Diploma Project

Graph-Theoretic Problems in Stochastic Planning

Sotiris Ektoros

University of Cyprus
Department of Computer Science

Department of Computer Science

January 2026

UNIVERSITY OF CYPRUS

Faculty of Pure and Applied Sciences

Department of Computer Science

Graph-Theoretic Problems in Stochastic Planning

Sotiris Ektoros

Advisor:

Marios Mavronicolas

The Thesis was submitted in partial fulfillment of the requirements for obtaining the Computer Science degree of the Department of Computer Science of the University of Cyprus

January 2026

Abstract

This thesis studies graph-theoretic aspects of stochastic planning under behavioral biases, starting from the sunk cost model of Kleinberg, Oren, Raghavan, and Sklar [1]. In this model, an agent plans in a directed acyclic graph with probabilistic transitions, pays edge costs along a path, and receives a fixed reward only upon reaching a designated target state. Sunk cost bias is captured by a parameter $\lambda \geq 0$, if the agent abandons after incurring total cost C_{sunk}, it experiences an additional psychological loss of $\lambda \cdot C_{\text{sunk}}$. The performance of a biased agent is measured by the loss in expected payoff relative to an optimal, unbiased agent facing the same stochastic environment. Kleinberg et al. established general upper bounds on this loss for *sophisticated* agents who correctly anticipate their own future sunk cost distortions and provided evidence that fan graphs might be the worst-case for the gap between sophisticated and optimal payoffs.

The main contribution of this thesis is to resolve that open question in the negative. Working in the same stochastic sunk cost framework, we introduce a new class of *Layered Graph With Skips* and construct within it a family of instances in which the payoff gap between the sophisticated and optimal agents is strictly larger than in fan graphs. At the technical level, we develop a collection of structural lemmas that track how sunk costs accumulate along multiple paths in a layered stochastic environment, and how a sophisticated agent's forward-looking continuation decisions interact with this accumulation. These lemmas yield a general upper bound on sophisticated loss of the form

$$\pi_o(s) - \pi_\sigma(s) \;\leq\; \lambda\, a\, b\, m R \left(1 - \frac{1}{k}\right)^k,$$

where a, b, and m are structural parameters of the graph, R is the reward obtained upon reaching the terminal node, π_o denotes the expected payoff of the optimal agent, π_σ the expected payoff of the sophisticated agent, and s the starting node. Moreover, we exhibit an asymptotically tight worst-case family within this class for which the loss approaches $\frac{\lambda abmR}{e}$, and in which the product abm can exceed 1.This shows that fan graphs are not the worst-case and that more damaging network topologies exist.

Beyond the sunk cost setting, the thesis also extends the Layered Graph With Skips construction to *present biased, reward-seeking agents* who overweight immediate rewards relative to future ones. We analyze the behavior of sophisticated present biased agents and optimal agents on our layered family and show that it generates an exponential gap in reward, matching the asymptotic behavior of fan graphs up to constant factors. Taken together, the results suggest general principles for designing environments that either amplify or mitigate the impact of behavioral biases, and they illustrate how a common structural template Layered Graph With Skips can support worst-case analyses for multiple distortions, including sunk cost bias and present bias.

Contents

Chapter 1

Introduction

1.1 Motivation and Purpose of the Thesis

Human decision making is frequently skewed by *sunk cost bias* once resources have been invested in a project, individuals are inclined to persist even when a fresh evaluation would recommend stopping. This phenomenon has been documented in diverse settings, from personal finance to large-scale organizational decisions, and it becomes especially subtle when decisions unfold over time under uncertainty. In sequential environments, an agent repeatedly decides whether to continue or abandon a project, facing stochastic transitions between states and uncertain future costs and rewards. Understanding how sunk cost bias interacts with such uncertainty is therefore important both for behavioral economics and for the algorithmic study of planning problems.

Recent theoretical work has introduced a formal model of sunk cost bias in stochastic planning, due to Kleinberg, Oren, Raghavan, and Sklar [1]. In this model, an agent plans in a directed acyclic graph with probabilistic transitions, pays costs along a path, and receives a fixed reward only upon reaching a designated target state. Sunk cost bias is captured by a parameter $\lambda \geq 0$, if the agent abandons after incurring total cost C_{sunk}, it experiences an additional psychological loss of $\lambda \cdot C_{\text{sunk}}$. The benchmark is an optimal unbiased agent that maximizes expected payoff in the same stochastic environment. A central focus of this line of work is *sophisticated* agents, who anticipate their own future sunk cost bias and plan accordingly. The performance of such an agent is measured by the gap between its expected payoff and that of the optimal unbiased agent.

Kleinberg et al. obtained general upper bounds on the payoff loss of sophisticated sunk cost agents and identified *fan graphs* as natural candidates for worst-case instances in these graphs, the gap between sophisticated and optimal payoffs can be made on the order of $\frac{\lambda R}{e}$, where R is the terminal reward. They also provided evidence that fan graphs might in fact be worst-case, but stopped short of a definitive characterization. This leaves open a fundamental structural question among all stochastic planning environments of this form, which network topologies yield the largest possible performance loss for sophisticated agents with sunk cost bias?

This thesis addresses that question. Working in the same stochastic sunk cost framework, we introduce an alternative class of environments based on *Layered Graphs with Skips* and develop an analytical framework for sophisticated agents in these graphs. Conceptually, the classical fan graph construction fits naturally into this framework. A fan graph can be viewed as a degenerate Layered Graph With Skips the states lie on a single chain of layers, forward transitions move to the next layer, and each "fan" edge corresponds to a skip directly to the terminal layer. In this sense, fan graphs are a highly restricted special case of Layered Graphs With Skips, with only one node per layer and skips only to the final layer. Our model strictly generalizes this template by allowing multiple states per layer and skip transitions between arbitrary layers. Within this framework, we prove a general upper bound on the loss of a sophisticated agent in terms of structural parameters of the instance, and we construct a family of Layered Graphs with Skips in which the sophisticated-optimal gap is strictly larger than in fan graphs. In particular, we show that in this family, the loss can approach

$$\frac{\lambda\, abmR}{e},$$

where a, b, and m are structural parameters describing the graph, and it is possible to have $abm > 1$. This demonstrates that fan graphs are not worst-case instances; there exist network topologies that generate an even greater degradation in expected payoff due to sunk cost bias.

