Commitment and Welfare*

Frank N. Caliendo and T. Scott Findley

This Version: February 13, 2015. First Version: May 28, 2014.

Abstract

How should welfare analysis be conducted when individuals have dynamically inconsistent preferences? Behavioral economists typically treat the preferences of the time-zero self as the welfare criterion. Critics argue that this practice is unfair to the later time-dated selves of an individual. In this paper we establish conditions under which the time-zero welfare criterion is consistent with a multiself Pareto criterion. We find that these two criteria are consistent as long as the number of time-dated selves (decision nodes) in the choice problem exceeds a specific threshold. This threshold can be as small as 3 selves, depending on the application. Thus, the time-zero welfare criterion prescribes an allocation that may very well leave all selves better off than the equilibrium of the intra-personal conflict. This result helps to at least partly overcome common concerns about the welfare criterion that is commonly used when individuals have dynamically inconsistent preferences.

Key words: Behavioral Welfare Analysis, Commitment, Time-Zero Preferences, Dynamic Inconsistency, Hyperbolic Discounting.

JEL Classification: D03, D91, C61.

*Utah State University, frank.caliendo@usu.edu and tscott.findley@usu.edu. We are especially grateful for insightful suggestions from three anonymous referees and Erzo Luttmer. We also thank Madhav Chandrasekher, Jim Feigenbaum, Aspen Gorry, Fabian Herweg, Botond Kőszegi, Andrew Samwick, Heiner Schumacher, Charlie Sprenger, and participants at the BABEE Workshop at Stanford University, the QSPS Workshop at Utah State University, and the 9^{th} Nordic Conference on Behavioral and Experimental Economics at Aarhus University.

1 Introduction

Beginning with the work of Strotz (1956), models with dynamically inconsistent preferences have been developed to help explain various economic data that are hard to reconcile with standard models in which preferences are dynamically consistent. However, a long-standing concern is that models with dynamically inconsistent preferences are ill-suited for welfare analysis. The fundamental question in welfare economics—How should scarce resources be allocated?—seems hard to answer because it is unclear how to define a welfare improvement when all of the time-dated selves of a single individual disagree on how resources should be allocated over time.

As a practical way forward, O'Donoghue and Rabin (1999, 2000, 2001, 2003, 2007) and other behavioral economists typically assume that the goal of a policymaker should be to maximize the utility of the time-zero self.[1] This approach is based on the idea of helping individuals reach the goals that they themselves initially had in mind before they fell victim to self-control problems. And any deviation from these goals is considered an error in decision making.

However, critics argue that there isn't any particularly good reason to equate welfare with time-zero preferences. They argue that this practice arbitrarily favors the preferences of the first self at the expense of all the other selves—why should a policymaker want to hold all of the later selves captive to the desires of the first self? Essentially, the concern is that committing individuals to their initial goals unfairly forces later selves to do something that they consider to be suboptimal.[2]

In this paper we provide an assessment of the practice of using time-zero preferences as the welfare criterion in models with dynamically inconsistent preferences. We are motivated by Bryan, Karlan, and Nelson (2010) and others who believe that addressing this controversy should be at the top of the research agenda in behavioral economics.[3] We focus on hyperbolic discounting with sophisticated decision making over a sequence of consumption choices. Our main finding is this: the time-zero welfare criterion is consistent with a multiself Pareto

[1]See Brocas, Carrillo, and Dewatripont (2004) and Bryan, Karlan, and Nelson (2010) for surveys.

[2]For instance, Gul and Pesendorfer (2004, p.263) point out that welfare analysis based on time-zero preferences "has the planner forever guarding the perceived interests of the nonexistent former selves," which they later describe as "odd" (Gul and Pesendorfer (2008, p.38)). Likewise, Rubinstein (2006, p.248) states "One criticism made of behavioral economics is the arbitrariness of the welfare criterion...why should the utility of the first self be the basis for welfare considerations?" And Brocas, Carrillo, and Dewatripont (2004, p.51) state that "there is no normative foundation" for equating welfare with time-zero preferences. See Bernheim and Rangel (2009) for additional discussion.

[3]Bryan, Karlan, and Nelson (2010, p.694) state that "settling on a particular approach and providing empirical support or clear philosophical arguments for [that approach] are the hard questions....[that deserve] more thought and research."

criterion as long as the number of time-dated selves (i.e., decision nodes) exceeds a specific threshold. That is, with enough decision nodes in the choice problem, the consumption sequence or allocation that maximizes time-zero utility (i.e., the "commitment allocation") is viewed by *all* time-dated selves as a strictly better outcome than the allocation that emerges as the equilibrium of the intra-personal conflict.[4] If this condition is met, then commitment cannot hurt the selves who follow after the first self, even though these later selves consider the commitment allocation to be suboptimal.

The intuition for why the existence of Pareto gains (in moving from the equilibrium to the commitment allocation) hinges on the number of selves is as follows. If there is a small number of selves, then a given self may have the power to significantly influence the equilibrium allocation. In this case, the equilibrium allocation may be relatively close to what he wants. However, if there is a large number of selves, then the power to influence the equilibrium allocation is diffuse. And in this case, the equilibrium allocation is relatively dissimilar to the desired allocation of a given self and therefore it becomes possible to find allocations that Pareto dominate the equilibrium. In other words, although a large number of selves means that there are many conflicting points of view on how resources should ideally be allocated, it also means that the equilibrium is far away from the wishes of all of the selves, and this creates space for a Pareto improvement.[5]

To illustrate our claims, we consider two classic dynamic programming problems: eating fruit from a tree (renewable resource problem) and eating a cake (nonrenewable resource problem). These specific examples are just to fix ideas; the general lessons extend to settings in which there is a binary choice to take an action now or later that repeats itself over and over (as in the fruit-eating example) and to settings in which a fixed quantity of a resource is depleted over time (as in the cake-eating example). Taken together, these two examples span a range of economic settings in which dynamically inconsistent preferences are often considered.

In the first example, an individual plants a tree at time zero and the tree bears fruit at

[4] We assume individuals are sophisticated and are therefore fully aware of the choices that future selves will make, rather than oblivious (naive) to the choices of future selves. We do this because naiveté presents a conceptual challenge that we are not prepared to deal with. In addition to the usual complication that different selves of a single individual disagree, naiveté creates the added complication that a given self anticipates a different future consumption allocation than he actually experiences. It then becomes philosophically difficult to define the welfare of a given self, and we do not have anything to add to that particular debate.

[5] The only way to break our threshold result is to assume that each time-dated self only cares about himself. In this case, the equilibrium is Pareto efficient no matter how many selves there are, and any movement away from the equilibrium allocation would hurt at least one of the selves. But in the more general case in which each self potentially values the consumption of selves that come before and after him, appropriately discounted to his present vantage point (as in Strotz (1956) and Caplin and Leahy (2004)), the equilibrium allocation is not necessarily Pareto efficient and the door is open for a Pareto improvement.

each subsequent node. All nodes after time zero are called decision nodes. Eventually the tree dies. The individual faces a repeated choice: pick the fruit now and eat it when it tastes good, or leave the fruit on the tree for one extra node to fully ripen and eat it when it tastes great. The fruit spoils completely if left on the tree for more than one node. We focus on the popular "once-off" quasi-hyperbolic discount function $\{1, \beta, \beta, \beta, ...\}$. The individual applies this sequence to the future, and he applies a potentially different sequence $\{1, \gamma, \gamma, \gamma, ...\}$ to his valuation of the past.

We analytically prove four fundamental points. All of these points are robust to the particular value of γ. First, the commitment allocation will never multiself Pareto dominate the equilibrium allocation if there are only two decision nodes. Second, as long as there are more than two decision nodes, there is a non-empty parameter space over β for which the commitment allocation multiself Pareto dominates the equilibrium allocation. Third, this non-empty parameter space expands as the number of decision nodes increases, and this space ultimately subsumes the entire β-space as the model approaches infinitely many (countable) decision nodes. In other words, if the tree lives forever, then the commitment allocation multiself Pareto dominates the equilibrium allocation for any parameterization. Fourth, and most importantly, the commitment allocation will multiself Pareto dominate the equilibrium allocation if and only if the number of decision nodes exceeds a threshold. This threshold is decreasing in β and increasing in the relative utility of immediate consumption (i.e., the relative utility of eating unripe fruit).

We provide numerical calculations to show the relevance of these analytical results. For instance, suppose fruit tastes twice as good if it fully ripens. The time-zero self wants to be patient at each node and wait for the fruit to ripen. But the individual discounts the future by too much to be patient, say $\beta = 0.4$. In this case, in equilibrium the individual always eats unripe fruit from the tree at every decision node. And yet, if there are just 3 or more decision nodes, then all of the selves would prefer the commitment allocation of fully ripe fruit over the equilibrium.[6] These thresholds hold for any assumptions about the degree of backward discounting of past consumption.

Notice the surprise in these results. If there are only two selves that follow after the time-zero self in the choice problem, then committing both of these later selves to being patient and eating fully ripe fruit will not make both of them happy. But simply adding one more self or decision node to the lineup, for a total of three time-dated selves to follow after the time-zero self, can completely change the welfare conclusion because now all three of these later selves agree that commitment dominates the equilibrium. Almost paradoxically, as the

[6] If we instead assume that fruit tastes only 33% better if left to fully ripen, then the threshold number of decision nodes increases to 6.

number of selves increases and there are more conflicting points of view on how resources should ideally be allocated over time, it becomes easier for the selves to reach a unanimous agreement that the commitment allocation beats the equilibrium allocation.

In the second example, an individual eats an infinitely divisible cake over a specified number of decision nodes. The cake does not spoil, nor does it grow. This is a classic first model in textbooks on dynamic programming because it captures essential trade-offs that are common across many dynamic problems. We consider both the "once-off" quasi-hyperbolic discount function $\{1, \beta, \beta, \beta, ...\}$ as well as the typical quasi-hyperbolic discount function $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, ...\}$. The individual values the past according to either $\{1, \gamma, \gamma, \gamma, ...\}$ or $\{1, \gamma\eta, \gamma\eta^2, \gamma\eta^3, ...\}$. In all cases, the results align with our findings from the fruit tree example: if the number of decision nodes exceeds a specific threshold, then the commitment allocation multiself Pareto dominates the equilibrium allocation. Here commitment implies a gradual rate of consumption, whereas the rate of consumption is more rapid in equilibrium, leaving very little cake for later selves.

We consider a range of assumptions about the degree of forward and backward discounting. For example, in the case of once-off quasi-hyperbolic discounting of both the past and the future, the threshold number of decision nodes ranges from a low of 4 to a high of 9 for the rectangle defined by $\beta \in [0.2, 0.8]$ and $\gamma \in [0, 1]$. Any number of decision nodes above these threshold values leads to the conclusion that commitment represents a multiself Pareto improvement over the equilibrium. In other words, simply expanding the number of opportunities that the individual has to eat a portion of the fixed quantity of cake (i.e., expanding the number of time-dated selves of a single individual) brings the time-zero welfare criterion into alignment with the multiself Pareto criterion.

