

Procrastination with Variable Present Bias

NICK GRAVIN, Massachusetts Institute of Technology
NICOLE IMMORLICA, Microsoft Research
BRENDAN LUCIER, Microsoft Research
EMMANOUIL POUNTOURAKIS, Northwestern University

Individuals working towards a goal often exhibit time inconsistent behavior, making plans and then failing to follow through. One well-known model of such behavioral anomalies is present-bias discounting: individuals over-weight present costs by a bias factor. This model explains many time-inconsistent behaviors, but can make stark predictions in many settings: individuals either follow the most efficient plan for reaching their goal or procrastinate indefinitely.

We propose a modification in which the present-bias parameter can vary over time, drawn independently each step from a fixed distribution. Following Kleinberg and Oren (2014), we use a weighted *task graph* to model task planning, and measure the cost of procrastination as the relative expected cost of the chosen path versus the optimal path. We use a novel connection to optimal pricing theory to describe the structure of the worst-case task graph for any present-bias distribution. We then leverage this structure to derive conditions on the bias distribution under which the worst-case ratio is exponential (in time) or constant. We also examine conditions on the task graph that lead to improved procrastination ratios: graphs with a uniformly bounded distance to the goal, and graphs in which the distance to the goal monotonically decreases on any path.

1. INTRODUCTION

Intertemporal tradeoffs – in which an individual incurs short-term costs to achieve long-term goals – have significant implications for health, wealth, and educational outcomes. A student may complete a homework assignment due next week instead of partying with friends. A marathoner may decide to go for a morning training run instead of sleeping in. A skier with a season pass may choose to buy skies instead of renting them each weekend, if it is less costly to do so. Long-term goals like these often involve many sacrifices over time, requiring the individual to form a consistent plan of future action. Anecdotally, however, individuals often have trouble sticking to such plans. This is supported empirically by the observation that experimental subjects discount future payoffs at a non-constant rate. When asked how much money a subject would require in a month/year/decade to offset \$15 today, the median response was \$20/\$50/\$100 [Thaler 1981]. Since these responses are not consistent with any constant discount factor, a subject with these preferences might plan today to invest \$20 next month to gain \$100 in a decade, but then fail to follow through next month.

These observations have led to an extensive line of work in behavioral economics, proposing behavior models that allow divergence between plan and action. The economic theory of *hyperbolic discounting* is one such model developed to explain these time inconsistencies. A hyperbolic discount function assigns weights to future costs.

Author's addresses:

N. Gravin: CSAIL MIT, email: ngravin@mit.edu; N. Immorlica, B. Lucier: Microsoft Research New England, email: brlucier@microsoft.com, nicimm@gmail.com; Emmanouil Pountourakis: EECS Northwestern University, manolis@u.northwestern.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

EC'16, July 24–28, 2016, Maastricht, The Netherlands.

ACM 978-1-4503-3936-0/16/07.

<http://dx.doi.org/10.1145/XXXXXXX.XXXXXXX>

At each point in time, an individual forms a plan which minimizes total costs calculated according to the hyperbolic discounting function. A special case of hyperbolic discounting, developed by Akerloff [Akerlof 1991], is *present-bias discounting*. In its simplest form, present-bias discounting predicts that individuals will inflate any costs incurred at the present moment by a model parameter $b > 1$, and leave all future costs unweighted.

As an example illustrating how present bias can lead to procrastination, consider a student who must complete a homework assignment.¹ The assignment, which is based on today’s lecture, is due next week. The homework takes one night to complete, and there is a cost (in mental effort and missed social opportunities) to complete the assignment immediately after the lecture. Each day that goes by, the student’s recollection of lecture becomes dimmer, and the cost for completing the assignment increases. In this scenario, it is clearly optimal, in terms of minimizing induced costs, to complete the homework assignment the day of the lecture. However, if the student inflates present-day costs, he may *perceive* that it is less costly to complete the assignment tomorrow than today. Facing a similar decision tomorrow, he again perceives that procrastinating is less costly than completing the assignment immediately. This procrastination behavior persists, and the student doesn’t complete the assignment until the last day, spending substantially more than the optimal plan.

The preceding example contains a very stark prediction: the student will either procrastinate indefinitely, if present-biased, or not at all. As noted in prior work of Kleinberg and Oren [Kleinberg and Oren 2014], this dichotomy persists even in the face of exponentially increasing costs, causing a present-biased student to spend exponentially more than necessary. In reality, one should expect a more nuanced behavior – any given individual might procrastinate for some amount of time, but eventually even the most lackadaisical students will exert effort to avoid future penalties. Such variation in procrastination behavior can arise from variability in the present-bias parameter. In this work, we assume that the agent has a present-bias parameter b_i drawn on each day i independently at random from a distribution \mathcal{F} supported on $[0, 1)$. For the preceding example, if the student’s present bias parameter has a constant probability of taking value 1, then the student will delay for only a constant number of days, in expectation, before completing their assignment. We ask: how does variability in present-bias change procrastination behavior? Are there structured choices, e.g., homework late policies, that guide individuals away from costly procrastination?

To study these questions, we adopt a graph-theoretic framework for task completion proposed by Kleinberg and Oren [Kleinberg and Oren 2014]. In this framework, an agent with a goal n days in the future traverses through a weighted task graph G starting from node s and ending at node t . On each day i , if the agent’s current present bias is b_i and current state is v , then the weight of the first edge along a $v - t$ path is multiplied by an extra factor of b_i , while all the remaining edges are evaluated according to their true weights. The agent chooses a minimal-weight path, under this distortion, and follows the first edge of that path. For a given task graph G and present-bias distribution \mathcal{F} , the *procrastination ratio* is the ratio between the expected total weight of the traversed path to the weight of the initial shortest path. We are interested in bounding the procrastination ratio as a function of n , the number of days until the deadline for the task.²

¹This example is a minor reformulation of an example due to Akerloff.

²Kleinberg and Oren [Kleinberg and Oren 2014] bounded the procrastination ratio with respect to the total number of nodes in the graph, rather than the path length n , but these are essentially equivalent; see Section 2.

Under this model, we can state our question more concretely: are there natural conditions that guarantee a low procrastination ratio? Recall that in the student example, if there is a constant probability of bias $b_i = 1$, the student will complete the assignment in a constant number of days in expectation. As it turns out, this is not sufficient to guarantee a sub-exponential procrastination ratio (see Section 2 for more details). Our first result shows that, for any distribution, the graph with maximal procrastination ratio has exactly this form: on any day, the agent may either complete the task for a growing cost or procrastinate for free. The proof of this result involves a connection to optimal pricing theory: the problem of constructing a worst-case graph can be reduced to the problem of designing a revenue-optimal auction. For example, the fact that on any day there is only one costly edge is a consequence of the fact that the optimal pricing menu for a single-parameter agent has a single deterministic option [Myerson 1981].

We can leverage this description of the worst-case task graph to develop bounds on the procrastination ratio as a function of the present-bias distribution. In optimal pricing theory, a special subclass of Pareto distributions, often referred to as *equal-revenue* distributions, are useful for worst-case examples. For a given present-bias distribution, we calculate the “smallest” equal-revenue distribution that stochastically dominates it. As worst-case procrastination ratios grow with stochastic dominance, this gives us an upper bound on the worst-case procrastination ratio of any given present-bias distribution. One can similarly derive lower bounds for any distribution that is *not* dominated by a particular equal-revenue distribution. One implication of our analysis is that present-bias distributions can be roughly divided into two categories: those with light tails tend to have low (linear or even constant) worst-case procrastination ratios, whereas heavy-tailed distributions tend to have worst-case procrastination ratios that are exponential in n .

These worst-case examples lead us to ask whether there are natural conditions on the task graph, as well as the present-bias distribution, that lead to smaller procrastination ratios. Intuitively, the key feature of the worst-case graph discussed above is that sub-optimal planning can cause an agent to not only incur extra costs today, but also reach a state from which the agent is strictly worse off; i.e., the cost of completing the task can be higher than it was initially. We say that a graph has the *bounded distance property* if the weight of the shortest path from any node v to the target t is at most the weight of the shortest path from the initial node s to t . This condition captures scenarios, like training for a marathon, in which each day of training improves preparedness whereas procrastination diminishes it, but not below the initial base level. We show for any distribution \mathcal{F} , the procrastination ratio of a bounded shortest-path graph is at most linear in n , and that this bound is tight even for distributions that have a constant probability of $b_i = 1$.

We next consider a stronger condition under which the procrastination ratio is constant. We say a graph has the *monotone distance property* if, for any $s - t$ path ($s = v_0, v_1, \dots, v_{k-1}, t = v_k$), the shortest path from v_i to t is decreasing in i . Roughly speaking, this condition captures scenarios in which progress made by an individual is not lost: regardless of an agent’s action in the present round, the total cost required to complete the task has not increased. For example, consider a skier deciding whether to rent or buy skies each weekend of the skiing season: as the cost of buying skies remains constant over time, the total cost required to ski the remainder of the season does not grow. We show that if the distribution over present-bias parameters has sufficient mass at biases close to 1,³ then the procrastination ratio of graphs with the

³See Section 5.2 for a formal statement of the condition on the distribution. For example, the condition is satisfied whenever there is a positive probability that the bias is equal to 1.

monotone distance property is bounded by a constant. Moreover, the condition on the present-bias distribution is necessary for this result: if the agent has constant bias $b > 1$, there exist monotone-shortest-path graphs for which the procrastination ratio is $\Omega(n)$.⁴

1.1. Related work

Behavioral Economics. Behavioral economics and game theory study models that can explain anomalies in human behavior which are consistently observed in experimental data, see for example [Colin 2003]. Our work adopts the model of [Akerlof 1991] which aims to explain individual’s procrastination that violates classical assumption of utility-maximizing individual’s behavior. Closely related to procrastination are issues of abandonment of long-time projects [O’Donoghue and Rabin 2008] and benefit of imposing a deadline in the context of task completion [Ariely and Wertenbroch 2002]. Other examples include models addressing attentiveness issues, e.g., where reduction of choice among the options available to an agent [O’Donoghue and Rabin 1999], or delayed notification messages [Ely 2015] may improve agent’s performance.