The analysis proceeds via a collection of structural lemmas about biased planning in layered stochastic graphs. These lemmas track how sunk costs accumulate along multiple paths and how a sophisticated agent's forward-looking continuation decisions interact with this accumulation. They yield new upper and lower bounds on the agent's expected payoff and identify asymptotically tight worst-case instances within the class of Layered Graph With Skips.

Beyond the sunk cost setting, We also extend the Layered Graph With Skips construction to *present biased, reward-seeking agents* who overweight immediate rewards relative to future ones. We analyze the behavior of a sophisticated present biased agent and an optimal agent on the same layered family and show that it produces an exponential gap in reward, matching the asymptotic behavior of the classical fan graph constructions up to constant factors. This illustrates that Layered Graph With Skips form a flexible structural template for worst-case analysis across multiple behavioral distortions, including sunk cost bias and present bias.

1.2 Related Literature

This thesis builds directly on the stochastic sunk-cost framework of Kleinberg, Oren, Raghavan, and Sklar [1]. More broadly, a line of work has developed graph-theoretic models of planning to study how different behavioral distortions affect sequential decision-making. These include deterministic planning models for present-biased sophisticated agents [2], extensions in which multiple behavioral biases interact, notably present bias and sunk-cost bias [4], and principal–agent formulations where incentives are shaped through subgraph design for present-biased agents [3]. In addition, related graph-based planning frameworks have been studied in other contexts, such as stochastic planning against prophet benchmarks [5] and planning models capturing projection bias [6].

1.3 Structure of the Thesis

Chapter 2 presents the formal stochastic planning model, introduces Layered Graph With Skips, and sets up the notation and fundamental quantities used throughout the thesis.

Chapter 3 develops a general upper bound on the loss of sophisticated sunk cost agents in Layered Graph With Skips. The chapter proves structural lemmas relating costs, transition probabilities, and optimal payoffs, and then combines them into a bound of the form

$$\pi_o(s) - \pi_\sigma(s) \leq \lambda\, abmR \left(1 - \frac{1}{k}\right)^k,$$

where a, b, and m are structural parameters of the graph, π_o denotes the expected payoff of the optimal agent, π_σ the expected payoff of the sophisticated agent, and s the starting node. Moreover, we exhibit an asymptotically tight worst-case family within the class of Layered Graph With Skips. It shows that the upper bound can be achieved up to lower-order terms, that the loss can scale like $\frac{\lambda abmR}{e}$, and that there exist instances with $abm > 1$. The chapter concludes by comparing these layered constructions with the fan graph examples of Kleinberg et al. [1] and showing that fan graphs are not the worst-case.

Chapter 4 extends the framework to present biased, reward-seeking agents. It adapts the Layered Graph With Skips construction to a reward-based model with present bias, analyzes the behavior of sophisticated and optimal agents on this family, and compares the resulting exponential reward gap to known present bias results on fan graphs.

Chapter 5 concludes the thesis. It summarizes the main contributions, discusses the broader implications for the design and analysis of stochastic environments with behavioral agents, and outlines directions for extending the approach to heterogeneous structural settings, paths of non-uniform lengths, and interactions with other behavioral distortions.

Chapter 2

The Stochastic Planning Model and Foundational Properties for Sunk Cost Bias

2.1 The Layered Graph with Skip Edges: Model Definition

A Layered Graph With Skips contains various paths of differing lengths that a single agent may traverse from one node to another. The graph consists of L layers, with each layer containing a finite number of nodes, which may vary across layers. For simplicity, we presume that all pathways terminating at the node of interest (namely, node u_k) possess identical lengths. An illustration of a representative path is the path $(s = u_1, u_2, \ldots, u_n, t)$, where n denotes the number of nodes the agent traverses before reaching the target, and where for each integer i such that $1 \leq i < n$, the transition probability is defined as $p(u_i, u_{i+1}) = 1/j$ for a fixed integer $j \geq 1$, with $p(u_n, t) = 1$. Traversing from a node u_i incurs a cost $c(u_i) = c_i$, with c_n denoting the cost incurred at the final non-terminal node. A depiction of a Layered Graph With Skips is presented in Figure 2.1.

Layered Graphs With Skips can represent numerous circumstances. For example, they can illustrate a project development process wherein, at each stage, one must select an investment level that enables the project to proceed; if an erroneous choice is made, the project ceases.