We view this paper as a step toward a better understanding of when the time-zero welfare criterion that has been popularized by O'Donoghue and Rabin and others can be justified on Pareto grounds. We claim that it may be enough to know the total number of decision nodels (time-dated selves) in a choice problem in order to determine whether this practice has a Pareto foundation. While it may not be possible to *completely* generalize our claim to every conceivable setting, knowing the number of selves is indeed enough in the two settings (renewable and nonrenewable resources) that we consider.

1.1 Related Literature

Our finding that commitment may be Pareto improving is related to a series of papers by David Laibson. Laibson (1996) shows that the saving rate that the time-zero individual would commit himself to follow in an infinite-horizon setting is Pareto superior to the equilibrium saving rate (also see Goldman (1979)). Similarly, Laibson, Repetto, and Tobacman (1998) consider the welfare gains of commitment to the plan that a dynamically-consistent (exponential) individual follows, and they quantify these gains from the perspective of various selves in a finite-horizon life-cycle economy. Likewise, Laibson (1997) computes the welfare gains associated with "partial commitment"—the gains that result from a lack of access to instantaneous credit and a lack of ability to immediately liquidate asset holdings—and Cropper and Laibson (1999) compute the optimal capital subsidy that would perfectly replicate the consumption-saving allocation preferred by the time-zero self. While these studies are related to our paper, they do not uncover the fundamental connection between the number of decision nodes and the appropriateness of the time-zero welfare criterion.[7]

İmrohoroğlu, İmrohoroğlu, and Joines (2003) perform two related sets of welfare experiments (full commitment and partial commitment). Their first set is very similar to Laibson, Repetto, and Tobacman (1998) in that they also consider the welfare gains that accrue to each self from commitment to the plan that a dynamically-consistent (exponential) individual would follow in a finite-horizon life-cycle economy. Their second set of experiments is similar to Laibson (1997) in that both studies consider the gains that accrue to each self from partial commitment, though the papers differ in that partial commitment in Laibson's paper relates to credit market frictions while social security is the partial commitment device in İmrohoroğlu, İmrohoroğlu, and Joines. Once again, the role of the frequency of choice (i.e., the number of decision nodes) is not the focus of their paper.

Finally, we conclude our introductory remarks with a subtle point concerning *choice* versus *paternalism*. Although it may seem paternalistic to equate welfare with time-zero preferences, there is a sense in which doing so is choice-based rather than paternalistic. It is choice-based in the sense that Bernheim and Rangel (2009) define choice-based welfare criteria: the commitment allocation would be *chosen* over the equilibrium allocation by all the selves if such a choice were always available. In other words, if all the time-dated selves of a single individual could step outside of time and coordinate, we cannot say what allocation would result from such coordination in general, but we can say that they would all agree that the commitment allocation dominates the non-cooperative equilibrium allocation.

[7] In each of these studies, Laibson and his coauthors focus on the special case of infinite backward discounting (no weight is given to past consumption), whereas we generalize the multiself welfare calculations to allow for a range of assumptions about backward discounting. By doing this, we address Bernheim and Rangel's (2009) concern that typical multiself Pareto analysis suffers from the "conceptual deficiency" that individuals are assumed to derive no utility at all from past consumption. Bernheim and Rangel argue that there is no empirical basis for such an assumption, nor is there a solid basis for making any other *specific* assumption about backward discounting "because we lack critical information [about] backward looking preferences" (p.89). Consequently, they advocate, in part, that we "adopt a notion of multiself Pareto efficiency that is robust with respect to a wider range of possibilities [about the nature of backward discounting]" (p.88). For this reason, we consider a range of assumptions about backward discounting.

2 Notation and Definitions

Time is discrete and indexed by a finite number of selves or nodes $t = 0, 1, 2, ..., T$. Nothing happens at $t = 0$; no decisions are made and no economic activities occur (e.g., no consumption, no income received, etc.). It is an inaction node. It is there simply to allow us to consider what self 0 would like his future selves to do. All the other nodes $t > 0$ are action nodes or decision nodes.

Definition 1 *An* ***allocation*** *is a vector of consumption decisions* $\mathbf{c} = \{c_1, c_2, ..., c_T\}$. *The set of* ***feasible allocations*** *is* S.

Definition 2 *Following Caplin and Leahy (2004),* ***lifetime utility*** *is a mapping* $U(t, \mathbf{c}) : \mathbb{R}^T \mapsto \mathbb{R}$ *that depends on the vantage point* $t \in [0, T]$. *Note that this definition is general enough to include any assumption about how the individual values past consumption, including the common special case in which he places no value on past consumption.*

Definition 3 *Following the terminology of Bernheim and Rangel (2009), an allocation* $\mathbf{c}' \in S$ ***strictly multiself Pareto dominates*** *another allocation* $\mathbf{c}'' \in S$ *if and only if*

$$U(t, \mathbf{c}') > U(t, \mathbf{c}'') \text{ for all } t \in [0, T].$$

If the inequality holds for a wide range of parameterizations of a given model, then we say that $\mathbf{c}'$ ***strictly and robustly multiself Pareto dominates*** $\mathbf{c}''$.

Definition 4 *The* ***commitment allocation*** $\mathbf{c^0}$ *is the optimal allocation from the vantage point of self 0,*

$$\mathbf{c^0} = \arg\max_{\mathbf{c} \in S} U(0, \mathbf{c}).$$

Definition 5 *The* ***equilibrium allocation*** $\mathbf{c}^*$ *is the allocation that actually materializes from the internal conflict among the many time-dated selves who potentially each have a different view on optimal decision making. We assume individuals are* ***sophisticated*** *and are therefore fully aware of the choices that future selves will make, rather than oblivious (****naive****) to the choices of future selves.*

Definition 6 *The* ***Pareto set****,* $P = \{\mathbf{c} \in S\colon \mathbf{c} = \arg\max U(t, \mathbf{c}) \text{ for at least one } t \in [0, T]\}$. *The elements of* P *are* ***Pareto optima****. Note that* $\mathbf{c^0}$ *is one such element.*

Definition 7 ***Dynamic inconsistency*** *is a situation in which* $\mathbf{c^0} \neq \mathbf{c}^*$.

The thrust of this paper is to understand the conditions under which $\mathbf{c^0}$ strictly multiself Pareto dominates $\mathbf{c}^*$. Of course, even though $\mathbf{c^0}$ is an element of the Pareto set and $\mathbf{c}^*$ may not be an element of the Pareto set, there is no guarantee that $\mathbf{c^0}$ strictly multiself Pareto dominates $\mathbf{c}^*$ because such requires that all selves be made better off in moving from $\mathbf{c}^*$ to $\mathbf{c^0}$.

3 Dynamic Program Part I: Eating Fruit from a Tree

An individual plants a tree at $t = 0$. At each subsequent node (beginning at $t = 1$) the tree bears exactly one piece of new fruit. The fruit may be consumed immediately or it may be left on the tree an additional period to fully ripen. Either way, the fruit tastes good, but it tastes better if it is left to fully ripen. The fruit is totally rotten if it is left on the tree for more than one additional period. The last piece of new fruit is produced at $t = T - 1$, which may be consumed immediately or consumed one period later at $t = T$, after which the tree dies and no consumption takes place beyond T. Throughout the paper we refer to T as the number of decision nodes.

In this example we assume utility is linear to facilitate a variety of analytical results (in the next section utility is concave). Therefore, at each age $t \in [1, T-1]$ the individual faces a simple choice. Take a small amount of utility now c^- or a larger amount c^+ one period later. The choice repeats itself over and over, for all $t \in [1, T-1]$.

Following Laibson (2003) and many others, a given self t applies the following forward discount function to future period utility that is s periods from the present[8]

$$F(s) = \begin{cases} 1 \text{ for } s = 0 \\ \beta \text{ for } s > 0 \end{cases}$$

with $\beta < 1$, and he applies the following backward discount function to past utility that is s periods from the present

$$B(s) = \begin{cases} 1 \text{ for } s = 0 \\ \gamma \text{ for } s > 0 \end{cases}$$

with $\gamma < 1$.

From the perspective of $t = 0$, the individual wants to always be patient and eat ripe fruit, c^+ at $t \in [2, T]$, because there is no perceived cost of waiting for the fruit to ripen.

[8]This is a simple way to capture "present-biased" dynamically-inconsistent preferences. This is a typical calibration (Laibson (2003)) of quasi-hyperbolic discounting, and its simplicity allows us to derive a variety of analytical results. We will pursue the more generic $\beta\delta$ form in the next section of the paper.

Hence the ideal allocation of self 0 is

$$\begin{aligned} c_1 &= 0, \\ c_t &= c^+ \text{ for } t \in [2, T]. \end{aligned}$$

Likewise, from the perspective of any $t > 0$, the individual would want future selves to be patient at all nodes $t+1$ and beyond, but he will be impatient at the current time if

$$c^- > \beta c^+ \Longrightarrow \beta < \frac{c^-}{c^+}.$$

Let us assume this condition is met so that we can have a meaningful discussion about dynamic inconsistency. Hence, self t will be impatient at the current moment and eat unripened fruit. And there is no way for him to impose on later selves his preference for patience at later dates, nor do past choices affect the options currently available to him. Hence, he recognizes that, in equilibrium, all selves will choose to eat the fruit immediately rather than let it ripen; the equilibrium allocation therefore is

$$\begin{aligned} c_t &= c^- \text{ for } t \in [1, T-1] \\ c_T &= 0. \end{aligned}$$

3.1 Welfare

The lifetime utility of the individual at age t is

$$U(t, \mathbf{c}) = \begin{cases} \beta \sum_{s=1}^{T} c_s & t = 0 \\ \gamma \sum_{s=1}^{t-1} c_s + c_t + \beta \sum_{s=t+1}^{T} c_s & t \in [1, T] \end{cases}$$

where we are handling the boundaries $t = 1$ and $t = T$ by making use of the summation notation for an *empty sum* $\sum_{i=a}^{b} x_i = 0$ when $b < a$.

3.2 $T = 2$

The equilibrium allocation is $\mathbf{c}^* = (c^-, 0)$, which confers utility $U(1, \mathbf{c}^*) = c^-$ to self 1. On the other hand the commitment allocation $\mathbf{c^0} = (0, c^+)$ confers utility $U(1, \mathbf{c^0}) = \beta c^+$, which is less than $U(1, \mathbf{c}^*)$ because we have already assumed $\beta < c^-/c^+$. The equilibrium allocation dominates the commitment allocation from the perspective of self 1 and hence commitment cannot represent a Pareto improvement (regardless of how self 2 feels about the commitment allocation over the equilibrium allocation).

