Present-Biased Optimization

Fedor V. Fomin,[1] **Pierre Fraigniaud,** [2] **Petr A. Golovach** [1]

[1] Department of Informatics, University of Bergen, Norway
[2] IRIF, Université de Paris and CNRS, France
fedor.fomin@uib.no, pierre.fraigniaud@irif.fr, petr.golovach@uib.no

Abstract

This paper explores the behavior of *present-biased* agents, that is, agents who erroneously anticipate the costs of future actions compared to their real costs. Specifically, the paper extends the original framework proposed by Akerlof (1991) for studying various aspects of human behavior related to time-inconsistent planning, including procrastination, and abandonment, as well as the elegant graph-theoretic model encapsulating this framework recently proposed by Kleinberg and Oren (2014). The benefit of this extension is twofold. First, it enables to perform fine grained analysis of the behavior of present-biased agents depending on the *optimisation task* they have to perform. In particular, we study covering tasks vs. hitting tasks, and show that the ratio between the cost of the solutions computed by present-biased agents and the cost of the optimal solutions may differ significantly depending on the problem constraints. Second, our extension enables to study not only underestimation of future costs, coupled with minimization problems, but also all combinations of minimization/maximization, and underestimation/overestimation. We study the four scenarios, and we establish upper bounds on the cost ratio for three of them (the cost ratio for the original scenario was known to be unbounded), providing a complete global picture of the behavior of present-biased agents, as far as optimisation tasks are concerned.

Introduction

Present bias is the term used in behavioral economics to describe the gap between the anticipated costs of future actions and their real costs. A simple mathematical model of present bias was suggested by Akerlof 1991. In this model the cost of an action that will be perceived in the future is assumed to be β times smaller than its actual cost, for some constant $\beta < 1$, called the *degree of present bias*. The model was used for studying various aspects of human behavior related to *time-inconsistent* planning, including *procrastination*, and *abandonment*.

Kleinberg and Oren 2014; 2018 introduced an elegant graph-theoretic model encapsulating Akerlof's model. The approach of Kleinberg and Oren is based on analyzing how an agent navigates from a source s to a target t in a directed edge-weighted graph G, called *task graph*. At any step, the agent chooses the next edge to traverse from the current vertex v thanks to an estimation of the length of the shortest path from v to t passing through each edge outgoing from v. A crucial characteristic of the model is that the estimation of the path lengths is *present-biased*. More specifically, the model of Kleinberg and Oren includes a positive parameter $\beta < 1$, the degree of present bias, and the length of a path $x_0, \ldots, x_k$ from $x_0 = v$ to $x_k = t$ in G is evaluated as $\omega_0 + \beta \sum_{i=1}^{k-1} \omega_i$ where ω_i denotes the weight of edge (x_i, x_{i+1}), for every $i \in \{0, \ldots, k-1\}$. As a result, the agent may choose an outgoing edge that is not on any shortest path from v to t, modeling *procrastination* by underestimating the cost of future actions to be performed whenever acting now in some given way. With this effect cumulating along its way from s to t, the agent may significantly diverge from shortest s-t paths, which demonstrates the negative impact of procrastination. Moreover, the *cost ratio*, which is the ratio between the cost of the path traversed by an agent with present bias and the cost of a shortest path, could be arbitrarily large. An illustrating example is depicted on Fig. 1, borrowed from (Kleinberg and Oren 2018), and originally due to Akerlof 1991. Among many results, Kleinberg and Oren showed that any graph in which a present-biased agent incurs significantly more cost than an optimal agent must contain a large specific structure as a minor. This structure, called *procrastination structure*, is specifically the one depicted on Fig. 1.

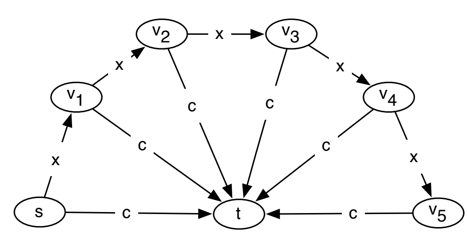

Figure 1: Procrastination structure as displayed in (Kleinberg and Oren 2018); Assuming $\mathsf{x} + \beta\mathsf{c} < \mathsf{c}$, the path followed by the agent is $\mathsf{s}, \mathsf{v}_1, \ldots, \mathsf{v}_5, \mathsf{t}$; The ratio between the length of the path followed by the agent and the shortest s-t path can be made arbitrarily large by adding more nodes v_k with $k \geq 5$.

In this paper, we are interested in understanding what kind of *tasks* performed by the agent result in large cost ratio. Let us take the concrete example of an agent willing to acquire the knowledge of a set of scientific concepts, by reading books. Each book covers a certain number of these concepts, and the agent's objective is to read as few books as possible. More generally, each book could also be weighted according to, say, its accessibility to a general reader, or its length. The agent's objective is then to read a collection of books with minimum total weight. Both the weight and the collection of concepts covered by each book are known to the agent a priori. This scenario is obviously an instance of the (weighted) *set-cover* problem. Let us assume, for simplicity, that the agent has access to a *time-biased* oracle providing it with the following information. Given the subset of concepts already acquired by the agent when it queries the oracle, the latter returns to the agent a set $\{b_0, \dots, b_{k-1}\}$ of books minimizing $\omega_0 + \beta \sum_{i=1}^{k-1} \omega_i$ among all sets of books covering the concepts yet to be acquired by the agent, where $\omega_0 \leq \omega_1 \leq \cdots \leq \omega_{k-1}$ are the respective weights of the books $b_0, \dots, b_{k-1}$. This corresponds to the procrastination scenario in which the agent picks the easiest book to read now, and underestimates the cost of reading the remaining books later. Then the agent moves on by reading b_0, and querying the oracle for figuring out the next book to read for covering the remaining uncovered concepts after having read book b_0.

