arise when it aggregates the preferences of heterogeneous citizens, even if they are all time consistent (as in Jackson and Yariv (2015)), or because of uncertainty in the political turnover (as in Aguiar and Amador (2011), for example). Anticipating its inconsistency, ex ante the government may commit to some fiscal rules. Since information is not persistent, considering fiscal rules that bind only for one year is without loss of generality (Amador et al. (2003); Halac and Yared (2014)). To map this setting into the previous model, we can assume an exogenous upper bound on borrowing $Z < +\infty$, let $x_0^t = -z^t/\rho$ and $I(x_0^t) = \tau + \rho x_0^t$, and define $v(x_0^t) = \hat{v}(-\rho x_0^t)$. The feasibility constraint becomes $\sum_{i=1}^n x_i^t + x_0^t \leq I(x_0^{t-1})$. At the beginning of each year t , given x_0^{t-1} the government can design fiscal rules that restrict the implementable allocations to some $D \subset B(x_0^{t-1})$.

¹⁸This property is arguably strong, but it is consistent with some empirical evidence. For example, Peltzman (1992) finds that U. S. voters penalize federal spending growth, but its composition seems irrelevant.

4 Tractable Delegation Policies: Caps and Floors

In principle, we would like to find the planner’s best policy among all possible commitment rules, that is, all $D \subset B$. The usual mechanism-design approach would rely on the revelation principle to turn the problem of finding an optimal *set* $D \subset B$ into the equivalent problem of finding an optimal, incentive compatible, and resource-feasible allocation *function* of the state. Searching among such functions is usually easier than searching among sets, but in the present setting it remains an intractable problem as explained in Section 2.

This calls for a different approach. As Holmström (1977) noted, “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 setting where the agent’s decision is unidimensional, 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.”¹⁹ 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.” With regard to the consumption-savings application, Thaler and Shefrin (1981) argue that commitment “rules by nature must be simple.”²⁰ For these reasons, this paper focuses on the class of policies that correspond to the multidimensional version of Holmström’s intervals, that is, rectangles. Of course, one may want to consider other classes of policies. However, it is difficult to identify classes of sets which are sensible for our setting and have enough structure to render the problem tractable.

The class of rectangle policies implies that the planner can commit to imposing a cap or a floor on how much the doer is allowed to save and to spend on each consumption good. Formally, given $\mathbf{f}, \mathbf{c} \in [0, 1]^{n+1}$ that satisfy $f_i \leq c_i$ for all i and $\sum_{i=0}^n f_i \leq 1$, a rectangle policy is defined by

$$D_{\mathbf{f}, \mathbf{c}} = \{(\mathbf{x}, x_0) \in B : f_i \leq x_i \leq c_i \text{ for all } i\}.$$

Denote the collection of all such policies by \mathcal{R} . Given $D_{\mathbf{f}, \mathbf{c}}$, some floors and caps may never affect the doer’s decision and others may constrain it only in some states. Therefore, when describing a policy from the ex-ante viewpoint, I will call a floor (or cap) *binding* if it constrains the doer in a set of states with strictly positive probability. When considering policies in \mathcal{R} , I will also leave \mathbf{f} and \mathbf{c} implicit unless required by the circumstances.

The planner has to choose $D \in \mathcal{R}$ so as to maximize

$$\mathcal{U}(D) = \int_S U(\mathbf{a}(\mathbf{s}); \mathbf{s}) dG \tag{3}$$

subject to

$$\mathbf{a}(\mathbf{s}) \in \arg \max_{(\mathbf{x}, x_0) \in D} V(\mathbf{x}, x_0; \mathbf{s}), \quad \mathbf{s} \in S. \tag{4}$$

¹⁹For Holmström’s (1977) settings, the literature has recently identified conditions for intervals to be optimal among all possible delegation policies (see Footnote 3). Similar conditions are not available, however, for multidimensional settings like in the present paper.

²⁰Benhabib and Bisin (2005) provide a rationale for why people may prefer simple commitment rules based on higher psychological costs of complying with complex rules.

This problem has a solution. The proof of this result as well as all the others appear in the Appendix.

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

Given this, we can turn to characterizing the optimal policies. To this end, we will first examine the effects of restricting savings and spending on each consumption good in isolation. This will provide useful insights for understanding how she combines caps and floors, thus paving the way for the main result of the paper in Section 5.

4.1 Restricting Conflict Dimensions

The planner and the doer disagree on how much they value present vs. future utility and hence on how income should be allocated between consumption and savings. Since present bias leads to overconsumption, it seems intuitive that the planner wants to restrict the total amount that the doer spends on consumption. A simple way to implement such a constraint is to impose a floor on savings. When binding, the floor prevents the doer from splurging, but never affects how he divides spendable income across goods, a decision which raises no conflict.

The following preliminary result shows that, *if the planner can only restrict* how much the doer is allowed to save, then a binding floor f_0 strictly improves on the full-discretion policy ($D = B$). Also, the optimal f_0 is *strictly* higher than the lowest first-best savings level ($f_0 > \underline{p}_0^*$). This result relies only on monotonicity and concavity in \mathbf{x} of the function \hat{u} in (1). Hence, it continues to hold if we allow for general interactions across goods and dependence on \mathbf{s} .²¹

Lemma 3. *When the only available delegation policies involve a floor on x_0 , it is optimal to set f_0 strictly between \underline{p}_0^* and \bar{p}_0^* .*

In some settings, for practical reasons it may be possible to restrict only x_0 . For example, in the application to research vs. teaching in academia, a university can easily request and monitor that a professor allocates at least f_0 hours per week to teaching, but may not be able to restrict the time that she spends on each of her research projects.

The intuition for Lemma 3 is simple. On the one hand, under full discretion the doer saves strictly less than \underline{p}_0^* for some states, which is never justifiable from the planner's viewpoint. On the other hand, fixing any x_0 , both parties always agree on how to divide $1 - x_0$ across consumption goods; therefore, even when f_0 is binding, it never distorts the final consumption bundle. These two observations imply that the planner always wants to set $f_0 \geq \underline{p}_0^*$. To see why this inequality must be strict, note that for states in which the planner would save $x_0 > \underline{p}_0^*$, she strictly prefers setting $f_0 = x_0$ than $f_0 = \underline{p}_0^*$, which leads to overconsumption. By contrast, for states in which the planner would save $x_0 = \underline{p}_0^*$, setting $f_0 = \underline{p}_0^*$ already induces the doer to choose the first-best allocation. Therefore, a marginal increase in f_0 above \underline{p}_0^* has a first-order positive

²¹Proposition 6 will imply that it is never optimal to impose a binding cap on savings (that is, $c_0 < \bar{a}_0^*$).

effect and only a second-order negative effect on the planner's payoff. A similar logic explains the inequality $f_0 < \bar{p}_0^*$.

Lemma 3 is reminiscent of the results in AWA, but differs in several respects. In AWA's setting both consumption and information are unidimensional, which allows AWA to consider *all* subsets of the resource constraint as feasible delegation policies and show that the optimal one must always involve a binding minimum-savings rule. AWA derive their result through a clever application of mechanism-design techniques. As is well known,²² similar techniques are not available for the present multidimensional setting, which significantly complicates the problem of finding the optimal policy among all sets $D \subset B$. Therefore, compared to AWA, this paper restricts attention to policies which can involve only a floor or cap on savings, and to prove the optimality of a binding floor, it relies on different techniques.

The key step is to show that the planner's payoff is differentiable in the floor f_0 . If we let D_{f_0} be the policy which involves only f_0 , Lemma 10 in the Appendix shows that $\frac{d}{df_0}\mathcal{U}(D_{f_0})$ exists and provides a simple expression for it. Of course, at an interior optimum we must have $\frac{d}{df_0}\mathcal{U}(D_{f_0}) = 0$. The expression of $\frac{d}{df_0}\mathcal{U}(D_{f_0})$ shows that this first-order condition captures the following trade-off. Consider $f_0 \in (\underline{p}_0^*, \bar{p}_0^*)$ and the consequences of increasing it marginally. In some states the planner would save more than f_0 (that is, $p_0^*(\mathbf{s}) > f_0$), and hence benefits from increasing the doer's savings when f_0 binds. In other states the planner would save less than f_0 (that is, $p_0^*(\mathbf{s}) < f_0$), and hence loses by inducing the doer to save even more than in the first best. The first-order condition says that, when f_0 affects the doer's decision, the expected benefit for the states that demand higher savings should be equal to the expected loss for the states that demand lower savings.

To derive the main result below, it will be useful to know how the optimal floor changes as the doer becomes more biased. Everything else equal, the planner should tighten f_0^* . Indeed, it turns out that $\frac{d}{df_0}\mathcal{U}(D_{f_0})$ decreases in b , and hence $\mathcal{U}(D_{f_0})$ is a submodular function of (f_0, b) . Intuitively, as the bias worsens, the doer penalizes savings more; so any f_0 is more likely to bind. This strengthens the expected benefit of raising f_0 for the states where the planner would save more than f_0 , but does *not* change the expected cost of raising f_0 for the states where she would save less than f_0 : In such states f_0 binds for any bias, as the doer always prefers to save less than does the planner.

Lemma 4. *The set of optimal floors, denoted by $F(b)$, is decreasing in b in the strong set order.²³ In particular, $\max F(b)$ is decreasing in b and converges to \underline{p}_0^* as $b \uparrow 1$. Moreover, there exists $\underline{b} > 0$ such that $F(b) = \{\bar{f}_0\}$ for all $b \leq \underline{b}$, where \bar{f}_0 satisfies*

$$\mathcal{U}(D_{\bar{f}_0}) = \max_{f_0 \in [\underline{p}_0^*, \bar{p}_0^*]} \left\{ v(f_0) + \int_S \hat{u}(\mathbf{x}^{f_0}(\mathbf{s}); \mathbf{s}) dG \right\}$$

and

$$\mathbf{x}^{f_0}(\mathbf{s}) = \arg \max_{\{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i \leq 1 - f_0\}} \hat{u}(\mathbf{x}; \mathbf{s}), \quad \mathbf{s} \in S.$$

²²See, for example, Rochet and Choné (1998) and their discussion on direct and dual approaches to screening problems.

²³Given two sets F and F' in \mathbb{R} , $F \geq F'$ in the strong set order if, for every $f \in F$ and $f' \in F'$, $\min\{f, f'\} \in F'$ and $\max\{f, f'\} \in F$ (Milgrom and Shannon (1994)).

Thus, when the bias is sufficiently strong, the planner sacrifices entirely the option of letting the doer adjust savings to information. This happens even though the doer does care about the future and hence would adjust x_0 to the state.

4.2 Restricting Agreement Dimensions

Besides by imposing a savings floor, the planner can limit the doer’s splurging by restricting consumption directly. When there is only one consumption good, a saving floor is always equivalent to a consumption cap, and hence nothing else can be done within the class of rectangle policies. This is no longer true when consumption involves multiple goods. A savings floor is equivalent to a limit on *total* consumption (that is, $\sum_{i=1}^n x_i \leq 1 - f_0$), but there are many ways to enforce this limit using good-specific caps. Moreover, the planner can impose such caps *in addition to* an aggregate limit on consumption (via f_0). For example, she can require that consumption never exceed 80% of her income and that “going out” alone never exceed 10%, or “gifts” never exceed 5%.” Since the planner and the doer have the same preference when it comes to dividing money across goods, we may expect specific caps to be useless, or even harmful. This section examines this conjecture, starting with two observations.

First of all, good-specific caps cannot implement the same allocations as a policy which relies only on a binding savings floor—unless information affects only the intertemporal utility trade-off and the component \mathbf{r} is known from the outset. A formal statement of this property appears in the Appendix (Lemma 11). To see the intuition, consider a binding floor f_0 . Since \mathbf{r} can vary, the doer does not always choose the same consumption bundle while f_0 binds: As the marginal utility of some good i rises (higher r_i), he spends on it a larger share of $1 - f_0$. Thus, if the planner wants to let him respond to \mathbf{r} as under f_0 , she has to set a cap on each good at least as large as the highest amount of that good consumed under f_0 . By the previous observation, however, the sum of such caps must exceed $1 - f_0$, and hence cannot ensure that the doer saves at least f_0 . Therefore, to ensure this minimum savings, some caps must be set strictly lower and hence prevent the doer from responding to \mathbf{r} as under f_0 . Since the doer always chooses an efficient consumption bundle under f_0 —that is, marginal utilities coincide across all goods—it follows that trying to implement f_0 using good-specific caps must induce the doer to choose inefficient bundles.

The second observation is that, when binding, a good-specific cap mitigates the doer’s *aggregate* overconsumption, but without other constraints it also exacerbates overconsumption in all other goods. Lemma 12 in the Appendix states this formally. To see the logic, note that since income is fixed, overspending on good j comes at the cost of subtracting money from savings, which the doer undervalues, or from other goods like i , which he values on par with j . When good i is already capped, however, the second cost decreases, inducing the doer to overspend on j even more. Note the key role that the resource constraint plays in this logic, which is further highlighted by the absence of any utility interaction across goods.

Given these observations, one may wonder whether the planner can ever benefit from

imposing specific caps. To answer this question, we will use the next result, which shows that capping even only one good strictly dominates granting the doer full discretion.

Lemma 5. *Fix $i \neq 0$ and consider policies $D_{\mathbf{0},\mathbf{c}}$ with $c_j = 1$ for all $j \neq i$. There exists $c_i < \bar{a}_i^*$ such that the planner strictly benefits from it, that is, $\mathcal{U}(D_{\mathbf{0},\mathbf{c}}) > \mathcal{U}(B)$.*

To gain intuition, start from $c_i = \bar{a}_i^*$ and imagine lowering it a bit. On the one hand, when binding, the cap distorts the allocation across consumption goods. This reduces the planner's expected payoff, but this loss is initially of second-order importance. The reason is that, under full discretion, the doer's allocation across \mathbf{x} is always efficient; moreover, both parties have the same preference regarding \mathbf{x} . Hence, marginal distortions in \mathbf{x} do not change the planner's payoffs. On the other hand, the cap induces the doer to save more with strictly positive probability. Since the doer undersaves from the planner's viewpoint, this reallocation causes a first-order gain in her payoff. Overall the cap should then benefit the planner, *provided* that the doer does not reallocate money to the unrestricted goods at a much faster rate than to savings, which is not obvious. This key property is guaranteed by the additive structure of preferences. It should continue to hold if the goods are complements: Capping one of them renders all the others less valuable and hence should incentivize the doer to save even more. In this case, the thrust of the paper does not change.

Lemma 5 suggests that the planner may always benefit by combining a savings floor with good-specific caps. Perhaps surprisingly, this depends on the strength of the doer's bias and the nature of his information: Whether it only affects the intertemporal utility trade-off (θ) or also the intratemporal trade-offs across goods (\mathbf{r}). Before deriving these results in the next section, Lemma 6 shows that binding caps on savings or floors on consumption goods are never part of the planner's optimal policy.

Lemma 6. *For any $D_{\mathbf{f},\mathbf{c}} \in \mathcal{R}$, let $D_{f_0,\mathbf{c}_{-0}}$ be the policy obtained by removing the savings cap and all good-specific floors. Then $\mathcal{U}(D_{f_0,\mathbf{c}_{-0}}) \geq \mathcal{U}(D_{\mathbf{f},\mathbf{c}})$, where the inequality is strict if under $D_{\mathbf{f},\mathbf{c}}$ either c_0 or f_i for some $i \neq 0$ binds with strictly positive probability.*

5 The Main Result: Restrictions on both Conflict and Agreement Dimensions

This section examines the leading case with information on both the intratemporal and intertemporal trade-offs. It provides conditions for the optimal delegation policy within \mathcal{R} to involve good-specific caps and to involve *only* a savings floor. It also considers how reducing the doer's idiosyncratic information on consumption goods renders the second policy more likely to be optimal, thereby highlighting the role of that kind of information. Note that these results hold under the very weak assumptions on the distribution of information introduced in Section 3.

Proposition 1 shows that, everything else equal, there always exists a sufficiently *weak* degree of the doer's bias such that every optimal policy must involve good-specific caps.

Proposition 1. *There exists $b^* \in (0, 1)$ such that, if $b > b^*$, then every optimal $D \in \mathcal{R}$ must involve binding good-specific caps.²⁴*

Proposition 1 relies on the following properties of the first-best and full-discretion allocations (\mathbf{p}^* and \mathbf{a}^*). They hold because the planner and the doer want to reallocate resources to good i when its marginal utility rises *relative* to all other goods and savings.

Lemma 7. *The allocations \mathbf{p}^* and \mathbf{a}^* satisfy the following properties:*

- $p_0^* = p_0^*(\bar{\theta}, \bar{\mathbf{r}}) < p_0^*(\bar{\theta}, \bar{\mathbf{r}}_i, \underline{r}_{-i})$ and $a_0^* = a_0^*(\bar{\theta}, \bar{\mathbf{r}}) < a_0^*(\bar{\theta}, \bar{\mathbf{r}}_i, \underline{r}_{-i})$,
- $p_i^*(\bar{\theta}, \bar{\mathbf{r}}) < p_i^*(\bar{\theta}, \bar{\mathbf{r}}_i, \underline{r}_{-i}) = \bar{p}_i^*$ and $a_i^*(\bar{\theta}, \bar{\mathbf{r}}) < a_i^*(\bar{\theta}, \bar{\mathbf{r}}_i, \underline{r}_{-i}) = \bar{a}_i^*$ for every $i \neq 0$.

To gain intuition, suppose that there are only two goods and the model is fully symmetric with regard to them. Figure 1 shows the surface of the resource constraint B as the two-dimensional simplex. We can focus on this surface due to monotonicity of preferences and the fact that it is never optimal to use a savings cap or good-specific floors (Lemma 6). We can represent savings floors as horizontal lines and good-specific caps as lines parallel to the oblique edges of the simplex. The doer can choose only allocations above all these lines; the higher the line, the tighter the corresponding constraint. Figure 1 also represents the range of \mathbf{p}^* and \mathbf{a}^* as the areas inside the dotted and solid lines, respectively. Both the planner and the doer want to spend more on good i as its value relative to good j or savings rises. Hence, the states in which their optimal consumption of good 1 (respectively 2) is highest are *not* the states in which their optimal savings level is lowest (Lemma 7). In Figure 1, the former states map to the light-shaded areas, the latter to the dark-shaded areas.

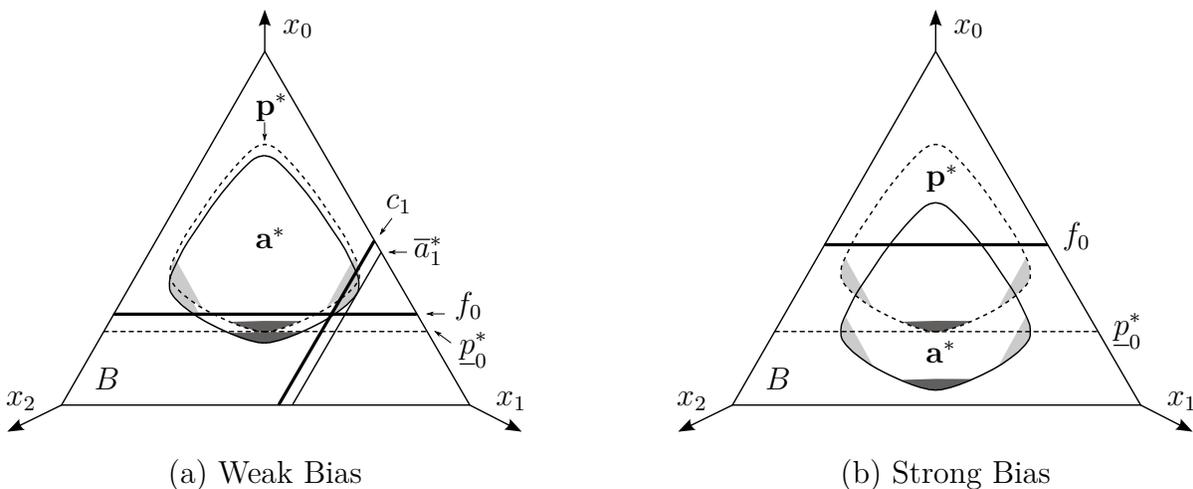


Figure 1: Optimal Delegation Policy – Intuition

A savings floor primarily targets the doer's decisions in the dark-shaded areas, but may have no effect in the light-shaded areas. Yet, in these states the doer continues to

²⁴The Appendix describes a procedure to calculate the threshold b^* .

overconsume and the planner would like to intervene. To see how, recall our preliminary results in Sections 4.1 and 4.2. By Lemma 4, if policies can involve only a savings floor, the planner sets it lower and lower as the doer's bias weakens; thus the floor becomes less and less likely to affect the states in which, say, good 1 is very valuable but good 2 is not. We can see this by comparing the levels of f_0 in panels (a) and (b) of Figure 1. To address those states, the planner does not want to use the floor when the bias is weak, but she can add a cap on good 1 that binds only when the floor does not (as $c_1 < \bar{a}_1^*$ in Figure 1(a)). By Lemma 5, such a cap will mitigate overconsumption when the allocation to good 1 is high, and despite its distorting effects, it strictly benefits the planner.