3.3 $T > 2$: The Case of $\gamma = \beta$

An allocation $\mathbf{c}$ confers lifetime utility from the perspective of vantage point $t > 0$ according to

$$U(t, \mathbf{c}) = \beta \sum_{s=1}^{t-1} c_s + c_t + \beta \sum_{s=t+1}^{T} c_s \text{ for } t \in [1, T].$$

Self 1

$$U(1, \mathbf{c^0}) = 0 + \beta \sum_{s=2}^{T} c^+ = \beta(T-1)c^+$$

$$U(1, \mathbf{c}^*) = c^- + \beta \sum_{s=2}^{T-1} c^- = c^- + \beta(T-2)c^-.$$

He prefers commitment over the equilibrium, $U(1, \mathbf{c^0}) > U(1, \mathbf{c}^*)$, if

$$\beta(T-1)c^+ > c^- + \beta(T-2)c^-$$

or

$$\beta > \frac{c^-}{T(c^+ - c^-) + 2c^- - c^+} \equiv \bar{\beta}(T).$$

Self $t \in [2, T-1]$

$$U(t, \mathbf{c^0}) = c^+ + \beta \sum_{s=3}^{T} c^+ = c^+ + \beta(T-2)c^+$$

$$U(t, \mathbf{c}^*) = c^- + \beta \sum_{s=3}^{T} c^- = c^- + \beta(T-2)c^-$$

and hence

$$U(t, \mathbf{c^0}) > U(t, \mathbf{c}^*) \text{ for all } t \in [2, T-1].$$

Self T

$$U(T, \mathbf{c^0}) = c^+ + \beta \sum_{s=3}^{T} c^+ = c^+ + \beta(T-2)c^+$$

$$U(T, \mathbf{c}^*) = \beta \sum_{s=2}^{T} c^- = \beta(T-1)c^- = \beta c^- + \beta(T-2)c^-$$

and hence

$$U(T, \mathbf{c^0}) > U(T, \mathbf{c}^*).$$

Notice that self 1 is the only potential holdup in drawing the conclusion that the commitment allocation Pareto dominates the equilibrium allocation. Also recall that we assumed that $\beta < c^-/c^+$ so that there is enough discounting to ensure disagreement between the multiple selves (otherwise preferences are time consistent and all the selves choose to be patient, just like self 0 would like). Hence, to get time-inconsistent preferences and a Pareto role for commitment, we need

$$\beta \in \left(\bar{\beta}(T), \frac{c^-}{c^+}\right).$$

With some algebra we can write

$$\frac{c^-}{c^+} - \bar{\beta}(T) = \frac{c^-}{c^+}\left(\frac{(T-2)(c^+ - c^-)}{(T-2)(c^+ - c^-) + c^+}\right) > 0 \text{ for all } T > 2.$$

This proves that there always exists a non-empty range of values for β for which the commitment allocation Pareto dominates the equilibrium allocation as long as $T > 2$. Moreover the larger is T, the larger is the parameter space that can deliver a Pareto gain from commitment. In fact,

$$\lim_{T\to\infty} \bar{\beta}(T) = 0,$$

which means that the commitment allocation Pareto dominates the equilibrium allocation for *any* $\beta < c^-/c^+$ if the tree lives forever.

Figure 1 is a graph of the parameter space $\beta \in \left(\bar{\beta}(T), c^-/c^+\right)$ as a function of T. Note the convexity of the lower bound: $\bar{\beta}(T)$ falls rapidly as T increases, which means the parameter space opens up rapidly as T increases. This is true for any assumptions about the other parameters (c^- and c^+).

A little more understanding can be gleaned from inverting the condition $\beta > \bar{\beta}(T)$, while recognizing that T is an integer,

$$T \geq \bar{T}\left(2 + \frac{c^- - \beta c^+}{\beta(c^+ - c^-)}\right),$$

where we define $\bar{T}(x)$ as a mapping of real number x to the smallest integer that is greater than x. This condition ensures that the commitment allocation multiself Pareto dominates the equilibrium allocation. We know that $\bar{T} > 2$ because we have already assumed that the decision problem is dynamically inconsistent ($\beta < c^-/c^+$). And we also note that

$$\frac{\partial}{\partial \beta}\left(2 + \frac{c^- - \beta c^+}{\beta(c^+ - c^-)}\right) < 0$$

$$\frac{\partial}{\partial c^-}\left(2 + \frac{c^- - \beta c^+}{\beta(c^+ - c^-)}\right) > 0.$$

As β increases, the necessary number of decision nodes either stays the same or decreases. Likewise, as c^- increases (i.e., as the utility from eating unripened fruit increases), the necessary number of decision nodes either stays the same or increases. The change in β and c^- must be large enough to cross a jump discontinuity in the function $\bar{T}(x)$, in order to trigger a change in the number of decision nodes that are necessary for our Pareto result.

To get a feel for the magnitudes, suppose $c^- = 1/2$ and $c^+ = 1$ (the fruit tastes twice as good if left to ripen). If the individual discounts the future extremely heavily (say $\beta = 0.1$) then there must be 11 or more nodes in order for the holdup self (self 1) to prefer to always be patient and eat ripe fruit. But at more modest levels of discounting (say $\beta = 0.4$), then just 3 nodes is enough to ensure that the holdup self prefers to always be patient and wait for the fruit to ripen. That is, in this particular parameterization, just 3 nodes is enough to ensure that all selves would prefer the commitment allocation of always waiting to let the fruit ripen over the equilibrium allocation of always eating unripe fruit.

Alternatively, suppose $c^- = 3/4$ and $c^+ = 1$. Then there must be 29 or more decision nodes when $\beta = 0.1$ and there must be 6 or more decision nodes when $\beta = 0.4$.

3.4 $T > 2$: The Case of $\gamma = 0$

An allocation $\mathbf{c}$ confers lifetime utility from the perspective of vantage point $t > 0$ according to

$$U(t, \mathbf{c}) = c_t + \beta \sum_{s=t+1}^{T} c_s \text{ for } t \in [1, T].$$

Self 1

$$U(1, \mathbf{c^0}) = 0 + \beta \sum_{s=2}^{T} c^+ = \beta(T-1)c^+$$

$$U(1, \mathbf{c}^*) = c^- + \beta \sum_{s=2}^{T-1} c^- = c^- + \beta(T-2)c^-.$$

He prefers commitment over the equilibrium, $U(1, \mathbf{c^0}) > U(1, \mathbf{c}^*)$, if

$$\beta(T-1)c^+ > c^- + \beta(T-2)c^-$$

or

$$\beta > \frac{c^-}{T(c^+ - c^-) + 2c^- - c^+} \equiv \bar{\beta}(T).$$

Self $t \in [2, T-1]$

$$U(t, \mathbf{c^0}) = c^+ + \beta \sum_{s=t+1}^{T} c^+ = c^+ + \beta(T-t)c^+$$

$$U(t, \mathbf{c}^*) = c^- + \beta \sum_{s=t+1}^{T-1} c^- = c^- + \beta(T-t-1)c^-.$$

And hence

$$U(t, \mathbf{c^0}) > U(t, \mathbf{c}^*) \text{ for all } t \in [2, T-1].$$

Self T

$$U(T, \mathbf{c^0}) = c^+$$

$$U(T, \mathbf{c}^*) = 0$$

and hence

$$U(T, \mathbf{c^0}) > U(T, \mathbf{c}^*).$$

Once again, we arrive at the same result: self 1 is the only potential holdup in drawing the conclusion that the commitment allocation Pareto dominates the equilibrium allocation, and self 1 prefers commitment under the same condition as above

$$\beta > \frac{c^-}{T(c^+ - c^-) + 2c^- - c^+} \equiv \bar{\beta}(T).$$

3.5 Intuition and Summary of Findings

We know self 1 is the holdup. Let's take a deeper look at why the time grid matters. Consider how his ideal allocation $\mathbf{c^1}$ compares to $\mathbf{c^0}$ and $\mathbf{c}^*$. For $T > 2$,

$$\begin{aligned}\mathbf{c^0} &= (0, c^+, ..., c^+) \\ \mathbf{c^1} &= (c^-, 0, c^+, ..., c^+) \\ \mathbf{c}^* &= (c^-, ..., c^-, 0).\end{aligned}$$

Self 1 disagrees with Self 0's ideal allocation only at nodes $t = 1$ and $t = 2$. On the other hand, self 1 disagrees with the equilibrium allocation at every node except $t = 1$. In other words, as T increases, the equilibrium allocation starts to look less and less like the desires of self 1 and the commitment allocation starts to look more and more like the desires of self 1. This pattern can be seen in Figures 2 and 3, which show the specific cases of $T = 2$ and $T = 6$ to illustrate this point.

Formally,

$$\begin{aligned}U(1, \mathbf{c^0}) &= \beta(T-1)c^+ \\ U(1, \mathbf{c^1}) &= c^- + \beta(T-2)c^+ \\ U(1, \mathbf{c}^*) &= c^- + \beta(T-2)c^-.\end{aligned}$$

Compute forward first differences D in the T dimension

$$\begin{aligned}D[U(1, \mathbf{c^1}) - U(1, \mathbf{c^0})] &= \left[c^- + \beta(T-1)c^+ - \beta T c^+\right] \\ &\quad - \left[c^- + \beta(T-2)c^+ - \beta(T-1)c^+\right] \\ &= 0\end{aligned}$$

$$\begin{aligned}D[U(1, \mathbf{c^1}) - U(1, \mathbf{c}^*)] &= [c^- + \beta(T-1)c^+ - c^- - \beta(T-1)c^-] \\ &\quad -[c^- + \beta(T-2)c^+ - c^- - \beta(T-2)c^-] \\ &= \beta(c^+ - c^-) > 0.\end{aligned}$$

Hence, as T increases, the utility gap between $\mathbf{c^1}$ and $\mathbf{c^0}$ stays fixed, while the utility gap between $\mathbf{c^1}$ versus $\mathbf{c}^*$ increases without bound. In other words, while self 1 may prefer $\mathbf{c}^*$ over $\mathbf{c^0}$ at very low T, this preference is guaranteed to eventually reverse itself as T increases.

In sum, when there is a small number of decision nodes (i.e., a small number of selves), a given self may have the power to significantly influence the equilibrium allocation. If so, then the equilibrium allocation may not be so bad from his perspective; in fact, it may be relatively close to what he wants. However, when there is a large number of selves, then the power to influence the equilibrium allocation is diffuse among selves. No one self has much power to influence the equilibrium outcome and hence the equilibrium is relatively dissimilar to the desires of a given self. And in such an environment where power is diffuse and no self gets what he wants, it is possible to find allocations that Pareto dominate the equilibrium allocation. In other words, just because the many selves disagree on the ideal allocation of resources doesn't mean that they can't all agree that certain allocations are better than others; and, the further the equilibrium allocation gets from any one self's ideal allocation, the more room there is for such an agreement.