The question is: by how much the agent eventually diverges from the optimal set of books to be read? This set-cover example fits with the framework of Kleinberg and Oren, by defining the vertex set of the task graph as the set of all subsets of concepts, and placing an edge (u, v) of weight ω from u to v if there exists a book b of weight ω such that v is the union of u and the concepts covered by b. In this setting, the agent needs to move from the source $s = \varnothing$ to the target t corresponding to the set of all the concepts to be acquired by the agent. Under this setting, the question can be reformulated as: under which circumstances the set-cover problem yields a large cost ratio?

More generally, let us consider a *minimization* problem where, for every feasible solution S of every instance of the problem, the cost $c(S)$ of S can be expressed as $c(S) = \sum_{x \in S} \omega(x)$ for some weight function ω. This includes, e.g., set-cover, min-cut, minimum dominating set, feedback-vertex set, etc. We then define the biased cost c_β as

$$c_\beta(S) = \omega(x^*) + \beta\, c(S \smallsetminus \{x^*\}), \qquad (1)$$

where $x^* = \arg\min_{x \in S} \omega(x)$. Given an instance I of the minimization problem at hand, the agent aims at finding a feasible solution $S \in I$ minimizing $c(S)$. It does so using the following present-biased planning, where $I_0 = I$.

Minimization scenario: For $k \geq 0$, given an instance I_k, the agent computes the feasible solution S_k with minimum cost $c_\beta(S_k)$ among all feasible solutions for I_k. Let $x_k^* = \arg\min_{x \in S_k} w(x)$. The agent stops whenever $\{x_0^*, x_1^*, \dots, x_k^*\}$ is a feasible solution for I. Otherwise, the agent moves to $I_{k+1} = I_k \smallsetminus x_k^*$, that is, the instance obtained from I_k when one assumes x_k^* selected in the solution.

This general scenario is indeed captured by the Kleinberg and Oren model, by defining the vertex set of the graph task graph as the set of all "sub-instances" of the instance I at hand, and placing an edge (u, v) of weight w from u to v if there exists an element x of weight ω such that v results from u by adding x to the current solution. The issue is to analyze how far the solution computed by the present-biased agent is from the optimal solution. The first question addressed in this paper is therefore the following.

Question 1. *For which minimization tasks a large cost ratio may appear?*

In the models of Akerlof 1991 and Kleinberg and Oren 2018 the degree β of present bias is assumed to be less than one. However, there are natural situations where underestimating the future costs does not hold. For example, in their influential paper, Loewenstein, O'Donoghue, and Rabin 2003 gave a number of examples from a variety of domains demonstrating the prevalence of *projection bias*. In particular, they reported an experiment by Jepson, Loewenstein, and Ubel 2001 who "asked people waiting for a kidney transplant to predict what their quality of life would be one year later if they did or did not receive a transplant, and then asked those same people one year later to report their quality of life. Patients who received transplants predicted a higher quality of life than they ended up reporting, and those who did not predicted a lower quality of life than they ended up reporting". In other words, there are situations in which people may also *overestimate* the future costs. In the model of Kleinberg and Oren 2018 overestimation bias corresponds to the situation of putting the degree of present bias $\beta > 1$. This brings us to the second question.

Question 2. *Could a large cost ratio appear for minimization problems when the degree of present bias β is more than 1?*

Reformulating the analysis of procrastination, as stated in Question 1, provides inspiration for tackling related problems. As a matter of fact, under the framework of Kleinberg and Oren, procrastination is a priori associated to *minimization* problems. We also investigate *maximization* problems, in which a present-biased agent aims at, say, maximizing its revenue by making a sequence of actions, each providing some immediate gain that the agent maximizes while underestimating the incomes resulting from future actions. As a concrete example, let us consider an instance of Knapsack. The agent constructs a solution gradually by picking the item x_0 of highest value $\omega(x_0)$ in a feasible set $\{x_0, \dots, x_{k-1}\}$ of items that is maximizing $\omega(x_0) + \beta \sum_{i=1}^{k-1} \omega(x_i)$ for the current sub-instance of Knapsack. In general, given an instance I of a maximisation problem, we assume that the agent applies the following present-biased planning, with $I_0 = I$:

Maximization scenario: Given an instance I_k for $k \geq 0$, the agent computes the feasible solution S_k with maximum cost $c_\beta(S_k)$ among all feasible solutions for I_k — where the definition of x^* in Eq. (1) is replaced by $x^* = \arg\max_{x \in S} w(x)$. With $x_k^* = \arg\max_{x \in S_k} w(x)$, the agent stops whenever $\{x_0^*, x_1^*, \ldots, x_k^*\}$ is an inclusion-wise maximal feasible solution for I, and moves to $I_{k+1} = I_k \setminus x_k^*$ otherwise.