Although Proposition 1 says that good-specific caps must be part of an optimal policy, it does not say that they are always combined with a savings floor. Indeed, it is possible to have situations in which the planner combines some caps with the floor as well as situations in which she uses only caps. Section 5.1 illustrates this.

Does the optimal policy always rely on some good-specific cap? The answer is no. There always exists a sufficiently strong degree of the doer's bias such that, to be optimal, a policy should impose only a savings floor.

Proposition 2. *There exists $b_* \in (0, 1)$ such that, if $b < b_*$, then every optimal $D \in \mathcal{R}$ involves only a binding savings floor. Moreover, if $\underline{\mathbf{r}}' \geq \underline{\mathbf{r}}$ and $\bar{\mathbf{r}}' \leq \bar{\mathbf{r}}$ with $\underline{\mathbf{r}}' \neq \underline{\mathbf{r}}$ and $\bar{\mathbf{r}}' \neq \bar{\mathbf{r}}$, then the corresponding b'_* and b_* satisfy $b'_* > b_*$.²⁵*

Proposition 2 uses the following lemma, which shows that every optimal policy sets an effective lower bound on savings which is at least as high as the lowest first-best level \underline{p}_0^* .

Lemma 8. *For every $b \in (0, 1)$, if $D \in \mathcal{R}$ is optimal, then $\max\{f_0, 1 - \sum_{i=1}^n c_i\} \geq \underline{p}_0^*$.*

If D lets the doer save $x_0 < \underline{p}_0^*$, the planner realizes that no state justifies such a low x_0 . By increasing f_0 up to \underline{p}_0^* , she uniformly improves her payoff with regard to savings. Moreover, as a consequence of the lower level of spendable income, caps (if any) become less likely to bind—recall that all goods are normal—and hence less likely to distort the consumption bundle. Thus, the planner also benefits on this front.

We can now see the intuition behind Proposition 2. When b is very small, the doer wants to save much less than \underline{p}_0^* , no matter what information he observes. By Lemma 8, however, the planner never allows him to save less than \underline{p}_0^* . Hence, when a cap forces the doer to consume less of good i , he reallocates all the unspent money across the other goods, but not to savings. Since binding caps distort the chosen consumption bundle, the planner cannot benefit from imposing them if they do not improve savings. This reasoning also leads to the following simple observation, which may be useful to discard certain policies.

Remark 1. Suppose that $D \in \mathcal{R}$ involves binding caps but always induces the same level of savings, say x'_0 . Then D cannot be optimal. The planner can strictly improve on D by imposing only a floor $f_0 = x'_0$.

²⁵The Appendix shows how to calculate the threshold b_* ; importantly, b_* does not depend on the distribution G .

Proposition 2 also gives an idea of how reducing the doer’s idiosyncratic information on each consumption good may affect the optimal commitment policy. By shrinking the range of such information—without changing the information on the intertemporal trade-off (θ)—it becomes more likely that the simple policy with only a savings floor is optimal for a fixed degree of the doer’s bias. Given this, one may wonder what happens in the limit when the doer’s information is only about θ , but consumption still involves multiple goods. Does the optimal policy always involve only a savings floor? If not, which conditions ensure this? Section 6 will provide the answers.

The weakest bias for which optimal policies include good-specific caps obviously depends on the details of the setting at hand. Intuitively, as b falls below b^* , for any policy D it increases the probability that the doer ends up in a state where D ’s effective lower bound on savings, denoted by \underline{x}_0 , binds. Since in these states binding caps only create inefficiencies, their appeal for the planner falls accordingly. How the planner balances the inefficiencies in those states with the benefits that a cap can yield in other states ultimately depends on their distribution G . Nonetheless, since she can always set $f_0 = \underline{x}_0$, for biases below some level $\hat{b} \geq b_*$ every optimal policy will involve only a savings floor.

Overall Propositions 1 and 2 suggest that richer commitment policies involving many rules may in fact prevail when the individual has weaker self-control problems, whereas simple policies may prevail when such problems are stronger. This prediction perhaps goes against an initial intuition that if the planner adds good-specific restrictions on consumption to an aggregate one when the conflict with the doer is weak, then *a fortiori* she should do so when the conflict becomes stronger. Put differently, at first glance one might think that when the conflict is weak more rules are actually *less* valuable than discretion: As the doer bias weakens, the planner cares relatively more about allowing him to act on his information, especially along the dimensions for which their preferences agree.

5.1 Caps and Floor or Only Caps?

This section shows that there exist both cases in which the planner combines binding good-specific caps with a savings floor and cases in which she uses only the caps. We will first consider a setting with three states to develop the key intuitions. We will then show that the two types of optimal policies can exist with a continuum of states. Throughout this section, we will focus on the following symmetric setting: Let $n = 2$, $u^1(x; r) = u^2(x; r) = r \ln(x)$, $\underline{r}_1 = \underline{r}_2 = \underline{r} > 0$, $\bar{r}_1 = \bar{r}_2 = \bar{r} > \underline{r}$, and $v(x) = \ln(x)$.²⁶

Suppose that the possible states are $\mathbf{s}^0 = (\bar{\theta}, \bar{r}_1, \bar{r}_2)$, $\mathbf{s}^1 = (\underline{\theta}, \bar{r}_1, \underline{r}_2)$, and $\mathbf{s}^2 = (\underline{\theta}, \underline{r}_1, \bar{r}_2)$; let their distribution be $(g, \frac{1}{2}(1 - g), \frac{1}{2}(1 - g))$, where g is the probability of \mathbf{s}^0 . By Lemma 1 and symmetry, we have that

$$a_0^*(\mathbf{s}^0) < a_0^*(\mathbf{s}^1) = a_0^*(\mathbf{s}^2), \quad a_1^*(\mathbf{s}^2) = a_2^*(\mathbf{s}^1) < a_1^*(\mathbf{s}^1) = a_2^*(\mathbf{s}^2), \quad a_1^*(\mathbf{s}^0) = a_2^*(\mathbf{s}^0);$$

²⁶The function $\ln(x)$ violates the continuity and differentiability assumptions of Section 3 at $x = 0$, but this is irrelevant for the analysis. On the other hand, this function will greatly simplify the analysis.

similar properties hold for \mathbf{p}^* . Relying on continuity, we can always find $b < 1$ sufficiently high so that $a_0^*(\mathbf{s}^1) = a_0^*(\mathbf{s}^2) > p_0^*(\mathbf{s}^0)$; also, we can always find $\underline{\theta} < \bar{\theta}$ sufficiently close to $\bar{\theta}$ so that $p_1^*(\mathbf{s}^1) > p_1^*(\mathbf{s}^0)$ and $p_2^*(\mathbf{s}^2) > p_2^*(\mathbf{s}^0)$. Figure 2(a) represents such a situation. Concretely, we can think about this situation in the following terms. Imagine an individual, Bob, who enjoys going out for dinner (x_1) and attending live-music events (x_2). In a given period, his best friend Ann may visit him ($\bar{\theta}$) or not ($\underline{\theta}$). If on his own, depending on the mood Bob prefers to either go to a fancy restaurant with a piano bar (\mathbf{s}^1) or grab a quick sandwich and attend a great concert (\mathbf{s}^2). By contrast, when Ann is in town, Bob prefers to combine a good restaurant with a good concert (\mathbf{s}^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 $D \in \mathcal{R}$ satisfies $f_0 = p_0^*(\mathbf{s}^0)$, $c_1 = p_1^*(\mathbf{s}^1)$, and $c_2 = p_2^*(\mathbf{s}^2)$.*

The intuition is as follows. If Bob knew for sure that Ann was not planning to visit him, he could impose a savings floor that prevents him from splurging in both \mathbf{s}^1 and \mathbf{s}^2 . However, such a floor is too stringent if Ann happens to visit (see Figure 2(a)). Therefore, if he thinks that Ann's visit is sufficiently likely, in expectation Bob views raising f_0 enough to influence his choices in \mathbf{s}^1 and \mathbf{s}^2 as too costly, and hence prefers to set $f_0 = p_0^*(\mathbf{s}^0)$. Such a floor grants full discretion in \mathbf{s}^1 and \mathbf{s}^2 . In these states, however, the logic of Lemma 5 applies and Bob can again limit the consequences of his bias by using good-specific caps; moreover, here he can do so without affecting his choice in \mathbf{s}^0 . The fact that the optimal caps coincide with the first-best allocation to the respective good is just a consequence of the logarithmic payoffs.

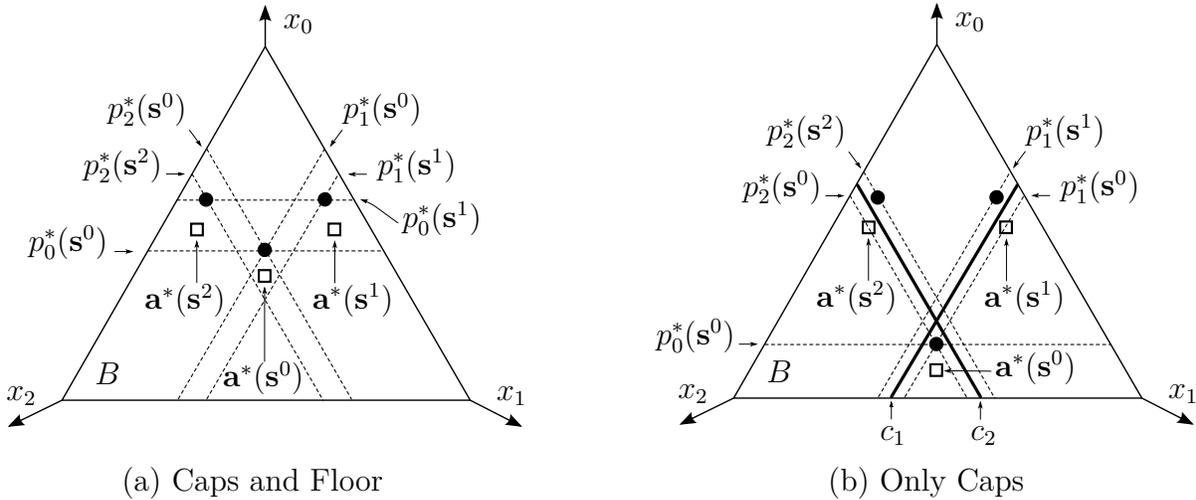


Figure 2: Three-State Example

A simple change of the previous three-state setting suffices to show that the optimal policy can involve only good-specific caps. Fix $g > g^*$ and all the other parameters of the model, except $\bar{\theta}$. If we increase $\bar{\theta}$, the planner and the doer want to consume

more of each good in \mathbf{s}^0 . This eventually leads to a situation as in Figure 2(b), where $p_1^*(\mathbf{s}^0) > p_1^*(\mathbf{s}^1)$ and $p_2^*(\mathbf{s}^0) > p_2^*(\mathbf{s}^2)$. Continuing our previous story, we can interpret this case as a situation in which, being a spendthrift, Ann always insists on choosing fancy restaurants and attending the best concerts.

Proposition 4. *There exists $\bar{\theta}'$ such that in the optimal policy $D \in \mathcal{R}$ both caps bind, but the floor never binds. In particular, $c_1 = c_2$ and $p_i^*(\mathbf{s}^0) < c_i < p_i^*(\mathbf{s}^i)$ for every $i = 1, 2$.*

Figure 2 helps us see the intuition. In this case, Bob would want to spend more both on restaurants and on concerts when Ann is in town than when she is not. Therefore, the good-specific caps that Bob would set to curb his splurging in \mathbf{s}^1 and \mathbf{s}^2 already create an aggregate limit on consumption which binds in \mathbf{s}^0 . As a result, he is willing to relax such caps; and if his first-best consumption is not too high in \mathbf{s}^0 , he may keep the caps sufficiently low so as to still curb overspending in \mathbf{s}^1 and \mathbf{s}^2 . However, since these caps already push savings above the first-best level in \mathbf{s}^0 , Bob would be strictly worse off by also imposing a savings floor sufficiently high to bind.

It remains to show that the qualitative properties of the policies in Propositions 3 and 4 can also arise in settings with a continuum of states. Intuitively, this should be the case if the planner assigns sufficiently high probability to states that induce similar trade-offs as do \mathbf{s}^0 , \mathbf{s}^1 , and \mathbf{s}^2 . To formalize this idea, we can consider distributions that are sufficiently concentrated on \mathbf{s}^0 , \mathbf{s}^1 , and \mathbf{s}^2 . One way to do this is the following. Let G^{fc} be a distribution over $(\mathbf{s}^0, \mathbf{s}^1, \mathbf{s}^2)$ that leads to Proposition 3 and \bar{G} the uniform distribution over $[\underline{\theta}, \bar{\theta}] \times [\underline{r}, \bar{r}]^2$. Similarly, let G^c be a distribution that leads to Proposition 4 and \bar{G}' the uniform distribution over $[\underline{\theta}, \bar{\theta}'] \times [\underline{r}, \bar{r}]^2$, where $\bar{\theta}'$ is as in Proposition 4. Finally, define

$$G_\alpha^{\text{fc}} = \alpha G^{\text{fc}} + (1 - \alpha) \bar{G} \quad \text{and} \quad G_\alpha^c = \alpha G^c + (1 - \alpha) \bar{G}', \quad \alpha \in [0, 1].$$

Corollary 1.

- (1) *There exists $\alpha \in (0, 1)$ such that, given G_α^{fc} , every optimal $D \in \mathcal{R}$ involves a binding savings floor as well as binding caps for both goods.*
- (2) *There exists $\alpha' \in (0, 1)$ such that, given $G_{\alpha'}^c$, for every optimal $D \in \mathcal{R}$ both good-specific caps bind, but the savings floor never binds.*

Note that, given this $G_{\alpha'}^c$, for a set of states with strictly positive probability the doer's choices are unaffected by the optimal D and a savings floor would curb his overconsumption. However, imposing such a floor requires forcing the doer to save even more in \mathbf{s}^0 than the level induced by the good-specific caps, which already exceeds the first-best level. Hence, when the planner cares enough about \mathbf{s}^0 , she will not use any savings floor.

5.2 Discussion: Consumer Behavior and Demand of Commitment Devices

The previous results may be of interest with regard to the consumption-savings behavior of individuals with self-control problems for several reasons.

The first is that they shed light on a phenomenon called “mental budgeting.” It has often been observed that some individuals earmark their income according to multiple

spending categories—sometimes by dividing it into use-specific envelopes or “tin cans,” and more recently by setting up category-specific budgets via services like Mint.com, Quicken.com, or StickK.com. Individuals set up these budgets with the goal of controlling their spending on certain categories, similarly to how good-specific caps work in our model. Heath and Soll (1996) argue that “descriptions of consumers over the past 50 years indicate [that budgeting] is a pervasive part of consumer behavior” and discuss a body of evidence supporting this view. According to Bénabou and Tirole (2004), “mental accounts and other personal rules [...] appear to be common in economic decisions;” their analysis provides an answer, based on the idea of self-reputation, to the question of what renders such rules effective. Ameriks et al. (2003) found that 37% of households in their TIAA-CREF survey sets detailed spending budgets for themselves, and that “a substantial minority of them [45%] agrees that their budgeting help them refrain their spending.” Antonides et al. (2011) show that “mental budgeting appears to be widely practiced in the Netherlands,” where they conducted their study.

Budgeting has important implications for consumer behavior—beyond satiation and income effects (Heath and Soll (1996)): Most notably, it violates fungibility of money, which has far-reaching consequences for how firms market their products or how governments design policies that regulate saving and borrowing devices. The literature has informally suggested or assumed that budgeting represents how individuals deal with self-control problems which lead them to overconsume and undersave (Thaler (1985, 1999), Heath and Soll (1996), Prelec and Loewenstein (1998), Antonides et al. (2011)).

The present paper offers the first (to the best of my knowledge) explicit foundation of budgeting. Using a standard consumption-savings model, it shows that a well-studied cause of self-control problems, namely present bias, can induce individuals to optimally limit how much income they can allocate to some consumption goods, besides savings.²⁷ The theory is consistent with the existing evidence. For instance, in Heath and Soll’s (1996) study, when a budget is not binding people tend to overconsume in that category; however, after previous expenses reduce the balance of a budget, people tend to underconsume in that category. Antonides et al. (2011) show that having saving goals has a positive effect on budgeting, which is consistent with the reason why the planner imposes caps in our model.

The theory also suggests some qualifications of existing views on mental budgeting. First, Heath and Soll (1996) emphasize that their data show that budgets lead to underconsumption and in particular for “unobjectionable goods”—goods that one would not obviously classify as temptations. Given this, they conclude that the planner may impose budgets that are too stringent. According to our theory, this conclusion may be invalid: Optimal budgets have to bind with positive probability and are imposed on what we called agreement categories. Second, the theory predicts that, in fact, category-specific budgets are used only by individuals who have a *weak* present bias.²⁸ This, perhaps

²⁷In Brocas and Carillo’s (2008) model with two consumption goods, one of which has ex-ante uncertain utility, the optimal commitment policy is not a simple budgeting rule, but rather a rule that punishes higher spending in one good by requiring lower spending in the other. Also, since in their model the doer is fully myopic ($b = 0$), an optimal interval policy would involve only a savings floor.

²⁸This prediction refers to the ex-ante optimal commitment strategy and hence continues to hold for

counterintuitive, prediction is consistent with some evidence in Antonides et al. (2011): Their analysis shows that people who exhibit a short-term time orientation—which can be interpreted as a low b in our model—are less likely to set up budgets, whereas people who exhibit a long-term time orientation—that is, a b closer to 1—are more likely to set up budgets. Of course, further empirical tests of these predictions are needed.

The results in this paper may also be of interest for researchers eager to understand the demand of commitment devices. According to Bryan et al. (2010), “there is insufficient work to understand [this] demand,” especially for what they call “soft commitments” (like mental budgeting). Since the seminal work of Thaler and Shefrin (1981) and Laibson (1997), the previous literature on consumption and savings has focused on the individuals’ problem of curbing their undersaving, often stressing the key role played in this by external commitment devices like illiquid assets. More recently, Amador et al. (2006) reached the conclusion that, under weak conditions, a minimum-savings rule coincides with the optimal commitment and argued that illiquid assets may suffice to implement it (see also the next section). However, we saw that in a world with multiple consumption goods, minimum-savings rules alone may be strictly dominated by commitment strategies involving good-specific budgets. Such budgets cannot be implemented using illiquid assets, which therefore no longer allow individuals to achieve their desired form of commitment. This can explain why, besides possibly adopting internal rules, some individuals demand services that allow them to budget expenses by categories, such as those offered by companies like Mint.com, Quicken.com, or StickK.com. Such services should have no value according to the existing theory, which therefore cannot explain their active market.

6 No Information on Intratemporal Trade-offs

Motivated in part by Proposition 2, this section considers settings in which consumption continues to involve multiple goods, but information is only about the intertemporal utility trade-off (θ). Even though it may seem unrealistic or special that at the time of committing the planner faces no uncertainty at all about the trade-offs across goods, studying this case will provide further understanding of the leading case considered before. We will show that most of the time the optimal commitment policy within \mathcal{R} will involve a savings floor but *no* good-specific caps. This result further highlights that the uncertainty on intratemporal trade-offs plays a key role in the usefulness of good-specific caps as commitment rules.

To capture the settings of interest in this section with the formalism used so far, let $r_i = \bar{r}_i$ for all $i = 1, \dots, n$. Given this, G now denotes the distribution of $\theta \in [\underline{\theta}, \bar{\theta}]$; for this section, assume that G has a density function g which is strictly positive and continuous on $[\underline{\theta}, \bar{\theta}]$. Since \mathbf{r} is fixed, hereafter we will omit it from the consumption utility \hat{u} . As will become clear, the results in this section do not hinge at all on the separability of \hat{u} across goods.

partially naive individuals who incorrectly think, ex ante, that their present bias is weak (O’Donoghue and Rabin (2001)).

Given that now the doer’s information is only about the intertemporal trade-off, does the multidimensionality of consumption still play any role in the planner’s problem? Note that the same level of total consumption expenditure y yields different utilities $\hat{u}(\mathbf{x})$ depending on how y is divided across goods. This fact can be exploited to curb the doer’s tendency to overconsume. For instance, by requiring (perhaps via caps) that y be allocated in a distorted way which does not yield much more utility than some $y' < y$, a policy curbs the doer’s willingness to spend y in states where the planner prefers y' .