We can summarize the result thus far as follows:

Result 1 *When there are only two decision nodes, the commitment allocation will never multiself Pareto dominate the equilibrium allocation.*

Result 2 *When the number of decision nodes is greater than two, there is a non-empty parameter space over β for which the commitment allocation multiself Pareto dominates the equilibrium allocation.*

Result 3 *The non-empty parameter space grows wider as the number of decision nodes increases, and this space ultimately subsumes the entire parameter space as the model approaches infinitely many (countable) decision nodes.*

Result 4 *The commitment allocation will multiself Pareto dominate the equilibrium allocation if and only if the number of decision nodes exceeds a threshold. This threshold is decreasing in the forward discount factor β and increasing in the utility of immediate consumption c^-.*

3.6 $T > 2$: The Generic Case (Unrestricted γ)

Finally, up to this point we have focused on two special cases of $\gamma = \beta$ and $\gamma = 0$. The purpose of this section is to illustrate that we have not lost any generality in doing so, though the more general case is more cumbersome algebraically. All the same results hold when we allow for any value of $\gamma < 1$: the commitment allocation Pareto dominates the equilibrium allocation if and only if self 1 prefers commitment, and this happens if $\beta > \bar{\beta}(T)$ as in Figure 1. All the lessons and intuition from the previous special cases therefore carry over to the case of generic γ. Hence, this subsection is presented for technical completeness. Readers who are less interested in the technical details can skip to the cake-eating problem.

An allocation $\mathbf{c}$ confers lifetime utility from the perspective of vantage point $t > 0$ according to

$$U(t, \mathbf{c}) = \gamma \sum_{s=1}^{t-1} c_s + c_t + \beta \sum_{s=t+1}^{T} c_s \text{ for } t \in [1, T].$$

Self 1

$$U(1, \mathbf{c^0}) = 0 + \beta \sum_{s=2}^{T} c^+ = \beta(T-1)c^+$$

$$U(1, \mathbf{c}^*) = c^- + \beta \sum_{s=2}^{T-1} c^- = c^- + \beta(T-2)c^-.$$

He prefers commitment over the equilibrium, $U(1, \mathbf{c^0}) > U(1, \mathbf{c}^*)$, if

$$\beta(T-1)c^+ > c^- + \beta(T-2)c^-$$

or

$$\beta > \frac{c^-}{T(c^+ - c^-) + 2c^- - c^+} \equiv \bar{\beta}(T)$$

which is just the same parameter space as in the previous cases (it doesn't matter how he discounts past consumption because there is no past consumption).

But the results for self $t \in [2, T-1]$ are more complicated than in the previous special cases

$$U(t, \mathbf{c^0}) = \gamma \sum_{s=2}^{t-1} c^+ + c^+ + \beta \sum_{s=t+1}^{T} c^+ = \gamma(t-2)c^+ + c^+ + \beta(T-t)c^+$$

$$U(t, \mathbf{c}^*) = \gamma \sum_{s=1}^{t-1} c^- + c^- + \beta \sum_{s=t+1}^{T-1} c^- = \gamma(t-1)c^- + c^- + \beta(T-t-1)c^-.$$

And hence

$$U(t, \mathbf{c^0}) > U(t, \mathbf{c}^*) \text{ for all } t \in [2, T-1]$$

if

$$\beta > \frac{\gamma\left[(t-1)c^- - (t-2)c^+\right] + c^- - c^+}{(T-t)c^+ - (T-t-1)c^-} \equiv \hat{\beta}(\gamma, T, t) \text{ for all } t \in [2, T-1].$$

Self T

$$U(T, \mathbf{c^0}) = c^+ + \gamma \sum_{s=3}^{T} c^+ = c^+ + \gamma(T-2)c^+$$

$$U(T,\mathbf{c}^*) = \gamma \sum_{s=2}^{T} c^- = \gamma(T-1)c^- = \gamma c^- + \gamma(T-2)c^-$$

and hence it is always the case that

$$U(T,\mathbf{c}^{\mathbf{0}}) > U(T,\mathbf{c}^*).$$

In summary, for a given number of nodes T, the commitment allocation dominates the equilibrium allocation from the perspective of self 1 if

$$\beta \in \left(\bar{\beta}(T), \frac{c^-}{c^+}\right).$$

Likewise, for a given number of nodes T, and for a given backward discount factor γ, the commitment allocation dominates the equilibrium allocation from the perspective of self $t \in [2, T-1]$ if

$$\beta \in \left(\hat{\beta}(\gamma, T, t), \frac{c^-}{c^+}\right).$$

However, it can be show that

$$\hat{\beta}(\gamma, T, t) < \bar{\beta}(T) \text{ for all } t \in [2, T-1] \text{ and for all } \gamma < 1$$

which means that if the commitment allocation dominates the equilibrium allocation from the perspective of self 1, then the commitment allocation will multiself Pareto dominate the equilibrium allocation.

To see this, note that

$$\frac{\partial \hat{\beta}(\gamma, T, t)}{\partial t} = (c^- - c^+) \times \left(\frac{\gamma(T-2)\,(c^+ - c^-) + c^+ - c^-}{[(T-t)c^+ - (T-t-1)c^-]^2}\right) < 0 \text{ for all } t \geq 2$$

which means that

$$2 = \arg\max_t \hat{\beta}(\gamma, T, t)$$

$$\hat{\beta}(\gamma, T, 2) = \frac{\gamma c^- + c^- - c^+}{(T-2)c^+ - (T-3)c^-}.$$

Further note that $\hat{\beta}(\gamma, T, 2)$ is maximized in the γ dimension as $\gamma \to 1$ (we are assuming γ cannot exceed 1). Rewrite

$$\bar{\beta}(T) = \frac{c^-}{(T-2)(c^+ - c^-) + c^+}$$

$$\hat{\beta}(1,T,2) = \frac{2c^- - c^+}{(T-2)(c^+ - c^-) + c^-}$$

and compute

$$\bar{\beta}(T) - \hat{\beta}(1,T,2) = \frac{c^-[(T-2)(c^+ - c^-) + c^-] - (2c^- - c^+)[(T-2)(c^+ - c^-) + c^+]}{[(T-2)(c^+ - c^-) + c^+][(T-2)(c^+ - c^-) + c^-]}.$$

The denominator is positive, and with some algebra the numerator can be rewritten as

$$\left(c^+ - c^-\right)^2 (T-1) > 0$$

and hence

$$\bar{\beta}(T) > \hat{\beta}(1,T,2) > \hat{\beta}(\gamma,T,t), \text{ for all } t \in [2, T-1] \text{ and for all } \gamma < 1.$$

Hence, as in the special cases in the previous subsections, self 1 is the holdup. The commitment allocation Pareto dominates the equilibrium allocation if and only if self 1 prefers commitment, and this happens if $\beta > \bar{\beta}(T)$ as in Figure 1, and therefore all the lessons and intuition from the previous special cases carry over to the case of generic γ.

4 Dynamic Programming Part II: Cake-Eating

The cake-eating problem is a first model in most lectures and textbooks on dynamic programming. There is an infinitely divisible quantity of cake C. At $t = 0$ the individual orders a cake that will arrive at $t = 1$, and hence no cake can be eaten at $t = 0$ but cake may be eaten at all the other T decision nodes. The cake doesn't spoil (nor does it grow). The flow of consumption of the cake at node t is c_t. Any cake not yet consumed is available for consumption at future decision nodes. The period utility function $u(c_t)$ satisfies $u' > 0, u'' < 0$.[9]

Self t applies the following forward discount function $F(s)$ to future period utility that is s years away, with $F(0) = 1$. He likewise applies the backward discount function $B(s)$ to past period utility that is s years away, with $B(0) = 1$.

[9] Unlike the fruit tree example in which we had the luxury of assuming linear utility, which helped us derive a variety of analytical results, here we must assume utility is strictly concave. Without concavity, there is no solution to the time-zero cake-eating problem under once-off discounting $\{1, \beta, \beta, \beta, ...\}$ because the time-zero self doesn't care when the cake is eaten. And under quasi-hyperbolic discounting $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, ...\}$ with linear utility, then there is a well-defined solution to the time-zero problem but it is the same as the equilibrium allocation: eat all the cake at $t = 1$, and in this case, commitment has nothing to offer.

4.1 The Case of $T = 2$

Consumption takes place at $t = 1$ and $t = 2$ but not at $t = 0$. From the perspective of $t = 0$, the optimal allocation solves

$$\max : F(1)u(c_1) + F(2)u(c_2), \quad \text{s.t. } c_1 + c_2 = C,$$

which gives the first order condition

$$\frac{u'(c_1)}{u'(c_2)} = \frac{F(2)}{F(1)}.$$

However, the preferences of self 1 will prevail and actual consumption solves

$$\max : u(c_1) + F(1)u(c_2), \quad \text{s.t. } c_1 + c_2 = C,$$

which has the first order condition

$$\frac{u'(c_1)}{u'(c_2)} = F(1).$$

Note that the first order conditions will differ, except in the special case of exponential discounting $F(2) = (F(1))^2$. Hence, time-inconsistent preferences creates disagreement about which consumption allocation is optimal, and in a three period setting this necessarily implies that the commitment allocation cannot multi-self Pareto dominate the actual allocation, because the actual allocation is the optimal allocation of self 1. However, as with the previous example, there is a deeper point at work here: this logic is specific to the low frequency setting. If we add more decision nodes, we will see that even though preferences are non-stationary and each self has a different perspective on what is optimal, it is possible to construct examples where the commitment allocation dominates the actual allocation from the perspective of all the selves. When there are more selves (i.e., more decision nodes) it is easier to identify allocations that make everyone better off.

4.2 The Case of $T \geq 2$

Now it becomes a numerical issue to determine how many decision nodes are needed to generate a Pareto argument for commitment. We assume log utility.