We are interested in analyzing how far the solution computed by the present-biased agent is from the optimal solution. More generally even, we aim at revisiting time-inconsistent planning by considering both cases $\beta < 1$ and $\beta > 1$, that is, not only scenarios in which the agent underestimates the cost of future actions, but also scenarios in which the agent *overestimates* the cost of future actions. The last, more general question addressed in this paper is therefore the following.

Question 3. *For which optimization tasks, and for which time-inconsistency planning (underestimation, or overestimation of the future actions), the solutions computed by a present-biased agent are far from optimal, and for which they are close?*

For all these problems, we study the *cost ratio* $\varrho = \frac{c(S)}{\text{OPT}}$ (resp., $\varrho = \frac{\text{OPT}}{c(S)}$) where S is the solution returned by the present-biased agent, and $\text{OPT} = c(S_{\text{OPT}})$ is the cost of an optimal solution for the same instance of the considered minimization (resp., maximization) problem.

Our Results

Focussing on agents aiming at solving tasks, and not just on agents aiming at reaching targets in abstract graphs, as in the generic model in (Kleinberg and Oren 2018), allows us not only to refine the worst-case analysis of present-biased agents, but also to extend this analysis to scenarios corresponding to overestimating the future costs to be incurred by the agents (by setting the degree β of present bias larger than 1), and to maximisation problems.

Minimization & underestimation. In the original setting of minimization problems, with underestimation of future costs (i.e., $\beta < 1$), we show that the cost ratio ϱ of an agent performing k steps, that is, computes a feasible solution $\{x_1^*, \ldots, x_k^*\}$, satisfies $\varrho \leq k$. This is in contrast to the general model in (Kleinberg and Oren 2018), in which an agent can incur a cost ratio exponential in k when returning a k-edge path from the source to the target. Hence, in particular, our minimization scenarios do not produce the worst cases examples constructed in (Kleinberg and Oren 2018), i.e., obtained by considering travels from sources to targets in arbitrary weighted graphs.

On the other hand, we also show that a "minor structure" bearing similarities with the one identified in (Kleinberg and Oren 2018) can be identified. Namely, if an agent incurs a large cost ratio, then the minimization problem addressed by the agent includes a large instance of a specific form of minimization problem.

Min/maximization & under/overestimation. Interestingly, the original setting of minimization problems, with underestimation of future costs, is far from reflecting the whole nature of the behavior of present-biased agents. Indeed, while minimization problems with underestimation of future costs may result in unbounded cost ratios, the worst-case cost ratios corresponding to the three other settings can be upper bounded, some by a constant independent of the task at hand. Specifically, we show that:

- For any minimization problem with $\beta > 1$, the cost ratio is at most β;
- For any maximization problem with $\beta < 1$, the cost ratio is at most $\frac{1}{\beta}$;
- For any maximization problem with $\beta > 1$, the cost ratio is at most β^c, where $c \leq \text{OPT}$ is the cost of a solution constructed by the agent.

Our results are summarized in Table 1.

	minimization	maximization
$\beta < 1$	∞ (Kleinberg and Oren 2018)	$1/\beta$ [Thm 5(i)]
$\beta > 1$	β [Thm 4]	$(1 + \log \beta)\,\frac{\text{OPT}}{\log \text{OPT}}$ [Cor 1]

Table 1: Upper bounds on the worst case ratio between the solution cost returned by the present-biased agent and the optimal solution OPT. The symbol ∞ means that the cost ratio can be arbitrarily large, independently of the values of β, and OPT.

Let us remark that, for minimization problems with $\beta > 1$, as well as for maximization problems with $\beta < 1$, we have that the cost ratio is bounded by a constant. However, for maximization problems with $\beta > 1$, the cost ratio can be exponential in the cost of the computed solution. We show that this exponential upper bound is essentially tight.

Approximated evaluations. Actually, in many settings, discrete optimization problems are hard. Therefore, for evaluating the best feasible solution according to the biased cost function c_β, an agent may have to solve computationally intractable problems. Thus, in a more realistic scenario, we assume that, instead of computing an optimal solution for c_β at every step, the agent computes an α-approximate solution.

Fine-grained analysis. In contrast to the general model in (Kleinberg and Oren 2018), the refined model of this paper enables fine-grain analysis of the agents' strategies, that is, it enables identifying different behaviors of the agents as a function of the considered optimisation problems. Specifically, there are natural minimization problems for which specific bounds on the cost ratio can be established.

To illustrate the interest of focusing on optimisation tasks, we study two tasks in detail, namely *set-cover* and *hitting set*, and show that they appear to behave quite differently. For set-cover, we show that the cost ratio is at most $d \cdot \text{OPT}$, where d is the maximum size of the sets. For hitting set, we show that the cost ratio is at most $d!\,(\frac{1}{\beta}\text{OPT})^d$, again for d equal to the maximum size of the sets.

Finally, we identify a simple restriction of the agent's strategy, which guarantees that the cost of the solution computed by the agent is not more than β times the cost of an optimal solution.