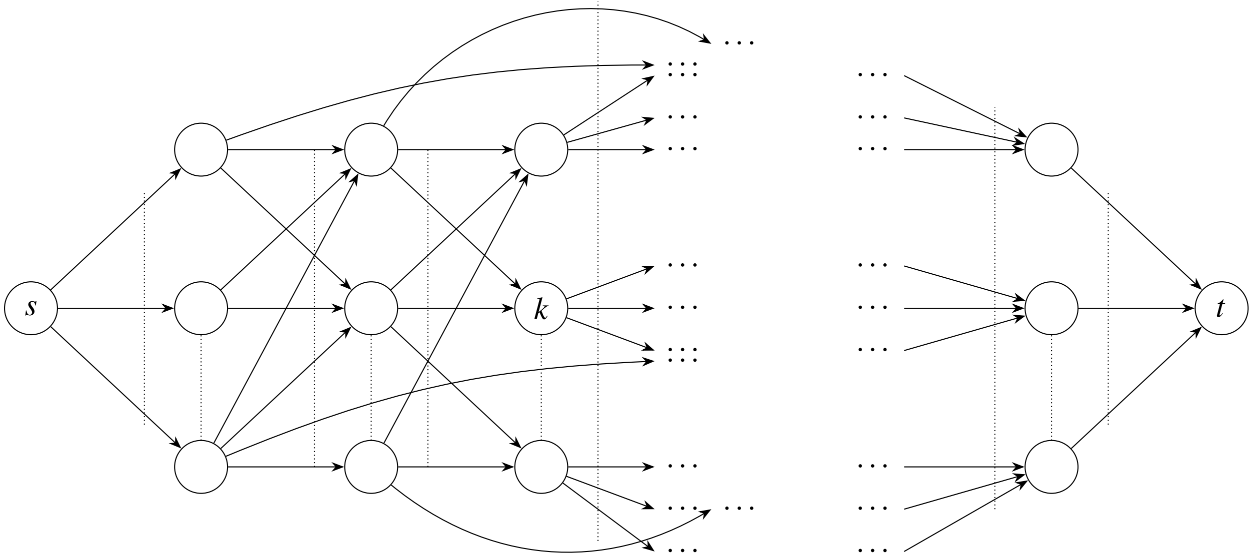

Figure 2.1: Layered Graph with Skip Edges

2.2 Fundamental Quantities and Structural Properties

We collect here the notation and quantities that will be used throughout the analysis.

L**:** Number of layers of the graph.

j**:** Number of next nodes that a non-terminal node has (a layer may contain more or fewer than j nodes.

$1/j$**:** Uniform transition probability from a node to each of its j next nodes.

c_i**:** Cost incurred when traversing from a node in layer i to one of its next nodes.

s**:** Starting node of the graph.

t**:** Terminal (target) node where the agent receives the reward R.

$\pi_o(u)$**:** The expected payoff when the optimal agent starts from node u:

$$\pi_o(u) = \max \left\{ \sum_{v \in N(u)} p(u,v)\, \pi_o(v) - c(u),\, 0 \right\}.$$

$\pi_\sigma(u)$**:** The expected payoff when the sophisticated agent starts from node u

u_k**:** Node at which, by assumption, the optimal agent stops.

S**:** Event that the optimal agent eventually reaches the target node t. The probability of failure is therefore $1 - p(S)$.

m**:** Total number of distinct paths that reach node u_k.

$E[C \mid \overline{S}]$**:** Expected total cost conditional on the event that the agent stops before reaching the target.

Chapter 3

Bounding Sophisticated Loss in Layered Graphs With Skips for Sunk Cost Bias

In this chapter, we prove an upper bound on the loss of the sophisticated agent on Layered Graphs with Skips. The proof is based on a sequence of structural lemmas that control the costs, probabilities, and structural parameters of the graph, and then we combine them to obtain the desired bound on $\pi_\sigma(u)$ in terms of $\pi_o(u)$.

3.1 The General Upper-Bound Theorem

We are now ready to state and prove the main result of this chapter.

Theorem 1. *There are Layered Graphs With Skips such that* $\pi_\sigma(s) \geq \pi_o(s) - \lambda\, abmR \left(1-\frac{1}{k}\right)^k$.

Proof. By Corollary 3.5 of Kleinberg et al. [1], the sophisticated agent's payoff satisfies $\pi_\sigma \geq \pi_o - \lambda(1-p(S))\,\mathbb{E}[C \mid \overline{S}]$, where S is the event that the optimal agent reached the target and $\mathbb{E}[C \mid \bar{S}]$ is the expected cost of the optimal agent for paths in which it stopped before the target. To bound the payoff of the sophisticated agent, we bound $(1-p(S))\,\mathbb{E}[C \mid \overline{S}]$. Assume without loss of generality that the optimal agent stopped traversing at u_k. In the Layered Graph With Skips, many paths can lead to u_k (we assume for our convenience that all these paths have the same length). Thus,

$$(1-p(S))\cdot\mathbb{E}[C \mid \bar{S}] = m\left(\prod_{i=1}^{k-1}(p_i)\right)\cdot\left(\sum_{i=1}^{k-1}c_i\right).$$

We now use the fact that the optimal agent traverses the graph till u_k to get an upper bound on $(1-p(S))\,\mathbb{E}[C \mid \overline{S}]$.

□

3.2 Structural Lemmas for Costs and Continuation

We begin with a lemma relating the costs and the continuation probabilities along an optimal path.

Lemma 1. *If the optimal agent reaches u_k along a path, then $\sum_{i=1}^{k-1} c_i \leq \sum_{i=1}^{k-1} \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i\, \pi_o(u)$.*

Proof. Define $\pi_o(u_j, u_k)$ as the expected reward of the optimal agent navigating the graph from the node u_j. We demonstrate using backward induction that for all j≥1, $\pi_o(u_j, u_k) \leq \sum_{i=j}^{k-1} \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i\, \pi_o(u) - c_i$. In the base case, we note that when the agent proceeds from node u_{k-1}, it holds true that $\pi_o(u_{k-1}, u_k) = \sum_{u \in next(k-1)} p_i\, \pi_o(u) - c_{k-1} = p_{k-1} \pi_o(u_k) + \sum_{\substack{u \in next(k-1) \\ u \neq k}} p_i\, \pi_o(u) - c_{k-1} = \sum_{\substack{u \in next(k-1) \\ u \neq k}} p_i\, \pi_o(u) - c_{k-1}$. In the induction stage, we assume the validity of u_{j+1} and demonstrate it for an agent navigating the graph commencing from u_j. It is noted that as the agent moves across the graph from u_j to u_k, the equation is expressed as $\pi_o(u_j, u_k) = p_j \pi_o(u_{j+1}) + \sum_{\substack{u \in next(j) \\ u \neq j+1}} p_i\, \pi_o(u) - c_j$. According to the induction hypothesis, it follows that $\pi_o(u_{j+1}, u_k) \leq \sum_{i=j+1}^{k-1} \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i\, \pi_o(u) - c_i$ By consolidating this information, we derive that $\pi_o(u_j, u_k) \leq \sum_{i=j}^{k-1} \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i\, \pi_o(u) - c_i$, as necessitated. The expected payoff for the optimal agent attaining u_k is no greater than $\sum_{i=1}^{k-1} \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i\, \pi_o(u) - c_i$. Given that this amount must be non-negative, it follows that $\sum_{i=1}^{k-1} c_i \leq \sum_{i=1}^{k-1} \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i\, \pi_o(u)$. □