4.2.1 The Time-Zero Allocation (Commitment Allocation)

Self 0 would like his future selves to obey

$$\max : \sum_{t=1}^{T} F(t) \ln c_t, \quad \text{s.t.} \ \sum_{t=1}^{T} c_t = C,$$

which has the following solution (commitment allocation)

$$c_t = \frac{CF(t)}{\sum_{s=1}^{T} F(s)}, \text{ for all } t > 0.$$

4.2.2 Non-Cooperative Equilibrium Allocation

We solve the problem recursively:

At vantage point $T-1$ he solves

$$\max : \ln c_{T-1} + F(1) \ln c_T, \quad \text{s.t.} \ c_{T-1} + c_T = C - \sum_{t=1}^{T-2} c_t.$$

The first order condition is

$$c_T = F(1) c_{T-1},$$

which implies

$$c_{T-1} = \frac{C - \sum_{t=1}^{T-2} c_t}{1 + F(1)}.$$

At vantage point $T-2$ he solves

$$\max : \ln c_{T-2} + F(1) \ln c_{T-1} + F(2) \ln c_T$$

s.t.[10]

$$\begin{aligned} c_{T-2} + c_{T-1} + c_T &= C - \sum_{t=1}^{T-3} c_t \\ c_{T-1} &= \frac{C - \sum_{t=1}^{T-2} c_t}{1 + F(1)} \\ c_T &= F(1) c_{T-1} \end{aligned}$$

or

[10] The first of these constraints—remaining consumption must equal remaining cake—is actually redundant because this constraint was already imposed in the first step of the recursive solution.

$$\max : \ln c_{T-2} + F(1)\ln\left(\frac{C-\sum_{t=1}^{T-2} c_t}{1+F(1)}\right) + F(2)\ln\left(F(1)\frac{C-\sum_{t=1}^{T-2} c_t}{1+F(1)}\right)$$

which is the same as maximizing

$$\max : \ln c_{T-2} + F(1)\ln\left(C-\sum_{t=1}^{T-3} c_t - c_{T-2}\right) + F(2)\ln\left(C-\sum_{t=1}^{T-3} c_t - c_{T-2}\right)$$

or

$$\max : \ln c_{T-2} + (F(1)+F(2))\ln\left(C-\sum_{t=1}^{T-3} c_t - c_{T-2}\right).$$

The first order condition is

$$\frac{1}{c_{T-2}} - \frac{F(1)+F(2)}{C-\sum_{t=1}^{T-3} c_t - c_{T-2}} = 0$$

$$c_{T-2} = \frac{C-\sum_{t=1}^{T-3} c_t}{1+F(1)+F(2)}.$$

At vantage point $T-3$ he solves

$$\max : \ln c_{T-3} + F(1)\ln c_{T-2} + F(2)\ln c_{T-1} + F(3)\ln c_T$$

s.t.[11]

$$\begin{aligned} c_{T-3}+c_{T-2}+c_{T-1}+c_T &= C-\sum_{t=1}^{T-4} c_t \\ c_{T-2} &= \frac{C-\sum_{t=1}^{T-3} c_t}{1+F(1)+F(2)} \\ c_{T-1} &= \frac{C-\sum_{t=1}^{T-2} c_t}{1+F(1)} \\ c_T &= F(1)c_{T-1}. \end{aligned}$$

Note that

$$c_{T-2} = \frac{C-\sum_{t=1}^{T-3} c_t}{1+F(1)+F(2)} = \frac{C-\sum_{t=1}^{T-4} c_t - c_{T-3}}{1+F(1)+F(2)}$$

[11] Again, the first constraint is redundant and can be dropped.

$$
\begin{aligned}
c_{T-1} &= \frac{C-\sum_{t=1}^{T-2} c_t}{1+F(1)} = \frac{C-\sum_{t=1}^{T-4} c_t - c_{T-3} - c_{T-2}}{1+F(1)} \\
&= \frac{C-\sum_{t=1}^{T-4} c_t - c_{T-3}}{1+F(1)} - \frac{1}{1+F(1)} \frac{C-\sum_{t=1}^{T-4} c_t - c_{T-3}}{1+F(1)+F(2)} \\
&= \left(C-\sum_{t=1}^{T-4} c_t - c_{T-3}\right)\left(\frac{1}{1+F(1)}\right)\left(\frac{F(1)+F(2)}{1+F(1)+F(2)}\right).
\end{aligned}
$$

Hence

$$
\begin{aligned}
\max \quad &: \quad \ln c_{T-3} + F(1)\ln\left(\frac{C-\sum_{t=1}^{T-4} c_t - c_{T-3}}{1+F(1)+F(2)}\right) \\
&+F(2)\ln\left(\left(C-\sum_{t=1}^{T-4} c_t - c_{T-3}\right)\left(\frac{1}{1+F(1)}\right)\left(\frac{F(1)+F(2)}{1+F(1)+F(2)}\right)\right) \\
&+F(3)\ln\left(F(1)\left(C-\sum_{t=1}^{T-4} c_t - c_{T-3}\right)\left(\frac{1}{1+F(1)}\right)\left(\frac{F(1)+F(2)}{1+F(1)+F(2)}\right)\right)
\end{aligned}
$$

which is the same as maximizing

$$
\begin{aligned}
\max \quad &: \quad \ln c_{T-3} + F(1)\ln\left(C-\sum_{t=1}^{T-4} c_t - c_{T-3}\right) + F(2)\ln\left(C-\sum_{t=1}^{T-4} c_t - c_{T-3}\right) \\
&+F(3)\ln\left(C-\sum_{t=1}^{T-4} c_t - c_{T-3}\right)
\end{aligned}
$$

or

$$
\max : \ln c_{T-3} + (F(1)+F(2)+F(3))\ln\left(C-\sum_{t=1}^{T-4} c_t - c_{T-3}\right)
$$

which has the first order condition

$$
\frac{1}{c_{T-3}} - \frac{F(1)+F(2)+F(3)}{C-\sum_{t=1}^{T-4} c_t - c_{T-3}} = 0
$$

which implies

$$
c_{T-3} = \frac{C-\sum_{t=1}^{T-4} c_t}{1+F(1)+F(2)+F(3)}.
$$

Note the pattern that has emerged, standing n nodes back from the final node

$$
c_{T-n} = \frac{C-\sum_{t=1}^{T-(n+1)} c_t}{1+\sum_{t=1}^{n} F(t)}.
$$

For convenience, change the index dummy from t to s

$$c_{T-n} = \frac{C - \sum_{s=1}^{T-(n+1)} c_s}{1 + \sum_{s=1}^{n} F(s)}$$

and note that standing n nodes back from the last node is the same as standing at node $t = T - n$, so

$$c_t = \frac{C - \sum_{s=1}^{t-1} c_s}{1 + \sum_{s=1}^{T-t} F(s)}, \text{ for } t > 0,$$

and hence we have the following recursion

$$\begin{aligned} c_{t+1} &= \frac{C - \sum_{s=1}^{t} c_s}{1 + \sum_{s=1}^{T-t-1} F(s)} \\ &= \frac{C - \sum_{s=1}^{t-1} c_s - c_t}{1 + \sum_{s=1}^{T-t-1} F(s)} \\ &= \frac{C - \sum_{s=1}^{t-1} c_s}{1 + \sum_{s=1}^{T-t-1} F(s)} - \frac{c_t}{1 + \sum_{s=1}^{T-t-1} F(s)} \\ &= c_t \frac{1 + \sum_{s=1}^{T-t} F(s)}{1 + \sum_{s=1}^{T-t-1} F(s)} - \frac{c_t}{1 + \sum_{s=1}^{T-t-1} F(s)} \\ &= c_t \left(\frac{\sum_{s=1}^{T-t} F(s)}{1 + \sum_{s=1}^{T-t-1} F(s)} \right) \\ &< c_t. \end{aligned}$$

4.2.3 Special Case: Once-Off Discounting

For a delay of length s,

$$\begin{aligned} F(s) &= \begin{cases} 1 \text{ for } s = 0 \\ \beta \text{ for } s > 0 \end{cases} \\ B(s) &= \begin{cases} 1 \text{ for } s = 0 \\ \gamma \text{ for } s > 0. \end{cases} \end{aligned}$$

Self 0 would like his future selves to obey (commitment allocation)

$$c_t = C/T, \text{ for all } t.$$

However, the equilibrium consumption recursion satisfies

$$c_{t+1} = c_t \left(\frac{\beta(T-t)}{1+\beta(T-t-1)} \right)$$

which has the closed form solution

$$c_t = c_1 \prod_{s=1}^{t-1} \left(\frac{\beta(T-s)}{1+\beta(T-s-1)} \right), \text{ for } t > 0,$$

where we recall the notation for an *empty product* $\prod_{i=a}^{b} x_i = 1$ when $b < a$, and

$$c_1 = \frac{C}{1+\beta(T-1)}.$$

For self $t > 0$, a consumption allocation $\mathbf{c}$ confers utility

$$U(t, \mathbf{c}) = \gamma \sum_{s=1}^{t-1} \ln c_s + \ln c_t + \beta \sum_{s=t+1}^{T} \ln c_s.$$

Table 1 reports the number of decision nodes to that are needed to ensure that the commitment allocation is a multiself Pareto improvement over the equilibrium allocation. In other words, Table 1 reports the number of decision nodes that are needed to ensure that $U(t, \mathbf{c^0}) > U(t, \mathbf{c}^*)$ for all t.

Table 1. Necessary and Sufficient Decision Nodes for Pareto Improvement

	$\beta = 0.2$	$\beta = 0.3$	$\beta = 0.4$	$\beta = 0.5$	$\beta = 0.6$	$\beta = 0.7$	$\beta = 0.8$
$\gamma = 1$	$T \geq 9$	$T \geq 8$	$T \geq 8$	$T \geq 8$	$T \geq 8$	$T \geq 8$	$T \geq 8$
$\gamma = \beta$	$T \geq 6$	$T \geq 5$	$T \geq 5$	$T \geq 4$	$T \geq 4$	$T \geq 4$	$T \geq 4$
$\gamma = 0$	$T \geq 6$	$T \geq 5$	$T \geq 5$	$T \geq 4$	$T \geq 4$	$T \geq 4$	$T \geq 4$

Note: β is the forward discount factor and γ is the backward discount factor, and T is the number of nodes that ensure $U(t, \mathbf{c^0}) > U(t, \mathbf{c}^*)$ for all t.

There are three basic results in Table 1.

Result 5 *If the number of decision nodes is equal to or greater than a given threshold, then the commitment allocation multiself Pareto dominates the equilibrium allocation.*

When there is a small number of decision nodes (i.e., a small number of selves), a given self may have the power to significantly influence the equilibrium allocation. In this case, the equilibrium allocation may not be so bad from his perspective. However, when there is a large number of selves, then the power to influence the equilibrium allocation is diffuse among selves; no one self has much power, and hence the equilibrium is relatively dissimilar to the desires of a given self. Table 1 reports the threshold at which power becomes diffuse enough that all the selves are far enough from what they want that they all prefer the commitment allocation to the equilibrium allocation.

Result 6 *The necessary and sufficient number of decision nodes (to ensure that the commitment allocation multiself Pareto dominates the equilibrium allocation) is increasing in the degree of forward discounting.*

Considering an extreme case helps to highlight the intuition for this result. As $\beta \to 0$, it is impossible for the commitment allocation to Pareto dominate the equilibrium allocation. To see this, it is enough to show that at least one self prefers the actual allocation. Consider self 1. The allocation that maximizes his utility is $\lim_{\beta \to 0} c_1 = C$ with $\lim_{\beta \to 0} c_t = 0,$ for $t > 1$. The individual wants everything now and wants future selves to have nothing. And in this special case, the individual is able to perfectly enforce his will because eating all the cake at $t = 1$ will eliminate all opportunities for future decision making. Hence it is impossible for the commitment allocation (which sets consumption to C/T for all t) to dominate the actual allocation for any finite T because the first self prefers the equilibrium allocation.