However, if the goal is only to lower the utility that the doer can derive from the income that he does *not* save, a simpler method may be to prevent him from spending all of it in the first place. This method, usually called “money burning,” assumes that the planner can force the doer to literally “throw away” part of what he does not save.²⁹ It is easy to see that any utility level \hat{u} achieved by spending $y = 1 - x_0$ inefficiently can also be achieved by burning part of $1 - x_0$ and spending the rest to buy an *efficient* consumption bundle. Indeed, for any $y \in [0, 1]$ and $\mathbf{x}' \in \mathbb{R}_+^n$ that satisfy $\sum_{i=1}^n x'_i = y$, the utility $\hat{u}(\mathbf{x}')$ belongs to the interval $[\hat{u}(\mathbf{0}), u^*(y)]$, where

$$u^*(y) = \max_{\{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i \leq y\}} \hat{u}(\mathbf{x}). \quad (5)$$

Since u^* is strictly increasing and continuous and $u^*(0) = \hat{u}(\mathbf{0})$, there always exists $y' \leq y$ such that $u^*(y') = \hat{u}(\mathbf{x}')$.

These observations suggest that it should be possible to restrict attention to delegation policies that regulate only savings and *total* consumption expenditures. This is true, provided that we allow money burning to freely depend on the level of savings—which requires to go beyond simple rectangle policies. Formally, let B^{ac} be the set of feasible allocations defined only in terms of total consumption y and savings x_0 :

$$B^{\text{ac}} = \{(y, x_0) \in \mathbb{R}_+^2 : y + x_0 \leq 1\}.$$

Given *any* $D^{\text{ac}} \subset B^{\text{ac}}$, for each θ the doer’s problem becomes to maximize $\theta \hat{u}(\mathbf{x}) + bv(x_0)$ subject to $\sum_{i=1}^n x_i \leq y$ and $(y, x_0) \in D^{\text{ac}}$.

Lemma 9. *Suppose information affects only the intertemporal utility trade-off. There exists an optimal $D \subset B$ with $\mathcal{U}(D) = \mathcal{U}^*$ if and only if there exists an optimal $D^{\text{ac}} \subset B^{\text{ac}}$ with $\mathcal{U}(D^{\text{ac}}) = \mathcal{U}^*$.*

Remark 2. If money burning cannot be tailored to the level of savings, Lemma 9 need not hold. Without money burning, for instance, constraints on x_0 translate one-to-one into constraints on y . Yet, using the richer policies $D \subset B$, the planner can still regulate how the doer is allowed to divide any amount y across goods, thereby affecting his resulting consumption utility (see also Proposition 6 below).

Lemma 9 allows us to recast our problem into AWA’s analysis. Since the function \hat{u} is strictly increasing, the constraint $\sum_{i=1}^n x_i \leq y$ will always bind when the doer faces D^{ac} . Using (5), we can express the planner’s problem as choosing $D^{\text{ac}} \subset B^{\text{ac}}$ so as to maximize

$$\int_{\underline{\theta}}^{\bar{\theta}} [\theta u^*(a_y(\theta)) + v(a_0(\theta))] g(\theta) d\theta$$

²⁹One way to do this might be committing to some charitable donations conditional on the level of savings. Note, however, that such donations qualify as “money burning” only if the individual derives no utility *directly* from them (for example, due to altruism).

subject to

$$(a_y(\theta), a_0(\theta)) \in \arg \max_{(y, x_0) \in D^{\text{ac}}} \{\theta u^*(y) + bv(x_0)\}, \quad \theta \in [\underline{\theta}, \bar{\theta}]. \quad (6)$$

AWA show that this problem is equivalent to another problem in which the planner's objective can be written as

$$\int_{\underline{\theta}}^{\bar{\theta}} H(\theta) u^*(a_y(\theta)) g(\theta) d\theta + k,$$

where $k \in \mathbb{R}$ and

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

To understand what $H(\theta)$ captures, ignore feasibility for the moment. Suppose that the planner allows the doer to save a bit less in state θ , without changing his total payoff—so that he does not select other allocations. Doing so requires inducing the doer to consume a bit more. Overall this adjustment harms the planner when θ occurs, because she cares discretely more about savings. This explains the negative term $-(1 - b)\theta g(\theta)$. The adjustment, however, also renders the allocation chosen in θ more attractive for the doer in the states where he values present utility more: in all $\theta' > \theta$, which have mass $1 - G(\theta)$. Thus, the planner can induce the doer to save more in these states, which is exactly what she wants. This explains the positive term $1 - G(\theta)$.

AWA's main result shows that the solution to the planner's problem crucially depends on the following condition, where θ^* is defined by

$$\theta^* = \min \left\{ \theta \in [\underline{\theta}, \bar{\theta}] : \int_{\theta'}^{\bar{\theta}} H(\hat{\theta}) d\hat{\theta} \leq 0 \text{ for all } \theta' \geq \theta \right\}. \quad (7)$$

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

Also, let

$$D^{\text{ac}}(\theta^*) = \{(y, x_0) \in B^{\text{ac}} : x_0 \geq a_0^*(\theta^*)\},$$

which is essentially a rectangle policy which involves only the savings floor $f_0 = a_0^*(\theta^*)$.

Proposition 5 (Amador et al. (2006)). *The policy $D^{\text{ac}}(\theta^*)$ is optimal among all subsets of B^{ac} if and only if Condition 1 holds.*³⁰

As AWA noted, Condition 1 is satisfied for all $b \in [0, 1]$ by many distributions, especially those commonly used in applications. More generally, since G is strictly increasing, Condition 1 is more likely to hold when the doer's bias is weak (that is, b is close to 1). Recall that AWA proved the optimality of policies that use only a savings floor for environments with one consumption good. Thus, together with Lemma 9, Proposition 5 implies that the multiplicity of consumption goods adds no useful tool to achieve superior commitment policies when information is only about the intertemporal trade-off.

By Proposition 1, however, AWA's result does not extend to settings in which information also affects the intratemporal trade-offs. In these settings, for any distribution \hat{G} with full support on S —even if $\hat{G}(\cdot, \mathbf{r})$ satisfies Condition 1 for all \mathbf{r} —we can always find

³⁰It can be easily checked that, when the function \hat{u} is strictly increasing, concave, and continuously differentiable, then the function u^* in (5) satisfies the same properties, as assumed in AWA.

a sufficiently weak bias that renders policies involving only a savings floor strictly dominated. Note that, by contrast, a weak bias characterizes exactly those settings where Condition 1 almost certainly holds if information is only about θ , and hence those policies would be optimal.

Together with Lemma 9, Proposition 5 provides a sufficient condition for a policy D_{f_0} which involves only a savings floor to be optimal within the class of rectangle policies.

Corollary 2. *Define θ^* as in (7). If Condition 1 holds, then D_{f_0} with $f_0 = a_0^*(\theta^*)$ is optimal within \mathcal{R} .*

In general, Condition 1 is not necessary for the conclusion of Corollary 2. This is because policies that improve on D_{f_0} may lie outside \mathcal{R} (see AWA for an example). Hence, D_{f_0} can be optimal for an even larger family of distributions.

Even though most distributions satisfy Condition 1, what can we say about the optimal policies when it fails? AWA argue that in this case an optimal policy may have to rely on money burning. This reopens the door for the planner to benefit from the multidimensionality of consumption. By forcing the doer to choose inefficient bundles based on the level of savings, the planner can achieve the same curbing effect on the doer's tendency to overconsume with strictly less (possibly no) money burning. This highlights a possible limitation of treating consumption as a monolithic entity, even if information is only about θ . To state the result, define

$$u_*(y) = \min_{\{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i = y\}} \hat{u}(\mathbf{x}), \quad y \in [0, 1].$$

Note that, since \hat{u} is strictly concave, $u_*(y) < u^*(y)$ for all $y > 0$.

Proposition 6. *Suppose that the optimal $D^{\text{ac}} \subset B^{\text{as}}$ induces an allocation \mathbf{a} in (6) which exhibits positive consumption and money burning over some set $\Theta \subset [\underline{\theta}, \bar{\theta}]$ (that is, $a_y(\theta) > 0$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$ and $a_y(\theta) < 1 - a_0(\theta)$ for all $\theta \in \Theta$).*

(1) *There exists $D' \subset B$ that satisfies $\mathcal{U}(D') = \mathcal{U}(D^{\text{ac}})$ and involves less money burning: The induced allocation \mathbf{a}' in (4) satisfies $a'_0(\theta) = a_0(\theta)$ and $\sum_{i=1}^n a'_i(\theta) \geq a_y(\theta)$ for all θ , with strict inequality over Θ .*

(2) *If $u_*(1 - a_0(\theta)) \leq u^*(a_y(\theta))$ for all $\theta \in \Theta$, then D' can be chosen so that money burning never occurs: $\sum_{i=0}^n a'_i(\theta) = 1$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$.*

In words, the condition in part (2) means the following: In every state in which the planner wants to burn some money, she can instead let the doer spend all the unsaved income, but in such an inefficient way that the extra money does not yield higher utility than the efficient bundle induced with money burning. This condition holds if, for instance, zero consumption of some good is extremely inefficient and leads to a very low present utility. Examples of such goods may include one's favorite drink or food, or going out with friends.³¹ Reducing the share of $1 - x_0$ allocated to these goods can

³¹In other applications, this condition may be even easier to satisfy. For instance, in the fiscal-constitution and public-finance applications where the agreement dimensions are public services, the payoff of society may be very low if the government allocates no resources to services like national defense, law enforcement, or criminal detention.

replicate the effect of money burning. In these settings, if money burning is infeasible, the planner may achieve strictly higher expected payoffs by again imposing distortions along dimensions which cause *no* conflict of interest with the doer. Finally, note that one way to implement these distortions is again to use good-specific caps and floors, which however may now have to vary based on how much the doer saves.

7 Implications for Other Applications

This section discusses the implications of our analysis for the other applications outlined in Section 3.1.

Corporate governance. Our results suggest that to best incentivize a manager who undervalues R&D, the owner of a multi-product company may have to impose caps on how much can be spent each year on promoting specific products, possibly in addition to requiring a minimum investment in R&D. Due to the caps, the manager may end up promoting too little some products and others too much from the owner’s viewpoint. This, however, is a risk she 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 enough with the manager on how important R&D is for the company. This may be true, for instance, if the manager himself has significant stakes in the company. Otherwise, the owners can simply demand only a minimum investment in R&D, or even impose a fix one if she thinks that the manager is seriously biased. Of course, in this case she may consider addressing the manager’s bias by changing his compensation scheme, or hiring a new manager.

Fiscal-constitution design. When designing a fiscal constitution behind the veil of ignorance, society can limit the consequences of the government’s tendency to overspend on public services either by imposing an aggregate cap on public spending (via a floor on private consumption), or by setting specific caps on how much the government can allocate to some services. The government’s spending bias may depend on the incentives created by the country’s political institutions, for instance. The analysis shows that if these institutions lead to a sufficiently weak bias, then an optimal fiscal constitution must involve service-specific caps. The analysis highlights that such caps are “no free lunch,” in the sense that they cause inefficiencies in the composition of public spending. However, these inefficiencies are more than compensated by the resulting higher level of private consumption. On the other hand, if the political institutions lead to governments that are strongly biased, the constitution should involve only an aggregate spending cap; any binding specific cap distorts public spending without producing enough benefits in terms of private consumption.

Research vs. teaching in academia. Research universities usually specify a minimum amount of time that professors have to allocate each week to their teaching duties (lectures, office hours, etc.), granting them discretion on, say, course-preparation time. However, we do not observe universities restricting how much time professors should spend on each of their research projects. This is consistent with our theory if we believe that universities’ view professors as severely biased against teaching, so that they always

allocate to it only the required minimum time—which may be plausible. Alternatively, the theory suggests that existing practices may leave room for improvements. Of course, other reasons can explain the absence of restrictions on research: Monitoring how much time a professor spends on a project may simply be infeasible, or it may be in conflict with some principle that one should not interfere with the creative process of research. In either case, the theory identifies when minimum teaching requirements can still achieve the optimum.

Public finance. When setting the fiscal rules for the coming year, a government may realize that, due to its present bias, it will tend to borrow excessively against future tax revenues. It is then intuitive that, *ex ante*, the government benefits from committing to a cap on how much it will be allowed to borrow (recall that this is equivalent to a floor on $x_0^t = -z^t/\rho$, where z^t is the amount borrowed at time t). We often observe such a rule in reality in the form of budget-deficit ceilings. The paper shows, however, when and how the government can easily improve on a policy that imposes only a deficit ceiling. Although present bias never interferes with how the government trades off public services within a period, introducing specific caps on how much it will be allowed to spend on some services can lead to a superior policy. Such caps often appear in reality as part of fiscal budgets. The theory offers one explanation for why we observe both deficit ceilings and service-specific budgets; it also suggests that a government introduces specific caps not when it anticipates to be severely biased, but in fact when it anticipates to be mildly biased. Another point highlighted in the paper is again that service-specific caps are no free lunch: Even though they are useful to curb excessive borrowing, they also distort public spending.

Workers' time management. In this application, the worker overvalues his break time, x_0 , which causes the conflict with the employer ($b > 1$). An analysis similar to that above is possible and leads to results in which caps and floors swap roles. If the employer can only impose a cap or floor on how much time the worker spends on x_0 , she sets a cap c_0 which is strictly below the maximum time that she finds acceptable (analogously to Lemma 3). Moreover, c_0 becomes tighter as the worker's bias worsens (analogously to Lemma 4). Regarding task-specific rules, a floor on task i induces the worker to allocate less time to x_0 , but also less time to the other tasks. Nonetheless, the employer would strictly benefit from a single binding floor on any task relative to granting the worker full discretion (analogously to Lemma 5).

The main result of the paper (Propositions 1 and 2) changes as follows. If the worker's tendency to indulge in breaks is sufficiently strong, the employer should impose only a cap on x_0 . By contrast, if that tendency is sufficiently weak, an optimal delegation policy must involve binding task-specific floors. In practice, it may be impossible to monitor the worker's breaks; nonetheless, c_0 can always be implemented by setting an overall minimum for the time the worker has to allocate to his tasks, which are easier to monitor. The main result highlights that the possibility of monitoring each task individually may allow the employer to design strictly superior policies by adding specific floors on top of an aggregate one.

8 Concluding Remarks

This paper examines a new, broad class of principal-agent delegation problems which arise in many economic settings, from individual commitment problems to public finance, corporate governance, and workforce management. In such problems, the agent controls how to allocate finite resources (money, time, etc.) across multiple categories, having better information on their returns than the principal but pursuing different goals from hers.

The paper characterizes how optimal delegation policies trade off rules and discretion and how they depend on the degree of conflict between parties as well as the nature of the agent's information. Perhaps counterintuitively, it can be optimal for the principal to impose distorting restrictions on categories for which there is no conflict of interest with the agent, so as to curb how the conflict along other categories affects his overall resource allocation. Moreover, such restrictions are more likely to be optimal when the conflict of interest is weaker and when the agent's information is about the specific value of categories causing no conflict. The paper also shows that requiring distorted allocations across these categories can reduce or even eliminate the need for money burning as a way to manage the agent's incentives. By considering a tractable class of simple delegation policies, this paper offers insights that can be easily applied to many concrete problems, which existing models cannot handle.

A Appendix: Proofs

A.1 Proof of Lemma 2

Each $D \in \mathcal{R}$ is defined by a vector $(\mathbf{f}, \mathbf{c}) \in \mathbb{R}_+^{2(n+1)}$. Without loss we can restrict attention to the following compact subset of $\mathbb{R}_+^{2(n+1)}$:

$$\mathcal{FC} = \{(\mathbf{f}, \mathbf{c}) \in [0, 1]^{2(n+1)} : \mathbf{f} \leq \mathbf{c}, \sum_{i=0}^n f_i \leq 1\}.$$

Thus, we can think that the planner chooses $(\mathbf{f}, \mathbf{c}) \in \mathcal{FC}$.

Given any such (\mathbf{f}, \mathbf{c}) , let $\mathbf{a}(\mathbf{s}|\mathbf{f}, \mathbf{c})$ be the doer's optimal allocation in state \mathbf{s} from the compact set $D_{\mathbf{f}, \mathbf{c}}$. Since $D_{\mathbf{f}, \mathbf{c}}$ is convex (Theorem 2.1 in Rockafellar (1997)), $\mathbf{a}(\mathbf{s}|\mathbf{f}, \mathbf{c})$ is unique for every $\mathbf{s} \in S$ by strict concavity of $V(\cdot; \mathbf{s})$. Clearly, the correspondence that for each $(\mathbf{f}, \mathbf{c}) \in \mathcal{FC}$ maps to $D_{\mathbf{f}, \mathbf{c}}$ is non-empty, compact valued, and continuous. It follows from the Maximum Theorem that $\mathbf{a}(\mathbf{s}; \cdot, \cdot)$ is continuous for every $\mathbf{s} \in S$.

We can now show that the planner's payoff is continuous in (\mathbf{f}, \mathbf{c}) . For each $(\mathbf{f}, \mathbf{c}) \in \mathcal{FC}$, let

$$\mathcal{U}(\mathbf{f}, \mathbf{c}) = \int_S U(\mathbf{a}(\mathbf{s}|\mathbf{f}, \mathbf{c}); \mathbf{s}) dG.$$

Since $U(\mathbf{a}(\mathbf{s}|\mathbf{f}, \mathbf{c}); \mathbf{s})$ is continuous in (\mathbf{f}, \mathbf{c}) for every $\mathbf{s} \in S$ and is uniformly bounded over 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.

A.2 Lemma 10

For any floor $f_0 \in [\underline{a}_0^*, \bar{p}_0^*]$,³² for simplicity denote by D_{f_0} the corresponding policy in \mathcal{R} .

Lemma 10. Define $\bar{S}(f_0) = \{\mathbf{s} \in S : a_0^*(\mathbf{s}) \leq f_0\}$ and

$$\mathbf{x}^{f_0}(\mathbf{s}) = \arg \max_{\{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i \leq 1 - f_0\}} \hat{u}(\mathbf{x}; \mathbf{s}), \quad \mathbf{s} \in S.$$

The payoff $\mathcal{U}(D_{f_0})$ is differentiable in f_0 over the domain $[\underline{a}_0^*, \bar{p}_0^*]$ with

$$\frac{d}{df_0} \mathcal{U}(D_{f_0}) = \int_{\bar{S}(f_0)} \left[v'(f_0) - \frac{\partial}{\partial x_i} \hat{u}(\mathbf{x}^{f_0}(\mathbf{s}); \mathbf{s}) \right] dG,$$

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

Proof. For simplicity, drop the subscript 0 from f_0 and let $\Psi(f) = \mathcal{U}(D_f)$. Also, we will consider only $f \in [\underline{a}_0^*, \bar{p}_0^*]$ without specifying this every time. Given f and any \mathbf{s} , define

$$\tilde{u}(f; \mathbf{s}) \equiv \hat{u}(\mathbf{x}^f(\mathbf{s}); \mathbf{s}) = \max_{\{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i \leq 1 - f\}} \hat{u}(\mathbf{x}; \mathbf{s}). \quad (8)$$

and $\tilde{U}(f; \mathbf{s}) = \tilde{u}(f; \mathbf{s}) + v(f)$. We first want to show that $\tilde{U}(f; \mathbf{s})$ is strictly concave in f for every \mathbf{s} . Take any f, f' , and $\zeta \in (0, 1)$. We have

$$\begin{aligned} \tilde{u}(\zeta f + (1 - \zeta)f'; \mathbf{s}) + v(\zeta f + (1 - \zeta)f') &\geq \hat{u}(\zeta \mathbf{x}^f(\mathbf{s}) + (1 - \zeta)\mathbf{x}^{f'}(\mathbf{s}); \mathbf{s}) \\ &\quad + v(\zeta f + (1 - \zeta)f') \\ &> \zeta \left[\hat{u}(\mathbf{x}^f(\mathbf{s}); \mathbf{s}) + v(f) \right] \\ &\quad + (1 - \zeta) \left[\hat{u}(\mathbf{x}^{f'}(\mathbf{s}); \mathbf{s}) + v(f') \right] \\ &= \zeta [\tilde{u}(f; \mathbf{s}) + v(f)] \\ &\quad + (1 - \zeta) [\tilde{u}(f'; \mathbf{s}) + v(f')], \end{aligned} \quad (9)$$

where the weak inequality follows because $\sum_i x_i^f(\mathbf{s}) \leq f$ and $\sum_i x_i^{f'}(\mathbf{s}) \leq f'$ implies

$$\sum_i \left[\zeta x_i^f(\mathbf{s}) + (1 - \zeta)x_i^{f'}(\mathbf{s}) \right] \leq \zeta f + (1 - \zeta)f',$$

and the strict inequality follows from strict concavity of $\hat{u}(\cdot, \mathbf{s})$ and $v(\cdot)$.