Next, we state separate upper and lower bounds on the optimal payoff.

Lemma 2. *The upper bound on the optimal payoff is at most R.*

Proof. Assume for the sake of contradiction that $\pi_o(u) > R$. However, this contradicts Proposition 3.3 in the stochastic sunk cost model [1], that proves $\pi_o(u) = p(S) \cdot R \;-\; p(S) \cdot E[C \mid S] \;-\; (1-p(S)) \cdot E[C \mid \overline{S}]$, thus the payoff can never be bigger than the reward but it can be at most equal to it. □

Lemma 3. *The lower bound on the optimal payoff is at least 0.*

Proof. Assume for the sake of contradiction that $\pi_o(u) < 0$. However, this contradicts the optimal expected payoff equation that states that the optimal expected payoff is at least zero, $\pi_o(u) = \max\left\{\sum_{v \in N(u)} p(u,v)\, \pi_o(v) - c(u),\, 0\right\}$, thus the minimum optimal expected payoff is zero. □

Combining the previous two lemmas, we can normalize the optimal payoff as a fraction of the reward.

Lemma 4. *There exists a parameter $b \in [0,1]$ such that $\pi_o(u) \;=\; bR$.*

Proof. Assume for the sake of contradiction that there is not a $b \in [0,1]$ such that $\pi_o(u) = bR$. From Lemma 2 and Lemma 3, we proved that $\pi_o(u) \in [0,R]$ for all the nodes of the Graph. We observe that we can alter the equation $\pi_o(u) = bR$ to $b = \frac{\pi_o(u)}{R}$. After this alteration, we use the $\pi_o(u) \in [0,R]$, we see that $0 \leq \pi_o(u) \leq R$, now we divide this with R and we observe that now we have $0 \leq \frac{\pi_o(u)}{R} \leq \frac{R}{R}$ and this is $0 \leq \frac{\pi_o(u)}{R} \leq 1$. However, this $\frac{\pi_o(u)}{R}$ is equal to b and leads to a contradiction, and thus there is a $b \in [0,1]$ such that $\pi_o(u)$. □

We also isolate the structural parameter a describing the fraction of non-terminal next nodes among the next nodes of a node.

Lemma 5. *Let a given node have j next nodes in total, of which $j-1$ are not part of the path, and of which d are non-terminal of the $j-1$ next nodes not in the path. Define $a = \frac{d}{j-1}$. Then $a \in [0,1]$.*

Proof. Assume for the sake of contradiction that there is not a $a \in [0,1]$ such that $a = \frac{d}{j-1}$. We know that d is the total of non-terminal next nodes from the $j-1$ next nodes of the node, that is, the total next nodes that have their payoff greater than zero. We know that $0 \leq d \leq j-1$ from the subset rule, now, we divide by $j-1$, and thus we have $0 \leq \frac{d}{j-1} \leq \frac{j-1}{j-1}$. However this $\frac{d}{j-1}$ is equal to a and leads to contradiction, and thus there is a $a \in [0,1]$ such that $a = \frac{d}{j-1}$. □

3.3 Parameter Interpretation and Upper-Bound Construction

We now describe how the parameters a, b, and m arise from the Layered Graph with Skips, and how they interact with sunk cost bias.

Let π_o denote the expected payoff of the optimal agent, and let π_σ be the expected payoff of a sophisticated agent with sunk cost parameter λ.

- The optimal payoff can be written as $\pi_o(s) = bR$ for some $b \in [0,1]$ (Lemma 4).
- The structural parameter $a \in [0,1]$ describes the fraction of non-terminal next nodes among the next nodes of a typical node (Lemma 5).
- The quantity m denotes the total number of distinct paths that reach a given critical node (for example, u_k) in the layered graph.

Using Lemma 1, we can upper bound the total cost incurred along an optimal path by the continuation probabilities and the optimal payoffs of the off-path next nodes. Recall that Lemma 1 states

$$\sum_{i=1}^{k-1} c_i \leq \sum_{i=1}^{k-1} \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i \pi_o(u).$$

Substituting this inequality into the expression for the conditional expected sunk cost, we obtain

$$(1-p(S)) \cdot \mathbb{E}[C \mid \bar{S}] \leq m \left(\prod_{i=1}^{k-1} (p_i) \right) \cdot \left(\sum_{i=1}^{k-1} \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i \pi_o(u) \right).$$

We aim to eradicate the reliance on $\pi_o(u)$ and articulate this bound only in terms of the reward R and the graph's structural properties. This is the application of Lemma 4 and Lemma 5. Lemma 4 indicates that for each node u, we may express $\pi_o(u) = bR$ for some $b \in [0,1]$, whereas Lemma 5 represents the proportion $a \in [0,1]$ of non-terminal offspring among the $j-1$ off-path next nodes of a typical node. In our stratified Layered Graph With Skips construction, we select the case to ensure uniformity of parameters throughout the graph, each pertinent off-path kid possesses a reward $\pi_o(u) = bR$, and at each layer, the proportion of non-terminal off-path next nodes is precisely a. Thus, we can regard a and b as structural constants of the instance, thereby constraining each $\pi_o(u)$ by bR and limiting the number of pertinent next nodes at a layer by a factor of a, resulting in

$$(1-p(S)) \cdot \mathbb{E}[C \mid \bar{S}] \;\leq\; m\left(\prod_{i=1}^{k-1}(p_i)\right) \cdot \left(\sum_{i=1}^{k-1} ab \sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i R\right).$$

In particular, this shows that the entire quantity $(1-p(S)) \cdot \mathbb{E}[C \mid \bar{S}]$ can be bounded in terms of the continuation probabilities p_i, the reward R, and the structural parameters a, b, and m.