Result 7 *The necessary and sufficient number of decision nodes is decreasing in the degree of backward discounting.*

We now quantify the welfare gains (or losses) from commitment. Let $\mathbf{c^0}(C)$ denote the commitment allocation as a function of the size of the cake, and let $\mathbf{c}^*(C)$ denote the actual allocation as a function of the cake. Define Δ_t as the solution to

$$U(t, \mathbf{c^0}(C\Delta_t)) = U(t, \mathbf{c}^*(C)).$$

Note that

$$U(t, \mathbf{c^0}(C\Delta_t)) = U(t, \mathbf{c^0}(C)) + \gamma(t-1)\ln\Delta_t + \ln\Delta_t + \beta(T-t)\ln\Delta_t$$

and hence

$$\Delta_t = \exp\left(\frac{U(t, \mathbf{c}^*(C)) - U(t, \mathbf{c^0}(C))}{\gamma(t-1) + 1 + \beta(T-t)}\right), \text{ for } t > 0.$$

The function $1 - \Delta_t$ tells us what fraction of the cake self t would be willing to give up in order to consume the commitment allocation $(\Delta_t C/T)$. Figures 4, 5, and 6 plot $1 - \Delta_t$

for each self $t > 0$, for various assumptions about the total number of nodes T and discount factors β and γ. In generating these figures we have assumed $\beta = 0.5$ in order to give a feel for the magnitude of the potential welfare gains from commitment. The overall visual appearance of these graphs is not terribly sensitive to the precise value of β.

4.2.4 Special Case: Quasi-hyperbolic Discounting

For a delay of length s,

$$F(s) = \begin{cases} 1 \text{ for } s = 0 \\ \beta\delta^s \text{ for } s > 0 \end{cases}$$

$$B(s) = \begin{cases} 1 \text{ for } s = 0 \\ \gamma\eta^s \text{ for } s > 0. \end{cases}$$

Self 0 would like his future selves to obey (commitment allocation)

$$c_t = \frac{C\delta^t}{\sum_{s=1}^{T} \delta^s} = C\delta^t \left(\frac{1-\delta}{\delta - \delta^{T+1}} \right), \text{ for all } t > 0.$$ [12]

However, the equilibrium consumption allocation is

$$c_1 = \frac{C}{1 + \beta \sum_{s=1}^{T-1} \delta^s} = \frac{C}{1 + \beta \left(\frac{\delta - \delta^T}{1-\delta} \right)}$$

$$c_{t+1} = c_t \left(\frac{\beta \sum_{s=1}^{T-t} \delta^s}{1 + \beta \sum_{s=1}^{T-t-1} \delta^s} \right) = \beta c_t \left(\frac{\frac{\delta - \delta^{T-t+1}}{1-\delta}}{1 + \beta \left(\frac{\delta - \delta^{T-t}}{1-\delta} \right)} \right), \text{ for } t > 0.$$

For self $t > 0$, a consumption allocation $\mathbf{c}$ confers utility

$$U(t, \mathbf{c}) = \gamma \sum_{s=1}^{t-1} \eta^{t-s} \ln c_s + \ln c_t + \beta \sum_{s=t+1}^{T} \delta^{s-t} \ln c_s.$$

As above, let $\mathbf{c^0}(C)$ denote the commitment allocation as a function of the size of the cake, and let $\mathbf{c}^*(C)$ denote the actual allocation as a function of the cake. Define Δ_t as the solution

[12] Recall that a one-dollar annuity that lasts for T periods at rate r per period has a present value $\sum_{s=1}^{T} 1/(1+r)^s = (1 - 1/(1+r)^T)/r$. Setting $\delta = 1/(1+r)$ gives the reduced form shown in the equation above.

to

$$U(t,\mathbf{c^0}(C\Delta_t)) = U(t,\mathbf{c}^*(C)).$$

Note that

$$U(t,\mathbf{c^0}(C\Delta_t)) = U(t,\mathbf{c^0}(C)) + \left(\gamma\sum_{s=1}^{t-1}\eta^{t-s} + 1 + \beta\sum_{s=t+1}^{T}\delta^{s-t}\right)\ln\Delta_t$$

and hence

$$\begin{aligned}
\Delta_t &= \exp\left(\frac{U(t,\mathbf{c}^*(C)) - U(t,\mathbf{c^0}(C))}{\gamma\sum_{s=1}^{t-1}\eta^{t-s} + 1 + \beta\sum_{s=t+1}^{T}\delta^{s-t}}\right) \\
&= \exp\left(\frac{U(t,\mathbf{c}^*(C)) - U(t,\mathbf{c^0}(C))}{\gamma\sum_{s=1}^{t-1}\eta^{s} + 1 + \beta\sum_{s=1}^{T-t}\delta^{s}}\right) \\
&= \exp\left(\frac{U(t,\mathbf{c}^*(C)) - U(t,\mathbf{c^0}(C))}{\gamma\left(\frac{\eta-\eta^t}{1-\eta}\right) + 1 + \beta\left(\frac{\delta-\delta^{T-t+1}}{1-\delta}\right)}\right), \text{ for } t > 0.
\end{aligned}$$

We will avoid duplicating too many tables and figures because it is sufficient to say that the principle lesson continues to go through: for a given parameterization (now there are four parameters instead of two, $\beta, \delta, \gamma, \eta$) there is a threshold number of decision nodes, below which the commitment allocation does not make all the selves better off relative to the equilibrium, and above which the commitment allocation multiself Pareto dominates the equilibrium.

For example, consider a fairly typical calibration of the quasi-hyperbolic forward discount function $\beta = 0.8, \delta = 0.95$, and suppose the backward discount function takes the same shape $\gamma = 0.8, \eta = 0.95$. In this case, as long as there are 4 or more decision nodes ($T \geq 4$) then the commitment allocation multiself Pareto dominates the equilibrium allocation. Suppose $T = 7$; that is, suppose the cake can be eaten each day for a week. In this case we know that all 7 selves will prefer commitment because there are more than 4 decision nodes, and their willingness to pay $1 - \Delta_t$ for the commitment allocation is monotonically increasing in t (the age of the self). Self 1 would give up 0.85% of the cake to commitment himself to the time-zero allocation, while self 7 would give up 6.60% for commitment.

Suppose we keep the same assumptions about the forward discount function $\beta = 0.8, \delta = 0.95$, but now we assume the individual doesn't care about the past ($\gamma = 0$). As in the previous example, it still happens to be the case that as long as there are 4 or more decision nodes ($T \geq 4$) then the commitment allocation multiself Pareto dominates the equilibrium allocation. If we continue to assume the individual has one week to eat the cake ($T = 7$), then willingness to pay $1 - \Delta_t$ for the commitment allocation is monotonically increasing in t. Self 1 would give up 0.85% of the cake to commitment himself to the time-zero allocation, while self 7 would give up 33.66% for commitment. The oldest self is, of course, willing to pay a huge premium for commitment because commitment leaves him a bigger piece of cake at day 7 than the equilibrium outcome. Self 0 would like to leave 12.18% of the cake for day 7 while the equilibrium allocation only leaves 8.08% for day 7. The latter is 33.66% less than the former and hence this is self 7's willingness to pay for commitment. Of course, self 7's ideal allocation is to eat all of the cake on day 7, so neither the commitment allocation nor the equilibrium allocation come close to his first best. But our point is not that commitment helps all selves reach their first best, our point is that commitment can help all selves do better than the equilibrium if there are enough decision nodes (selves) in the choice problem.

5 Concluding Remarks

In this paper we have tried to make some progress toward laying a stronger philosophical foundation for the O'Donoghue-Rabin approach to behavioral welfare analysis (i.e., equating welfare with time-zero preferences). Critics of this approach worry that it is ad hoc or arbitrary. And our sense is that even among behavioral economists who support equating welfare with time-zero preferences on intuitive grounds, they too would like to see more formal work to strengthen the logical underpinnings of such an approach. Our finding that equating welfare with time-zero preferences is consistent with a multiself Pareto criterion, as long as the number of decision nodes exceeds a specific threshold, provides a little stronger foundation than has previously existed.

We conclude with the fundamental question that we opened with in the introduction: How should scarce resources be allocated when individuals have dynamically inconsistent preferences? Critics of behavioral economics may argue that meaningful welfare analysis is hopeless because the time-dated selves of a single individual disagree on how resources should be allocated over time. But just because the different selves each have different ideas on how resources should *ideally* be allocated over time doesn't mean that they can't at least agree that some allocations (commitment) are better than others (equilibrium). And our contribution is to show that as the number of decision nodes increases and therefore there are more conflicting points of view on how resources should ideally be allocated over time, *almost* paradoxically it becomes easier for the selves to reach a unanimous agreement that the commitment allocation beats the equilibrium allocation. The reason is that with many selves instead of just a few, the power to influence the equilibrium allocation becomes diffuse and the equilibrium therefore diverges from the ideal allocations of any of the selves. This gap opens the door for a Pareto improvement.

While we do not expect that our findings will necessarily extend to every conceivable economic setting for which time-inconsistent preferences could be studied, the fact that our findings persist across a broad space of model settings (renewable and nonrenewable resources) seems to be an important consideration in the ongoing welfare debate. At a minimum, we can safely conclude that the O'Donoghue-Rabin approach to behavioral welfare analysis (i.e., equating welfare with time-zero preferences) may have a firmer philosophical foundation than previously supposed.

There are at least two interesting extensions to this paper. First, rather than just comparing the time-zero allocation to the equilibrium allocation, it would be interesting to compare the ideal allocations of all the time-dated selves to the equilibrium allocation, and then see which of these Pareto optima tend to Pareto dominate the equilibrium. This will provide additional insights as to whether the time-zero allocation is just one of many Pareto optima that Pareto dominate the equilibrium, or whether it is instead one of just a few allocations that share this feature. Understanding the answer to this question would help to better understand the concern that the time-zero criterion is arbitrary. Second, it would also be interesting (and ambitious) to characterize the entire set of feasible allocations that Pareto dominate the equilibrium allocation. Note that this second set may include allocations that are not Pareto optimal. Characterizing the scope and boundaries of this set would provide further understanding of multiself welfare analysis and the validity of the time-zero welfare criterion.

References

Bernheim, Douglas B. and Antonio Rangel (2009). Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Analysis. *Quarterly Journal of Economics* 124(1), 51-104.

Brocas, Isabelle, Juan D. Carrillo, and Mathias Dewatripont (2004). Commitment Devices under Self-control Problems: An Overview. In *The Psychology of Economic Decisions: Volume II: Reasons and Choices*, edited by Isabelle Brocas and Juan D. Carrillo. Oxford: Oxford University Press, 49-65.

Bryan, Gharad, Dean Karlan, and Scott Nelson (2010). Commitment Devices. *Annual Review of Economics* 2, 671-698.

Caplin, Andrew and John Leahy (2004). The Social Discount Rate. *Journal of Political Economy* 112(6), 1257-1268.

Cropper, Maureen and David Laibson (1999). The Implications of Hyperbolic Discounting for Project Evaluation. In *Discounting and Intergenerational Equity,* edited by Paul R. Portney and John P. Weyant. Resources for the Future, 1999.