Some behavioral models study various levels of agents’ rationality. For example Kaur, Kremer, and Mullainathan [Kaur et al. 2010] study sophisticated customers who are aware of their possible future procrastination, and Wright and Leyton-Brown [Wright and Leyton-Brown 2010, 2012, 2014] model agents’ behavior via quantal cognitive hierarchies in single-round simultaneous-move games where agents are empirically observed not to follow Nash equilibrium strategies.

Hyperbolic Discounting. A significant amount of work in economics literature has been devoted to the study of *quasi-hyperbolic discounting*, see [Frederick et al. 2002] for a survey. This form of discounting function generalizes Akerlof’s model by modeling agent’s behavior with two parameters: present-bias factor b and discounting factor $\delta \leq 1$. In this model the agent at every step scales up the immediate cost by b , and discounts every cost t steps away into the future by a factor of δ^t . This form of discounting has been used in many areas including addiction [Beshears et al. 2015], self-control [Laibson 2015], and health and pre-retirement savings [Choi et al. 2004].

Graphical Model. Our work uses the model of procrastination proposed by Akerlof [Akerlof 1991] and adapts the graphical framework of Kleinberg and Oren [Kleinberg and Oren 2014]. The latter work studies graphical properties of the network which may cause high procrastination rates, and identifies a graph-minor condition on the task graph that implies a bound on the procrastination ratio. The follow-up work of Tang et al. [Tang et al. 2014] provides improved bounds on the procrastination ratio, again as a function of this graph-minor condition.

Techniques. We model an individual as an agent performing graph traversal under uncertainty. Related models also appear in the literature on route planning under uncertainty [Nikolova and Karger 2008], or risk-aversion in routing and congestion games [Nikolova and Moses 2011, 2014, 2015] but under different objectives.

Some of our technical results exploit a connection between agent behavior under hyperbolic discounting and auction theory. In particular, we use tools characterizing the space of optimal single-item auctions under a distribution of agent values, as explored by Myerson [Myerson 1981].

⁴Akerlof’s original example of mailing a package is such an example: it satisfies the monotone distance property and has linear procrastination ratio. See Kleinberg and Oren [Kleinberg and Oren 2014] for a graph-theoretic formulation of that example.

2. MODEL

We consider a setting in which an agent must complete a task over a period of n days. We track the progress of the agent towards task completion using states of intermediate progress, which are nodes of a task graph $G = (\cup_{i=1}^{n+1} V_i, E)$ with edge weights $w : E \rightarrow \mathbb{R}^+$. The set of nodes V_i are the possible states of the agent on day i . We assume without loss of generality that $V_1 = \{s\}$, the *start* node, and $V_{n+1} = \{t\}$, the *end* node. All edges occur between nodes on consecutive days, i.e., $v_i v_j \in E$ if and only if $v_i \in V_i$ and $v_j \in V_{i+1}$. The weights represent the cost of transitioning between states: i.e., an agent in state v_i wishing to transition to a state v_{i+1} incurs a cost of $w(v_i v_{i+1})$. The agent starts from state $s = v_1$ on day 1, follows a path $(v_1 v_2, \dots, v_n v_{n+1})$ in G , and ends in state $t = v_{n+1}$ on day $n + 1$, for a total cost of $\sum_{i=1}^n w(v_i v_{i+1})$.⁵ For convenience we will refer to the shortest path between two nodes v_i and v_j in G as the distance, denoted $d(v_i, v_j)$. Thus the minimum cost way to complete the task is $d(s, t)$.

We are interested in the inefficiency of the traversed path due to the variable present-bias of the agent. As such, for each day i we introduce a present-bias factor b_i drawn i.i.d. from a given distribution \mathcal{F} supported on $[1, \infty)$. When calculating costs on day i , the agent inflates costs incurred on that day by b_i while leaving remaining costs unweighted. Thus an agent currently in state v_i on day i will choose a transition $v_i v_{i+1}$ minimizing

$$b_i \cdot w(v_i v_{i+1}) + d(v_{i+1}, t).$$

Note that on the following day $i+1$ the agent again chooses an edge using a (potentially different) present-bias factor b_{i+1} , and may therefore choose a state v_{i+2} different from his planned path on day i .

We consider the ratio of the weight of the chosen path $v_1 v_2, \dots, v_n v_{n+1}$ with the weight of the shortest $s - t$ path:

$$\frac{\sum_{i=1}^n w(v_i v_{i+1})}{d(s, t)}.$$

The *procrastination ratio* is then the expectation (over the choice of the present bias factors) of this ratio.

In summary, an instance of our problem consists of a weighted task graph G and present-bias distribution \mathcal{F} . We would like to understand the structure of worst-case instances, i.e., the distributions and tasks graphs that achieve the maximum procrastination ratio. We would also like to describe natural classes of instances with small procrastination ratios. To this end, it is useful to introduce the following restrictions on task graphs.

Definition 2.1 (Bounded Distance). A weighted task graph G satisfies *bounded distance*, if $\forall i$, and $\forall v_i \in V_i$, $d(v_i, t) \leq d(s, t)$.

Intuitively, the bounded distance property captures scenarios in which one’s current state, in terms of the optimal cost to reach the goal, is never worse than the initial state. For example, this includes scenarios in which “starting over” is always a free and feasible option: i.e., from each vertex there is a 0-cost edge to a vertex whose distance from t is at most $d(v_i, t)$.

Definition 2.2 (Monotone Distance). A weighted graph G satisfies *monotone distance* if $\forall i$ and $\forall v_i v_{i+1} \in E$ with $v_i \in V_i, v_{i+1} \in V_{i+1}$, $d(v_i, t) \geq d(v_{i+1}, t)$.

⁵Note we do not assume the agent completes the task; task abandonment can be represented by having a costly edge from a state $v_n \in V_n$ to t .

This captures scenarios where the agent always makes progress towards the end state, i.e., no transition causes the agent to lose any of his prior accomplishments. We note that the latter restriction implies the former one, in the sense that any graph with monotone distances also has bounded distances.

Remark. Kleinberg and Oren [Kleinberg and Oren 2014] studied task planning in generic weighted directed acyclic graphs $G' = (V', E')$ with a designated start node s and end node t , and presented the procrastination ratio as a function of the number of nodes in G' , say N . To transform their setup to ours, simply add a self-loop to t of weight 0, create N sets V_i where V_i contains copies of all nodes $v \in V'$ that can be reached from s via a path of exactly $i - 1$ edges in G' , and create weighted edges between nodes of V_i and V_{i+1} if such edges exist in E' . Note that under this transformation, the number of steps n is at most N , and the total number of vertices in the constructed graph is $O(N^2)$.

3. EXAMPLES

To illustrate the model and definitions introduced in Section 2, we work through instantiations of the scenarios discussed in the introduction. Throughout the course of this discussion, we will assume an agent with a present-bias factor b_i drawn from distribution \mathcal{F} which takes value 1 with probability $1/3$ and value 3 with probability $2/3$. Note that this distribution is “close to rational” in the sense that the agent has a constant probability of $b_i = 1$. In each example, we will normalize the shortest $s - t$ path to be 1 and so the procrastination ratio is just the expected cost induced by the agent.

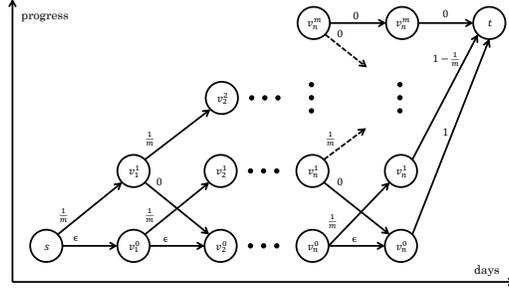
3.1. Students Preparing Homework

Our first scenario in the introduction concerned a student assigned homework, where the student’s memory of the requisite material diminishes as time passes. Suppose the homework is assigned on day 1 and due on day $n + 1$. We construct a task graph $G = (\cup_{i=1}^{n+1} V_i, E)$ with $V_1 = \{s_1\}$ and $V_{n+1} = \{t_{n+1}\}$. On each day $i \in 2, \dots, n$, there are two possible states s_i and t_i corresponding to whether the student completed the assignment (s_i) or not (t_i). The edges are: $E = \{s_i s_{i+1} | i = 1 \dots n - 1\} \cup \{t_i t_{i+1} | i = 2 \dots n\} \cup \{s_i t_{i+1} | i = 1 \dots n\}$. The cost of transitioning from s_i to s_{i+1} or from t_i to t_{i+1} is zero. The cost of transitioning from s_i to t_{i+1} is $w_{s_i t_{i+1}} = 2^{i-1}$. This models the idea that the mental effort to complete the assignment grows by a factor of 2 each day as the student’s recollection of the material from lecture becomes dimmer. See Figure 1(b) for a representation of this task graph.