Related Work

Our work is directly inspired by the aforementioned contribution of Kleinberg and Oren 2014, which was itself motivated by the earlier work by Akerlof 1991. We refer to (Kleinberg and Oren 2014, 2018) for a survey of earlier work on time-inconsistent planning, with connections to procrastination, abandonment, and choice reduction. Hereafter, we discuss solely (Kleinberg and Oren 2014), and the subsequent work. Using their graph-theoretic framework, Kleinberg and Oren reasoned about time-inconsistency effects. In particular, they provided a characterization of the graphs yielding the worst-case cost-ratio, and they showed that, despite the fact that the degree β of present bias can take all possible values in $[0,1]$, it remains that, for any given digraph, the collection of distinct s-t paths computed by present-biased agents for all degrees of present bias is of size at most polynomial in the number of nodes. They also showed how to improve the behavior of present-biased agents by deleting edges and nodes, and they provided a characterization of the subgraphs supporting efficient agent's behavior. Finally, they analyzed the case of a collection of agents with different degrees of present bias, and showed how to divide the global task to be performed by the agents into "easier" sub-tasks, so that each agent performs efficiently her sub-tasks.

As far as we are aware of, all contributions subsequent to (Kleinberg and Oren 2014), and related to our paper, essentially remain within the same graph theoretic framework as (Kleinberg and Oren 2014), and focus on algorithmic problems related to this framework. In particular, Albers and Kraft 2019 studied the ability to place rewards at nodes for motivating and guiding the agent. They show hardness and inaproximability results, and provide an approximation algorithm whose performances match the inaproximability bound. The same authors considered another approach in (Albers and Kraft 2017a) for overcoming these hardness issues, by allowing not to remove edges but to increase their weight. They were able to design a 2-approximation algorithm in this context. Tang et al. 2017 also proved hardness results related to the placement of rewards, and showed that finding a motivating subgraph is NP-hard. Gravin et al. 2016a (see (Gravin et al. 2016b) for the full paper) extended the model by considering the case where the degree of present bias may vary over time, drawn independently at each step from a fixed distribution. In particular, they described the structure of the worst-case graph for any distribution, and derived conditions on this distribution under which the worst-case cost ratio is exponential or constant.

Kleinberg, Oren, and Raghavan 2016; 2017 revisited the model in (Kleinberg and Oren 2014). In (Kleinberg, Oren, and Raghavan 2016), they were considering agents estimating erroneously the degree β of present bias, either underestimating or overestimating that degree, and compared the behavior of such agents with the behavior of "sophisticated" agents who are aware of their present-biased behavior in future and take this into account in their strategies. In (Kleinberg, Oren, and Raghavan 2017), they extended the model by considering not only agents suffering from present-biases, but also from *sunk-cost* bias, i.e., the tendency to incorporate costs experienced in the past into one's plans for the future. Albers and Kraft 2017b considered a model with uncertainty, bearing similarities with (Kleinberg, Oren, and Raghavan 2016), in which the agent is solely aware that the degree of present bias belongs to some set $B \subset (0,1]$, and may or may not vary over time.

Procrastination Under Minimization Problems

This section includes a formal definition of inconsistent planning by present-biased agents, and describes two extreme scenarios: one in which a present-biased agent constructs worst case plannings, and one in which the plannings generated by a present-biased agent are close to optimal.

Model and Definition

We consider minimization problems defined as triples $(\mathcal{I}, F, c)$, where $\mathcal{I}$ is the set of instances (e.g., the set of all graphs), F is a function that returns the set $F(I)$ of feasible solutions for every instance $I \in \mathcal{I}$ (e.g., the set of all edge-cuts of any given graph), and c is a non-negative function returning the cost $c(I,S)$ of every feasible solution $S \in F(I)$ of every instance $I \in \mathcal{I}$ (e.g., the number of edges in a cut). We focus solely on optimization problems for which

(i) a finite ground set $\mathcal{S}_I \neq \varnothing$ is associated to every instance I,

(ii) every feasible solution for I is a set $S \subseteq \mathcal{S}_I$, and

(iii) $c(I,S) = \sum_{x \in S} \omega(x)$ where $\omega : \mathcal{S}_I \to \mathbb{N}$ is a weight function.

Moreover, we enforce two properties that are satisfied by classical minimization problems. Specifically we assume that:

- All considered problems are closed downward, that is, for every considered minimization problem $(\mathcal{I}, F, c)$, every $I \in \mathcal{I}$, and every $x \in \mathcal{S}_I$, the instance $I \smallsetminus \{x\}$ defined by the feasible solutions $S \smallsetminus \{x\}$, for every $S \in F(I)$, is in $\mathcal{I}$ with the same weight function ω as for I. This guarantees that an agent cannot be stuck after having performed some task x, as the sub-problem $I \smallsetminus \{x\}$ remains solvable for every x.
- All considered feasible solutions are closed upward, that is, for every minimization problem $(\mathcal{I}, F, c)$, and every $I \in \mathcal{I}$, $\mathcal{S}_I$ is a feasible solution, and, for every $S \in F(I)$, if $S \subseteq S' \subseteq \mathcal{S}_I$ then $S' \in F(I)$. This guarantees that an agent performing a sequence of tasks $x_0, x_1, \dots$ eventually computes a feasible solution.

Inconsistent planning can be rephrased in this framework as follows.