Goldman, Steven Marc (1979). Intertemporally Inconsistent Preferences and the Rate of Consumption. *Econometrica* 47(3), 621-626.

Gul, Faruk and Wolfgang Pesendorfer (2004). Self-Control, Revealed Preference, and Consumption Choice. *Review of Economic Dynamics* 7, 243-264.

Gul, Faruk and Wolfgang Pesendorfer (2008). The Case for Mindless Economics. In *Foundations of Positive and Normative Economics,* edited by Andrew Caplin and Andrew Schotter. Oxford University Press, 2008.

İmrohoroğlu, Ayşe, Selahattin İmrohoroğlu, and Douglas Joines (2003). Time-Inconsistent Preferences and Social Security. *Quarterly Journal of Economics* 118(2), 745–784.

Laibson, David I. (1996). Hyperbolic Discount Functions, Undersaving, and Savings Policy. NBER Working Paper 5635.

Laibson, David I. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics* 112(2), 443-478.

Laibson, David I. (2003). Intertemporal Decision Making. *Encyclopedia of Cognitive Science*. Reference code 708.

Laibson, David I., Andrea Repetto, and Jeremy Tobacman (1998). Self-Control and Saving for Retirement. Brookings Papers on Economic Activity 1, 91-172.

O'Donoghue, Ted and Matthew Rabin (1999). Doing It Now or Later. *American Economic Review* 89(1), 103-124.

O'Donoghue, Ted and Matthew Rabin (2000). The Economics of Immediate Gratification. *Journal of Behavioral Decision Making* 13, 233-250.

O'Donoghue, Ted and Matthew Rabin (2001). Choice and Procrastination. *Quarterly Journal of Economics* 11(1), 121-160.

O'Donoghue, Ted and Matthew Rabin (2003). Self-Awareness and Self-Control. In *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*, edited by George Loewenstein, Daniel Read, and Roy F. Baumeister. Washington DC: Russell Sage Foundation Publications, 217-243.

O'Donoghue, Ted and Matthew Rabin (2007). Incentives and Self Control. In *Advances in Economics and Econometrics: Theory and Applications*, edited by Richard Blundell, Whitney Newey, and Torsten Persson. Cambridge University Press.

Rubinstein, Ariel (2006). Discussion of "Behavioral Economics". In *Advances in Economics and Econometrics: Theory and Application*, Volume 2, edited by Richard Blundell, Whitney K. Newey, and Torsten Persson. New York: Cambridge University Press, 246-254.

Strotz, Robert H. (1956). Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies* 23(3), 165-180.

Tomiyama, Ken (1985). Two-Stage Optimal Control Problems and Optimality Conditions. *Journal of Economic Dynamics and Control* 9(3), 317-337.

Figure 1. Parameter Space where $U(t, \mathbf{c^0}) > U(t, \mathbf{c}^*)$ for all t

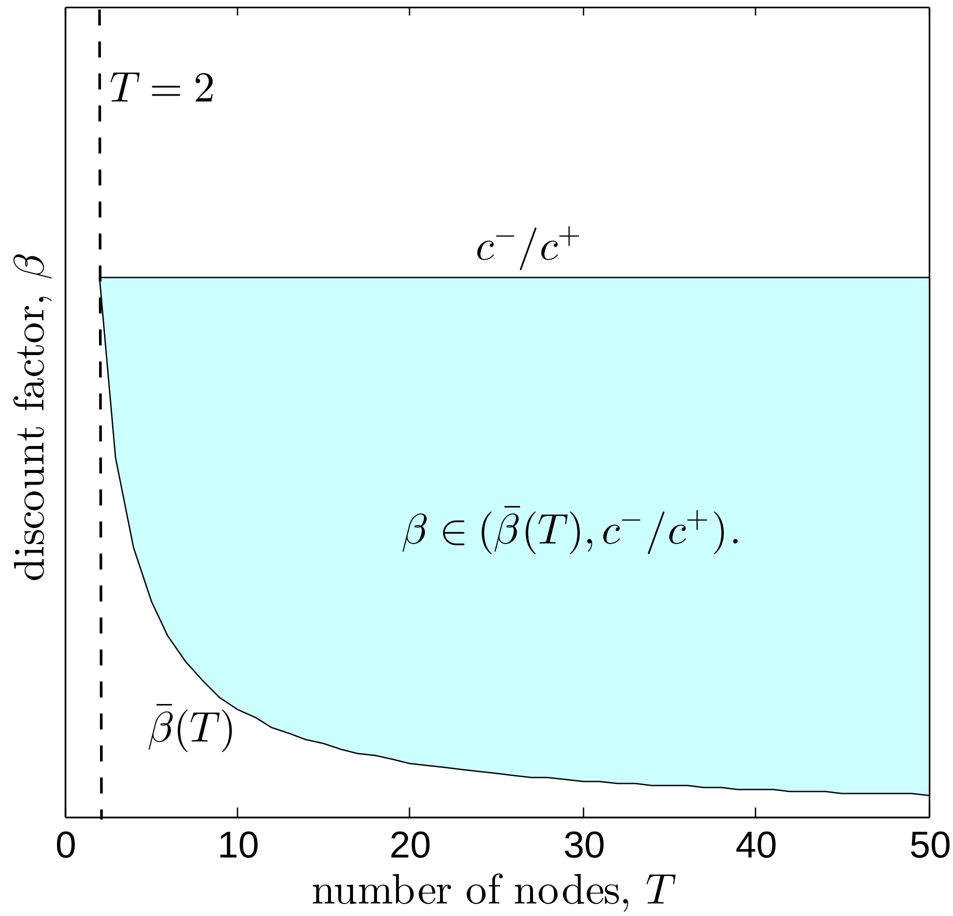

Note: $U(t, \mathbf{c^0}) > U(t, \mathbf{c}^*)$ for all t, if $\beta \in (\bar{\beta}(T), c^-/c^+)$.

Figure 2. 3 Allocations and 2 Decision Nodes ($T = 2$)

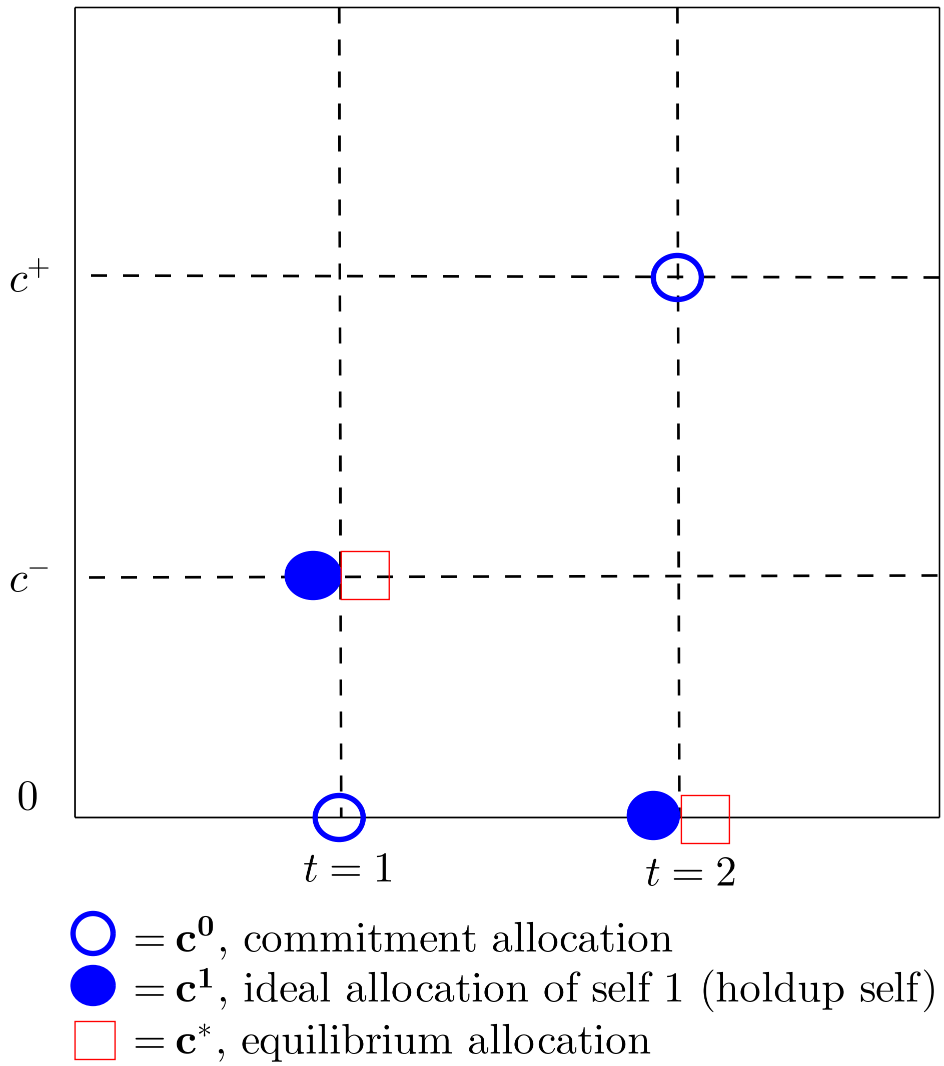

Figure 3. 3 Allocations and 6 Decision Nodes ($T = 6$)

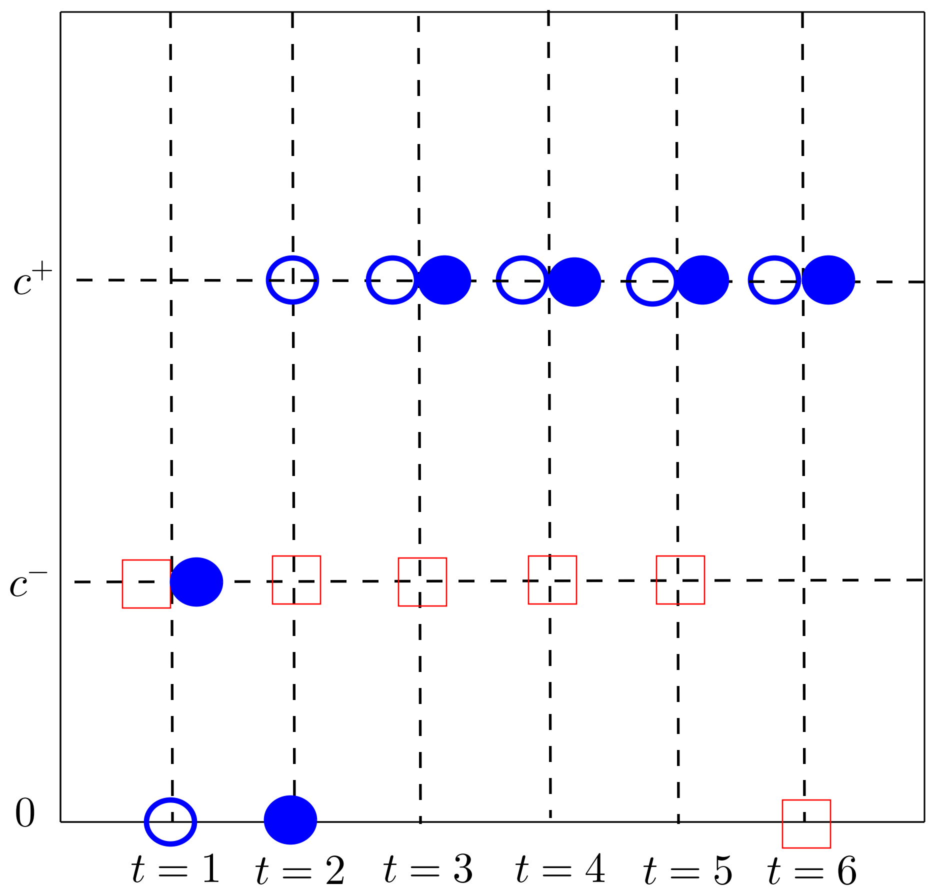

○ $= \mathbf{c^0}$, commitment allocation
● $= \mathbf{c^1}$, ideal allocation of self 1 (holdup self)
□ $= \mathbf{c}^*$, equilibrium allocation