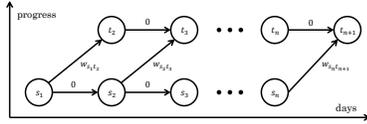
Note that this task graph does not have the bounded distance property. Accordingly, the procrastination ratio can be exponential, and indeed is exponential for the distribution \mathcal{F} . To see this, note on day i , the perceived cost of completing the assignment that day is $b_i \cdot 2^{i-1}$ while the cost of waiting to complete the homework on the next day is 2^i . Thus, if the instantiation of b_i is greater than 2, then the student will decide to complete the assignment on the following day. For the present-bias distribution \mathcal{F} described above, this event occurs with probability $2/3$, and so the total expected cost incurred by the student is $\sum_{i=1}^{n+1} (2/3)^i 2^{i-1}$, which is exponential in n .

3.2. Runners Training for a Marathon

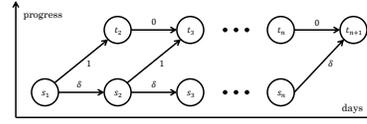
Our second scenario in the introduction concerned a runner training for a marathon. The runner’s fitness diminishes each day she sleeps in instead of going for a training run. Suppose the runner starts training on day 1 and the marathon is scheduled for day $n + 1$. The task graph $G = (\cup_{i=1}^{n+1} V_i, E)$ has $V_1 = \{s\}$ and $V_{n+1} = \{t\}$. The runner has a fitness level in $\{0, \dots, m\}$ where 0 is the base-level fitness and m is peak fitness. If the runner is at fitness level $j < m$ on day $i < n$, then the runner can pay $\frac{1}{m}$ to train



(a) Training for a marathon.



(b) Completing homework.



(c) Renting skis.

Fig. 1. Task completion graphs for the examples in Section 3.

and increase fitness to level $j + 1$ on day $i + 1$. Alternatively, if the runner is at fitness level $j > 0$ on day $i < n$, then the runner can sleep in, paying 0, and drop to fitness level $j - 1$ on day $i + 1$. We assume the runner can maintain fitness level m for 0 cost and level 0 at a small cost $\epsilon < \frac{2}{3m}$.⁶ On day n , if the runner is at fitness level j , then she pays $1 - \frac{j}{m}$ to run the marathon. See Figure 1(a) for a representation of this task graph.

Note that this task graph has the bounded distance property. As we show in Section 5.1, such graphs have at most a linear procrastination ratio, for any distribution over present-bias parameters. For $b_i \sim \mathcal{F}$, this task graph nearly achieves this bound. To see this, note that on day $i < n$, if the runner is at fitness level $0 < j < m$, then she perceives training as costing her $b_i \cdot \frac{1}{m} + (1 - \frac{j+1}{m})$ and not training as costing her $1 - \frac{j-1}{m}$. Therefore, she trains when $b_i = 1$ and sleeps in when $b_i = 3$ (ignoring boundary conditions at fitness level 0). This is a bounded random walk on the integers from $\{0, \dots, m\}$ with an absorbing state at m . The transition probabilities are biased toward lower fitness levels, and so the runner will spend most of her time transitioning between the low states, leading to a total cost of $\Theta(n/m)$ so long as m is sufficiently large that she never hits the absorbing state. Taking $m = \Omega(\log n)$ suffices, and this leads to a bound of $\Theta(n/\log(n))$. In Section 5.1, we show how to improve this example to get a linear bound for a broad class of present-bias distributions.

⁶At fitness level m , running ceases to become a chore. At fitness level 0, lack of exercise causes remorse and/or disutility from a low fitness level.

3.3. Skiers Renting or Buying Skies

Our third scenario in the introduction concerned a skier debating whether to rent or buy skies. The skier buys a season pass on day 1 and commits to skiing every weekend for the next n weeks. We construct a task graph $G = (\cup_{i=1}^{n+1} V_i, E)$ with $V_1 = \{s\}$ and $V_{n+1} = \{t\}$. On each day $i \in 2, \dots, n$, there are two possible states, **rent** and **own**, corresponding to whether the skier uses rental equipment or his own equipment. Transitioning from a **rent** state (or s) to an **own** state on the following weekend costs 1, the cost of buying skies. Transitioning from a **rent** state (or s) to another **rent** state (or t) costs δ , the cost of renting skies. Once the skier is in the **own** state, all future transitions are free. See Figure 1(c) for a representation of this task graph.

Note that this task graph has the monotone distance property. As we show in Section 5.2, if the present-bias is “close enough” to 1 sufficiently often, such graphs have a constant procrastination ratio. Indeed, for $\delta < 2/3$, we can bound the procrastination ratio by noting that the skier will choose to buy when $b_i = 1$ (except possibly towards the very end of the season) and rents otherwise. The cost is then at most $\sum_{i=1}^n (\frac{2}{3})^i (\delta \cdot i + 1)$, which is a constant. Note that if the present-bias was bounded away from 1, e.g., a point mass at 3, then the skier would rent for the entire season and pay a linear cost.

4. WORST-CASE PROCRASTINATION

For a fixed present-bias distribution, we would like to characterize the worst-case procrastination ratio over all task graphs with n days. To this end, we first derive the form of a task graph that achieves the worst-case procrastination ratio for an arbitrary fixed present-bias distribution. We then develop a natural parameterization of distributions and bound this worst-case ratio as a function of the distribution parameters.

4.1. Worst-Case Task Graphs

Given present-bias distribution \mathcal{F} , we will show that there is a worst-case task graph with a simple form reminiscent of that from the student/homework example. On each day there are only two possible states: the uncompleted state and the completed state. Moving between uncompleted states or completed states has zero cost while moving from an uncompleted state to a completed one has a cost which is a function of the distribution.

The proof uses the following theorem from *optimal pricing theory*: A buyer has a value for an item, drawn from a distribution, and faces a menu of options, each consisting of a probability of allocation and a price.⁷ The buyer, acting to maximize his expected utility (defined as value times allocation probability minus price), induces an expected allocation and price of the menu. The theorem states that for any linear function of expected allocation and price, there is a simple *single posted price* menu that optimizes that function.

THEOREM 4.1 ([MYERSON 1981]). *Let $M \subseteq [0, 1] \times \mathbb{R}$ be such that $(0, 0) \in M$. For any such M and $v \in \mathbb{R}^+$, write $(x_M(v), p_M(v)) = \operatorname{argmax}_{(x,p) \in M} [vx - p]$. Then for all distributions \mathcal{F} with support contained in \mathbb{R}^+ , and for all $\alpha, \beta \in \mathbb{R}$, there exists a p^* and corresponding $M^* = \{(0, 0), (1, p^*)\}$ such that*

$$\alpha E_{v \sim \mathcal{F}}[x_M(v)] + \beta E_{v \sim \mathcal{F}}[p_M(v)] \leq \alpha E_{v \sim \mathcal{F}}[x_{M^*}(v)] + \beta E_{v \sim \mathcal{F}}[p_{M^*}(v)].$$

We will construct the worst-case task graph recursively. From a node v_i , the subproblem requires us to construct neighbors v_{i+1} , edge weights $w(v_i v_{i+1})$, and the weights of the shortest paths $d(v_{i+1} t_{n+1})$ from these neighbors to the end state, all subject to

⁷The menu must contain the option of not buying for a cost of zero, i.e., the buyer is allowed to walk away.

the constraint that the weight of the shortest path $d(v_i, t_{n+1})$ from v_i to the end state is preserved. We solve this subproblem by relating the ratio $w(v_i v_{i+1})/d(v_i, t_{n+1})$ to the allocation probability and $d(v_{i+1}, t_{n+1})/d(v_i, t_{n+1})$ to the price of a menu in pricing theory.

THEOREM 4.2. *Given a present-bias distribution \mathcal{F} , there is a task graph over n days, $G_{\mathcal{F}}^n = (\cup_{i=1}^{n+1} V_i, E)$, achieving maximum procrastination ratio with the following form. Let $V_1 = \{s_1\}$, $V_{n+1} = \{t_{n+1}\}$, and $d(s_1, t_{n+1}) = 1$. Then*

- for $i \in 2, \dots, n$, $V_i = \{s_i, t_i\}$,
- for $i \in 1, \dots, n-1$, $w(s_i s_{i+1}) = 0$,
- for $i \in 2, \dots, n$, $w(t_i t_{i+1}) = 0$,
- and for $i \in 1, \dots, n$, $w(s_i t_{i+1}) = d(s_i, t_{n+1}) \leq d(s_{i+1}, t_{n+1})$,

where $d(s_i, t_{n+1})$ is chosen as a function of \mathcal{F} .

PROOF. First note the assumption $d(s_1, t_{n+1}) = 1$ is without loss of generality as we can always normalize all weights by the initial shortest path without changing the procrastination ratio. We prove the theorem by constructing a worst-case task graph recursively. In the i 'th iteration, for $i = 1 \dots n-1$, we will construct:

- the nodes of set V_{i+1} ,
- the weights of edges from V_i to V_{i+1} ,
- and the shortest continuation path weights $d(v_{i+1}, t_{n+1})$ for $v_{i+1} \in V_{i+1}$,

to maximize the expected procrastination ratio of the chosen path subject to constraints on the shortest paths $d(v_i, t_{n+1})$ for $v_i \in V_i$.