Under our assumptions, each non-terminal node has exactly j next nodes and the transition probabilities are uniform, so $p_i = 1/j$ for all layers i. Moreover, at each layer there are at most $a(j-1)$ non-terminal off-path next nodes, and for each such child we have $\pi_o(u) \leq bR$. Hence, for every layer i,

$$\sum_{\substack{u \in next(i) \\ u \neq i+1}} p_i \pi_o(u) \;\leq\; \frac{1}{j} \cdot a(j-1) \cdot bR \;=\; ab\,\frac{j-1}{j}R.$$

Summing over the $k-1$ layers on the optimal path and using $\prod_{i=1}^{k-1} p_i = (1/j)^{k-1}$, we obtain

$$(1-p(S)) \cdot \mathbb{E}[C \mid \bar{S}] \;\leq\; m\left(\frac{1}{j}\right)^{k-1} \cdot \left((k-1)\,ab\,\frac{j-1}{j}R\right) = mabR\,(k-1)\,\frac{j-1}{j^k}.$$

It is convenient to write this bound as

$$f(j) = mab\,(k-1)\,\frac{j-1}{j^k},$$

and view $j > 1$ as a real variable. Differentiating with respect to j, we get

$$\frac{\partial f}{\partial j} = mab(k-1)\,\frac{-kj+k+j}{j^{k+1}}.$$

The numerator vanishes when

$$-kj+k+j = 0 \iff (k-1)j = k \iff j = \frac{k}{k-1},$$

so at the critical point, we have

$$\frac{1}{j} = \frac{k-1}{k} = 1 - \frac{1}{k}.$$

Substituting this value of j into $f(j)$ yields

$$f(j) = mab\left(1 - \frac{1}{k}\right)^k.$$

Therefore,

$$(1-p(S))\cdot\mathbb{E}[C\mid\bar{S}]\;\le\; mabR\left(1-\frac{1}{k}\right)^{k}.$$

In particular, the bound in Theorem 1 depends on the factor $\left(1-\frac{1}{k}\right)^{k}$, whose limit as $k\to\infty$ is $\frac{1}{e}$; this asymptotic behavior will be exploited to obtain a worst-case family with loss on the order of $\frac{\lambda abmR}{e}$.

We now build on the general upper bound and show that, for an appropriate Family of Layered Graph With Skips, the sophisticated agent's loss approaches the asymptotic value suggested by that bound. Recall that Theorem 1 established

$$\pi_\sigma(s)\;\ge\;\pi_o(s)\;-\;\lambda\,abmR\left(1-\frac{1}{k}\right)^{k},$$

and the previous discussion highlighted that

$$\left(1-\frac{1}{k}\right)^{k}\longrightarrow\frac{1}{e}\quad\text{as } k\to\infty.$$

Our goal is to construct a family of Layered Graph With Skips whose loss asymptotically matches this bound, yielding a worst-case value on the order of $\frac{\lambda abmR}{e}$.

3.4 Tightness Theorem for Layered Graph With Skips

We are now ready to state the main result of this chapter.

Theorem 2. *There exists a family of instances in which λ is a function of n, such that as n goes to infinity*

$$\lim_{n\to\infty}\left(\pi_o-\frac{\lambda\,abmR}{e}\right)\;=\;\pi_\sigma.$$

Proof. We examine a Layered Graph With Skips where the transition probabilities and costs along the paths are uniform. For each $1\le i<n$, $p_i=p=\frac{1}{j}=\frac{n-1}{n}$ and $c_i=c=\frac{ab}{n}-\frac{ab}{n^2}$. We additionally establish $R=1$ and indicate that there exists a layer where all nodes possess the same value b. This graph is illustrated in Figure 3.1. The optimal agent navigates the graph until it arrives at u_n, as the expected reward of each step is unequivocally positive. The optimal agent would cease at u_n since the expected payoff of the final step is negative. The expected payoff of the optimal agent is

$$\pi_o=m\sum_{i=1}^{n-1}\left(\frac{1}{j}\right)^{i-1}\left(ab\frac{j-1}{j}-c\right)$$

$$=m\frac{j-\left(\frac{1}{j}\right)^{n-2}}{j-1}\left(ab\frac{j-1}{j}-c\right)$$

$$=m\left(\left(\frac{n}{n-1}\right)-\left(\frac{n-1}{n}\right)^{n-2}\right)\frac{n-1}{n^2}ab$$