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

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

By considering the FOC of the Lagrangian defining $\tilde{u}(f; \mathbf{s})$, we see that $\frac{\partial}{\partial x_i} \hat{u}(\mathbf{x}^f(\mathbf{s}); \mathbf{s}) = \lambda(\mathbf{s}; f)$ for any $i = 1, \dots, n$, where $\lambda(\mathbf{s}; f)$ is the Lagrange multiplier of the constraint $\sum_{i=1}^n x_i \leq 1 - f$. Since $\mathbf{x}^f(\mathbf{s})$ is continuous in f for every \mathbf{s} , so it $\lambda(\mathbf{s}; f)$ given our assumptions on \hat{u} . By Theorem 1, p. 222, of Luenberger (1969), for every $f', f'' \in [0, 1]$ we have

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

³²Any other floor is dominated by one in this range.

Continuity of $\lambda(\mathbf{s}; \cdot)$ then implies that $\tilde{u}'(f; \mathbf{s})$ exists for every f and equals $-\lambda(\mathbf{s}; f)$. Therefore,

$$\tilde{U}'(f; \mathbf{s}) = v'(f) - \frac{\partial}{\partial x_i} \hat{u}(\mathbf{x}^f(\mathbf{s}); \mathbf{s}) \quad \text{for all } \mathbf{s}. \quad (10)$$

For any f , denote by \mathbf{a}^f the doer's behavior as a function of \mathbf{s} under f . Note that $\mathbf{a}^f(\mathbf{s})$ is continuous in both f and \mathbf{s} by the Maximum Theorem. Since, given any choice of x_0 , the planner and the doer choose the same bundle \mathbf{x} in every state, by definition we have

$$\Psi(f) = \int_S \tilde{U}(a_0^f(\mathbf{s}); \mathbf{s}) dG.$$

Consider any $f > \hat{f}$. Recall that $\bar{S}(f) = \{\mathbf{s} : a_0^*(\mathbf{s}) \leq f\}$. Then,

$$\begin{aligned} \Psi(f) - \Psi(\hat{f}) &= \int_S \left[\tilde{U}(a_0^f(\mathbf{s}); \mathbf{s}) - \tilde{U}(a_0^{\hat{f}}(\mathbf{s}); \mathbf{s}) \right] dG \\ &= \int_{\bar{S}(f)} \left[\tilde{U}(f; \mathbf{s}) - \tilde{U}(a_0^{\hat{f}}(\mathbf{s}); \mathbf{s}) \right] dG \\ &= \int_{\bar{S}(f) \cap (\bar{S}(\hat{f}))^c} \left[\tilde{U}(f; \mathbf{s}) - \tilde{U}(a_0^{\hat{f}}(\mathbf{s}); \mathbf{s}) \right] dG \\ &\quad + \int_{\bar{S}(\hat{f})} \left[\tilde{U}(f; \mathbf{s}) - \tilde{U}(\hat{f}; \mathbf{s}) \right] dG. \end{aligned}$$

where the second equality follows because $a_0^f(\mathbf{s}) = a_0^{\hat{f}}(\mathbf{s})$ for $\mathbf{s} \notin \bar{S}(f)$ and $a_0^f(\mathbf{s}) = f$ for $\mathbf{s} \in \bar{S}(f)$. Dividing both sides by $f - \hat{f}$, we get

$$\begin{aligned} \lim_{f \downarrow \hat{f}} \frac{\Psi(f) - \Psi(\hat{f})}{f - \hat{f}} &= \lim_{f \downarrow \hat{f}} \int_{\bar{S}(f)} \frac{\tilde{U}(f; \mathbf{s}) - \tilde{U}(\hat{f}; \mathbf{s})}{f - \hat{f}} dG \\ &\quad + \lim_{f \downarrow \hat{f}} \int_{\bar{S}(f) \cap (\bar{S}(\hat{f}))^c} \frac{\tilde{U}(f; \mathbf{s}) - \tilde{U}(a_0^{\hat{f}}(\mathbf{s}); \mathbf{s})}{f - \hat{f}} dG. \end{aligned} \quad (11)$$

Consider the first limit. For all \mathbf{s} , we have

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

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

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

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

$$\max_{\{(f, \mathbf{s}) \in [a_0^*, \bar{p}_0^*] \times \mathcal{S}\}} \left| \tilde{U}'(f; \mathbf{s}) \right| = M < +\infty.$$

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

$$\lim_{f \downarrow \hat{f}} \int_{\bar{S}(f)} \frac{\tilde{U}(f; \mathbf{s}) - \tilde{U}(\hat{f}; \mathbf{s})}{f - \hat{f}} dG = \int_{\bar{S}(\hat{f})} \tilde{U}'(\hat{f}; \mathbf{s}) dG.$$

Consider now the second limit in (11). Again, by concavity of $\tilde{U}(\cdot; \mathbf{s})$ and since $a_0^f(\mathbf{s}) \in [\underline{a}_0^*, \bar{p}_0^*]$ for $f \in [\underline{a}_0^*, \bar{p}_0^*]$, we have that

$$\left| \frac{\tilde{U}(f; \mathbf{s}) - \tilde{U}(a_0^f(\mathbf{s}); \mathbf{s})}{f - a_0^f(\mathbf{s})} \right| \leq M.$$

Therefore,

$$\begin{aligned} \left| \int_{\bar{S}(f) \cap (\bar{S}(\hat{f}))^c} \frac{\tilde{U}(f; \mathbf{s}) - \tilde{U}(a_0^f(\mathbf{s}); \mathbf{s})}{f - \hat{f}} dG \right| &\leq \int_{\bar{S}(f) \cap (\bar{S}(\hat{f}))^c} \left| \frac{\tilde{U}(f; \mathbf{s}) - \tilde{U}(a_0^f(\mathbf{s}); \mathbf{s})}{f - \hat{f}} \right| dG \\ &\leq \int_{\bar{S}(f) \cap (\bar{S}(\hat{f}))^c} \left| \frac{\tilde{U}(f; \mathbf{s}) - \tilde{U}(a_0^f(\mathbf{s}); \mathbf{s})}{f - a_0^f(\mathbf{s})} \right| dG \\ &\leq M \int_{\bar{S}(f) \cap (\bar{S}(\hat{f}))^c} dG. \end{aligned}$$

Now, observe that $\bar{S}(f) \cap (\bar{S}(\hat{f}))^c = \{\mathbf{s} : \hat{f} < a_0^f(\mathbf{s}) \leq f\}$ 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 $\hat{f} \in [\underline{a}_0^*, \bar{p}_0^*)$, we have

$$\Psi'(f+) = \int_{\bar{S}(\hat{f})} \tilde{U}'(\hat{f}; \mathbf{s}) dG.$$

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

$$\begin{aligned} \Psi(f) - \Psi(\hat{f}) &= \int_S [\tilde{U}(a_0^f(\mathbf{s}); \mathbf{s}) - \tilde{U}(a_0^{\hat{f}}(\mathbf{s}); \mathbf{s})] dG \\ &= \int_{\bar{S}(\hat{f})} [\tilde{U}(a_0^f(\mathbf{s}); \mathbf{s}) - \tilde{U}(\hat{f}; \mathbf{s})] dG \\ &= \int_{\bar{S}(\hat{f})} [\tilde{U}(f; \mathbf{s}) - \tilde{U}(\hat{f}; \mathbf{s})] dG + \int_{\bar{S}(\hat{f})} [\tilde{U}(a_0^f(\mathbf{s}); \mathbf{s}) - \tilde{U}(f; \mathbf{s})] dG \\ &= \int_{\bar{S}(\hat{f})} [\tilde{U}(f; \mathbf{s}) - \tilde{U}(\hat{f}; \mathbf{s})] dG + \int_{\bar{S}(\hat{f}) \cap (\bar{S}(f))^c} [\tilde{U}(a_0^f(\mathbf{s}); \mathbf{s}) - \tilde{U}(f; \mathbf{s})] dG, \end{aligned}$$

where the second equality follows because $a_0^f(\mathbf{s}) = a_0^{\hat{f}}(\mathbf{s})$ for $\mathbf{s} \notin \bar{S}(\hat{f})$ and $a_0^f(\mathbf{s}) = \hat{f}$ for $\mathbf{s} \in \bar{S}(\hat{f})$, and the last equality follows because $a_0^f(\mathbf{s}) = f$ for $\mathbf{s} \in \bar{S}(f)$. By the same argument as before,

$$\begin{aligned} \lim_{f \uparrow \hat{f}} \int_{\bar{S}(\hat{f})} \frac{\tilde{U}(f; \mathbf{s}) - \tilde{U}(\hat{f}; \mathbf{s})}{f - \hat{f}} dG &= \int_{\bar{S}(\hat{f})} \tilde{U}'(\hat{f}; \mathbf{s}) dG, \\ \lim_{f \uparrow \hat{f}} \int_{\bar{S}(\hat{f}) \cap (\bar{S}(f))^c} \frac{\tilde{U}(a_0^f(\mathbf{s}); \mathbf{s}) - \tilde{U}(f; \mathbf{s})}{f - \hat{f}} dG &= 0. \end{aligned}$$

We conclude that for every $\hat{f} \in (\underline{a}_0^*, \bar{p}_0^*]$, we have

$$\Psi'(f-) = \int_{\bar{S}(\hat{f})} \tilde{U}'(\hat{f}; \mathbf{s}) dG.$$

Hence, $\Psi(f)$ is differentiable over the restricted domain $[\underline{a}_0^*, \bar{p}_0^*]$.

□

A.3 Proof of Lemma 3

For simplicity, drop the subscript from f_0 . We shall show that $\Psi'(f) > 0$ for all $f \in (\underline{a}_0^*, \underline{p}_0^*]$ and $\Psi'(f-) < 0$ for $f = \bar{p}_0^*$. Recall that $a^f(\mathbf{s})$ is continuous in f for every \mathbf{s} and therefore $\Psi(f)$ is continuous in f . These observations imply that the optimal $f^* \in (\underline{p}_0^*, \bar{p}_0^*)$.

For any $f \in (\underline{a}_0^*, \bar{p}_0^*]$, define

$$S^+(f) = \{\mathbf{s} : p_0^*(\mathbf{s}) > f\} \quad \text{and} \quad S^-(f) = \{\mathbf{s} : p_0^*(\mathbf{s}) \leq f\}.$$

For $\mathbf{s} \in S^+(f)$, consider the the following problem:

$$\max \hat{u}(\mathbf{x}; \mathbf{s}) + v(x_0)$$

subject to $\sum_i x_i \leq 1$ and $x_0 \leq f$. The associated Lagrangian is

$$\hat{u}(\mathbf{x}; \mathbf{s}) + v(x_0) + \mu(\mathbf{s}) \left[1 - \sum_{i=0}^n x_i \right] + \phi^+(\mathbf{s})[f - x_0].$$

Hence, the FOC are³³

$$v'(x_0) = \mu(\mathbf{s}) + \phi^+(\mathbf{s}) \quad \text{and} \quad \frac{\partial}{\partial x_i} \hat{u}(\mathbf{x}; \mathbf{s}) = \mu(\mathbf{s}) \quad \text{for all } i.$$

Clearly, the constraint $x_0 \leq f$ must be binding for $\mathbf{s} \in S^+(f)$, which implies that $x_0 = f$ and $\phi^+(\mathbf{s}) > 0$. Also, conditional on choosing $x_0 = f$, both the planner and the doer choose the same \mathbf{x} in state \mathbf{s} , which therefore equals $\mathbf{x}^f(\mathbf{s})$. Using (10), it follows that, for any i ,

$$\phi^+(\mathbf{s}) = v'(f) - \frac{\partial}{\partial x_i} \hat{u}(\mathbf{x}^f(\mathbf{s}); \mathbf{s}) = \tilde{U}'(f; \mathbf{s}) \tag{12}$$

when $\mathbf{s} \in S^+(f)$.

For $\mathbf{s} \in S^-(f)$, consider the following problem:

$$\max \hat{u}(\mathbf{x}; \mathbf{s}) + v(x_0)$$

subject to $\sum_i x_i \leq 1$ and $x_0 \geq f$. The associated Lagrangian is

$$\hat{u}(\mathbf{x}; \mathbf{s}) + v(x_0) + \mu(\mathbf{s}) \left[1 - \sum_{i=0}^n x_i \right] + \phi^-(\mathbf{s})[x_0 - f].$$

Hence, the FOC are

$$v'(x_0) = \mu(\mathbf{s}) - \phi^-(\mathbf{s}) \quad \text{and} \quad \frac{\partial}{\partial x_i} \hat{u}(\mathbf{x}; \mathbf{s}) = \mu(\mathbf{s}) \quad \text{for all } i,$$

Clearly, the constraint $x_0 \geq f$ must be binding for $\mathbf{s} \in S^-(f)$ except when $p_0^*(\mathbf{s}) = f$, which implies that $x_0 = f$ and $\phi^-(\mathbf{s}) \geq 0$. Also, conditional on choosing $x_0 = f$, both the planner

³³Here, as well as in the other proofs, the complementary slackness conditions are omitted for simplicity.

and the doer choose the same \mathbf{x} in state \mathbf{s} , which therefore equals $\mathbf{x}^f(\mathbf{s})$. Using (10), it follows that, for any i ,

$$\phi^-(\mathbf{s}) = \frac{\partial}{\partial x_i} \hat{u}(\mathbf{x}^f(\mathbf{s}); \mathbf{s}) - v'(f) = -\tilde{U}'(f; \mathbf{s})$$

when $\mathbf{s} \in S^-(f)$.

Consider any $f \in (\underline{a}_0^*, \underline{p}_0^*]$. Recall that $\bar{S}(f) = \{\mathbf{s} : a_0^*(\mathbf{s}) \leq f\}$. Using Lemma 10, we have

$$\begin{aligned} \Psi'(f) &= \int_{\bar{S}(f)} \tilde{U}'(f; \mathbf{s}) dG \\ &= \int_{\bar{S}(f) \cap S^+(f)} \tilde{U}'(f; \mathbf{s}) dG + \int_{\bar{S}(f) \cap S^-(f)} \tilde{U}'(f; \mathbf{s}) dG \\ &= \int_{\bar{S}(f) \cap S^+(f)} \phi^+(\mathbf{s}) dG, \end{aligned}$$

where the last equality follows because either $S^-(f) = \emptyset$ or $\phi^-(\mathbf{s}) = 0$ for $\mathbf{s} \in S^-(f)$. The function $\phi^+(\mathbf{s})$ is strictly positive over $\bar{S}(f) \cap S^+(f)$. We need to show that this set has strictly positive measure, which implies $\Psi'(f) > 0$. This is immediate if $f \in (\underline{a}_0^*, \underline{p}_0^*)$, because in this case $S^+(f) = S$. Consider $f = \underline{p}_0^*$. Clearly, $\bar{S}(\underline{p}_0^*) \cap S^+(\underline{p}_0^*)$ contains the open set

$$\bar{S}^\circ(\underline{p}_0^*) \cap S^+(\underline{p}_0^*) = \{\mathbf{s} : a_0^*(\mathbf{s}) < \underline{p}_0^* < p_0^*(\mathbf{s})\}.$$

If we can show that this set is non-empty, we are done because G assigns strictly positive probability to it. Both $\bar{S}^\circ(\underline{p}_0^*)$ and $S^+(\underline{p}_0^*)$ are nonempty. Suppose that there is no $\mathbf{s} \in S^+(\underline{p}_0^*)$ such that we also have $\mathbf{s} \in \bar{S}^\circ(\underline{p}_0^*)$. Then, it means that for every $\mathbf{s} \in S^+(\underline{p}_0^*)$, we have $a_0^*(\mathbf{s}) \geq \underline{p}_0^*$ and that $\bar{S}^\circ(\underline{p}_0^*) \subset S^-(\underline{p}_0^*) = \{\mathbf{s} : p_0^*(\mathbf{s}) = \underline{p}_0^*\}$. Now, consider $\hat{\mathbf{s}} \in \bar{S}^\circ(\underline{p}_0^*)$ and any sequence $\{\mathbf{s}_n\}$ is $S^+(\underline{p}_0^*)$ converging to $\hat{\mathbf{s}}$. We have that

$$\liminf_{\mathbf{s}_n \rightarrow \hat{\mathbf{s}}} a_0^*(\mathbf{s}_n) \geq \underline{p}_0^* > a_0^*(\hat{\mathbf{s}}).$$

But this contradicts the continuity of a^* and hence leads to a contradiction.

Now consider $f = \bar{p}_0^*$. Using again Lemma 10, we have

$$\Psi'(\bar{p}_0^* -) = \int_{\bar{S}(\bar{p}_0^*)} \tilde{U}'(\bar{p}_0^*; \mathbf{s}) dG = \int_S \tilde{U}'(\bar{p}_0^*; \mathbf{s}) dG = - \int_S \phi^-(\mathbf{s}) dG,$$

where $\phi^-(\mathbf{s}) > 0$ for all \mathbf{s} such that $p_0^*(\mathbf{s}) < \bar{p}_0^*$. Therefore, $\Psi'(\bar{p}_0^* -) < 0$.³⁴

A.4 Proof of Lemma 4

Fix $f_0 \in [\underline{a}_0^*, \bar{p}_0^*]$. Changes in b affect $\bar{S}(f_0)$ through the change in \mathbf{a}^* . By standard arguments, if $b < b' < 1$, then $a_0^*(\mathbf{s}; b) < a_0^*(\mathbf{s}; b')$ for every \mathbf{s} and hence $\bar{S}(f_0; b') \subset \bar{S}(f_0; b)$. On the other

³⁴It is easy to see that the optimal f satisfies $f \leq \bar{p}_0^*$. Suppose $f \in (\bar{p}_0^*, 1)$. Then, for all \mathbf{s} , the doer chooses $x_0(\mathbf{s}) = f$ and $\mathbf{x}(\mathbf{s}) = \mathbf{x}^f(\mathbf{s})$. Take any $f' \in (\bar{p}_0^*, f)$. Then, for every \mathbf{s} , $f' = \zeta(\mathbf{s})f + (1 - \zeta(\mathbf{s}))\bar{p}_0^*(\mathbf{s})$ for some $\zeta(\mathbf{s}) \in (0, 1)$. Therefore, for every \mathbf{s} , $\tilde{U}(f'; \mathbf{s}) > \tilde{U}(f; \mathbf{s})$ because $\tilde{U}(\bar{p}_0^*(\mathbf{s}); \mathbf{s}) > \tilde{U}(f; \mathbf{s})$ and $\tilde{U}(\cdot; \mathbf{s})$ is strictly concave. It follows that the planner's payoff is strictly larger under f' than under f .

hand, for every $b < 1$, we have $S^-(f_0) \subset \bar{S}(f_0; b)$ because $a_0^*(\mathbf{s}; b) < p_0^*(\mathbf{s})$ for every \mathbf{s} . So, if $b < b' < 1$, we have

$$\Psi'(f_0; b) - \Psi'(f_0; b') = \int_{(\bar{S}(f_0; b) \setminus \bar{S}(f_0; b')) \cap S^+(f_0)} \phi^+(\mathbf{s}) dG \geq 0,$$

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

Define $\bar{f}_0(b) = \max F(b)$. Since $\bar{f}_0(b) \geq \underline{p}_0^*$ for all b and $\bar{f}_0(\cdot)$ is decreasing, $\lim_{b \uparrow 1} \bar{f}_0(b)$ exists; denote it by $\bar{f}_0(1-) \geq \underline{p}_0^*$. Clearly, $\bar{f}_0(1) = \underline{p}_0^*$. Now suppose that $\bar{f}_0(1-) > \bar{f}_0(1)$. By a similar argument, for any $f_0 > \underline{p}_0^*$, $\lim_{b \uparrow 1} \Psi'(f_0; b)$ exists and satisfies

$$\lim_{b \uparrow 1} \Psi'(f_0; b) = - \int_{S^-(f_0)} \phi(\mathbf{s}) dG < 0.$$

This implies that for b close enough to 1, $\bar{f}_0(b) \geq \bar{f}_0(1-)$ cannot be optimal, a contradiction which implies that $\bar{f}_0(1-) = \bar{f}_0(1)$.