We first describe the behavior of the agent at node v_i for a fixed set of transition options V_{i+1} consisting of weighted edges $w(v_i v_{i+1})$ and continuation distances $d(v_{i+1}, t_{n+1})$ for $v_{i+1} \in V_{i+1}$. An agent with present-bias parameter b_i selects transition v_{i+1}^* minimizing $b_i w(v_i v_{i+1}) + d(v_{i+1}, t_{n+1})$. This minimization problem is equivalent to the following maximization problem, as can be seen by first changing the sign of the objective, then adding the constant $b_i \cdot d(v_i, t_{n+1})$, and then scaling by the constant $1/d(v_i, t_{n+1})$:

$$\max_{v_{i+1} \in V_{i+1}} \left(b_i \left(1 - \frac{w(v_i v_{i+1})}{d(v_i, t_{n+1})} \right) - \frac{d(v_{i+1}, t_{n+1})}{d(v_i, t_{n+1})} \right).$$

The above problem is the optimization problem faced by a buyer in a pricing menu where: the value of the buyer is b_i , and the allocation and price of menu item v_{i+1} are $x(v_{i+1}) = 1 - \frac{w(v_i v_{i+1})}{d(v_i, t_{n+1})}$ and $p(v_{i+1}) = \frac{d(v_{i+1}, t_{n+1})}{d(v_i, t_{n+1})}$, respectively. We argue this menu satisfies the conditions of Theorem 4.1. As $d(v_i, t_{n+1})$ is the weight of the shortest $v_i - t_{n+1}$ path, there must be an option v_{i+1} with $w(v_i v_{i+1}) \leq d(v_i, t_{n+1})$. Furthermore, any option v'_{i+1} with $w(v_i v'_{i+1}) > d(v_i, t_{n+1})$ will never be preferred to v_{i+1} and so we can eliminate such nodes from the graph. Therefore, the allocations $x(v_{i+1})$ are in $[0, 1]$. Furthermore, we can assume without loss of generality that there is a transition $t_{i+1} \in V_{i+1}$ with $w(v_i t_{i+1}) = d(v_i, t_{n+1})$ and $d(t_{i+1}, t_{n+1}) = 0$ as we are required to include a transition option with $w(v_i v_{i+1}) + d(v_{i+1}, t_{n+1}) = d(v_i, t_{n+1})$ and any such option weakly dominates t_{i+1} for any bias factor. This proves that the menu contains an option t_{i+1} with $x(t_{i+1}) = 0$ and $p(t_{i+1}) = 0$.

To apply Theorem 4.1, we must argue that the objective of maximizing the procrastination ratio is linear in the expected allocation and pricing of any menu V_{i+1} defined above. Let $\text{OPT}(i, d)$ denote the expected cost of the agent in the worst-case task graph with $n-i$ days and shortest path weight d . We can write $\text{OPT}(i, d(v_i, t_{n+1}))$, the worst-

case expected cost conditioned on being at v_i , recursively as:

$$\max_{V_{i+1}} \left[\mathbf{E}_{b_i \sim \mathcal{F}} [w(v_i v_{i+1}^*)] + \mathbf{E}_{b_i \sim \mathcal{F}} [\text{OPT}(i+1, d(v_{i+1}^*, t_{n+1}))] \right],$$

where v_{i+1}^* is the choice of the agent with bias b_i facing options V_{i+1} , and the base of the recursion $\text{OPT}(n, d) = d$. Using the linearity of $\text{OPT}(i, d)$ in its second argument and the menu pricing notation introduced in the preceding paragraph, we see the optimization problem is equivalent to:

$$\max_{V_{i+1}} \left(1 - E_{b_i \sim \mathcal{F}}[x(v_{i+1}^*)] + E_{b_i \sim \mathcal{F}}[p(v_{i+1}^*)] \text{OPT}(i+1, 1) \right) \cdot d(v_i, t_{n+1}). \quad (1)$$

Applying Theorem 4.1, we conclude that there is an optimal menu $V_{i+1} = \{s_{i+1}, t_{i+1}\}$ where $x(s_{i+1}) = 1$, $p(s_{i+1}) = p^* \geq 1$ (as the value $b_i \geq 1$), $x(t_{i+1}) = 0$, and $p(t_{i+1}) = 0$. Given a partially constructed task graph with $V_i = \{s_i, t_i\}$, we can therefore optimally extend it by defining $V_{i+1} = \{s_{i+1}, t_{i+1}\}$. The above discussion shows that the transitions from s_i to V_{i+1} should have edge weights $w(s_i s_{i+1}) = 0$ and $w(s_i t_{i+1}) = d(s_i, t_{n+1})$ and continuation shortest path weights $d(s_{i+1}, t_{n+1}) \geq d(s_i, t_{n+1})$ and $d(t_{i+1}, t_{n+1}) = 0$. To satisfy the constraint that $d(t_i, t_{n+1}) = 0$ we also add an edge from t_i to t_{i+1} of weight zero. This completes the i 'th iteration. To complete the construction, add edges $s_n t_{n+1}$ and $t_n t_{n+1}$ with weights $d(s_n, t_{n+1})$ and 0, respectively. \square

4.2. Bounding the Procrastination Ratio for a Bias Distribution

In Section 4.1 we characterized the task graph that maximizes the procrastination ratio for any given present-bias distribution. In this section we leverage this characterization to explore how properties of the present-bias distribution impacts the worst-case procrastination ratio. In particular, we are interested in determining features of the bias distribution that imply sub-exponential, or even constant, bounds on the procrastination ratio.

We first note that worst-case procrastination ratio can only be greater for distributions that generate higher present-bias parameters. Recall that distribution $\bar{\mathcal{F}}$ *stochastically dominates* distribution \mathcal{F} if, for all x , $\Pr_{b \sim \bar{\mathcal{F}}}[b > x] \geq \Pr_{b \sim \mathcal{F}}[b > x]$. That is, $\bar{\mathcal{F}}(x) \leq \mathcal{F}(x)$ for all x .

LEMMA 4.3. *Given two distributions \mathcal{F} and $\bar{\mathcal{F}}$ such that $\bar{\mathcal{F}}$ stochastically dominates \mathcal{F} , the worst-case procrastination ratio for $\bar{\mathcal{F}}$ is at least as large as the worst-case procrastination ratio for \mathcal{F} .*

The proof of Lemma 4.3 appears in Appendix 4.

Lemma 4.3 motivates us to categorize distributions into classes, where all distributions in a certain class stochastically dominate a family of distributions. Such a categorization, along with Lemma 4.3, will allow us to derive bounds on the procrastination ratio of any given distribution. Our bound will be with respect to the following parametrization of present-bias distributions, which is essentially the largest value z such that the distribution stochastically dominates the equal-revenue distribution with revenue z .⁸

Definition 4.4. Given a distribution \mathcal{F} with support contained in $[1, \infty)$, write $z(\mathcal{F}) \equiv \max(1 - \mathcal{F}(b))b$. Given $z \in (0, \infty)$, we will write $\mathfrak{F}(z)$ for the family of distributions \mathcal{F} with $z(\mathcal{F}) = z$.

Note that any distribution \mathcal{F} with bounded support belongs to some family $\mathfrak{F}(z)$, potentially with $z = \infty$. For $\mathcal{F} \in \mathfrak{F}(\infty)$, i.e., \mathcal{F} has infinite support and $\mathcal{F} \notin \mathfrak{F}(z)$ for any

⁸The equal-revenue distribution with revenue z satisfies $F(x) = 1 - \frac{z}{x}$ for all $x \geq z$.

$z > 0$, then there are task graphs for \mathcal{F} with procrastination ratio growing faster than any exponential function in n . We will largely focus on distributions for which $z(\mathcal{F})$ is finite.

To build some intuition for the classes $\mathfrak{F}(z)$, we give below a few examples.

- (1) The uniform distribution over the interval $[1, 3]$ belongs to the family $\mathfrak{F}(1.125)$, with $(1 - \mathcal{F}(1.5)) \cdot 1.5 = 1.125$.
- (2) The uniform distribution over the interval $[1, 2]$ belongs to the family $\mathfrak{F}(1)$.
- (3) Suppose \mathcal{F} is the upper half of a normal distribution with mean 1. That is, $b = \max(\xi, 1)$ where $\xi \sim N(1, 1)$ is a normal random variable with mean and normal deviation equal to 1. In this case, we can maximize $b(1 - \mathcal{F}(b))$ numerically to find that $\mathcal{F}(b) \in \mathcal{F}(z)$ for $z \approx 0.507$.
- (4) Suppose $\mathcal{F}(x) = 1 - \frac{1}{2\sqrt{x}}$ for $x \in [1, 100)$, and $\mathcal{F}(100) = 1$. This is a heavy-tailed distribution, which has an atom at $b = 1$ with probability $\frac{1}{2}$. This distribution belongs to $\mathfrak{F}(5)$.

We are now ready to bound the worst-case procrastination ratio for any given distribution \mathcal{F} , as a function of $z(\mathcal{F})$. The following two theorems tell us whether a given distribution has an exponential, linear, or constant procrastination ratio, in the worst case over task graphs. Theorem 4.5 provides an upper bound on the procrastination ratio, which is tight for the family $\mathfrak{F}(z)$. Theorem 4.6 provides a lower bound for any given distribution. The proofs appear in Appendix 4.2.