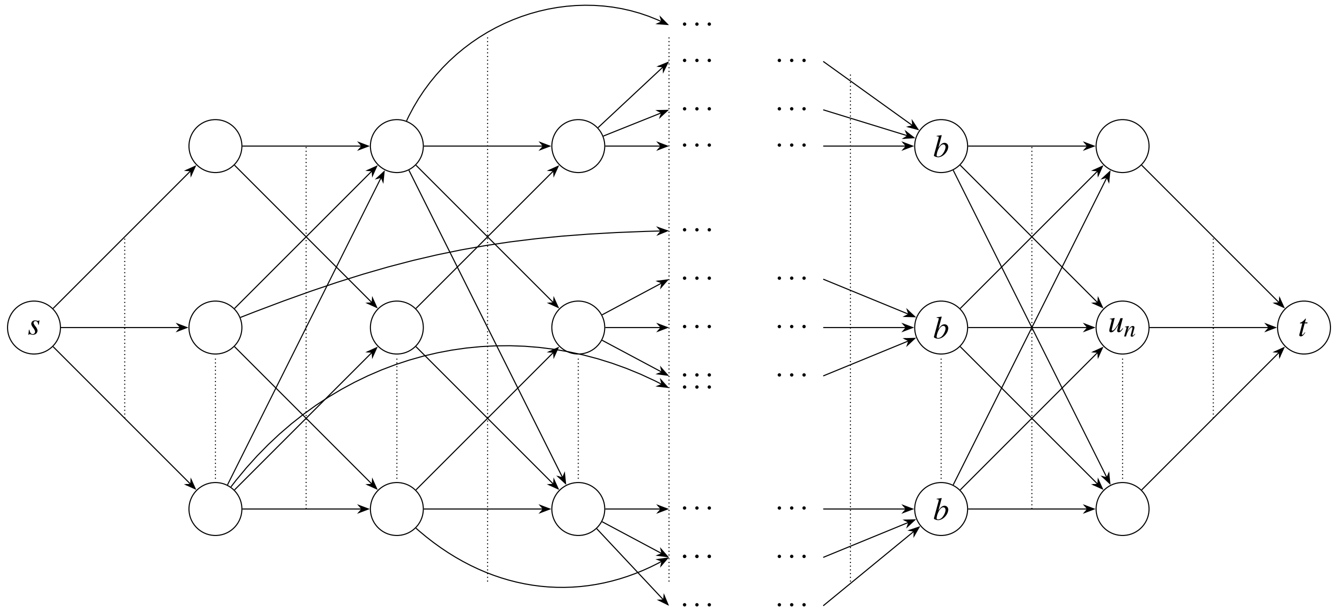

Figure 3.1: Layered Graph with Skip Edges with constant b and a

Next we let $\lambda = \frac{\pi_o}{m(\frac{1}{j})^{n-1}(n-1)c}$. We show that it is indeed the case when $\lambda \leq 1$.

Lemma 6. *If* $\pi_o = m\left(\left(\frac{n}{n-1}\right) - \left(\frac{n-1}{n}\right)^{n-2}\right)\frac{n-1}{n^2}ab$*, then for* $\lambda = \frac{\pi_o}{m\left(\frac{1}{j}\right)^{n-1}(n-1)c}$*, we have* $\lambda \leq 1$.

Proof. By plugging the values of $\pi_o, \frac{1}{j}$ and c we get

$$\frac{\left(\left(\frac{n}{n-1}\right) - \left(\frac{n-1}{n}\right)^{n-2}\right)\frac{1}{n}}{\left(\frac{n-1}{n}\right)^{n}}$$

To show that $\lambda \leq 1$, it is sufficient to show that

$$1 \leq \left(\frac{n-1}{n}\right)^{n+1} n + \left(\frac{n-1}{n}\right)^{n-1}$$

Let $f(n) = \left(\frac{n-1}{n}\right)^{n+1} n + \left(\frac{n-1}{n}\right)^{n-1}$. Observe that for $f(3) = \frac{28}{27}$. Thus showing that $f(n)$ is increasing will complete the proof.

$$f'(n) = \frac{\left(\frac{n-1}{n}\right)^{n}\left((n^2-n+1)\ln\left(\frac{n-1}{n}\right) + 2n - 1\right)}{n-1}.$$

and by using calculus, one can show that it is indeed the case that $f'(n) > 0$ for any $n > 2$, which completes the proof.

Note that $\lim_{n\to\infty} m\left(\frac{1}{j}\right)^{n-1}(n-1)c = \frac{mab}{e}$. Consequently, when n approaches infinity, π_o approaches $\frac{mab}{e}$. To finalize the proof, we demonstrate that for this value of λ, $\pi_\sigma = 0$. Consequently, we demonstrate that if the agent begins to navigate the Graph, it will reach t, resulting in an expected reward of zero. Note that the expected reward of a stochastic agent progressing from node u_k to time t is

$$\pi_\sigma(u_k) = \pi_o(u_k) - m\lambda\left(\frac{1}{j}\right)^{n-k}(n-1)c.$$

The initial term represents the reward of the sophisticated agent reaching t prior to u_n, equivalent to that of the optimal agent, while the subsequent term denotes the expense incurred by the sophisticated agent during the final transition from u_n to t.The expected payment of an astute agent commencing at s and ceasing immediately upon reaching t is

$$\pi_s = \pi_o - m\lambda\left(\frac{1}{j}\right)^{n-1}(n-1)c.$$

This is zero for the selection of λ.
We now prove that the sophisticated agent that starts traversing the Graph will go all the way to t. We give a proof by backward induction. That is, we assume that if the agent arrives at u_{k+1} from u_k, then it will go all the way to t. For the base case, observe that if the agent arrives at u_n, then the payoff for abandoning is $-\lambda\, c\,(n-1)$. On the other hand, moving to t incurs the same payoff $R-(\lambda\, c\,(n-1)+R) = -\lambda\, c\,(n-1)$ which is the same as quitting. For ease in presentation, we assume that the agent breaks ties in favor of continuing, and hence the agent will choose to continue. For the induction step, we assume correctness for u_{k+1} and prove for the agent traversing the Graph from u_{k+1} till t then its expected payoff for continuing is $\pi_\sigma(u_k)$. To complete the proof, we show that the expected payoff is greater than or equal to the agent's sunk cost in claim 1.