It is easy to see that, for all $\mathbf{s} \in S$, $a_0^*(\mathbf{s}; b) \rightarrow 0$ as $b \downarrow 0$. Therefore, $\bar{a}_0^*(b) = \max_S a_0^*(\mathbf{s}; b)$ also decreases monotonically to 0 as $b \downarrow 0$. Let $\underline{b} = \max\{b \in [0, 1] : \bar{a}_0^*(b) \leq \underline{p}_0^*\}$ which is strictly positive because $\underline{p}_0^* > 0$. Then, $\bar{S}(f_0) = S$ for all $b \leq \underline{b}$ and $f_0 \in [\underline{p}_0^*, \bar{p}_0^*]$, which implies that

$$\Psi(f_0; b) = v(f_0) + \int_S \hat{u}(\mathbf{x}^{f_0}(\mathbf{s}); \mathbf{s}) dG. \quad (13)$$

From the proof of Lemma 10, we have that $\hat{u}(\mathbf{x}^{f_0}(\mathbf{s}); \mathbf{s}) = \tilde{u}(f_0; \mathbf{s})$ is strictly concave in f_0 for all $\mathbf{s} \in S$. This implies that the maximizer of (13) is unique. From the proof of Lemma 3, we know that the derivative of (13) is negative at \bar{p}_0^* and hence $\bar{f}_0 < \bar{p}_0^*$.

A.5 Lemma 11

Lemma 11. *Fix $f_0 > \underline{a}_0^*$ and let $D_{\mathbf{0}, \mathbf{c}}$ be any policy that satisfies $\sum_{i=1}^n c_i = 1 - f_0$ and $c_0 = 1$. If \mathbf{r} is constant, then there exists a $D_{\mathbf{0}, \mathbf{c}}$ that implements the same allocations as D_{f_0} . If \mathbf{r} is not constant, then every $D_{\mathbf{0}, \mathbf{c}}$ implements allocations that differ with positive probability from those implemented by D_{f_0} .*

Proof.

Part 1: Fix $f_0 > \underline{a}_0^*$. Suppose \mathbf{r} is constant. For every $\theta \in [\underline{\theta}, \bar{\theta}]$, the doer maximizes $\theta \sum_{i=1}^n u^i(x_i; r_i) + bv(x_0)$ subject to $(\mathbf{x}, x_0) \in B$ and $x_0 \geq f_0$, which leads to the unique optimal allocation $\mathbf{a}(\theta)$. By strong monotonicity of preferences, $\sum_{i=1}^n a_i(\theta) = 1 - a_0(\theta)$ for every θ . By standard arguments, each $a_i(\cdot)$ is a strictly increasing, continuous function of θ for $i = 1, \dots, n$. Let $c_i = \max_{\theta \in [\underline{\theta}, \bar{\theta}]} a_i(\theta)$. Clearly, $1 - \sum_{i=1}^n c_i = \min_{\theta \in [\underline{\theta}, \bar{\theta}]} a_0(\theta)$. Since f_0 has to be binding for some θ , $\min_{\theta \in [\underline{\theta}, \bar{\theta}]} a_0(\theta) = f_0$. It is clear that if we replace the floor f_0 with the caps $\{c_i\}_{i=1}^n$, the doer's choices across θ 's do not change.

Part 2: Fix $f_0 > \underline{a}_0^*$. Suppose \mathbf{r} is not constant, i.e., $\underline{r}_i < \bar{r}_i$ for some $i = 1, \dots, n$. Let $\bar{\mathbf{r}} = (\bar{r}_1, \dots, \bar{r}_n)$. It is easy to see that $a_0^*(\bar{\theta}, \bar{\mathbf{r}}) = \underline{a}_0^*$. Therefore, f_0 must be binding in state $\bar{\mathbf{s}} = (\bar{\theta}, \bar{\mathbf{r}})$. Since \mathbf{a}^* is continuous in \mathbf{s} , there exists $\varepsilon > 0$ such that, if $|\mathbf{r} - \bar{\mathbf{r}}| < \varepsilon$, then $a_0^*(\bar{\theta}, \mathbf{r}) > f_0$ and hence the floor is still binding. When f_0 binds, the doer's allocation \hat{a}_{-0} must

maximize $\theta \sum_{i=1}^n u^i(x_i; r_i)$ subject to $\sum_{i=1}^n x_i \leq 1 - f_0$. So, for all \mathbf{r} with $|\mathbf{r} - \bar{\mathbf{r}}| < \varepsilon$, we must have

$$u_1^i(\hat{a}_i(\bar{\theta}, \mathbf{r}); r_i) = u_1^j(\hat{a}_j(\bar{\theta}, \mathbf{r}); r_j) \quad \text{for all } i, j.$$

It follows that there exists \mathbf{r}' with $|\mathbf{r}' - \bar{\mathbf{r}}| < \varepsilon$ such that $\hat{a}_{-0}(\bar{\theta}, \mathbf{r}') \neq \hat{a}_{-0}(\bar{\theta}, \bar{\mathbf{r}})$. Since $\sum_{i=1}^n \hat{a}_i(\bar{\theta}, \mathbf{r}') = \sum_{i=1}^n \hat{a}_i(\bar{\theta}, \bar{\mathbf{r}}) = 1 - f_0$, there exists $i \neq j$ such that $\hat{a}_i(\bar{\theta}, \mathbf{r}') > \hat{a}_i(\bar{\theta}, \bar{\mathbf{r}})$ and $\hat{a}_j(\bar{\theta}, \mathbf{r}') < \hat{a}_j(\bar{\theta}, \bar{\mathbf{r}})$. Now let $S(f_0)$ be the set of states for which $\hat{a}_0(\mathbf{s}) = f_0$. By the previous argument, \hat{a}_i and \hat{a}_j cannot be constant over $S(f_0)$.

For each $k = 1, \dots, n$, let $\hat{c}_k = \max_S \hat{a}_k(\mathbf{s})$. When $\hat{a}_i(\mathbf{s}) = \hat{c}_i$, we must have $\hat{a}_j(\mathbf{s}) < \hat{c}_j$, and when $\hat{a}_j(\mathbf{s}) = \hat{c}_j$, we must have $\hat{a}_i(\mathbf{s}) < \hat{c}_i$. Therefore, $\sum_{i=1}^n \hat{c}_i > 1 - f_0$. It follows that any collection of caps $\mathbf{c}_{-0} = \{c_i\}_{i=1}^n$ satisfying $\sum_{i=1}^n c_i = 1 - f_0$ must involve $c_i < \hat{c}_i$ for some $i = 1, \dots, n$. So, when the doer faces \mathbf{c}_{-0} , for some i and state \mathbf{s} , $a_i(\mathbf{s}) \leq c_i$ for all states \mathbf{s} such that the doer chooses $x_i > c_i$ under f_0 . Since $\hat{\mathbf{a}}$ is continuous in \mathbf{s} , the set $S(\mathbf{c}_{-0}) = \{\mathbf{s} : \hat{a}_i(\mathbf{s}) > c_i\}$ is open and hence it has strictly positive probability under G . \square

A.6 Lemma 12

Lemma 12. Fix $i \neq 0$ and consider $D_{\mathbf{0}, \mathbf{c}} \in \mathcal{R}$ with $c_j = 1$ for all $j \neq i$. In any state \mathbf{s} , if $c_i < a_i^*(\mathbf{s})$, then the doer chooses $x_0 > a_0^*(\mathbf{s})$, but also $x_j > a_j^*(\mathbf{s})$ for all $j \neq i$.

Proof. Without loss, let $i = 1$ and take any $c_1 \in (0, a_1^*(\mathbf{s}))$. Consider the doer's problem in state \mathbf{s} subject to c_1 :

$$\max_{\{(\mathbf{x}, x_0) \in B : x_1 \leq c_1\}} \hat{u}(\mathbf{x}, \mathbf{s}) + bv(x_0).$$

The first-order conditions of the associated Lagrangian are

$$\begin{aligned} bv'(a_0(\mathbf{s})) &= \mu(\mathbf{s}), \\ \theta u_1^1(a_1(\mathbf{s}); r_1) &= \mu(\mathbf{s}) + \lambda_1(\mathbf{s}), \\ \theta u_1^j(a_j(\mathbf{s}); r_j) &= \mu(\mathbf{s}) \quad \text{for all } j \neq 0, 1, \end{aligned}$$

where $\mu(\mathbf{s}) \geq 0$ and $\lambda_1(\mathbf{s}) \geq 0$ are the Lagrange multipliers for constraints $\sum_{i=1}^n x_i \leq 1$ and $x_1 \leq c_1$.

Suppose $a_0(\mathbf{s}) \leq a_0^*(\mathbf{s})$. Since $a_1(\mathbf{s}) = c_1 < a_1^*(\mathbf{s})$ and $\sum_j a_j(\mathbf{s}) = \sum_j a_j^*(\mathbf{s}) = 1$ by strong monotonicity of preferences, $a_j(\mathbf{s}) > a_j^*(\mathbf{s})$ for some $j \neq 0, 1$. By strict concavity of u_j and v ,

$$\theta u_1^j(a_j(\mathbf{s}); r_j) < \theta u_1^j(a_j^*(\mathbf{s}); r_j) = bv'(a_0^*(\mathbf{s})) \leq bv'(a_0(\mathbf{s})).$$

This violates the first-order conditions for $a(\mathbf{s})$. So we must have $a_0(\mathbf{s}) > a_0^*(\mathbf{s})$. This in turn implies that for $j \neq i$

$$\theta u_1^j(a_j(\mathbf{s}); r_j) = bv'(a_0(\mathbf{s})) < bv'(a_0^*(\mathbf{s})) = \theta u_1^j(a_j^*(\mathbf{s}); r_j).$$

By concavity, we have $a_j(\mathbf{s}) > a_j^*(\mathbf{s})$ for $j \neq 0, 1$. \square

A.7 Proof of Lemma 5

Fix $i = 1$ and consider any $c_1 \leq \bar{a}_1^*$. Let \mathbf{a}^{c_1} describe the doer's choices under cap c_1 . Then, let

$$\Phi(c_1) = \int_S U(\mathbf{a}^{c_1}(\mathbf{s}); \mathbf{s}) dG.$$

Let $S(c_1) = \{\mathbf{s} : a_1^*(\mathbf{s}) > c_1\}$. Note that for any $c_1 < \bar{a}_1^*$, since a_1^* is continuous, $S(c_1)$ is non-empty and open and hence has strictly positive probability under G . We have

$$\begin{aligned} \Phi(c_1) - \Phi(\bar{a}_1^*) &= \int_{S(c_1)} [U(\mathbf{a}^{c_1}(\mathbf{s}); \mathbf{s}) - U(\mathbf{a}^*(\mathbf{s}); \mathbf{s})] dG \\ &= (1 - b) \int_{S(c_1)} [v(a_0^{c_1}(\mathbf{s})) - v(a_0^*(\mathbf{s}))] dG \\ &\quad + \int_{S(c_1)} [\tilde{V}(a_1^{c_1}(\mathbf{s}); \mathbf{s}) - \tilde{V}(a_1^*(\mathbf{s}); \mathbf{s})] dG \end{aligned}$$

where

$$\tilde{V}(\hat{c}_1; \mathbf{s}) = V(\mathbf{a}^{\hat{c}_1}(\mathbf{s}); \mathbf{s}) = \max_{\{(\mathbf{x}, x_0) \in \mathbb{R}_+^{n+1} : \sum_{j=1}^n x_j \leq 1, x_1 \leq \hat{c}_1\}} \{\hat{u}(\mathbf{x}; \mathbf{s}) + bv(x_0)\}.$$

Clearly, $\tilde{V}(a_1^*(\mathbf{s}); \mathbf{s}) \geq \tilde{V}(c_1; \mathbf{s})$ for every \mathbf{s} . From the first-order conditions of the Lagrangian defining $V(a^{\hat{c}_1}(\mathbf{s}); \mathbf{s})$, we have $\lambda_1(\mathbf{s}; \hat{c}_1) = \theta u_1^1(a_1^{\hat{c}_1}(\mathbf{s}); \mathbf{r}) - bv'(a_0^{\hat{c}_1}(\mathbf{s}))$, where $\lambda_1(\mathbf{s}; \hat{c}_1)$ is the Lagrange multiplier on the constraint $x_1 \leq \hat{c}_1$. Since $\mathbf{a}^{\hat{c}_1}(\mathbf{s})$ is continuous in \hat{c}_1 as well as \mathbf{s} , so is $\lambda_1(\mathbf{s}; \hat{c}_1)$. Relying again on Theorem 1, p. 222, of Luenberger (1969), we conclude that $\tilde{V}'(\hat{c}_1; \mathbf{s})$ exists for every \hat{c}_1 and equals $\lambda_1(\mathbf{s}; \hat{c}_1)$. It follows that $\tilde{V}'(a_1^*(\mathbf{s}); \mathbf{s}) = 0$ for every \mathbf{s} by the definition of a^* . Therefore, by the Mean Value Theorem,

$$\tilde{V}(a_1^{c_1}(\mathbf{s}); \mathbf{s}) - \tilde{V}(a_1^*(\mathbf{s}); \mathbf{s}) = \tilde{V}'(\chi(\mathbf{s}); \mathbf{s})(a_1^{c_1}(\mathbf{s}) - a_1^*(\mathbf{s})),$$

$$v(a_0^{c_1}(\mathbf{s})) - v(a_0^*(\mathbf{s})) = v'(\xi(\mathbf{s}))(a_0^{c_1}(\mathbf{s}) - a_0^*(\mathbf{s})),$$

where $\chi(\mathbf{s}) \in [a_1^{c_1}(\mathbf{s}), a_1^*(\mathbf{s})]$ and $\xi(\mathbf{s}) \in [a_0^*(\mathbf{s}), a_0^{c_1}(\mathbf{s})]$.

Let $c_1^\varepsilon = \bar{a}_1^* - \varepsilon$ for some small $\varepsilon > 0$. Fix any $\mathbf{s} \in S(c_1^\varepsilon)$ and, for now, suppress the dependence on \mathbf{s} for simplicity. Recall that $\sum_i a_i^{c_1^\varepsilon} = \sum_i a_i^* = 1$. Since $a_0^{c_1^\varepsilon} > a_0^*$ for any $\varepsilon > 0$, we can write

$$-\frac{a_1^{c_1^\varepsilon} - a_1^*}{a_0^{c_1^\varepsilon} - a_0^*} = 1 + \sum_{j \neq 0,1} \frac{a_j^{c_1^\varepsilon} - a_j^*}{a_0^{c_1^\varepsilon} - a_0^*}.$$

Now, for any c_1^ε , the following first order condition must hold for every $j \neq 1$:

$$bv'(a_0) - \theta u_{11}^j(a_j; r_j) = 0.$$

This defines an implicit function $a_j(a_0)$ and, by the Implicit Function Theorem,

$$\frac{d}{da_0} a_j(a_0) = \frac{bv''(a_0)}{\theta u_{11}^j(a_j(a_0); r_j)}.$$

Since $u_{11}^j < 0$, $v'' < 0$, $\theta > 0$, we have $\frac{d}{da_0} a_j > 0$ everywhere. Moreover, u_{11}^j and v'' are continuous and we can restrict attention to a_0 and a_j that take values in the compact set

$[\underline{a}_0^*, 1] \times [\underline{a}_j^*, 1]$ where $\underline{a}_0^* > 0$ and $\underline{a}_j^* > 0$. Therefore, $\frac{d}{da_0} a_j$ is bounded above by some finite $k_j > 0$ for all $\mathbf{s} \in S(c_1^\varepsilon)$. Hence, for any $\varepsilon > 0$, $a_j^{c_1^\varepsilon} - a_j^* \leq k_j(a_0^{c_1^\varepsilon} - a_0^*)$. Letting $K = \sum_{j \neq 0,1} k_j$, we then have

$$-\frac{a_1^{c_1^\varepsilon} - a_1^*}{a_0^{c_1^\varepsilon} - a_0^*} \leq 1 + K \quad \Rightarrow \quad a_0^{c_1^\varepsilon} - a_0^* \geq \frac{a_1^* - a_1^{c_1^\varepsilon}}{1 + K}.$$

Using these observations, we have that $\Phi(c_1^\varepsilon) - \Phi(\bar{a}_1^*)$ is bounded below by

$$\int_{S(c_1^\varepsilon)} \left[\frac{1-b}{1+K} v'(\xi(\mathbf{s})) - \tilde{V}'(\chi(\mathbf{s}); \mathbf{s}) \right] (a_1^*(\mathbf{s}) - c_1^\varepsilon) dG. \quad (14)$$

Since v' is continuous and strictly positive everywhere and $\xi(\mathbf{s}) \in [\underline{a}_0^*, 1]$ with $\underline{a}_0^* > 0$ for all $\mathbf{s} \in S(c_1^\varepsilon)$, there exists a finite $\kappa > 0$ such that $v'(\xi(\mathbf{s})) \geq \kappa$ for all $\mathbf{s} \in S(c_1^\varepsilon)$.

Next let $\bar{S}(c_1^\varepsilon) = \{\mathbf{s} : a_1^*(\mathbf{s}) \geq c_1^\varepsilon\}$ which is a closed and bounded set by continuity of a_1^* and hence is compact. As a function of c_1^ε , the correspondence $\bar{S}(\cdot)$ is continuous by continuity of a_1^* . Note that, if $a_1^*(\mathbf{s}) = c_1^\varepsilon$, then $\tilde{V}'(\chi(\mathbf{s}); \mathbf{s}) = \tilde{V}'(a_1^*(\mathbf{s}); \mathbf{s}) = 0$. We have

$$\sup_{\mathbf{s} \in S(c_1^\varepsilon)} \tilde{V}'(\chi(\mathbf{s}); \mathbf{s}) = \sup_{\mathbf{s} \in \bar{S}(c_1^\varepsilon)} \tilde{V}'(\chi(\mathbf{s}); \mathbf{s}) \leq \max_{c_1^\varepsilon \leq \zeta \leq \bar{a}_1^*, \mathbf{s} \in \bar{S}(c_1^\varepsilon)} \tilde{V}'(\zeta; \mathbf{s}) \equiv \kappa(c_1^\varepsilon).$$

Clearly, $\kappa(c_1^\varepsilon) \geq 0$ for any $\varepsilon > 0$, $\varepsilon' > \varepsilon > 0$ implies that $\kappa(c_1^{\varepsilon'}) \leq \kappa(c_1^\varepsilon)$, and $\lim_{\varepsilon \rightarrow 0} \kappa(c_1^\varepsilon) = 0$ because $\kappa(\cdot)$ is also continuous. Therefore, there exists $\varepsilon^* > 0$ such that

$$\kappa(c_1^{\varepsilon^*}) < \kappa \frac{1-b}{1+K}.$$

It follows that for ε^* , expression (14) is strictly positive and hence $\Phi(c_1^{\varepsilon^*}) > \Phi(\bar{a}_1^*)$. This also holds for all $\varepsilon \in (0, \varepsilon^*)$.

A.8 Proof of Lemma 6

Let $\hat{\mathbf{a}}$ and $\hat{\mathbf{a}}'$ describe the doer's choices across states under $D_{\mathbf{f}, \mathbf{c}}$ and $D_{f_0, \mathbf{c}_{-0}}$. Then, $\mathcal{U}(D_{f_0, \mathbf{c}_{-0}}) - \mathcal{U}(D_{\mathbf{f}, \mathbf{c}})$ equals

$$\begin{aligned} \int_S [U(\hat{\mathbf{a}}'(\mathbf{s}); \mathbf{s}) - U(\hat{\mathbf{a}}(\mathbf{s}); \mathbf{s})] dG &= \int_S (1-b) [v(\hat{a}'_0(\mathbf{s})) - v(\hat{a}_0(\mathbf{s}))] dG \\ &\quad + \int_S [V(\hat{a}'(\mathbf{s}); \mathbf{s}) - V(\hat{a}(\mathbf{s}); \mathbf{s})] dG \\ &= \int_S (1-b) [v(\hat{a}'_0(\mathbf{s})) - v(\hat{a}_0(\mathbf{s}))] dG \\ &\quad + \int_S [\hat{V}(D_{f_0, \mathbf{c}_{-0}}; \mathbf{s}) - \hat{V}(D_{\mathbf{f}, \mathbf{c}}; \mathbf{s})] dG, \end{aligned} \quad (15)$$

where for any $(\tilde{\mathbf{f}}, \tilde{\mathbf{c}})$

$$\hat{V}(D_{\tilde{\mathbf{f}}, \tilde{\mathbf{c}}}; \mathbf{s}) = \max_{\{(\mathbf{x}, x_0) \in B: \hat{\mathbf{f}} \leq (\mathbf{x}, c_0) \leq \tilde{\mathbf{c}}\}} V(\mathbf{x}, x_0; \mathbf{s}).$$

Clearly, for every \mathbf{s} , $\hat{V}(D_{f_0, \mathbf{c}_{-0}}; \mathbf{s}) \geq \hat{V}(D_{\mathbf{f}, \mathbf{c}}; \mathbf{s})$. Moreover, the inequality is strict in states in which either c_0 or f_i are binding for the doer, given the strict concavity of the doer's payoff

function and convexity of the feasible set for the doer under both $D_{f_0, c_{-0}}$ and $D_{\mathbf{f}, \mathbf{c}}$. Therefore, if any of them binds with strictly positive probability, the second integral in (15) is strictly positive.