Figure 4. Willingness to Pay for Commitment: The Case of $\beta = 0.5$, $\gamma = 1$

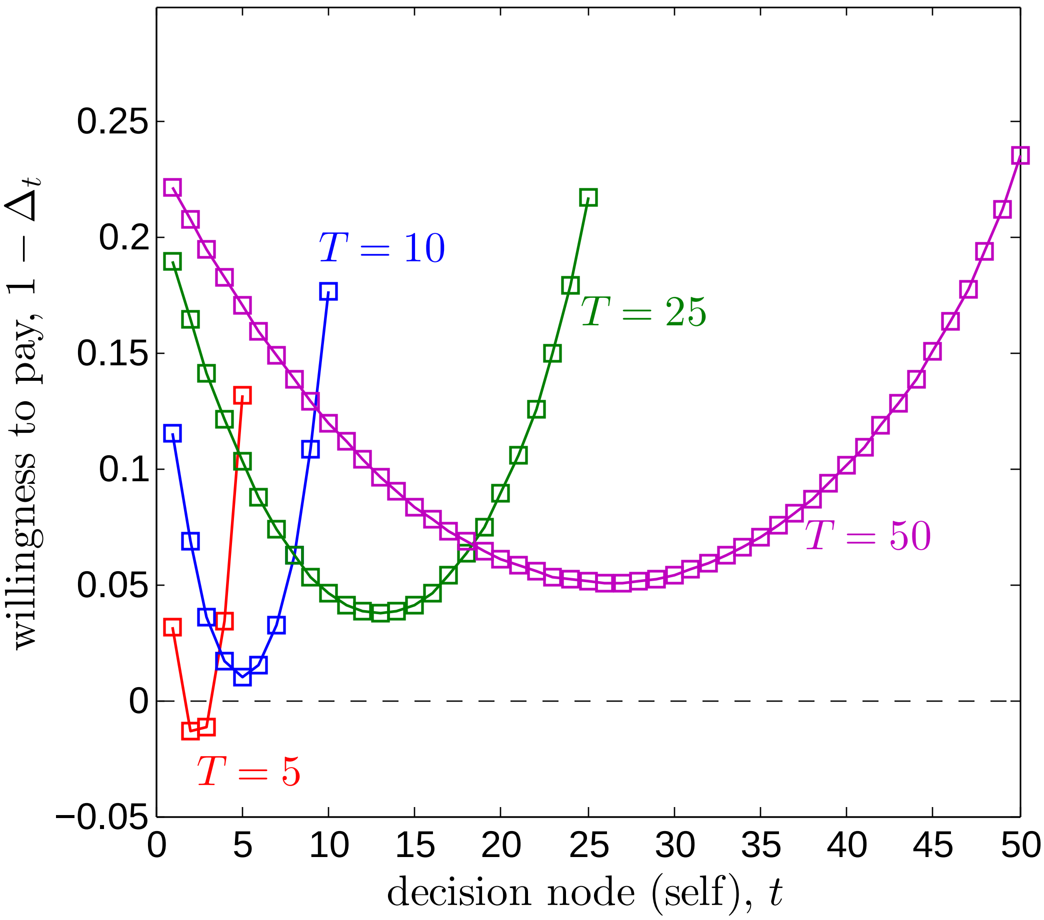

Note: $1 - \Delta_t$ is the fraction of cake self t would give up; β and γ are the forward and backward discount factors.

Figure 5. Willingness to Pay for Commitment: The Case of $\beta = \gamma = 0.5$

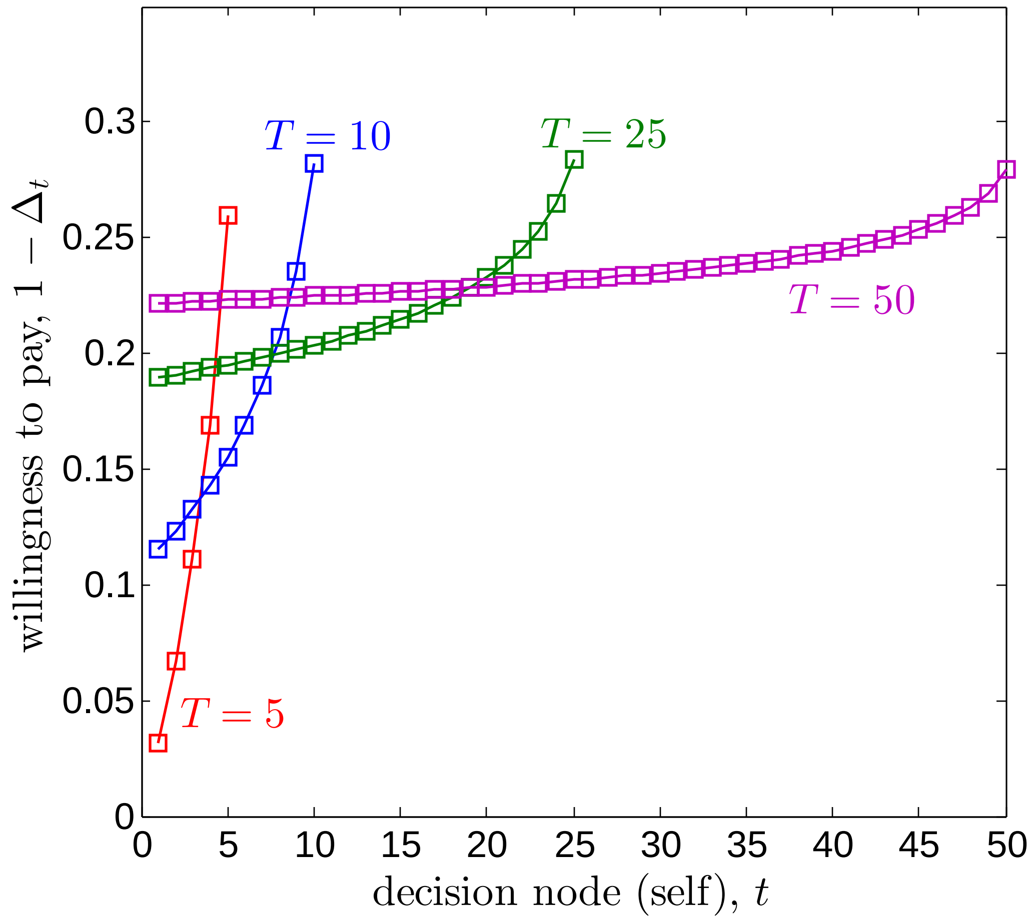

Note: $1 - \Delta_t$ is the fraction of cake self t would give up;
β and γ are the forward and backward discount factors.

Figure 6. Willingness to Pay for Commitment: The Case of $\beta = 0.5$, $\gamma = 0$

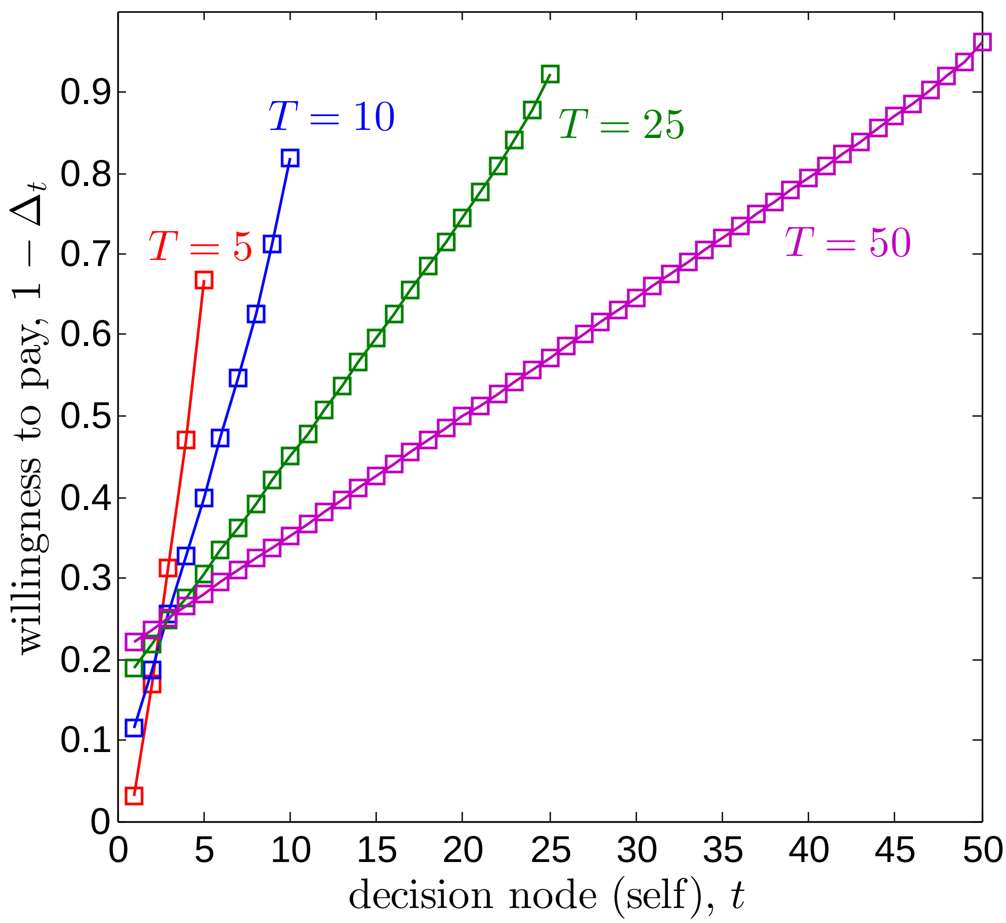

Note: $1 - \Delta_t$ is the fraction of cake self t would give up; β and γ are the forward and backward discount factors.