Inconsistent planning. Let $\beta < 1$ be a positive constant. Given a minimization problems $(\mathcal{I}, F, c)$, the biased cost c_β satisfies $c_\beta(S) = \omega(x) + \beta\, c(S \setminus \{x\})$ for every feasible solution S of every instance $I \in \mathcal{I}$, where $x = \arg\min_{y \in S} \omega(y)$. Given an instance I, the agent aims at finding a feasible solution $S \in I$ by applying a present-based planning defined inductively as follows. Let $I_0 = I$. For every $k \geq 0$, given the instance I_k, the agent computes a feasible solution S_k with minimum cost $c_\beta(S_k)$ among all feasible solutions for I_k. Let $x_k = \arg\min_{y \in S_k} \omega(y)$. The agent stops whenever $\{x_0, x_1, \ldots, x_k\}$ is a feasible solution for I. Otherwise, it carries on the construction of the solution by considering $I_{k+1} = I_k \setminus \{x_k\}$.

Observe that inconsistent planning terminates. Indeed, since the instances of the considered problem $(\mathcal{I}, F, c)$ are closed downwards, $I_k = I \setminus \{x_0, \ldots, x_{k-1}\} \in \mathcal{I}$ for every $k \geq 0$, i.e., inconsistent planning is well defined. Moreover, since the feasible solutions are closed upward, there exists $k \geq 0$ such that $\{x_0, x_1, \ldots, x_k\}$ is a feasible solution for I.

The *cost* of inconsistent planning is defined as the ratio $\varrho = \frac{c(S)}{\text{OPT}}$ where $S = \{x_0, x_1, \ldots, x_k\}$ is the solution returned by the agent, and $\text{OPT} = c(S_{\text{OPT}})$ is the cost of an optimal solution S_{OPT} for the same instance of the considered minimization problem.

Approximated evaluation. It can happen that the considered minimization problem is computationally hard, say NP-hard, and the agent is unable to compute a feasible solution S of minimum cost $c_\beta(S)$ exactly. Then the agent can pick an approximate solution instead. For this situation, we modify the above strategy of the agent as follows. Assume that the agent has access to an α-approximation algorithm $\mathcal{A}$ that, given an instance I, computes a feasible solution S^* to the instance such that $c_\beta(S^*) \leq \alpha \min c_\beta(S)$, where minimum is taken over all feasible solution S to I. For simplicity, we assume throughout the paper that $\alpha \geq 1$ is a constant, but our results can be generalized for the case, where α is a function of the input size or OPT.

Again, the agent uses an inductive scheme to construct a solution. Initially, $I_0 = I$. For every $k \geq 0$, given the instance I_k, the agent computes a feasible solution S_k of cost at most $\alpha \min c_\beta(S)$, where the minimum is taken over all feasible solutions S of I_k. Then, exactly as before, the agent finds $x_k = \arg\min_{y \in S_k} \omega(y)$. If $\{x_0, x_1, \ldots, x_k\}$ is a feasible solution for I, then the agent stops. Otherwise, we set $I_{k+1} = I_k \setminus \{x_k\}$ and proceed. The *α-approximative cost* of inconsistent planning is defined as the ratio $\varrho_\alpha = \frac{c(S)}{\text{OPT}}$ where $S = \{x_0, x_1, \ldots, x_k\}$. Clearly, the 1-approximative cost coincides with ϱ.

Worst-Case Present-Biased Planning

We start with a simple observation. Given a feasible solution S for an instance I of a minimization problem, we say that $x \in S$ is *superfluous* in S if $S \setminus \{x\}$ is also feasible for I. The ability for the agent to make superfluous choices yields trivial scenarios in which the cost ratio ϱ can be arbitrarily large. This is for instance the case of an instance of set-cover, defined as one set $y = \{1, \ldots, n\}$ of weight $c > 1$ covering all elements, and n sets $x_i = \{i\}$, each of weight 1, for $i = 1, \ldots, n$. Every solution $S_i = \{x_i, y\}$ is feasible, for $i = 1, \ldots, n$, and satisfies $c_\beta(S_i) = 1 + \beta c$. As a result, whenever $1 + \beta c < c$, the present-biased agent constructs the solution $S = \{x_1, \ldots, x_n\}$, which yields a cost ratio $\varrho = n/c$, which can be made arbitrarily large as n grows. Instead, if the agent is bounded to avoid superfluous choices, that is, to systematically choose *minimal* feasible solutions, then only the feasible solutions $\{y\}$ and $\{x_1, \ldots, x_n\}$ can be considered. As a result, the agent will compute the optimal solution $S_{\text{OPT}} = \{y\}$ if $c < 1 + \beta(n-1)$.

Unfortunately, bounding the agent to systematically choose *minimal* feasible solutions, i.e., solutions with no superfluous elements, is not sufficient to avoid procrastination. That is, it does not prevent the agent from computing solution with high cost ratio. This is for instance the case of another instance of set-cover, that we denote by $I_{SC}^{(n)}$ for further references.

Set-cover instance $I_{SC}^{(n)}$: specified by $2n$ subsets of $\{1, \ldots, n\}$ defined as $x_i = \{i\}$ with weights 1, and $y_i = \{i, \ldots, n\}$ with weight $c > 1$, for $i = 1, \ldots, n$.

The minimal feasible solutions of $I_{SC}^{(n)}$ are $\{y_1\}$ of weight c, $\{x_1, \ldots, x_i, y_{i+1}\}$ of weight $i + c$ for $i = 1, \ldots, n-1$, and $\{x_1, \ldots, x_n\}$ of weight n. Whenever $1 + \beta c < c$, a time-biased agent bounded to make non-superfluous choices only yet constructs the solution $\{x_1, \ldots, x_n\}$ which yields a cost ratio $\varrho = n/c$, which can be made arbitrarily large as n grows. We need the following lemma about biased solutions for minimization problems.