Now consider the first integral, if we can show that $\hat{a}'_0(\mathbf{s}) \geq \hat{a}_0(\mathbf{s})$ for every \mathbf{s} , we are done. To show this, we proceed in steps, removing one constraint from $D_{\mathbf{f}, \mathbf{c}}$ at a time. Consider first removing only c_0 which leads to an intermediate behavior of the doer described by the function \mathbf{a}^0 . If c_0 is never binding for the doer, then it does not affect his choices and hence $a_0^0(\mathbf{s}) = \hat{a}_0(\mathbf{s})$ for every \mathbf{s} . In any state \mathbf{s} in which c_0 is binding, removing only this cap cannot decrease $\hat{a}_0(\mathbf{s})$ because the doer could have decreased it when the cap was in place. Thus, $a_0^0(\mathbf{s}) \geq \hat{a}_0(\mathbf{s})$ for every \mathbf{s} . Note that, once we remove the cap on x_0 , for all \mathbf{s} we must have $\sum_i x_i = 1$ because v is strictly increasing.

Now consider removing one floor f_i for $i \neq 0$ at a time. Fix any state \mathbf{s} and suppress the dependence on it for simplicity. The Lagrangian of the doer's problem after we remove only c_0 is

$$\theta \sum_{i=1}^n u^i(x_i; r_i) + bv(x_0) + \mu \left[1 - \sum_{i=0}^n x_i \right] + \sum_{i=1}^n \gamma_i [c_i - x_i] + \sum_{i=0}^n \phi_i [x_i - f_i].$$

Hence, the first-order necessary and sufficient conditions are

$$\theta u_1^i(a_i^0; r_i) - \mu^0 + \phi_i^0 - \gamma_i^0 = 0 \quad \text{for } i = 1, \dots, n,$$

$$bv'(a_0^0) - \mu^0 + \phi_0^0 = 0,$$

with the usual complementary-slackness conditions. Without loss, start by removing f_1 , thus obtaining \mathbf{a}^1 . First, if $a_0^0 = f_0$, then $a_0^1 \geq a_0^0$. So suppose that $a_0^0 > f_0$ so that $\phi_0^0 = 0$. If $\phi_1^0 = 0$, then removing f_1 has no effect and hence again $a_0^1 \geq a_0^0$. So suppose that $\phi_1^0 > 0$; since $c_1 \geq f_1 = a_0^0$, it follows that $\gamma_1^0 = 0$ without loss of generality.³⁵ After removing f_1 only, the new conditions are

$$\theta u_1^i(a_i^1; r_i) - \mu^1 + \phi_i^1 - \gamma_i^1 = 0 \quad \text{for } i = 1, \dots, n,$$

$$bv'(a_0^1) - \mu^1 + \phi_0^1 = 0.$$

Clearly, at the resulting \mathbf{a}^1 , we must have $a_1^1 < a_0^1$ because the opposite choice was feasible for the doer before removing f_1 . Suppose $a_0^1 < a_0^0$. Then, we must have $a_j^1 > a_j^0$ for some $j \neq 0, 1$, because $\sum_{i=0}^n a_i^0 = \sum_{i=0}^n a_i^1 = 1$, and hence $u_1^j(a_j^1; r_j) < u_1^j(a_j^0; r_j)$ by strict concavity. To see that this leads to a contradiction, first observe that we must have $\gamma_j^0 = 0$, because if $\gamma_j^0 > 0$, then $a_j^0 = c_j \geq a_j^1$. Given this, then

$$bv'(a_0^1) + \gamma_j^1 = \theta u_1^j(a_j^1; r_j) < \theta u_1^j(a_j^0; r_j) = bv'(a_0^0) - \phi_j^0,$$

but this condition cannot hold because $v'(a_0^1) > v'(a_0^0)$ for $a_0^1 < a_0^0$ by our starting assumption. We conclude that $a_0^1 \geq a_0^0$.

Continuing in this way, we can remove every f_i for $i = 2, \dots, n$, obtaining at each step that $a_i^i \geq a_i^{i-1}$. Since $a_0^n = \hat{a}'_0$, by transitivity we get $\hat{a}'_0 \geq \hat{a}_0$. Since this steps assumed an arbitrary \mathbf{s} , we have that $\hat{a}'_0(\mathbf{s}) \geq \hat{a}_0(\mathbf{s})$ for every \mathbf{s} as desired.

³⁵Recall that, by Lagrange Duality, γ_1^0 is the result of a minimization of the Lagrangian at \mathbf{a}^0 .

A.9 Proof of Lemma 8

By Lemma 6, we can focus on policies $D \in \mathcal{R}$ that satisfy $c_0 = 1$. For such policies, define

$$\underline{x}_0 = \max\{f_0, 1 - \sum_{i=1}^n c_i\}$$

Given D , we have that $\sum_{i=0}^n a_i(\mathbf{s}) = 1$ and hence $a_0(\mathbf{s}) \geq \underline{x}_0$ for all $\mathbf{s} \in S$. Without loss of generality, we can let $\underline{x}_0 = \min_S a_0(\mathbf{s}) = \underline{a}_0$.³⁶

Now fix $b \in (0, 1)$. Suppose D' is optimal, but $\underline{x}'_0 < \underline{p}_0^*$. Consider $D'' \in \mathcal{R}$ identical to D' , except that $f''_0 = \underline{p}_0^*$. We claim that $\mathcal{U}(D'') > \mathcal{U}(D')$, which contradicts the optimality of D' . Since D' is convex and compact, the ensuing allocation \mathbf{a}' is a continuous function of \mathbf{s} . Hence, the set $S(\underline{p}_0^*) = \{\mathbf{s} \in S : a'_0(\mathbf{s}) < \underline{p}_0^*\}$ contains an open subset and hence has strictly positive probability under G .

Consider any $\mathbf{s} \in S(\underline{p}_0^*)$. Suppose the planner faces the following problem:

$$\max\{\hat{u}(\mathbf{x}; \mathbf{s}) + v(x_0)\}$$

subject to $(\mathbf{x}, x_0) \in \mathbb{R}_+^{n+1}$, $x_i \leq c'_i$, and $x_0 \leq f_0$. For any $f_0 < \underline{p}_0^*$, the latter constraint must bind for the planner because, by the same logic of Lemma 12, she would choose $p_0(\mathbf{s}) \geq p_0^*(\mathbf{s}) \geq \underline{p}_0^*$ if facing only the constraints $x_i \leq c'_i$ for $i = 1, \dots, n$. Therefore, the planner's payoff from this fictitious problem is strictly increasing in f_0 for $f_0 \leq \underline{p}_0^*$. When the doer faces D'' , the constraint $x_0 \geq \underline{p}_0^*$ must bind. Hence, his allocation $\mathbf{a}''(\mathbf{s}) = (a''_{-0}(\mathbf{s}), \underline{p}_0^*)$ solves $\max \hat{u}(\mathbf{x}; \mathbf{s})$ subject to $\mathbf{x} \in \mathbb{R}_+^n$, $x_i \leq c'_i$, and $\sum_{i=1}^n x_i \leq 1 - \underline{p}_0^*$. This allocation coincides with the planner's allocation under the fictitious problem with $f_0 = \underline{p}_0^*$. Hence, in \mathbf{s} , $\mathbf{a}''(\mathbf{s})$ is strictly better for the planner than $\mathbf{a}'(\mathbf{s})$.

We conclude that, for all $\mathbf{s} \in S(\underline{p}_0^*)$, the planner's payoff is strictly larger under D'' than under D' . Since for $\mathbf{s} \notin S(\underline{p}_0^*)$ the doer's allocation is unchanged, we must have $\mathcal{U}(D'') > \mathcal{U}(D')$.

A.10 Proof of Proposition 1

By Proposition 4, $\bar{f}_0(b) = \max F(b)$ decreases monotonically to \underline{p}_0^* when $b \uparrow 1$. Also, for every $i = 1, \dots, n$, we have that $a_0^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i}; b)$ increases monotonically to $p_0^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i})$ as $b \uparrow 1$. By Lemma 7, $p_0^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i}) > \underline{p}_0^*$. Given this, define

$$b^* = \inf\{b \in (0, 1) : \bar{f}_0(b) < \max_i a_0^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i}; b)\}.$$

Clearly, $b^* < 1$ and for every $b > b^*$ we have $a_0^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i}; b) > \bar{f}_0(b)$ for at least some $i = 1, \dots, n$. Hereafter, fix $b > b^*$ and any i that satisfies this last condition.

For $\varepsilon \geq 0$, consider $c_i^\varepsilon = \bar{a}_i^* - \varepsilon$ as in Proposition 5 where $\bar{a}_i^* = a_i^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i})$ by Lemma 7. Let $\Phi(c_i^\varepsilon, \bar{f}_0)$ be the planner's expected payoff from adding c_i^ε to the existing optimal floor \bar{f}_0 . We will show that there exists $\varepsilon > 0$ such that $\Phi(c_i^\varepsilon, \bar{f}_0) > \Phi(c_i^0, \bar{f}_0)$ where $\Phi(c_i^0, \bar{f}_0) = \Psi(\bar{f}_0)$ in Section (4.1). To do so, for any $\varepsilon \geq 0$, let \mathbf{a}^ε be the doer's allocation function under $(c_i^\varepsilon, \bar{f}_0)$ and $S(c_i^\varepsilon) = \{\mathbf{s} \in S : a_i^0(\mathbf{s}) > c_i^\varepsilon\}$. Then,

$$\Phi(c_i^\varepsilon, \bar{f}_0) - \Phi(c_i^0, \bar{f}_0) = \int_{S(c_i^\varepsilon)} [U(\mathbf{a}^\varepsilon(\mathbf{s}); \mathbf{s}) - U(\mathbf{a}^0(\mathbf{s}); \mathbf{s})] dG.$$

³⁶If $\underline{a}_0 > \underline{x}_0$, we could simply raise f_0 to \underline{a}_0 and nothing would change.

Note that, if there exists $\bar{\varepsilon} > 0$ such that for all $0 < \varepsilon < \bar{\varepsilon}$ we have $\mathbf{a}^0(\mathbf{s}) = \mathbf{a}^*(\mathbf{s})$ for all $\mathbf{s} \in S(c_i^\varepsilon)$, then for such ε 's the previous difference equals $\Phi(c_i^\varepsilon) - \Phi(\bar{a}_i^*)$ in the proof of Proposition 5. The conclusion of that proof then implies that there exists $\varepsilon^{**} \in (0, \bar{\varepsilon})$ such that $\Phi(c_i^{\varepsilon^{**}}, \bar{f}_0) > \Phi(c_i^0, \bar{f}_0)$.

Thus we only need to prove the existence of $\bar{\varepsilon}$. Let $\bar{S}(\bar{f}_0) = \{\mathbf{s} \in S : a_0^*(\mathbf{s}) \leq \bar{f}_0\}$, which is compact by continuity of \mathbf{a}^* . Define $\tilde{a}_i = \max_{\bar{S}(\bar{f}_0)} a_i^0(\mathbf{s})$ which is well defined by continuity of \mathbf{a}^0 . Since $a_0^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i}) > \bar{f}_0$, it follows that $(\bar{\theta}, \bar{r}_i, \underline{r}_{-i}) \notin \bar{S}(\bar{f}_0)$ and hence $a_i^0(\bar{\theta}, \bar{r}_i, \underline{r}_{-i}) = a_i^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i})$ where $a_i^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i}) = \bar{a}_i^*$ by Lemma 7. We must also have $\tilde{a}_i < \bar{a}_i^*$: for all $\mathbf{s} \in \bar{S}(\bar{f}_0)$, optimality requires

$$\theta u_1^i(a_i(\mathbf{s}); r_i) = v'(\bar{f}_0) + \lambda_0(\mathbf{s}) > v'(a_0^*(\bar{\theta}, \bar{r}_i, \underline{r}_{-i})) = \bar{\theta} u_1^i(\bar{a}_i^*; \bar{r}_i),$$

where $\lambda_0(\mathbf{s}) \geq 0$ is the Lagrange multiplier for constraint $x_0 \geq \bar{f}_0$. If $\mathbf{s} \in S$ is such that $a_i^0(\mathbf{s}) > \tilde{a}_i$, then $\mathbf{s} \notin \bar{S}(\bar{f}_0)$ —otherwise it would contradict the definition of \tilde{a}_i —and therefore $a^0(\mathbf{s}) = \mathbf{a}^*(\mathbf{s})$. Now define $\bar{\varepsilon} = \bar{a}_i^* - \tilde{a}_i > 0$. By construction for any $\varepsilon \in (0, \bar{\varepsilon})$, $a_i^0(\mathbf{s}) > c_i^\varepsilon$ implies that $a^0(\mathbf{s}) = \mathbf{a}^*(\mathbf{s})$, as desired.

A.11 Proof of Proposition 2

We first show that there exists $b_{**} > 0$ such that, if $b < b_{**}$, then for any $D \in \mathcal{R}$ with $\underline{x}_0 \geq \underline{p}_0^*$ the resulting \mathbf{a} satisfies $a_0(\mathbf{s}) = \underline{x}_0$ for all $\mathbf{s} \in S$. It is enough to show that $a_0(\underline{\mathbf{s}}) = \bar{a}_0 = \max_S a_0(\mathbf{s})$ must equal \underline{x}_0 . By strict concavity of v , $v'(\bar{a}_0) \leq v'(\underline{p}_0^*) < +\infty$ because $\underline{p}_0^* > 0$. By considering the Lagrangian of the doer's problem in state $\underline{\mathbf{s}}$ (see Proposition 6's proof), we have that $\mathbf{a}(\underline{\mathbf{s}})$ must satisfy

$$bv'(a_0(\underline{\mathbf{s}})) + \phi_0(\underline{\mathbf{s}}) + \gamma_i(\underline{\mathbf{s}}) = \theta u_1^i(a_i(\underline{\mathbf{s}}); r_i) \quad \text{for all } i = 1, \dots, n,$$

where $\phi_0(\underline{\mathbf{s}}) \geq 0$ and $\gamma_i(\underline{\mathbf{s}}) \geq 0$ are the Lagrange multipliers for constraints $x_0 \geq f_0$ and $x_i \leq c_i$. For every $i = 1, \dots, n$, since $a_i(\underline{\mathbf{s}}) \leq 1$ and $u^i(\cdot; r_i)$ is strictly concave, $u_1^i(a_i(\underline{\mathbf{s}}); r_i) \geq u_1^i(1; r_i) > 0$. Now let

$$b_{**} = \min_i \frac{\theta u_1^i(1; r_i)}{v'(\underline{p}_0^*)} > 0. \quad (16)$$

Then, for any $b < b_{**}$, we have $bv'(a_0(\underline{\mathbf{s}})) < \theta u_1^i(a_i(\underline{\mathbf{s}}); r_i)$ for all $i = 1, \dots, n$. Therefore, $\phi_0(\underline{\mathbf{s}}) + \gamma_i(\underline{\mathbf{s}}) > 0$ for all $i = 1, \dots, n$. Hence, either $\phi_0(\underline{\mathbf{s}}) > 0$, in which case $\bar{a}_0 = f_0 = \underline{x}_0$; or $\gamma_i(\underline{\mathbf{s}}) > 0$ for all $i = 1, \dots, n$, in which case $\bar{a}_0 = 1 - \sum_{i=1}^n a_i(\underline{\mathbf{s}}) = 1 - \sum_{i=1}^n c_i = \underline{x}_0$.

Finally, let $b < b_* = \min\{b, b_{**}\}$ where $b > 0$ was defined in Proposition 4. Let $D^b \in \mathcal{R}$ be an optimal policy for b . By Proposition 8, $\underline{x}_0^b \geq \underline{p}_0^*$. The previous result then implies that

$$\mathcal{U}(D^b) = v(\underline{x}_0^b) + \int_S \hat{u}(a_{-0}(\mathbf{s}); \mathbf{s}) dG.$$

Hence,

$$\mathcal{U}(D^b) \leq v(\underline{x}_0^b) + \int_S \hat{u}(\mathbf{x}^{\underline{x}_0^b}(\mathbf{s}); \mathbf{s}) dG \leq v(\bar{f}_0) + \int_S \hat{u}(\mathbf{x}^{\bar{f}_0}(\mathbf{s}); \mathbf{s}) dG = \mathcal{U}(D_{\bar{f}_0}),$$

where the first inequality follows since $\hat{u}(a_{-0}(\mathbf{s}); \mathbf{s}) \leq \max_{\{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i \leq \underline{x}_0^b\}} \hat{u}(\mathbf{x}; \mathbf{s}) = \hat{u}(\mathbf{x}^{\underline{x}_0^b}(\mathbf{s}); \mathbf{s})$ for all $\mathbf{s} \in S$ and from the definition of \bar{f}_0 in Proposition 4.1. It is immediate to see that if

D^b involves caps that bind for a set of states S' whose probability is strictly positive, then $\hat{u}(a_{-0}(\mathbf{s}); \mathbf{s}) < \hat{u}(\mathbf{x}^b(\mathbf{s}); \mathbf{s})$ for all $\mathbf{s} \in S'$, and hence $\mathcal{U}(D^b) < \mathcal{U}(D_{\bar{f}_0})$. Therefore, optimal policies can only involve a private-consumption floor.

Finally, let $\underline{\mathbf{r}}'$, $\underline{\mathbf{r}}$, $\bar{\mathbf{r}}'$, and $\bar{\mathbf{r}}$ satisfy the properties in the statement of Proposition 2. The corresponding states $\underline{\mathbf{s}}'$, $\underline{\mathbf{s}}$, $\bar{\mathbf{s}}'$, and $\bar{\mathbf{s}}$ satisfy the same properties. It follows that $\underline{p}_0^{*'} = p_0^*(\bar{\mathbf{s}}') \geq p_0^*(\bar{\mathbf{s}}) = \underline{p}_0^*$ with strict inequality if $\bar{\mathbf{s}} \neq \bar{\mathbf{s}}'$ (Lemma 7). Similarly, for each $b \in (0, 1)$, $\bar{a}_0^{*'}(b) = a_0^*(\underline{\mathbf{s}}'; b) \leq a_0^*(\underline{\mathbf{s}}; b) = \bar{a}_0^*(b)$ again with strict inequality if $\underline{\mathbf{s}}' \neq \underline{\mathbf{s}}$. Using the definition of b_{**} in (16), the strict concavity of the function v , and that $\underline{r}'_i \geq \underline{r}_i$, we have that $b_{**}' > b_{**}$. Using the definition of \underline{b} in the proof of Proposition 4 and that \bar{a}_0^* is strictly increasing in b , we have that $\underline{b}' > \underline{b}$. Therefore $b_*' > b_*$.

A.12 Proof of Proposition 3

By an argument similar to the proof of Lemma 2, we can conclude that an optimal policy $D^* \in \mathcal{R}$ 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 the floor f_0 , she sets $f_0 = p_0^*(\mathbf{s}^0)$.