Claim 1. $\pi_\sigma(u_k) \ \geq\ -\lambda\, c\,(k-1)$.

Proof. We need to show that:

$$\pi_o(u_k) - m\lambda\left(\frac{1}{j}\right)^{n-k}(n-1)c \ \geq\ -\lambda c(k-1).$$

By rearranging that we get:

$$\pi_o(u_k) \ \geq\ \lambda c\left(m\left(\frac{1}{j}\right)^{n-k}(n-1)-(k-1)\right).$$

Since $\lambda\, c > 0$, if $m\left(\frac{1}{j}\right)^{n-k}(n-1)-(k-1) \leq 0$ the lemma trivially holds. Else, assume that $m\left(\frac{1}{j}\right)^{n-k}(n-1)-(k-1) > 0$, hence we can divide by this and get that

$$\lambda \ \leq\ \frac{\pi_o(u_k)}{c\left(m\left(\frac{1}{j}\right)^{n-k}(n-1)-(k-1)\right)}.$$

By substituting for our choice of λ we get that:

$$\frac{\pi_o}{m(\frac{1}{j})^{n-1}(n-1)c} \ \leq\ \frac{\pi_o(u_k)}{c\left(m\left(\frac{1}{j}\right)^{n-k}(n-1)-(k-1)\right)}.$$

By rearranging, we get that:

$$\frac{1-\left(\frac{1}{j}\right)^{n-1}}{m\left(\frac{1}{j}\right)^{n-1}(n-1)} \leq \frac{1-\left(\frac{1}{j}\right)^{n-k}}{m\left(\frac{1}{j}\right)^{n-k}(n-1)-(k-1)}.$$

which implies that:

$$\underbrace{m\left(\frac{1}{j}\right)^{n-k}(n-1)-(k-1)+\left(\frac{1}{j}\right)^{n-1}(k-1)}_{f(k)} \leq m\left(\frac{1}{j}\right)^{n-1}(n-1).$$

Finally, we apply claim 2 to show that the above inequality holds. This is done by proving that $f(k)$ is bounded from above by $m\left(\frac{1}{j}\right)^{n-1}(n-1)$ for any $i \leq k \leq n$.

□

Claim 2. *The function* $f(x) = m\left(\frac{1}{j}\right)^{n-x}(n-1)-(x-1)+\left(\frac{1}{j}\right)^{n-1}(x-1)$, *where* $0 < \frac{1}{j} < 1$, *is bounded above by* $m\left(\frac{1}{j}\right)^{n-1}(n-1)$ *for each* $1 \leq x \leq n$.

Proof. In order to prove this claim we show that f is convex in $[1,n]$ and that $f(1) = f(n) = m\left(\frac{1}{j}\right)^{n-1}(n-1)$. Therefore, for each $x \in [1,n]$, $f(x) \leq m\left(\frac{1}{j}\right)^{n-1}(n-1)$.
Observe that indeed $f(1) = f(n) = m\left(\frac{1}{j}\right)^{n-1}(n-1)$. In addition:

$$f'(x) = -m\ln\left(\frac{1}{j}\right)(n-1)\left(\frac{1}{j}\right)^{n-x} + \left(\frac{1}{j}\right)^{n-1} - 1$$

$$f''(x) = m\ln^2\left(\frac{1}{j}\right)(n-1)\left(\frac{1}{j}\right)^{n-x}.$$

Thus, for $1 < x < n$ and $0 < \frac{1}{j} < 1$ and $m > 0$ we have that $f''(x) \geq 0$. Therefore,

$$f(x) \leq m\left(\frac{1}{j}\right)^{n-1}(n-1)$$

for each $x \in [1,n]$, as required. □

□

Corollary 1. *There exist Layered Graph With Skips instances for which the quantity abm is strictly greater than* 1. *In particular, fixing any constants* $a,b \in [0,1]$, *one can construct instances in which the number m of distinct paths from s to the critical node* u_k *is arbitrarily large (e.g., by adding parallel node-disjoint paths), while a and b remain unchanged. Choosing* $m > 1/(ab)$ *yields* $abm > 1$.

We conclude the chapter by briefly comparing our Layered Graph With Skips construction with the classical fan-graph instances in the original work of Kleinberg et al [1].

Recall that Theorem 1 implies, for any layered skip-edge instance,

$$\pi_\sigma(s) \geq \pi_o(s) - \lambda\, abmR\Big(1-\frac{1}{k}\Big)^k,$$

so for large k the loss satisfies

$$\pi_o(s) - \pi_\sigma(s) \;\leq\; \frac{\lambda\, abmR}{e}.$$

We constructed a family of Layered Graph With Skips graphs for which this bound is asymptotically tight, and we showed in Corollary 1 that there exist instances with $abm > 1$. Hence there are graphs in which

$$\pi_o(s) - \pi_\sigma(s) \;=\; \frac{\lambda\, abmR}{e} \;>\; \frac{\lambda\, R}{e}.$$

By contrast, in the classical fan-graph construction the worst-case loss is of order $\lambda R/e$ with an effective coefficient of 1. Thus the fan graph can be viewed as a special case of our framework with $abm = 1$, whereas our Layered Graph With Skips family realizes instances with strictly larger worst-case loss. In particular, this shows that fan graphs are not worst-case for sophisticated sunk cost agents.

Chapter 4

Present Biased Reward-Seeking Agents: An Extension

In this chapter, we adapt our Layered Graph With Skips construction to the setting of present biased, reward-seeking agents studied by Kleinberg, Oren, Raghavan, and Sklar [2], and we present an example of this adaptation. In particular, Figure 4.1 illustrates the base construction and its modification for the present-bias setting. In this model, the agent receives nonnegative rewards on edges, and at each decision point it overweights the immediate reward of the next step relative to all future rewards.