Lemma 1. *Let $\alpha \geq 1$ and let S^* be a feasible solution for minimization problem, satisfying $c_\beta(S^*) \leq \alpha \min c_\beta(S)$, where the minimum is taken over all feasible solutions. Then*

(i) $\omega(x) \leq \alpha \cdot \text{OPT}$ *for $x = \arg\min_{y \in S^*} \omega(y)$, and*

(ii) $c(S^*) \leq \frac{\alpha}{\beta} \text{OPT}$.

Proof. Let S be an optimum solution. As $\beta < 1$, it follows that $\omega(x) \leq \omega(x) + \beta \cdot \omega(S^* \setminus \{x\}) = c_\beta(S^*) \leq \alpha \cdot c_\beta(S) \leq \alpha \cdot c(S) = \alpha \cdot \text{OPT}$, and this proves (i). To show (ii), note that $c(S^*) = \omega(x) + \omega(S^* \setminus \{x\}) = \frac{1}{\beta}(\beta\omega(x) + \beta\omega(S^* \setminus \{x\}))$, from which it follows that $c(S^*) \leq \frac{1}{\beta}(\omega(x) + \beta\omega(S^* \setminus \{x\})) = \frac{1}{\beta} c_\beta(S^*) \leq \frac{\alpha}{\beta} c_\beta(S) \leq \frac{\alpha}{\beta} c(S) = \frac{\alpha}{\beta}\text{OPT}$, which completes the proof. □

Lemma 1 has a simple consequence that also can be derived from the results of Gravin et al. 2016b, Claim 5.1, that we state as a theorem despite its simplicity, as it illustrates one major difference between our model and the model in (Kleinberg and Oren 2018).

Theorem 1. *For every $\alpha \geq 1$ and every minimization problem, the α-approximative cost ratio ϱ_α cannot exceed $\alpha \cdot k$ where k is the number of steps performed by the agents to construct the feasible solution $\{x_1, \ldots, x_k\}$ by following the time-biased strategy.*

Proof. By Lemma 1(i), at any step $i \geq 1$ of the construction, the agent adds an element $x_i \in \mathcal{S}_I$ in the current partial solution, and this element satisfies $\omega(x_i) \leq \alpha\, c_\beta(S_{\text{OPT}}) \leq \alpha\, c(S_{\text{OPT}}) = \alpha \cdot \text{OPT}$. Therefore, if the agent computes a solution $\{x_1, \dots, x_k\}$, then the α-approximative cost ratio for this solution satisfies $\varrho_\alpha = \sum_{i=1}^{k} \omega(x_i)/\text{OPT} \leq \alpha\, k$, as claimed. □

Remark. The bound in Theorem 1 is in contrast to the general model in (Kleinberg and Oren 2018), in which an agent performing k steps can incur a cost ratio exponential in k. This is because the model in (Kleinberg and Oren 2018) enables to construct graphs with arbitrary weights. In particular, in a graph such as the one depicted on Fig. 1, one can set up weights such that the weight of (v_1, t) is a constant time larger than the weight of (s, t), the weight of (v_2, t) is in turn a constant time larger than the weight of (v_1, t), etc., and still a present-biased agent starting from s would travel via $v_1, v_2, \dots, v_k$ before reaching t. In this way, the sum of the weights of the edges traversed by the agent may become exponential in the number of traversed edges. This phenomenon does not occur when focussing on minimization tasks. Indeed, given a partial solution, the cost of completing this solution into a global feasible solution cannot exceed the cost of constructing a global feasible solution from scratch.

It follows from Theorem 1 that $I_{SC}^{(n)}$ is a worst-case instance. Interestingly, this instance fits with realistic procrastination scenarios in which the agent has to perform a task (e.g., learning a scientific topic T) by either energetically embracing the task (e.g., by reading a single thick book on topic T), or starting first by an easier subtask (e.g., by first reading a digest of a subtopic of topic T), with the objective of working harder later, but underestimating the cost of this postponed hard work. The latter strategy may result in procrastination, by performing a very long sequence of subtasks $x_1, x_2, \dots, x_n$.

In fact, $I_{SC}^{(n)}$ appears to be *the essence of procrastination* in the framework of minimization problems. Indeed, we show that if the cost ratio is large, then the considered instance I contains an instance of the form $I_{SC}^{(n)}$ with large n. More precisely, we say that an instance I contains an instance J as a *minor* if the ground set $\mathcal{S}_J$ associated to J is a collection of subsets of the ground set $\mathcal{S}_I$ associated to I, that is $\mathcal{S}_J \subseteq 2^{\mathcal{S}_I}$, and, for every $\bar{S} \subseteq \mathcal{S}_J$, $\bar{S}$ is feasible for J if and only if $S = \bigcup_{\bar{x} \in \bar{S}} \bar{x}$ is feasible for I. Moreover, the weight function $\bar{\omega}$ for the elements of $\mathcal{S}_J$ must be induced by the one for $\mathcal{S}_I$ as $\bar{\omega}(\bar{x}) = \sum_{x \in \bar{x}} \omega(x)$ for every $\bar{x} \in \mathcal{S}_J$. Let $J^{(n)}$ be any instance of a minimization problem such that its associated ground set is $\mathcal{S}_{J^{(n)}} = \{x_1, \dots, x_n\} \cup \{y_1, \dots, y_n\}$, and the set of feasible solutions for $J^{(n)}$ is $F(J^{(n)}) = \big\{\{y_1\}, \{x_1, y_2\}, \{x_1, x_2, y_3\}, \dots, \{x_1, \dots, x_{n-1}, y_n\}, \{x_1, \dots, x_n\}\big\}$. The following result sheds some light on *why* the procrastination structure of Fig. 1 pops up.