THEOREM 4.5. *Let $\mathcal{F} \in \mathfrak{F}(z)$ for some $z > 0$. Then the procrastination ratio of any task graph for \mathcal{F} is at most $\sum_{i=0}^{k-1} z^i$ and this bound is tight.*

THEOREM 4.6. *For arbitrary \mathcal{F} , if $(1 - \mathcal{F}(b_0))b_0 = z$ for some point $b_0 > 1$, then there is a task graph with procrastination ratio $\sum_{i=1}^{n-1} (1 - 1/b_0) \cdot z^{i-1} + z^{n-1}$.*

With these results in hand, we can bound the worst-case competitive ratios for the examples described above:

- (1) The uniform distribution over the interval $[1, 3]$ is in $\mathfrak{F}(1.125)$, so Theorem 4.5 implies that the worst-case procrastination ratio of any task graph for this distribution is not more than $O(1.125^n)$. Moreover, Theorem 4.6 applied at point $b_0 = 1.5$ implies that there are task graphs with procrastination ratio of $\Omega(1.125^n)$.
- (2) The uniform distribution over the interval $[1, 2]$ is in $\mathfrak{F}(1)$, and therefore has at most a linear procrastination ratio for any task graph (Theorem 4.5). On the other hand, one can obtain sublinear bound of $\Omega(\sqrt{n})$ in Theorem 4.6 by taking $b_0 = 1 + \varepsilon$ and $z = 1 - \varepsilon^2$ for any $\varepsilon > 0$.
- (3) For the upper half of a normal distribution with mean 1, since the distribution lies in $\mathfrak{F}(z)$ with $z \approx 0.507$, Theorem 4.5 implies that the procrastination ratio is at most 2.03.
- (4) For the distribution given by $\mathcal{F}(x) = 1 - \frac{1}{2\sqrt{x}}$ for $x \in [1, 100)$, and $\mathcal{F}(100) = 1$, since the distribution belongs to $\mathfrak{F}(5)$ we have that the procrastination ratio is bounded by $O(5^n)$, and invoking Theorem 4.6 with $b_0 = 100$ shows that it can be as large as $\Omega(5^n)$.

5. SPECIAL TASK GRAPHS

In this section we consider restrictions on the structure of a task graph. We first study the impact of bounding the length of the shortest path to the target node, from any vertex in the network. We then consider a stronger monotonicity property, which roughly states that following an edge in the graph can never increase the length of the shortest

path to the target. We will show that the former property implies a linear procrastination ratio, and this is tight, whereas the latter property implies a constant procrastination ratio for appropriate present-bias distributions.

5.1. Bounded Distance

Here we consider the class of task graphs such that distance from any given state to t is bounded by the initial distance $d(s, t)$. As it turns out, any such graph has at most a linear procrastination ratio, for any distribution \mathcal{F} .

CLAIM 5.1. *The procrastination ratio of any protocol with bounded distance does not exceed n .*

PROOF. Suppose the chosen path is v_1, v_2, \dots, v_n, t . By assumption, $d(v_i, t) \leq d(s, t)$ for all i . Since the agent chooses the edge $v_i v_{i+1}$ to minimize $b_i \cdot w(v_i v_{i+1}) + d(i+1, t)$ where $b_i \geq 1$, and since the first edge of the path realizing distance $d(v_i, t)$ is itself an option, it must be that $w(v_i v_{i+1}) \leq d(v_i, t) \leq d(s, t)$ for each i . The total cost of the chosen path is therefore at most $n \cdot d(s, t)$. \square

The bound in Claim 5.1 holds for arbitrary procrastination ratios. One might hope that if the present-bias distribution \mathcal{F} is sufficiently well-behaved, the procrastination ratio would improve. However, as we now show, this bound on the procrastination ratio is asymptotically tight for any $\mathcal{F} \in \mathfrak{F}(z)$, for any $z > 1$. In other words, one cannot hope to avoid a linear lower bound unless $\mathcal{F} \in \mathfrak{F}(z)$ for some $z \leq 1$, which we view as a particularly strict condition.

The rough idea for the construction is to simulate a random walk. Whenever the present-bias parameter is low the agent will incur a cost and reduce the length of the shortest path to the goal, but whenever the present-bias parameter is high the length of the shortest path will increase. Over the course of n steps, the agent will have repeatedly made and lost progress sufficiently often to have accrued a total cost that is linear in n , with high probability. The construction is formalized in Appendix B.

THEOREM 5.1. *If $\mathcal{F} \in \mathfrak{F}(z)$ for some $z > 1$, then there is a task graph with bounded shortest path that has procrastination ratio $\Omega(n)$.*

5.2. Monotone distance

We next consider the class of task graphs with the monotone distance property. Recall that for such task graphs, the distance from any given state to t is monotonically non-increasing along every edge in the graph.

We first note that if \mathcal{F} is the present-bias distribution, and there exists some $\delta > 0$ such that $\mathcal{F}(1 + \delta) = 0$ (that is, $1 + \delta$ is a lower bound on the present bias parameter), then there exist task graphs with the monotone distance property with procrastination ratio $\Omega(\delta n)$. Indeed, consider the ski-renting example from Section 3: if \mathcal{F} places all of its mass on values greater than $\frac{1}{1-\delta}$, the agent would choose to pay δ at each step. This leads to a total cost of $n\delta$, whereas a cost of 1 was possible.

This example motivates us to consider distributions \mathcal{F} that satisfy a mild condition, which essentially captures the requirement that there be sufficient mass in neighborhoods around $b = 1$. Specifically, we will require that there exist some constants $\beta, \delta > 0$ such that $\mathcal{F}(x) \geq \beta \cdot (x - 1)$ for all $x \in [1, 1 + \delta]$.⁹ We show that under this condition, the procrastination ratio is bounded by a constant.

⁹Sufficient conditions for our requirement including having a positive mass at $b = 1$, or having a constant lower bound on the density function in any neighborhood around $b = 1$.

THEOREM 5.2. *For any $\beta, \delta > 0$, if $\mathcal{F}(x) \geq \beta \cdot (x - 1)$ for all $x \in [1, 1 + \delta]$, then any task graph with monotone distances has procrastination ratio at most $\max(1 + \frac{1}{\beta}, \frac{1}{\beta\delta})$.*

PROOF. As in Theorem 4.2, we will characterize the task graphs with maximum procrastination ratio by way of an analogy with single-parameter auctions. Following the notation in the proof of Theorem 4.2, we have that the problem of maximizing the procrastination ratio can be expressed recursively (up to a constant $d(v_i, t_{n+1})$) as

$$\max_{V_{i+1}} \left(1 - E_{b_i \sim \mathcal{F}}[x(v_{i+1}^*)] + E_{b_i \sim \mathcal{F}}[p(v_{i+1}^*)] \text{OPT}(i+1, 1) \right).$$

The monotone distances property imposes the constraint that $d(v_{i+1}, t) \leq d(v_i, t)$ for all $v_{i+1} \in V_{i+1}$. Since $p(v_{i+1}) = \frac{d(v_{i+1}, t)}{d(v_i, t)}$, this translates to a constraint that $p(v_{i+1}^*) \in [0, 1]$. Unlike the proof of Theorem 4.2, we cannot employ Theorem 4.1, since the price suggested by Theorem 4.1 may be greater than 1. Instead, we will use the following variation, the proof of which appears in Appendix C.

CLAIM 5.2. *Let $M \subseteq [0, 1] \times [0, 1]$ be a menu of options (x, p) such that $(0, 0) \in M$. For any such M and $v \in \mathbb{R}_{\geq 1}$, write $(x_M(v), p_M(v)) = \arg\max_{(x, p) \in M} [vx - p]$. Then for all distributions \mathcal{F} with supported on $[1, \infty)$, and for all $\alpha \in \mathbb{R}, \beta \in \mathbb{R}_{\geq 0}$, there exists an $x^* \in [0, 1]$ and corresponding $M^* = \{(0, 0), (x^*, 1)\}$ such that*

$$\alpha E_{v \sim \mathcal{F}}[x_M(v)] + \beta E_{v \sim \mathcal{F}}[p_M(v)] \leq \alpha E_{v \sim \mathcal{F}}[x_{M^*}(v)] + \beta E_{v \sim \mathcal{F}}[p_{M^*}(v)].$$

Claim 5.2 implies that there is an optimal menu $V_{i+1} = \{v_{i+1}, u_{i+1}\}$ where $x(v_{i+1}) \leq 1, p(v_{i+1}) \leq 1, x(u_{i+1}) = 0$, and $p(u_{i+1}) = 0$. Given a partially constructed task graph with $V_i = \{v_i, u_i\}$, we can therefore optimally extend it by defining $V_{i+1} = \{v_{i+1}, u_{i+1}\}$. Applying the same reasoning as in the proof of Theorem 4.2, we can conclude that the transitions from v_i to V_{i+1} , and the continuation shortest path weights, will be as follows:

- (1) $w(v_i u_{i+1}) = d(v_i, t)$ and $d(u_{i+1}, t) = 0$, and
- (2) $w(v_i v_{i+1}) \leq d(v_i, t)$ and $d(v_{i+1}, t) \leq d(v_i, t)$.

Note that the task graph is now fully specified, except for the values of $w(v_i v_{i+1})$ and $d(v_{i+1}, t)$ for each $i \leq n$. (Recall that we must have $v_{n+1} = u_{n+1} = t$.) For any such graph and distribution \mathcal{F} , we can write q_i for the probability that the agent would choose u_{i+1} instead of v_{i+1} , from v_i . That is,

$$q_i = \Pr_{b_i \sim \mathcal{F}} \left[b_i \cdot \frac{w(v_i v_{i+1})}{d(v_i, t)} + \frac{d(v_{i+1}, t)}{d(v_i, t)} \geq b_i \cdot 1 \right].$$

We can write the following expression for the optimum at each state v_i scaled by $d(v_i, t)$, again following the notation from the proof of Theorem 4.2:

$$\text{OPT}(i, 1) = q_i \cdot 1 + (1 - q_i) \left(\frac{w(v_i v_{i+1})}{d(v_i, t)} + \frac{d(v_{i+1}, t)}{d(v_i, t)} \cdot \text{OPT}(i+1, 1) \right). \quad (2)$$

To complete the proof of Theorem 5.2, it suffices to show that $\text{OPT}(i, 1) \leq \max(1 + \frac{1}{\beta}, \frac{1}{\beta\delta})$ for each $1 \leq i \leq n$. We will show this by backward induction on i . The base case $i = n$ follows immediately from the fact that $\text{OPT}(n, 1) = 1 \leq \max(1 + \frac{1}{\beta}, \frac{1}{\beta\delta})$, since the agent has only one option: moving to vertex t .