Proof. If D can use only f_0 , we can clearly focus on $f_0 \in [a_0^*(\mathbf{s}^1), p_0^*(\mathbf{s}^1)] \cup \{p_0^*(\mathbf{s}^0)\}$. If $f_0 = p_0^*(\mathbf{s}^0)$, the planner's payoff is

$$gU(\mathbf{p}^*(\mathbf{s}^0); \mathbf{s}^0) + (1 - g)U(\mathbf{a}^*(\mathbf{s}^1); \mathbf{s}^1);$$

if $f_0 \in [a_0^*(\mathbf{s}^1), p_0^*(\mathbf{s}^1)]$, her payoff is

$$gU(\mathbf{x}^{f_0}(\mathbf{s}^0), f_0; \mathbf{s}^0) + (1 - g)U(\mathbf{x}^{f_0}(\mathbf{s}^1), f_0; \mathbf{s}^1),$$

where $\mathbf{x}^{f_0}(\mathbf{s})$ is defined in Lemma 10. Thus, $f_0 = p_0^*(\mathbf{s}^0)$ identifies the best policy that involves only f_0 if

$$\frac{g}{1 - g} > \max_{f_0 \in [a_0^*(\mathbf{s}^1), p_0^*(\mathbf{s}^1)]} \frac{U(\mathbf{x}^{f_0}(\mathbf{s}^1), f_0; \mathbf{s}^1) - U(\mathbf{a}^*(\mathbf{s}^1); \mathbf{s}^1)}{U(\mathbf{p}^*(\mathbf{s}^0); \mathbf{s}^0) - U(\mathbf{x}^{f_0}(\mathbf{s}^0), f_0; \mathbf{s}^0)} \geq 0. \quad (17)$$

The term on the right-hand side is well defined; also, for all $f_0 \in [a_0^*(\mathbf{s}^1), p_0^*(\mathbf{s}^1)]$ we have $U(\mathbf{p}^*(\mathbf{s}^1); \mathbf{s}^1) \geq U(\mathbf{x}^{f_0}(\mathbf{s}^1), f_0; \mathbf{s}^1) \geq U(\mathbf{a}^*(\mathbf{s}^1); \mathbf{s}^1)$ and $U(\mathbf{p}^*(\mathbf{s}^0); \mathbf{s}^0) > U(\mathbf{x}^{f_0}(\mathbf{s}^0), f_0; \mathbf{s}^0)$ because $a_0^*(\mathbf{s}^1) > p_0^*(\mathbf{s}^0)$. □

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

Claim 2. Suppose the planner knows that the state is \mathbf{s}^1 (resp. \mathbf{s}^2) and she can only impose a cap c_1 (resp. c_2). Then, it is optimal to set $c_1 = p_1^*(\mathbf{s}^1)$ (resp. $c_2 = p_2^*(\mathbf{s}^2)$).

Proof. Suppose the planner knows \mathbf{s}^1 —the other case is equivalent. Replicating the argument in the proof of Proposition 5, we can conclude that it is optimal to set $c_1 < a_1^*(\mathbf{s}^1)$. To find the optimal $c_1 \in (0, a_1^*(\mathbf{s}^1))$, consider first the doer's problem to maximize $\theta[\bar{r} \ln(x_1) + \underline{r} \ln(x_2)] + b \ln(x_0)$ subject to $x_0 + x_1 + x_2 \leq 1$ and $x_1 \leq c_1$. Since both constraints must bind, this problem becomes

$$\max_{x_0 \in [0, 1]} \{\theta \underline{r} \ln(1 - c_1 - x_0) + b \ln(x_0)\}.$$

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

$$x_0(c_1) = \frac{b}{\underline{\theta}\underline{r} + b}(1 - c_1) \quad \text{and} \quad x_2(c_1) = \frac{\underline{\theta}\underline{r}}{\underline{\theta}\underline{r} + b}(1 - c_1).$$

Given this, we can compute the planner's payoff in state \mathbf{s}^1 as a function of c_1 , which equals (up to a constant)

$$\underline{\theta}[\bar{r} \ln(c_1) + \underline{r} \ln(1 - c_1)] + \ln(1 - c_1). \quad (18)$$

The optimal c_1 is again characterized by the first-order condition, which leads to

$$c_1 = \frac{\underline{\theta}\bar{r}}{1 + \underline{\theta}(\underline{r} + \bar{r})}. \quad (19)$$

To complete the proof, we need to find $p_1^*(\mathbf{s}^1)$, which results from maximizing $\underline{\theta}[\bar{r} \ln(x_1) + \underline{r} \ln(x_2)] + \ln(x_0)$ subject to $x_0 + x_1 + x_2 \leq 1$. Substituting $x_0 = 1 - x_1 - x_2$, taking first-order conditions, and combining them, we get

$$p_1^*(\mathbf{s}^1) = \frac{\underline{\theta}\bar{r}}{1 + \underline{\theta}(\underline{r} + \bar{r})}.$$

□

Claim 3. Suppose the planner knows that the state is \mathbf{s}^1 (resp. \mathbf{s}^2). Then she strictly prefers to impose only c_1 (resp. c_2) than only c_2 (resp. c_1).

Proof. Suppose the planner knows \mathbf{s}^1 —the other case is equivalent. Mimicking the calculations in the proof of Claim 2, one can show that if the planner can impose only c_2 , then she sets

$$c_2 = \frac{\underline{\theta}\underline{r}}{1 + \underline{\theta}(\underline{r} + \bar{r})}. \quad (20)$$

We want to argue that her payoff in state \mathbf{s}^1 is strictly larger if she imposes only c_1 as in (19) than if she imposes only c_2 as in (20). Substituting the doer's allocations implied by c_1 and c_2 into the planner's utility function and simplifying the resulting expressions, one can show that c_1 in (19) is strictly better than c_2 in (20) if and only if

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

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

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

Therefore, $\varphi(b, \bar{r}) - \varphi(b, \underline{r})$ is strictly decreasing in b . Continuity gives the result.

□

Claim 4. If D is optimal, then f_0 can bind at most in \mathbf{s}^0 .

Proof. If f_0 binds in all states, then D is weakly dominated by a policy that involves only f_0 and no caps, as the caps distort consumption without improving savings. Given $g > g^*$, by Claim 4 the latter policy is strictly dominated by one imposing only the floor $p_0^*(\mathbf{s}^0)$. Clearly, if f_0 binds in \mathbf{s}^1 and \mathbf{s}^2 , then it must also bind in \mathbf{s}^0 .

Now suppose that f_0 binds only in \mathbf{s}^0 and another state, say, \mathbf{s}^1 —the same argument applies for \mathbf{s}^2 . There are two cases to consider:

Case 1: c_1 does not bind in state \mathbf{s}^2 . Then, removing c_1 leads to a weakly superior policy in which f_0 binds only in states \mathbf{s}^0 and \mathbf{s}^1 . Given $g > g^*$, however, the gain from raising f_0 above $p_0^*(\mathbf{s}^0)$ to improve the doer's allocation only in \mathbf{s}^1 does not justify the loss created in \mathbf{s}^0 . Therefore, D is again strictly dominated by the policy obtained if we remove c_1 and set the floor at $p_0^*(\mathbf{s}^0)$.

Case 2: c_1 binds also in state \mathbf{s}^2 . This implies that f_0 has to bind in all states. Indeed, since c_1 binds in both \mathbf{s}^1 and \mathbf{s}^2 , the doer chooses $x_1 = c_1$ in both states; moreover, since in \mathbf{s}^2 good 2 is more valuable than in \mathbf{s}^1 , he wants to allocate more income to good 2 than to savings relative to \mathbf{s}^1 and therefore f_0 also binds in \mathbf{s}^2 . However, we have already argued that such a policy is strictly dominated by one that imposes only the floor $p_0^*(\mathbf{s}^0)$. □

Claim 5. If D involves binding caps, then c_i can bind at most in \mathbf{s}^i for $i = 1, 2$.

Proof. Without loss, consider c_1 . Suppose first that c_1 binds in all states, which implies that $a_1(\mathbf{s}^i) = c_1$ for all $i = 0, 1, 2$. There are five cases to consider:

Case 1: Neither c_2 nor f_0 bind in any state. Since $\bar{\theta} > \underline{\theta}$, we have $a_2(\mathbf{s}^0) > a_2(\mathbf{s}^2)$. The policy cannot be optimal because, given c_1 , the planner would be strictly better off by adding a floor that binds only in \mathbf{s}^0 : Even if c_1 were binding for her in \mathbf{s}^0 , she would strictly prefer a level $x_2 < a_2(\mathbf{s}^0)$ of good 2.

Case 2: c_2 binds in all states. Then, $a_2(\mathbf{s}^i) = c_2$ and $a_0(\mathbf{s}^i) = 1 - c_1 - c_2$ for all $i = 0, 1, 2$. This policy is strictly dominated by a one that imposes only a floor equal to $1 - c_1 - c_2$ —because caps are distorting—which is in turn strictly dominated by the policy with only the floor $p_0^*(\mathbf{s}^0)$ given $g > g^*$.

Case 3: c_2 binds in no state. Then, as in case 1, for D to be optimal f_0 must bind at least in \mathbf{s}^0 and only in that state by Claim 4. Since by assumption c_1 binds in all states, it must be that $c_1 < p_1^*(\mathbf{s}^1)$. Indeed, if $c_1 \geq p_1^*(\mathbf{s}^1)$, the optimal f_0 equals $p_0^*(\mathbf{s}^0)$; since by assumption $p_1(\mathbf{s}^0) < p_1^*(\mathbf{s}^1)$, c_1 cannot bind in \mathbf{s}^0 . It follows that, with regard to \mathbf{s}^0 and \mathbf{s}^1 , the planner would be strictly better off replacing c_1 and f_0 with $\hat{c}_1 = p_1^*(\mathbf{s}^1)$ and $\hat{f}_0 = p_0^*(\mathbf{s}^0)$. With regard to \mathbf{s}^2 , the planner would be better off by replacing c_2 with $\hat{c}_2 = p_2^*(\mathbf{s}^2)$: By Claim 3, even if c_1 were perfectly tailored for \mathbf{s}^2 , it would be strictly dominated in that state by \hat{c}_2 .

Case 4: c_2 binds only in \mathbf{s}^0 . Since $a_2(\mathbf{s}^0) > a_2(\mathbf{s}^2)$ if the policy used only c_1 , it follows that the planner can obtain in all states the same allocations induced by D if she imposes a floor that binds only in \mathbf{s}^0 . Such a policy, however, is again strictly dominated for the same reasons as in case 3.

Case 5: c_2 binds in \mathbf{s}^0 and in \mathbf{s}^2 . Since $a_2(\mathbf{s}^0) > a_2(\mathbf{s}^2)$ if the policy used only c_1 , the planner could again obtain the same allocation in all states with a floor that binds only in \mathbf{s}^0 and \mathbf{s}^2 . By Claim 4, however, such a policy cannot be optimal.

Now suppose that c_1 binds in only two states. If c_1 binds only in \mathbf{s}^1 and in \mathbf{s}^0 , then by the same argument as in case 3 above the planner is strictly better off by replacing c_1 and f_0 with $\hat{c}_1 = p_1^*(\mathbf{s}^1)$ and $\hat{f}_0 = p_0^*(\mathbf{s}^0)$ as well as c_2 with $\hat{c}_2 = p_2^*(\mathbf{s}^2)$. If c_1 binds in \mathbf{s}^1 and \mathbf{s}^2 , then it

must also bind in \mathbf{s}^0 —which is the case we considered before. Indeed, if c_1 binds in \mathbf{s}^2 , then it will also bind at the fictitious state $(\underline{\theta}, \bar{r}, \bar{r})$ and hence in \mathbf{s}^0 where both consumption goods are more valuable. The case left is if c_1 binds only in \mathbf{s}^0 and \mathbf{s}^2 , but this is impossible: It would have to bind also in \mathbf{s}^1 , since in that state good 1 is more valuable than in \mathbf{s}^2 .

Finally, suppose that c_1 binds in only one state. We have just argued that if c_1 binds in \mathbf{s}^2 , then it must also bind in \mathbf{s}^1 . Thus, we only have to rule out the case in which c_1 binds only in \mathbf{s}^0 . This property is possible only if in \mathbf{s}^0 the cap c_2 also binds, inducing the doer to overconsume in good 1. However, such a c_2 must also bind in \mathbf{s}^2 ; hence, it cannot be part of an optimal D , because we just showed that a cap cannot bind in more than one state. \square

Combining Claims 1-5, we conclude that the optimal policy $D \in \mathcal{R}$ satisfies $f_0 = p_0^*(\mathbf{s}^0)$, $c_1 = p_1^*(\mathbf{s}^1)$, and $c_2 = p_2^*(\mathbf{s}^2)$.

A.13 Proof of Proposition 4

Start from the value of $\bar{\theta}$ which implies that $p_1^*(\mathbf{s}^1) > p_1^*(\mathbf{s}^0)$ and $p_2^*(\mathbf{s}^2) > p_2^*(\mathbf{s}^0)$ and hence leads to the optimal policy in Lemma 3. If we increase $\bar{\theta}$, both $p_1^*(\mathbf{s}^0)$ and $p_2^*(\mathbf{s}^0)$ increase continuously while always satisfying $p_1^*(\mathbf{s}^0) = p_2^*(\mathbf{s}^0)$. Therefore, there exists a unique $\bar{\theta}^\dagger$ such that, when $\bar{\theta} = \bar{\theta}^\dagger$, we have $p_1^*(\mathbf{s}^1) = p_1^*(\mathbf{s}^0)$ and $p_2^*(\mathbf{s}^2) = p_2^*(\mathbf{s}^0)$. For every $\bar{\theta} \leq \bar{\theta}^\dagger$, the optimal $D \in \mathcal{R}$ remains $c_1 = p_1^*(\mathbf{s}^1)$, $c_2 = p_2^*(\mathbf{s}^2)$, and $f_0 = p_0^*(\mathbf{s}^0)$, where the latter of course falls continuously as $\bar{\theta}$ rises towards $\bar{\theta}^\dagger$.

By the same logic of Proposition 6, we can focus on the class $\bar{\mathcal{R}} \subset \mathcal{R}$ that contains all policies that can use only f_0 , c_1 , and c_2 . Let $\mathcal{D}(\bar{\theta}) \subset \bar{\mathcal{R}}$ be the nonempty set of optimal policies as a function of $\bar{\theta}$. By Lemma 3 and the previous argument, $\mathcal{D}(\bar{\theta})$ is singleton for $\bar{\theta} \leq \bar{\theta}^\dagger$. Define the distance between any two policies D and D' as the Euclidean distance between the vector (f_0, c_1, c_2) describing D and the vector (f'_0, c'_1, c'_2) describing D' . By the Maximum Theorem, $\mathcal{D}(\bar{\theta})$ is upper hemicontinuous in $\bar{\theta}$.³⁷ Hence, by choosing $\bar{\theta} > \bar{\theta}^\dagger$ sufficiently close to $\bar{\theta}^\dagger$, we can render the distance between $D(\bar{\theta}^\dagger)$ and every $D \in \mathcal{D}(\bar{\theta})$ arbitrarily small. Thus, there exists $\varepsilon > 0$ such that, if $\bar{\theta} \in (\bar{\theta}^\dagger, \bar{\theta}^\dagger + \varepsilon)$, then for every $D \in \mathcal{D}(\bar{\theta})$ the following holds: (1) $c_i(\bar{\theta}) > a_i^*(\mathbf{s}^i)$ for $i = 1, 2$; and (2) $f_0(\bar{\theta})$ can bind neither in \mathbf{s}^1 nor in \mathbf{s}^2 . To see property (2), note that $\mathcal{D}(\bar{\theta}^\dagger)$ contains the policy defined by $c_i(\bar{\theta}^\dagger) = p_i^*(\mathbf{s}^i)$ for $i = 1, 2$ and $f_0(\bar{\theta}^\dagger) = p_0^*(\mathbf{s}^0)$, where $f_0(\bar{\theta}^\dagger) = 1 - c_1(\bar{\theta}^\dagger) - c_2(\bar{\theta}^\dagger)$ and hence f_0 is actually redundant. Thus, $\mathcal{D}(\bar{\theta})$ contains no policy with $f_0(\bar{\theta}) > 1 - c_1(\bar{\theta}) - c_2(\bar{\theta})$, because such policies are strictly dominated for the same argument that rules them out in the proof of Lemma 3. Since the largest value of $f_0(\bar{\theta})$ must be close to $f_0(\bar{\theta}^\dagger)$ for $\bar{\theta} \in (\bar{\theta}^\dagger, \bar{\theta}^\dagger + \varepsilon)$, it follows that $f_0(\bar{\theta})$ cannot bind in \mathbf{s}^1 and \mathbf{s}^2 as well.

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

³⁷Although 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 $D \in \mathcal{R}$ that forces 0 allocation to some dimension. Such a policy is always dominated by the optimal singleton D which gives the doer no discretion. Formally, there exists $\varepsilon > 0$ such that, if we required $f_i \leq 1 - \varepsilon$ and $c_i \geq \varepsilon$ for all $i = 0, 1, 2$, we would never affect the planner's problem.

Claim 6. For every $D \in \mathcal{D}(\bar{\theta})$, both $c_1(\bar{\theta})$ and $c_2(\bar{\theta})$ must bind in \mathbf{s}^0 —that is, $c_i(\bar{\theta}) = a_i(\mathbf{s}^0)$ for $i = 1, 2$. Given this, $a_0(\mathbf{s}^0) = 1 - c_1(\bar{\theta}) - c_2(\bar{\theta})$, and hence f_0 can be removed.

Proof. Note that the planner's objective in state \mathbf{s}^i as a function of c_i is strictly concave and decreasing for $c_i > p_i^*(\mathbf{s}^i)$ (see equation (18)). Thus, if for example $c_1(\bar{\theta})$ is not binding for the doer in state \mathbf{s}^0 —that is, $c_1(\bar{\theta}) > a_1(\mathbf{s}^0)$ —the planner can lower c_1 without affecting the doer's choice in \mathbf{s}^0 and \mathbf{s}^2 and strictly improve her payoff in \mathbf{s}^1 . Hence, the policy would not be optimal. □

Claim 7. $c_1(\bar{\theta}) = c_2(\bar{\theta})$ for every $D \in \mathcal{D}(\bar{\theta})$.

Proof. Without loss, suppose that $c_1(\bar{\theta}) > c_2(\bar{\theta})$. Note that $c_2(\bar{\theta}) < a_2^*(\mathbf{s}^0)$ because, otherwise, we would have $c_1(\bar{\theta}) > a_1^*(\mathbf{s}^0) = a_2^*(\mathbf{s}^0)$, which contradicts the previous point. Consider the alternative policy with $c_1^\delta = c_1(\bar{\theta}) - \delta$ and $c_2^\delta = c_2(\bar{\theta}) + \delta$, where $\delta > 0$. For δ sufficiently small, both c_1^δ and c_2^δ continue to be binding in \mathbf{s}^0 , and hence $1 - c_1^\delta - c_2^\delta = a_0(\mathbf{s}^0)$. In \mathbf{s}^0 , the planner's payoff is higher, because given $a_0(\mathbf{s}^0)$ 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 c_i in \mathbf{s}^i for $i = 1, 2$ (see (18)), we have that the decrease in her payoff in \mathbf{s}^2 resulting from the slacker c_2 is more than compensated by the increase in her payoff in \mathbf{s}^1 resulting from the tighter c_1 . Hence, overall the planner's payoff is strictly larger with (c_1^δ, c_2^δ) than with $(c_1(\bar{\theta}), c_2(\bar{\theta}))$, contradicts the optimality of the latter policy. □

Claim 8. $1 - c_1(\bar{\theta}) - c_2(\bar{\theta}) > p_0^*(\mathbf{s}^0)$ for every $D \in \mathcal{D}(\bar{\theta})$.

Proof. If $1 - c_1(\bar{\theta}) - c_2(\bar{\theta}) < p_0^*(\mathbf{s}^0)$, then the planner can set $f_0 = p_0^*(\mathbf{s}^0)$ and achieve a strictly higher payoff in \mathbf{s}^0 without affecting the doer's choices in \mathbf{s}^1 and \mathbf{s}^2 . If $1 - c_1(\bar{\theta}) - c_2(\bar{\theta}) = p_0^*(\mathbf{s}^0)$, then $c_i = p_i^*(\mathbf{s}^0)$ for $i = 1, 2$, which means that $\mathbf{a}(\mathbf{s}^0) = \mathbf{p}^*(\mathbf{s}^0)$. Therefore, it would be possible to lower both $c_1(\bar{\theta})$ and $c_2(\bar{\theta})$ by the same small amount δ , to induce a first-order gain in the planner's payoff for both \mathbf{s}^1 and \mathbf{s}^2 because $c_i(\bar{\theta}) > p_i^*(\mathbf{s}^i)$ for $i = 1, 2$, and to cause only a second-order loss in \mathbf{s}^0 . □

Claim 9. Every $D \in \mathcal{D}(\bar{\theta})$ is unique as far as c_1 and c_2 are concerned and satisfies the properties stated in Lemma 4.