Theorem 2 ($*$[1]). *Let I be an instance of a minimization problem for which the present-biased agent with parameter $\beta \in (0, 1)$ computes a solution for I with cost $\alpha \cdot \text{OPT}(I)$ for some $\alpha > 1$. Then I contains $J^{(n)}$ as a minor for some $n \geq \alpha$, and the present-biased agent with parameter β computes a solution for $J^{(n)}$ with cost $\alpha \cdot \text{OPT}(J^{(n)})$.*

[1]The proofs of the statements labled ($*$) are omitted.

Quasi-Optimal Present-Biased Planning

In the previous section, we have observed that forcing the agent to avoid superfluous choices, by picking minimal feasible solutions only, does not prevent it from constructing solutions that are arbitrarily far from the optimal. In this section, we show that, by enforcing consistency in the sequence of partial solutions constructed by the agent, such bad behavior does not occur. More specifically, given a feasible solution S for I, we say that x is *inconsistent* with S if $x \notin S$. The following result shows that inconsistent choices is what causes high cost ratio.

Theorem 3. *An agent using an α-approximation algorithm bounded to avoid inconsistent choices with respect to the feasible solutions used in the past for constructing the current partial solution returns an α/β-approximation of the optimal solution. This holds independently from whether the agent makes superfluous choices or not.*

Proof. Let I be an instance of a minimization problem $(\mathcal{I}, F, c)$. Let $S = \{x_0, \dots, x_k\}$ be the solution constructed by the agent for I, where x_i is the element computed by the agent at step i, for $i = 0, \dots, k$. Let S_i be the feasible solution of $I_i = I \smallsetminus \{x_0, \dots, x_{i-1}\}$ considered by the agent at step i. Since the agent is bounded to avoid any inconsistent choices with respect to the past, we have $x_i \in \cap_{j=0}^{i} S_j$ for every $i = 0, \dots, k$ because $x_i \notin S_j$ for some $j < i$ would be an inconsistent choice. It follows that $S \subseteq S_0$. Therefore, $c(S) \leq c(S_0)$. Since the agent uses an α-approximation algorithm, by Lemma 1(ii), $c(S_0) \leq \frac{\alpha}{\beta}\text{OPT}$ and the claim follows. □

Min/Maximization With Under/Overestimation

We first investigate the cost ratio for minimization problems for the case when $\beta > 1$. Similar bound was obtained by Kleinberg et al. (see (Kleinberg, Oren, and Raghavan 2016, Theorem 2.1)). However, their theorem is about sophisticated agents and cannot be applied in our case directly.

Theorem 4 ($*$). *Solutions computed by present-biased agents satisfy the following: For any minimization problem with $\beta > 1$, the cost ratio is at most β.*

Next, we consider maximization problems. The formalism for these variants can be set up in a straightforward manner by adapting the framework displayed in Section . We establish the following worst-case bounds.

Theorem 5 ($*$). *Solutions computed by present-biased agents satisfy the following:*

(i) For any maximization problem with $\beta < 1$, the cost ratio is at most $\frac{1}{\beta}$;

(ii) For any maximization problem with $\beta > 1$, the cost ratio is at most β^c, where $c \leq$ OPT is the cost of a solution constructed by the agent.

We also can write the bound for the cost ratio for $\beta > 1$ in the following form to obtain the upper bound that depends only on the value of OPT.

Corollary 1. *For any maximization problem with $\beta > 1$, the cost ratio is at most $(1 + \log \beta)\frac{\text{OPT}}{\log \text{OPT}}$.*

Proof. Let c be the cost of a solution constructed by the agent. By Theorem 5, $\text{OPT} \leq c\beta^c$. Therefore, $\log \text{OPT} \leq \log c + c \log \beta \leq (1 + \log \beta)c$, and $\frac{\text{OPT}}{c} \leq (1 + \log \beta)\frac{\text{OPT}}{\log \text{OPT}}$. $\square$

For minimization problems with $\beta > 1$, and maximization problems with $\beta < 1$, we have that the cost ratio is bounded by a constant. This differs drastically with the case of maximization problems with $\beta > 1$, when the cost ratio is still bounded but the bound is exponential. This exponential upper bound is however essentially tight, in the sense that the exponent cannot be avoided.

Theorem 6 ($*$)**.** *There are maximization problems for which a present-biased agent with $\beta > 1$ returns a solution whose cost ratio is at least $\frac{1}{c}\beta^{c-1}$, where c is the cost of the solution constructed by the agent.*

An example similar to the one in the proof of Theorem 6 can be constructed for the knapsack problem.

Covering and Hitting Problems

In Section , we have seen several instances of the set-cover problem whose cost ratio cannot be bounded by any function of OPT. The same obviously holds for the hitting-set problem. Recall that an instance of hitting-set is defined by a collection Σ of subsets of a finite set V, and the objective is to find the subset $S \subseteq V$ of minimum size, or minimum weight, which intersects (hits) every set in Σ. However, set-cover problems, and hitting set problems behave differently when the sizes of the sets are bounded. First, we consider the d-set cover problem.