Consider $i < n$. Let $\Delta \stackrel{\text{def}}{=} \max(1 + \frac{1}{\beta}, \frac{1}{\beta\delta})$. By (2), it is sufficient to show that $\Delta \geq q_i + (1 - q_i)(\frac{w(v_i v_{i+1})}{d(v_i, t)} + \frac{d(v_{i+1}, t)}{d(v_i, t)} \cdot \Delta)$, or equivalently that

$$\Delta q_i + \Delta(1 - q_i) \left(1 - \frac{d(v_{i+1}, t)}{d(v_i, t)}\right) \geq q_i + (1 - q_i) \frac{w(v_i v_{i+1})}{d(v_i, t)}. \quad (3)$$

We consider two cases based on the value of q_i .

Case 1: $q_i \geq \beta\delta$. Then $\Delta \cdot q_i \geq \frac{q_i}{\beta\delta} \geq 1$, and hence the left hand side of (3) is at least 1. On the other hand, since $\frac{w(v_i v_{i+1})}{d(v_i, t)} \leq 1$, the right hand side of (3) is at most 1.

Case 2: $q_i < \beta\delta$. We observe that $q_i = \mathcal{F}\left(\frac{d(v_{i+1}, t)}{d(v_i, t) - w(v_i v_{i+1})}\right)$ and by the condition of the theorem that $F(x) \geq \beta(x - 1)$ for $x \in [1, 1 + \delta]$ we therefore have $q_i \geq \beta \cdot \left(\frac{d(v_{i+1}, t)}{d(v_i, t) - w(v_i v_{i+1})} - 1\right)$. We therefore have

$$\begin{aligned} & \Delta q_i + \Delta(1 - q_i) \left(1 - \frac{d(v_{i+1}, t)}{d(v_i, t)}\right) \\ & \geq q_i + \frac{1}{\beta} q_i + (1 - q_i) \left(1 - \frac{d(v_{i+1}, t)}{d(v_i, t)}\right) \\ & \geq q_i + (1 - q_i) \left(\frac{1}{\beta} q_i + 1 - \frac{d(v_{i+1}, t)}{d(v_i, t)}\right) \\ & \geq q_i + (1 - q_i) \left(\frac{d(v_{i+1}, t)}{d(v_i, t) - w(v_i v_{i+1})} - 1 + 1 - \frac{d(v_{i+1}, t)}{d(v_i, t)}\right) \\ & = q_i + (1 - q_i) \frac{d(v_{i+1}, t)}{d(v_i, t) - w(v_i v_{i+1})} \frac{w(v_i v_{i+1})}{d(v_i, t)} \\ & \geq q_i + (1 - q_i) \frac{w(v_i v_{i+1})}{d(v_i, t)}, \end{aligned}$$

where the first inequality follows from $\Delta \geq 1 + \frac{1}{\beta}$, the second inequality we simply reduced the second term by $(1 - q_i)$; in the third inequality we applied our lower bound on q_i . The final inequality follows because $d(v_{i+1}, t) + w(v_i v_{i+1}) \geq d(v_i, t)$, which implies that $\frac{d(v_{i+1}, t)}{d(v_i, t) - w(v_i v_{i+1})} \geq 1$. We have therefore established (3), as desired. \square

6. CONCLUSIONS AND OPEN PROBLEMS

We have examined how variability in an individual's decision making process can affect behavior in time-inconsistent planning scenarios. We adopted the graphical framework of Kleinberg and Oren [Kleinberg and Oren 2014] and characterized worst-case task graphs for the agent. We showed that depending on the distribution of the present-bias parameter, the worst-case procrastination ratio is either bounded by constant, is at most linear, or grows at least exponentially with the number of agent's steps n . We also examined two natural families of tasks: (i) those in which the cost of reaching the goal is never more than at the outset, and (ii) those in which an agent can never lose progress toward the goal. We showed that in the first scenario the worst-case procrastination ratio is at most $O(n)$, and this is tight, and for the second scenario the procrastination ratio is $O(1)$ under some mild assumptions on the present-bias distribution.

Our work leaves open many avenues for future study. The following are a few concrete questions:

- (1) We always assume that $b \geq 1$. What could be the worst-case procrastination ratio if b_i can be smaller than 1? For example, an agent might receive a reminder about the task and becomes anxious to complete it as soon as possible.
- (2) We examine worst-case procrastination ratio for a fixed distribution of $b \sim \mathcal{F}$. Is there a bound on procrastination ratio for a given pair of the task graph and distribution \mathcal{F} ? In other words, are certain distributions better suited to certain tasks than others? What is an important intrinsic parameter of a weighted graph that can lead to large procrastination?
- (3) We examined the case of a simple, naïve agent who never anticipates their time-inconsistency. What can one say about agents who can reason about their future propensity to procrastinate?

REFERENCES

- George A. Akerlof. 1991. Procrastination and obedience. *American Economic Review* 81, 2 (May 1991), 1–19.
- Dan Ariely and Klaus Wertenbroch. 2002. Procrastination, Deadlines, and Performance: Self-Control by Precommitment. *Psychological Science* 13, 3 (May 2002), 219–224. DOI: <http://dx.doi.org/10.1111/1467-9280.00441>
- John Beshears, James J. Choi, David Laibson, Brigitte C. Madrian, and Katherine L. Milkman. 2015. The Effect of Providing Peer Information on Retirement Savings Decisions. *Journal of Finance* 70, 3 (2015), 1161–1201.
- James Choi, David Laibson, and Brigitte C. Madrian. 2004. Plan Design and 401(k) Savings Outcomes. *National Tax Journal* 57, 2 (2004), 275–298. <http://www.nber.org/digest/oct04/w10486.html>
- Camerer Colin. 2003. *Behavioral game theory: experiments in strategic interaction*. Vol. The Roundtable series in behavioral economics. Princeton University Press.
- Jeffrey C. Ely. 2015. Beeps. (2015). <http://jeffely.com/assets/beeps-final.pdf>
- Shane Frederick, George Loewenstein, and Ted O’Donoghue. 2002. Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature* 40, 2 (June 2002), 351–401.
- Supreet Kaur, Michael Kremer, and Sendhil Mullainathan. 2010. Self-Control and the Development of Work Arrangements. *American Economic Review* 100, 2 (2010), 624–28. DOI: <http://dx.doi.org/10.1257/aer.100.2.624>
- Jon M. Kleinberg and Sigal Oren. 2014. Time-inconsistent planning: a computational problem in behavioral economics. In *ACM Conference on Economics and Computation, EC ’14, Stanford, CA, USA, June 8-12, 2014*. 547–564. DOI: <http://dx.doi.org/10.1145/2600057.2602890>
- David Laibson. 2015. Why don’t present-biased agents make commitments? *American Economic Review Papers and Proceedings* 105, 5 (2015), 267–272.
- R. Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- Evdokia Nikolova and David R. Karger. 2008. Route Planning under Uncertainty: The Canadian Traveller Problem. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, July 13-17, 2008*. 969–974. <http://www.aaai.org/Library/AAAI/2008/aaai08-154.php>
- Evdokia Nikolova and Nicolás E. Stier Moses. 2011. Stochastic Selfish Routing. In *Algorithmic Game Theory, 4th International Symposium, SAGT 2011, Amalfi, Italy, October 17-19, 2011. Proceedings*. 314–325. DOI: http://dx.doi.org/10.1007/978-3-642-24829-0_28
- Evdokia Nikolova and Nicolás E. Stier Moses. 2014. A Mean-Risk Model for the Traffic Assignment Problem with Stochastic Travel Times. *Operations Research* 62, 2

- (2014), 366–382. DOI: <http://dx.doi.org/10.1287/opre.2013.1246>
- Evdokia Nikolova and Nicolás E. Stier Moses. 2015. The Burden of Risk Aversion in Mean-Risk Selfish Routing. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 489–506. DOI: <http://dx.doi.org/10.1145/2764468.2764485>
- Ted O’Donoghue and Matthew Rabin. 1999. Doing It Now or Later. *American Economic Review* 89, 1 (1999), 103–124. DOI: <http://dx.doi.org/10.1257/aer.89.1.103>
- Ted O’Donoghue and Matthew Rabin. 2008. Procrastination on long-term projects. *Journal of Economic Behavior & Organization* 66, 2 (May 2008), 161–175.
- Pingzhong Tang, Yifeng Teng, Ziheng Wang, Shenke Xiao, and Yichong Xu. 2014. Settling the open problems in Kleinberg-Oren time-inconsistent planning models. *CoRR* abs/1411.7472 (2014). <http://arxiv.org/abs/1411.7472>
- Richard Thaler. 1981. Some empirical evidence on dynamic inconsistency. *Economics Letters* 8, 3 (1981), 201 – 207. DOI: [http://dx.doi.org/10.1016/0165-1765\(81\)90067-7](http://dx.doi.org/10.1016/0165-1765(81)90067-7)
- James R. Wright and Kevin Leyton-Brown. 2010. Beyond Equilibrium: Predicting Human Behavior in Normal-Form Games. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*. <http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1946>
- James R. Wright and Kevin Leyton-Brown. 2012. Behavioral game theoretic models: a Bayesian framework for parameter analysis. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*. 921–930. <http://dl.acm.org/citation.cfm?id=2343828>
- James R. Wright and Kevin Leyton-Brown. 2014. Level-0 meta-models for predicting human behavior in games. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*. 857–874. DOI: <http://dx.doi.org/10.1145/2600057.2602907>

A. PROOFS FROM SECTION 4.2

In this section, we provide proofs for claims in Section 4.2, restated here for convenience. The following is a restatement of Lemma 4.3.