Proof. The planner's payoff in \mathbf{s}^1 and \mathbf{s}^2 is given by (18) up to a constant:

$$\underline{\theta}[\bar{r} \ln(c) + \underline{r} \ln(1 - c)] + \ln(1 - c).$$

Her payoff in \mathbf{s}^0 is given, up to a constant, by

$$2\bar{\theta} \bar{r} \ln(c) + \ln(1 - 2c).$$

Therefore, the optimal c maximizes

$$(1 - g) \{ \underline{\theta}[\bar{r} \ln(c) + \underline{r} \ln(1 - c)] + \ln(1 - c) \} + g \{ 2\bar{\theta} \bar{r} \ln(c) + \ln(1 - 2c) \}.$$

Since this function is strictly concave, there is a unique optimal c . To see that $p_i^*(\mathbf{s}^i) > c_i > p_i^*(\mathbf{s}^0)$ for every $i = 1, 2$, consider the following observations. Note that $c_i > p_i^*(\mathbf{s}^i)$ would be strictly dominated by $c_i = p_i^*(\mathbf{s}^i)$ for every i , because this is the optimal level of the cap in the corresponding state. Consequently, we must have $c_i < p_i^*(\mathbf{s}^i)$ because by assumption $1 - p_1^*(\mathbf{s}^1) - p_2^*(\mathbf{s}^2) > p_0^*(\mathbf{s}^0)$ for $\bar{\theta} > \bar{\theta}^\dagger$, and hence reducing c_i below $p_i^*(\mathbf{s}^i)$ by the same small amount for all $i = 1, 2$ causes a first-order gain in \mathbf{s}^0 and only a second-order loss in \mathbf{s}^1 and \mathbf{s}^2 . \square

A.14 Proof of Corollary 1

By Proposition 6, we can focus on the class $\bar{\mathcal{R}} \subset \mathcal{R}$ that contains all policies which can use only f_0 , c_1 , and c_2 . Let $\mathcal{R}_{f_0} \subset \bar{\mathcal{R}}$ contain all policies that can use only f_0 , $\mathcal{R}_{\mathbf{c}} \subset \bar{\mathcal{R}}$ contain all policies that can use both c_1 and c_2 , and $\mathcal{R}_{c_i} \subset \bar{\mathcal{R}}$ contain all policies that can use only c_i for $i = 1, 2$. To indicate that the planner's expected payoff from D is computed using some distribution \hat{G} , we will use the notation $\mathcal{U}(D; \hat{G})$.

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

$$\mathcal{U}(D; G_\alpha^{\text{fc}}) = \alpha \mathcal{U}(D; G^{\text{fc}}) + (1 - \alpha) \mathcal{U}(D; \bar{G}).$$

Now define

$$\mathcal{W}_{f_0}^{\text{fc}}(\alpha) = \max_{D \in \mathcal{R}_{f_0}} \mathcal{U}(D; G_\alpha^{\text{fc}}) \quad \text{and} \quad \mathcal{W}_{\mathbf{c}}^{\text{fc}}(\alpha) = \max_{D \in \mathcal{R}_{\mathbf{c}}} \mathcal{U}(D; G_\alpha^{\text{fc}}), \quad \alpha \in [0, 1].$$

Both $\mathcal{W}_{f_0}^{\text{fc}}$ and $\mathcal{W}_{\mathbf{c}}^{\text{fc}}$ are well defined by the same argument as in the proof of Lemma 2; moreover, by the Maximum Theorem, they are continuous functions of α .³⁸ Let D^{fc} denote the optimal policy in Lemma 3. Note that $\mathcal{U}(D^{\text{fc}}; \bar{G})$ is finite since the doer's resulting allocations are bounded away from 0 in all dimensions. We have that $\lim_{\alpha \uparrow 1} \mathcal{U}(D^{\text{fc}}; G_\alpha^{\text{fc}}) - \mathcal{W}_j^{\text{fc}}(\alpha) > 0$ for both $j = f_0$ and $j = \mathbf{c}$. Therefore, there exists $\hat{\alpha} \in (0, 1)$ such that D^{fc} strictly dominates every $D \in \mathcal{R}_{f_0} \cup \mathcal{R}_{\mathbf{c}}$ given the distribution $G_{\hat{\alpha}}^{\text{fc}}$.

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

$$\mathcal{U}(D; G_\alpha^{\text{c}}) = \alpha \mathcal{U}(D; G^{\text{c}}) + (1 - \alpha) \mathcal{U}(D; \bar{G}').$$

Let D^{c} denote the optimal policy in Lemma 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 D^{c} strictly dominates every $D \in \mathcal{R}_{f_0} \cup \mathcal{R}_{\mathbf{c}} \cup \mathcal{R}_{c_1} \cup \mathcal{R}_{c_2}$ given the distribution G_α^{c} . It remains to show that there exists $\alpha' \in (\alpha'', 1)$ such that D^{c} strictly dominates every $D \in \bar{\mathcal{R}}$ given $G_{\alpha'}^{\text{c}}$.

To this end, define

$$\mathcal{D}(\alpha) = \arg \max_{D \in \bar{\mathcal{R}}} \mathcal{U}(D; G_\alpha^{\text{c}}).$$

Another application of the Maximum Theorem implies that $\mathcal{D}(\cdot)$ is upper hemicontinuous. Note that $\mathcal{D}(1)$ is characterized by vectors (f_0^*, c_1^*, c_2^*) such that c_1^* and c_2^* are unique and satisfy the properties in Lemma 4, and $f_0^* \in [0, \bar{f}_0]$ where $\bar{f}_0 = 1 - c_1^* - c_2^*$. Therefore, for every $\delta > 0$, there exists $\varepsilon > 0$ such that, if $\alpha \in (1 - \varepsilon, 1]$, then $f_0 \in [0, \bar{f}_0 + \delta]$, $c_1 \in (c_1^* - \delta, c_1^* + \delta)$, and $c_2 \in (c_2^* - \delta, c_2^* + \delta)$ for every (f_0, c_1, c_2) corresponding to some $D \in \mathcal{D}(\alpha)$. This means that, by

³⁸Recall Footnote 37.

choosing δ sufficiently small, we can ensure that for every $D \in \mathcal{D}(\alpha)$ the following holds: (1) $1 - c_1 - c_2 > p_0^*(\mathbf{s}^0)$; (2) removing f_0 leads to a policy such that both c_1 and c_2 bind in \mathbf{s}^0 ; and (3) f_0 cannot bind in \mathbf{s}^1 and \mathbf{s}^2 , since \bar{f}_0 is strictly smaller than the doer's choice of x_0 in those states under every policy in $\mathcal{D}(1)$.

Take any $D \in \mathcal{D}(\alpha)$ and fix its c_1 and c_2 . The f_0 that completes D must be optimally chosen given c_1 and c_2 . We claim that such an f_0 must satisfy $f_0 \leq 1 - c_1 - c_2 = \bar{k}$ for α sufficiently close to 1. Suppose this is not true and consider the gain in the planner's expected payoff from imposing $f_0 > \bar{k}$. Her gain in \mathbf{s}^0 would be

$$(1 - b)[v(f_0) - v(\bar{k})] + \bar{V}(f_0; \mathbf{s}^0) - \bar{V}(\bar{k}; \mathbf{s}^0), \quad (21)$$

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

$$\int_{S(f_0)} \{(1 - b)[v(f_0) - v(\hat{a}_0(\mathbf{s}))] + \bar{V}(f_0; \mathbf{s}) - \bar{V}(\hat{a}_0(\mathbf{s}); \mathbf{s})\} d\bar{G}', \quad (22)$$

where $S(f_0) \subset S' = [\theta, \bar{\theta}'] \times [r, \bar{r}]^2$ is the set of states in which f_0 affects the doer's choices, \hat{a} is the doer's allocation function under the policy that involves only c_1 and c_2 , and

$$\bar{V}(k; \mathbf{s}) = \max_{\{(\mathbf{x}, x_0) \in B: x_1 \leq c_1, x_2 \leq c_2, x_0 \geq k\}} \{\hat{u}(\mathbf{x}; \mathbf{s}) + bv(x_0)\}, \quad k \in [\bar{k}, 1], \mathbf{s} \in S'.$$

Note that $\bar{V}(f_0; \mathbf{s}) \leq \bar{V}(\hat{a}_0(\mathbf{s}); \mathbf{s})$ and $\hat{a}_0(\mathbf{s}) \geq \bar{k}$ for all $\mathbf{s} \in S'$; therefore, for every $f_0 \geq \bar{k}$, the quantity (22) is bounded above by

$$\int_{S(f_0)} (1 - b)[v(f_0) - v(\hat{a}_0(\mathbf{s}))] d\bar{G}' \leq (1 - b)[v(f_0) - v(\bar{k})].$$

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

Now focus on $\bar{V}(k; \mathbf{s}^0)$. For every $f_0 > \bar{k}$, the following holds: (1) f_0 always binds, because $\bar{k} > p_0^*(\mathbf{s}^0)$ and hence the doer wants to save strictly less than f_0 ; (2) only one cap can bind, because if both bind, then $a_0(\mathbf{s}^0) = \bar{k} < f_0$, which is impossible; (3) one cap never binds, because consumption goods are normal, so for every $f_0 > \bar{k}$ the doer's chooses $a_i(\mathbf{s}^0) < c_i$ for at least one $i = 1, 2$. Without loss, suppose that the cap that never binds is c_2 . Therefore, if we remove c_2 , $\bar{V}(k; \mathbf{s}^0)$ coincides with the doer's indirect utility under the policies defined by $k \in [\bar{k}, 1]$ and c_1 only, which we denote by $\bar{V}(k; \mathbf{s}^0, c_1)$. By the same argument as in the proof of Lemma 10, $\bar{V}(k; \mathbf{s}^0, c_1)$ is continuously differentiable in k for $k \in [0, 1]$ and $\bar{V}'(k; \mathbf{s}^0, c_1) = \lambda(\mathbf{s}^0; k)$, where $\lambda(\mathbf{s}^0; k)$ is the Lagrange multiplier associated to the constraint $x_0 \geq k$. Using the Lagrangian defining $\bar{V}(k; \mathbf{s}^0, c_1)$, we have that

$$\lambda(\mathbf{s}^0; k) = \bar{\theta}' u_1^2(a_2(\mathbf{s}^0; k); \bar{\mathbf{r}}) - bv'(k).$$

Note that $\lambda(\mathbf{s}^0; k) > 0$ for all $k \in [\bar{k}, 1]$, because such levels of the floor must always bind for the doer. Moreover, $\lambda(\mathbf{s}^0; k)$ is strictly increasing in $k \in [\bar{k}, 1]$ because $v'' < 0$, $u_{11}^i < 0$, and $a_2(\mathbf{s}^0; k)$ is non-increasing in k by normality of goods. We conclude that $\bar{V}'(k; \mathbf{s}^0) = -\lambda(\mathbf{s}^0; k)$ for every $k \in (\bar{k}, 1]$ and $\bar{V}'(\bar{k}+; \mathbf{s}^0) = \lambda(\mathbf{s}^0; \bar{k})$, where the plus denotes the right derivative.³⁹ Moreover, $\bar{V}'(k; \mathbf{s}^0)$ is strictly decreasing in k .

³⁹In fact, $\bar{V}(k; \mathbf{s}^0)$ is not differentiable at $k = \bar{k}$ since $\bar{V}(k; \mathbf{s}^0)$ is constant for $k < \bar{k}$ and hence $\bar{V}'(\bar{k}-; \mathbf{s}^0) = 0$.

Observe that

$$(1 - b)v'(\bar{k}) + \bar{V}'(\bar{k}; \mathbf{s}^0) = v'(\bar{k}) - \bar{\theta}' u_1^2(a_2(\mathbf{s}^0; \bar{k}); \bar{\mathbf{r}}), \quad (23)$$

which is strictly negative. This is because $c_1 < p_1^*(\mathbf{s}^0)$ and $c_2 < p_2^*(\mathbf{s}^0)$ by Lemma 4 since α is close to 1, which implies that both caps must bind for the planner; consequently, $f_0 = \bar{k}$ and c_1 must also bind for the planner. The right-hand side of (23) coincides with the negative of the Lagrange multiplier associated with the constraint $x_0 \geq \bar{k}$ in the planner's problem that also includes the constraint $x_1 \leq c_1$.

Recall that \bar{k} depends on α —hence denote it by \bar{k}_α —and consider the quantity

$$g\bar{V}'(\bar{k}_\alpha; \mathbf{s}^0) + [\alpha g + (1 - \alpha)](1 - b)v'(\bar{k}_\alpha). \quad (24)$$

This quantity is strictly negative for $\alpha = 1$, which corresponds to $\bar{k}_1 = 1 - c_1^* - c_2^*$. By continuity of (24) as a function of (α, k) and upper hemicontinuity of $\mathcal{D}(\alpha)$, there exists $\varepsilon > 0$ such that (24) remains strictly negative for all $\alpha \in (1 - \varepsilon, 1]$. Given the monotonicity properties of v' and $\bar{V}'(\cdot; \mathbf{s}^0)$, (24) is strictly decreasing for all $k \geq \bar{k}_\alpha$.

Finally, for every $\alpha \in (1 - \varepsilon, 1]$ and $f_0 > \bar{k}_\alpha$, we have that

$$\begin{aligned} & [\alpha g + (1 - \alpha)](1 - b)[v(f_0) - v(\bar{k}_\alpha)] + g[\bar{V}(f_0; \mathbf{s}^0) - \bar{V}(\bar{k}_\alpha; \mathbf{s}^0)] \\ &= \int_{\bar{k}_\alpha}^{f_0} \left\{ [\alpha g + (1 - \alpha)](1 - b)v'(k) + g\bar{V}'(k; \mathbf{s}^0) \right\} dk \\ &< \left\{ [\alpha g + (1 - \alpha)](1 - b)v'(\bar{k}_\alpha) + g\bar{V}'(\bar{k}_\alpha; \mathbf{s}^0) \right\} (f_0 - \bar{k}_\alpha) < 0. \end{aligned}$$

We conclude that the planner is strictly worse off by imposing a binding savings floor in addition to the caps c_1 and c_2 , and hence every optimal policy must involve binding caps for both goods, but no binding floor on savings.

A.15 Proof of Lemma 9

Recall the definition of $\mathcal{U}(D)$ and $\mathbf{a}(\theta|D)$ in (3) and (4). There exists $D \subset B$ such that $\mathcal{U}(D) \geq \mathcal{U}(D')$ for all $D' \subset B$ if and only if there exist functions $\chi : [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+^n$ and $t : [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+$ that satisfy two conditions:

(1) for all $\theta, \theta' \in [\underline{\theta}, \bar{\theta}]$

$$\theta \hat{u}(\chi(\theta)) + bv(t(\theta)) \geq \theta \hat{u}(\chi(\theta')) + bv(t(\theta'))$$

and

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

(2) the pair (χ, t) maximizes

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

On the other hand, there exists $D^{\text{ac}} \subset B^{\text{as}}$ such that $\mathcal{U}(D^{\text{ac}}) \geq \mathcal{U}(\hat{D}^{\text{ac}})$ for all $\hat{D}^{\text{ac}} \subset B^{\text{as}}$ if and only if there exist functions $\varphi : [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+$ and $\tau : [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+$ that satisfy two conditions:

(1') for all $\theta, \theta' \in [\underline{\theta}, \bar{\theta}]$

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

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

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

(2') the pair (φ, τ) maximizes

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

Suppose (χ, t) that satisfies condition (1) and (2). Then, by our discussion on money burning before the statement of Lemma 9, there exists a function $\varphi : [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+$ such that $u^*(\varphi(\theta)) = \hat{u}(\chi(\theta))$ and $\varphi(\theta) \leq \sum_{i=1}^n \chi_i(\theta)$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$. Hence, letting $\tau \equiv t$, we have that (φ, τ) satisfy both (1') and (2').

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

$$\chi(\theta) = \arg \max_{\{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i \leq \varphi(\theta)\}} \hat{u}(\mathbf{x}).$$

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

A.16 Proof of Proposition 6

Let $D^{\text{ac}} \subset B^{\text{as}}$ satisfy the premise of Proposition 6. Then, as noted in the proof of Lemma 9, we can describe the doer's allocation from D^{ac} with the functions (φ, τ) that satisfy condition (1') and such that $0 < \varphi(\theta) < 1 - \tau(\theta)$ for all $\theta \in \Theta$ and

$$\mathcal{U}(D^{\text{ac}}) = \int_{\underline{\theta}}^{\bar{\theta}} [\theta u^*(\varphi(\theta)) + v(\tau(\theta))] g(\theta) d\theta.$$

Now, since \hat{u} is continuous and $E_y = \{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i = y\}$ is connected, $\hat{u}(E_y) = [u_*(y), u^*(y)]$. Since \hat{u} is strictly concave, $u_*(y) < u^*(y)$ for all $y > 0$. Since \hat{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 $\mathbf{x}(\theta) \in E_{y(\theta)}$ such that $\hat{u}(\mathbf{x}(\theta)) = u^*(\varphi(\theta))$. So, for every $\theta \in [\underline{\theta}, \bar{\theta}]$, define $t(\theta) = \tau(\theta)$ and

$$\chi(\theta) = \begin{cases} \mathbf{x}(\theta) & \text{if } \theta \in \Theta \\ \arg \max_{\{\mathbf{x} \in \mathbb{R}_+^n : \sum_{i=1}^n x_i \leq \varphi(\theta)\}} \hat{u}(\mathbf{x}) & \text{if } \theta \notin \Theta \end{cases}.$$

Then, by construction the pair (χ, t) satisfy conditions (1) and (2) in the proof of Lemma 9. Now, let $D' = \{(\mathbf{x}, x_0) \in \mathbb{R}_+^n : (\mathbf{x}, x_0) = (\chi(\theta), t(\theta)), \text{ for some } \theta \in [\underline{\theta}, \bar{\theta}]\}$. We have $D' \subset B$, $\mathcal{U}(D') = \mathcal{U}(D^{\text{ac}})$, and the doer's allocation satisfies $a'_{-0}(\theta) = \chi(\theta)$ and $a'_0(\theta) = \tau(\theta)$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$. By construction, \mathbf{a}' satisfies the stated relationship with \mathbf{a} .

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

References

- Aguiar, M. and M. Amador (2011). Growth in the shadow of expropriation. *Quarterly Journal of Economics* 126(2).
- Ali, S. N. (2011). Learning self-control. *The Quarterly Journal of Economics* 126(2), 857–893.
- Alonso, R., I. Brocas, and J. D. Carrillo (2013). Resource allocation in the brain. *The Review of Economic Studies*, rdt043.
- 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*, 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.
- Battaglini, M. (2005). Long-term contracting with markovian consumers. *American Economic Review*, 637–658.
- 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. *American Economic Review*, 1558–1590.
- Bernheim, B. D., D. Ray, and S. evin Yeltekin (2015). Poverty and self-control. *Econometrica* 83(5), 1877–1911.
- 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*, 148–167.
- Courty, P. and H. Li (2000). Sequential screening. *The Review of Economic Studies* 67(4), 697–717.
- Frankel, A. (2014). Aligned delegation. *The American Economic Review* 104(1), 66–83.
- Frankel, A. P. (2015). Delegating multiple decisions. *A EJ: Microeconomics forthcoming*.
- Fudenberg, D. and D. K. Levine (2006a). A dual-self model of impulse control. *The American Economic Review*, 1449–1476.
- Fudenberg, D. and D. K. Levine (2006b). A dual-self model of impulse control. *The American Economic Review*, 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*. *The Economic Journal* 121(557), 1252–1280.
- Halac, M. and P. Yared (2014). Fiscal rules and discretion under persistent shocks. *Econometrica* 82(5), 1557–1614.
- Heath, C. and J. B. Soll (1996). Mental budgeting and consumer decisions. *Journal of Consumer Research*, 40–52.
- Holmström, B. (1984). On the theory of delegation, 1 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.
- Jackson, M. O. and L. Yariv (2015). Collective dynamic choice: The necessity of time inconsistency. *American Economic Journal: Microeconomics* 7(4), 150–78.
- 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.
- Luenberger, D. G. (1969). *Optimization by vector space methods*. John Wiley & Sons.
- Manelli, A. M. and D. R. Vincent (2007). Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory* 137(1), 153–185.
- Milgrom, P. and C. Shannon (1994). Monotone comparative statics. *Econometrica*, 157–180.
- Niskanen, W. A. (1975). Bureaucrats and politicians. *Journal of law and economics*, 617–643.
- O'Donoghue, T. and M. Rabin (2001). Choice and procrastination. *Quarterly Journal of Economics*, 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. *The Quarterly Journal of Economics*, 327–361.
- Prelec, D. and G. Loewenstein (1998). The red and the black: Mental accounting of savings and debt. *Marketing Science* 17(1), 4–28.
- Quah, J. K.-H. (2007). The comparative statics of constrained optimization problems. *Econometrica* 75(2), 401–431.
- Rochet, J.-C. and P. Choné (1998). Ironing, sweeping, and multidimensional screening. *Econometrica*, 783–826.
- 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. *The Quarterly Journal of Economics*, 563–587.

Stole, L. and J.-C. Rochet (2003). The economics of multidimensional screening. *Advances in Economic Theory*, 7th World Congress, Mathias Dewatripont and Lars Hansen, editors, Cambridge University Press.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies* 23(3), 165–180.

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. *The Journal of Political Economy*, 392–406.