The d-set cover problem. Let d be a positive integer. The task of the d-set cover problem is, given a collection Σ of subsets with size at most d of a finite set V, and given a weight function $\omega\colon \Sigma \to \mathbb{N}$, find a set $S \subseteq \Sigma$ of minimum weight that covers V, that is, $\bigcup_{X \in S} X = V$.

Theorem 7 ($*$)**.** *Let $\alpha \geq 1$. For any instance of the d-set-cover problem, the α-approximative cost ratio is at most $\alpha \cdot d \cdot \text{OPT}$.*

The d-hitting set problem. Let d be a positive integer. We are given a collection Σ of subsets with size d of a finite set V, a weight function $\omega\colon V \to \mathbb{N}$. The task is to find a set $S \subseteq V$ of minimum weight that hits every set of Σ.

Theorem 8 ($*$)**.** *Let $\alpha \geq 1$. For any instance of the d-hitting-set problem, the α-approximative cost ratio is at most $\alpha d!\,(\frac{\alpha}{\beta}\text{OPT})^d$.*

Conclusion

We demonstrated that, by focussing on present-biased agents solving tasks, specific detailed analysis can be carried on for each considered task, which enables to identify very different agent's behavior depending on the tasks (e.g., set cover vs. hitting set). Second, focussing on present-biased agents solving tasks enables to generalize the study to over-estimation, and to maximization, providing a global picture of searching via present-biased agents. Yet, lots remain to be done for understanding the details of this picture.

In particular, efforts could be made for studying other specific classical problems in the context of searching by a present-biased agent. This includes classical optimization problems like traveling salesman (TSP), metric TSP, maximum matching, feedback vertex set, etc. Such study may lead to a better understanding of the class of problems for which present-biased agents are efficient, and the class for which they act poorly. And, for problems for which present-biased agents are acting poorly, it may be of high interest to understand what kind of restrictions on the agent's strategy may help the agent finding better solutions.

Another direction of research is further investigation of the influence of using approximation algorithms by agents, as it is natural to assume that the agents are unable compute the cost exactly. We made some initial steps in this direction, but it seems that this area is almost unexplored. For example, it can be noted that the upper bound for the cost ratio in Theorem 4 can be rewritten under the assumption that the agent uses an α-approximation algorithm. However, the bound gets blown-up by the factor α^s, where s is the size of the solution obtained by the agent (informally, we pay the factor α on each iteration). From the other side, the examples in Section show that this is not always so. Are there cases when this exponential blow-up unavoidable? The same question can be asked about maximization problems.

Acknowledgements

The paper received support from the Research Council of Norway via the project "MULTIVAL" (grant no. 263317), and from the French National Research Agency (ANR) via the project "DESCARTES".

References

Akerlof, G. A. 1991. Procrastination and obedience. *American Economic Review: Papers and Proceedings* 81(2): 1–19.

Albers, S.; and Kraft, D. 2017a. On the Value of Penalties in Time-Inconsistent Planning. In *44th International Colloquium on Automata, Languages, and Programming (ICALP)*, 10:1–10:12.

Albers, S.; and Kraft, D. 2017b. The Price of Uncertainty in Present-Biased Planning. In *13th Int. Conference on Web and Internet Economics (WINE)*, 325–339.

Albers, S.; and Kraft, D. 2019. Motivating Time-Inconsistent Agents: A Computational Approach. *Theory Comput. Syst.* 63(3): 466–487.

Gravin, N.; Immorlica, N.; Lucier, B.; and Pountourakis, E. 2016a. Procrastination with Variable Present Bias. In *ACM Conference on Economics and Computation (EC)*, 361.

Gravin, N.; Immorlica, N.; Lucier, B.; and Pountourakis, E. 2016b. Procrastination with variable present bias. *CoRR* abs/1606.03062.

Jepson, C.; Loewenstein, G.; and Ubel, P. 2001. Actual versus Estimated Differences in Quality of Life before and after Renal Transplant. Working Paper, Department of Social and Decision Sciences, Carnegie Mellon University.

Kleinberg, J. M.; and Oren, S. 2014. Time-inconsistent planning: a computational problem in behavioral economics. In *ACM Conference on Economics and Computation (EC)*, 547–564.

Kleinberg, J. M.; and Oren, S. 2018. Time-inconsistent planning: a computational problem in behavioral economics. *Commun. ACM* 61(3): 99–107.

Kleinberg, J. M.; Oren, S.; and Raghavan, M. 2016. Planning Problems for Sophisticated Agents with Present Bias. In *ACM Conference on Economics and Computation (EC)*, 343–360.

Kleinberg, J. M.; Oren, S.; and Raghavan, M. 2017. Planning with Multiple Biases. In *ACM Conference on Economics and Computation (EC)*, 567–584.

Loewenstein, G.; O'Donoghue, T.; and Rabin, M. 2003. Projection bias in predicting future utility. *The Quarterly Journal of economics* 118(4): 1209–1248.

Tang, P.; Teng, Y.; Wang, Z.; Xiao, S.; and Xu, Y. 2017. Computational issues in time-inconsistent planning. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI)*.