LEMMA A.1. *Given two distributions \mathcal{F} and $\bar{\mathcal{F}}$ such that $\bar{\mathcal{F}}$ stochastically dominates \mathcal{F} , the worst-case procrastination ratio for $\bar{\mathcal{F}}$ is at least as large as the worst-case procrastination ratio for \mathcal{F} .*

PROOF. Fix n and let $G_{\mathcal{F}}^n$ be a worst-case task graph for \mathcal{F} of the form specified in Theorem 4.2. We use cost_i to denote the expected cost paid by an agent with $b_i \sim \mathcal{F}$ on days i, \dots, n in the graph $G_{\mathcal{F}}^n$, conditional on being at node s_i on day i . Similarly, we use $\overline{\text{cost}}_i$ to denote the expected cost paid by an agent with $b_i \sim \bar{\mathcal{F}}$ on days i, \dots, n in the graph $G_{\mathcal{F}}^n$, conditional on being at node s_i on day i . In the following we show using backward induction on i that $\overline{\text{cost}}_i \geq \text{cost}_i$, implying that $\overline{\text{cost}}_0$, a lower bound on the worst-case procrastination ratio for $\bar{\mathcal{F}}$, is at least as large as cost_0 , the worst-case procrastination ratio for \mathcal{F} .

When $i = n$ we have $\overline{\text{cost}}_n = \text{cost}_n = d(s_n, t_{n+1})$. For $i < n$ note that by definition of $G_{\mathcal{F}}^n$, an agent with threshold bias $b_0 = d(s_{i+1}, t_{n+1})/d(s_i, t_{n+1})$ is indifferent between transitioning to s_{i+1} at a cost of zero, and transitioning to t_{i+1} at a cost of $d(s_i, t_{n+1})$ with zero future costs. Thus

$$\begin{cases} \text{cost}_i &= (1 - \mathcal{F}(b_0)) \cdot \text{cost}_{i+1} + \mathcal{F}(b_0) \cdot d(s_i, t_{n+1}) \\ \overline{\text{cost}}_i &= (1 - \bar{\mathcal{F}}(b_0)) \cdot \overline{\text{cost}}_{i+1} + \bar{\mathcal{F}}(b_0) \cdot d(s_i, t_{n+1}). \end{cases} \quad (4)$$

By induction, $\overline{\text{cost}}_{i+1} \geq \text{cost}_{i+1}$, and so

$$\overline{\text{cost}}_i - \text{cost}_i \geq (\mathcal{F}(b_0) - \bar{\mathcal{F}}(b_0)) \cdot (\text{cost}_{i+1} - d(s_i, t_{n+1})).$$

By stochastic dominance, $\mathcal{F}(b_0) \geq \bar{\mathcal{F}}(b_0)$, and by the construction of $G_{\mathcal{F}}^n$, $\text{cost}_{i+1} \geq d(s_{i+1}, t_{n+1}) \geq d(s_i, t_{n+1})$. As a result $\overline{\text{cost}}_i - \text{cost}_i \geq 0$, which completes the induction. \square

We next prove the following restatement of Theorem 4.5.

THEOREM A.2. *Let $\mathcal{F} \in \mathfrak{F}(z)$ for some $z > 0$. Then the procrastination ratio of any task graph for \mathcal{F} is at most $\sum_{i=0}^{k-1} z^i$ and this bound is tight.*

PROOF. To prove an upper bound on the procrastination ratio we analyze the following distribution \mathcal{F} with bounded support on $[1, C]$, where C is some sufficiently large constant:

$$\mathcal{F} : \begin{cases} \Pr[b = 1] &= 1 - z, & \text{if } z < 1 \\ \Pr[b \leq x] &= 1 - z/x, & x \in [\max(1, z), C] \\ \Pr[b = C] &= z/C \end{cases}$$

We note that \mathcal{F} stochastically dominates every distribution in $\mathfrak{F}(z)$ with support bounded by C .

Recall from Theorem 4.2 the form of the worst-case task graph for distribution \mathcal{F} . We briefly describe the form of the graph here. Given a current state v_i , the agent is offered two options: an edge with weight $d(v_i, t)$ (and no further costs thereafter), and an edge with weight 0 that transitions to a vertex v_{i+1} with increased distance to t . We claim that, in fact, $d(v_{i+1}, t) = C \cdot d(v_i, t)$ in the worst-case graph. To see why, note that according to (1) the worst-case task graph maximizes a linear function of $\mathbf{E}_{b_i \sim \mathcal{F}}[x(v_{i+1}^*)]$ and $\mathbf{E}_{b_i \sim \mathcal{F}}[p(v_{i+1}^*)]$. The former term does not depend on the set of offered price p , because $\mathbf{E}_{b_i \sim \mathcal{F}}[p(v_{i+1}^*)]$ corresponds in (1) to the expected revenue of a single item auction for the equal revenue distribution \mathcal{F} . Therefore, $\mathbf{E}_{b_i \sim \mathcal{F}}[p(v_{i+1}^*)] = z$ for any menu of options, and hence the optimization problem reduces to maximizing the term $1 - \mathbf{E}_{b_i \sim \mathcal{F}}[x(v_{i+1}^*)]$. It is maximized when the agent's probability of paying a positive cost at step i is maximized. This is maximized when agent can either increase the distance to t to $d(v_{i+1}, t) = C \cdot d(v_i, t)$, or pay the total amount $d(v_i, t)$ right away. Thus we obtain the worst-case task graph.

Now we calculate the procrastination ratio of this task graph. At step i , from a given vertex $v \in \{v_i, u_i\}$, the expected distance to t is $\mathbf{E}[d(v, t)] = z^{i-1}$. The probability that agent takes a non-zero edge is $1 - 1/C$. Thus at the steps $i = 1, \dots, n-1$ the agents pays a cost, in expectation, of $\mathcal{F}(C) \cdot \mathbf{E}[d(v, t)] = (1 - 1/C) \cdot z^{i-1}$ and $\mathbf{E}[d(v, t)] = z^{n-1}$ at the last step. As a result, the total expected cost paid by the agent is

$$\sum_{i=1}^{n-1} (1 - 1/C) \cdot z^{i-1} + z^{n-1}$$

As C tends to infinity this quantity converges to $\sum_{i=1}^n z^{i-1} = \frac{1-z^n}{1-z}$. This is an upper bound on the procrastination ratio for $\mathcal{F} \in \mathfrak{F}(z)$ by Lemma 4.3. Furthermore, this bound is tight for the class $\mathfrak{F}(z)$, since we have exhibited a distribution and a task graph that achieve this bound in the limit as $C \rightarrow \infty$. \square

Finally, we prove the following restatement of Theorem 4.6.

THEOREM A.3. *For arbitrary \mathcal{F} , if $(1 - \mathcal{F}(b_0))b_0 = z$ for some point $b_0 > 1$, then there is a task graph with procrastination ratio $\sum_{i=1}^{n-1} (1 - 1/b_0) \cdot z^{i-1} + z^{n-1}$.*

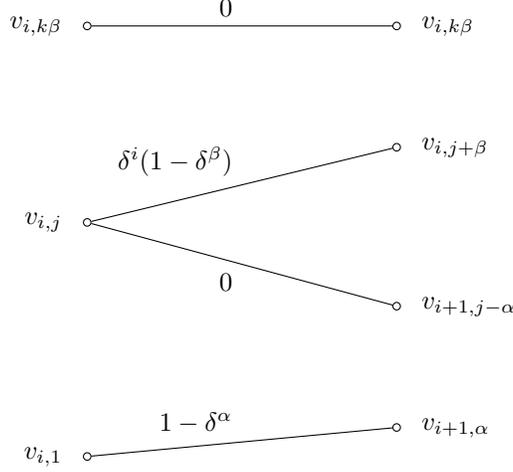


Fig. 2. Task completion graph for the protocol

PROOF. Indeed, we may assume that \mathcal{F} has finite support $\{1, b_0\}$, since every other distribution stochastically dominates this one and by revenue monotonicity (Lemma 4.3) the adversary may only get higher procrastination ratio. The task graph from the proof of Theorem 4.5 gives us the required bound. \square

B. PROOFS FROM SECTION 5.1

In this section, we prove Theorem 5.1, restated here for convenience.

THEOREM B.1. *If $\mathcal{F} \in \mathfrak{F}(z)$ for some $z > 1$, then there is a task graph with bounded shortest path that has procrastination ratio $\Omega(n)$.*

PROOF. Let $\delta \in (0, 1)$ and integers $\alpha, \beta \geq 1$ be constants that depend on \mathcal{F} , to be determined later. We construct a task graph as follows. Each layer V_i contains $n\beta$ vertices; say $V_i = \{v_{i,1}, \dots, v_{i,n\beta}\}$. Each vertex $v_{i,j}$ has out-degree at most 2. For $i < n$, starting from vertex $v_{i,j}$,

- if $\alpha \leq j \leq (n-1)\beta$, then the agent may go either to $v_{i+1,j-\alpha}$ at no cost, or to $v_{i+1,j+\beta}$ at a cost of $\delta^j - \delta^{j+\beta}$;
- if $j < \alpha$, then the agent must transition to $v_{i+1,\alpha}$ at a cost of $\delta^j - \delta^\alpha$.
- Otherwise, if $j > (n-1)\beta$, the agent must transition to $v_{i+1,j}$ at no cost.

For $i = n$, each vertex $v_{i,j}$ has only a single outgoing edge $(v_{i,j}, t)$, which has cost δ^j .

We claim that, for each vertex $v_{i,j}$, we have $d(v_{i,j}, t) = \delta^j$. This is achieved by the path in which the agent always chooses the option with positive cost, whenever multiple options are available. That is, the agent repeatedly chooses to take the route that increases the second coordinate of its vertex, if possible. If this path is $v_{i,j}, v_{i_2,j_2}, \dots, v_{i_\ell,j_\ell}, t$, then the total cost is

$$(\delta^j - \delta^{j_2}) + (\delta^{j_2} - \delta^{j_3}) + \dots + (\delta^{j_{\ell-1}} - \delta^{j_\ell}) + \delta^{j_\ell} = \delta^j,$$

via a telescoping sum. The agent starts at $v_{1,0}$, and $d(v_{1,0}, t) = 1$. We note that the task graph satisfies the bounded distance condition.

Let b^* be such that $(1 - \mathcal{F}(b^*))b^* = z > 1$; we know such a b^* exists by assumption. We will then choose α and β so that

$$\frac{1}{b^*} < \frac{\beta}{\alpha + \beta} < 1 - \mathcal{F}(b^*).$$

Note that the agent will select the 0-cost option, if multiple options are available, when

$$\delta^{j-\alpha} < b_i(\delta^j - \delta^{j+\beta}) + \delta^{j+\beta}.$$

We claim that this will occur for $b_i \geq b^*$, if δ is sufficiently close to 1. To see this, define $H(\delta) \stackrel{\text{def}}{=} b^* \cdot \delta^\alpha - \delta^{\alpha+\beta} \cdot (b^* - 1)$. Note that $H(1) = 1$. Furthermore, $H'(1) = b^* \cdot \alpha - (\alpha + \beta) \cdot (b^* - 1) = \alpha + \beta - \beta \cdot b^*$ is negative if $\frac{1}{b^*} < \frac{\beta}{\alpha + \beta}$. The claim follows, as $H(\delta) = H(1) - H'(1)(1 - \delta) + o(1 - \delta)$ for δ sufficiently close to 1. Choose δ to be sufficiently close to 1 for the claim to hold.

Given $i < n$ and j such that $\alpha \leq j \leq (n - 1)\beta$, let γ be the expectation, over the present-bias distribution, of the change in index of the current vertex when an agent makes a single move starting at vertex $v_{i,j}$. Then γ is at most $\beta\mathcal{F}(b^*) - \alpha(1 - \mathcal{F}(b^*))$, which is negative by our choice of α and β .

Note that since the agent started at $v_{1,0}$ it is impossible to reach a state $v_{i,j}$ where $j \geq (n - 1)\beta$. We now count the sum of all changes in vertex index, over the course of the agent's selected path, enumerated separately for two cases: (i) when $\alpha \leq j \leq (n - 1)\beta$, and when (ii) $j < \alpha$.

Let random variables s_1 and s_2 denote the sum of changes in vertex index for these two cases. We observe that the total sum must be non-negative, and hence

$$E[s_1] + E[s_2] \geq 0. \quad (5)$$

Also, by our construction, if $j < \alpha$ then the subsequent step will have $j \geq \alpha$, and hence the second case applies in at least half of the rounds. Therefore,

$$E[s_1] \leq \gamma \cdot n/2 \quad (6)$$

where recall that $\gamma < 0$. Equations (5) and (6) imply that

$$E[s_2] \geq -\gamma \cdot n/2. \quad (7)$$

Note that for each index-increasing step, the agent is reducing the shortest path and incurring an immediate cost equal to that reduction. The minimum cost reduction between two consecutive states is $\delta^{\alpha-1} - \delta^\alpha$, in the case when $j < \alpha$. We can therefore conclude that

$$\sum_{i=1}^n E[w_i] \geq -\gamma \cdot n/2 \cdot (\delta^{\alpha-1} - \delta^\alpha) = \Omega(n) \quad (8)$$

as required. \square

C. PROOFS FROM SECTION 5.2

In this section, we prove Claim 5.2, restated below.

CLAIM C.1. *Let $M \subseteq [0, 1] \times [0, 1]$ be a menu of options (x, p) such that $(0, 0) \in M$. For any such M and $v \in \mathbb{R}_{\geq 1}$, write $(x_M(v), p_M(v)) = \operatorname{argmax}_{(x,p) \in M} [vx - p]$. Then for all distributions \mathcal{F} supported on $[1, \infty)$, and for all $\alpha \in \mathbb{R}, \beta \in \mathbb{R}_{\geq 0}$, there exists an $x^* \in [0, 1]$ and corresponding $M^* = \{(0, 0), (x^*, 1)\}$ such that*

$$\alpha E_{v \sim \mathcal{F}}[x_M(v)] + \beta E_{v \sim \mathcal{F}}[p_M(v)] \leq \alpha E_{v \sim \mathcal{F}}[x_{M^*}(v)] + \beta E_{v \sim \mathcal{F}}[p_{M^*}(v)].$$

PROOF. For simplicity we will prove the claim under the assumption that menu M is finite. We note that the result directly extends to general M . Also, in the context in which we apply the claim (namely, when menu M corresponds to a task graph), M is finite.

Write $M = \{(x_1, p_1), \dots, (x_k, p_k)\}$. We will say that the item (x_i, p_i) is *chosen* by value v if $(x_M(v), p_M(v)) = (x_i, p_i)$. We will partition the set of possible values v (i.e., $\mathbb{R}_{\geq 1}$) according to their chosen lottery. Note that, for each i , the set of values that choose (x_i, p_i) forms an interval, say $[v_{i-1}, v_i]$. Without loss of generality, we can assume every element of M is chosen by some value. In this case, we can choose to index lotteries so that $1 = v_0 < v_1 < \dots < v_k = \infty$. In this case, we must have $0 = x_0 \leq x_1 < x_2 < \dots < x_k \leq 1$ and $p_1 < p_2 < \dots < p_k \leq 1$. We note that the latter can be derived from well-known payment identities of incentive compatible auctions [Myerson 1981]. Moreover, the fact that v_{i-1} is the threshold type between choosing (x_{i-1}, p_{i-1}) and (x_i, p_i) , and is therefore indifferent between these auctions, implies that

$$\begin{aligned} v_{j-1} \cdot (x_j - x_{j-1}) &= p_j - p_{j-1} & \text{if } j \in \{2, \dots, k\} \\ v_0 \cdot x_1 - p_1 &> 0 & \text{if } j = 1. \end{aligned} \tag{9}$$

We now consider a collection of two-element menus M_1, \dots, M_k , where $M_j = \{(0, 0), (\frac{1}{v_{j-1}}, 1)\}$. We will construct a distribution over these menus, and consider offering one of them at random: say menu M_j is chosen with probability m_j . We want to choose $\{m_1, \dots, m_k\}$ such that for every value $v \in [v_{j-1}, v_j]$, the expected allocation chosen by v under this distribution of menus is exactly x_j (the allocation under M), and the expected payment is at least p_j . Note then that this random menu, call it M^r , can only increase our objective of interest:

$$\alpha E_{v \sim \mathcal{F}}[x_M(v)] + \beta E_{v \sim \mathcal{F}}[p_M(v)] \leq \alpha E_{v \sim \mathcal{F}}[x_{M^r}(v)] + \beta E_{v \sim \mathcal{F}}[p_{M^r}(v)].$$

Since M^r is a convex combination of menus M_j , there must exist at least one M_j for which the objective is at least as high as for M^r . This M_j satisfies the requirements of the claim.

It remains to find appropriate probabilities m_1, \dots, m_k . We will choose $m_j = v_{j-1} \cdot (x_j - x_{j-1})$, where $j \in \{1, \dots, k\}$ and $x_0 = 0$.

We observe that each type $v \in (v_{i-1}, v_i)$ chooses the option $(\frac{1}{v_{j-1}}, 1)$ in M_j for every $j \in \{1, \dots, i\}$ and option $(0, 0)$ for the remaining j . Therefore, the allocation probability and expected payment of type v is

$$\begin{aligned} x(v) &= \sum_{j=1}^i \frac{1}{v_{j-1}} \cdot m_j = \sum_{j=1}^i x_j - x_{j-1} = x_i \\ p(v) &= \sum_{j=1}^i m_j = \sum_{j=1}^i v_{j-1} \cdot (x_j - x_{j-1}) = v_0 \cdot x_1 + \sum_{j=2}^i (p_j - p_{j-1}) = p_i - p_1 + x_1 \geq p_i, \end{aligned}$$

where the last inequality and previous two equalities follow from (9) and telescopic summation. Since $p(v_k) \leq 1$ we also get that $\sum_{j=1}^k m_j \leq 1$. If $\sum_{j=1}^k m_j$ is strictly less than 1, we can extend to a probability distribution by adding a dummy lottery $M_0 = \{(0, 0)\}$ that is offered with probability $1 - \sum_{j=1}^k m_j$. \square