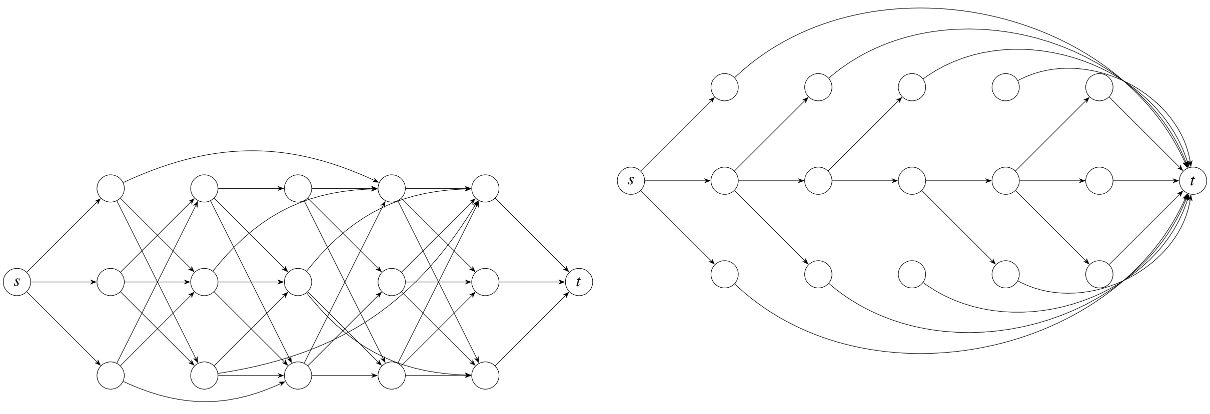

(a) Sunk cost base (b) Adoption to present bias

Figure 4.1: Base construction and its adaptation for present bias.

Our goal is to analyze the behavior of a sophisticated present biased agent on a Layered Graph With Skips, to compare it with the behavior of an optimal (unbiased) agent on the same graph, and to quantify the resulting in a performance loss. We show that our layered construction achieves an exponential gap in reward between the optimal and the sophisticated agent, matching the asymptotic behavior of the fan graphs in [2] up to a constant factor.

Throughout this chapter, we use two real parameters

$$n \geq 2 \quad \text{and} \quad 1 < c < b,$$

where $b > 1$ is the present bias factor and c determines the growth of edge rewards. This use of the symbol b is local to the present biased model and is independent of the parameter $b \in [0,1]$ that was used to normalize payoffs in Chapter 3.

4.1 The Present Biased Layered Graph with Skips: Model Definition

A present biased Layered Graph With Skips contains various paths of differing lengths that a reward-seeking agent may traverse from a starting node to a terminal node. As in Chapter 2, the nodes are arranged in layers and the graph is acyclic. In contrast to the sunk cost model, each edge now carries a nonnegative reward rather than a cost, and a path from the starting node to the terminal node yields a total reward equal to the sum of the rewards on its edges. The agent that traverses this graph will be present biased when choosing its next step, it will overweight the immediate reward on the outgoing edge by a factor $b > 1$, while evaluating all future rewards without this factor. We also fix a parameter $c > 1$ with $c < b$ that controls how the rewards grow across layers.

The graph is organized into $n+1$ layers of regular nodes and $n-1$ layers of special nodes, together with a terminal nodes. For each integer i such that $0 \leq i \leq n$, layer i contains a collection of regular nodes of the form $x_{i,k}$, where the index k ranges over a finite set (the number of such nodes may depend on i). The starting nodes is $s = x_{0,1}$. For each integer i such that $1 \leq i \leq n-1$, layer i also contains exactly two special nodes, denoted by $y_{i,1}$ and $y_{i,2}$. Finally, there is a single terminal nodes t.

Edges connect consecutive layers. For each integer i with $0 \leq i < n$, and for every pair of nodes k and k' in layers i and $i+1$ respectively, there is a directed edge from $x_{i,k}$ to $x_{i+1,k'}$. These edges represent the regular progression through the layers and have reward $r(x_{i,k}, x_{i+1,k'}) = 0$.

In addition to these regular edges, each regular nodes has skip edges to the special nodes in the next layer. For each integer i with $1 \leq i \leq n-1$ and every index k in layer $i-1$, there are directed edges from $x_{i-1,k}$ to $y_{i,1}$ and/or from $x_{i-1,k}$ to $y_{i,2}$. These skip edges carry reward $r(x_{i-1,k}, y_{i,l}) = c^i, l \in \{1,2\}$. Thus, a skip from layer $i-1$ into layer i yields an immediate reward of c^i.

Every special nodes leads directly to the terminal nodes. For each integer i with $1 \leq i \leq n-1$ and each $l \in \{1,2\}$, there is a directed edge from $y_{i,l}$ to t with reward $r(y_{i,l}, t) = 1$.

Finally, nodes in the last regular layer connect to the terminal nodes with a larger reward. For each index k in layer n, there is a directed edge from $x_{n,k}$ to t with reward $r(x_{n,k}, t) = c^n$.

There are no other edges in the graph. A path that moves only along edges of the form $(x_{i,k}, x_{i+1,k'})$ until it reaches some $x_{n,k}$, and then takes the edge from $x_{n,k}$ to t, collects a single reward of c^n at the terminal step. In contrast, a path that takes a skip edge $(x_{i-1,k}, y_{i,l})$ at some intermediate layer i obtains an immediate reward of c^i, then moves from $y_{i,l}$ to t and collects an additional reward of 1, and then terminates. The present biased agent that we study faces this structured choice between continuing along the regular layers to obtain a large reward c^n at the end, or taking earlier skip edges that offer tempting intermediate rewards of the form $c^i + 1$. A depiction of a present biased layered graph with skips is presented in Figure 4.2.

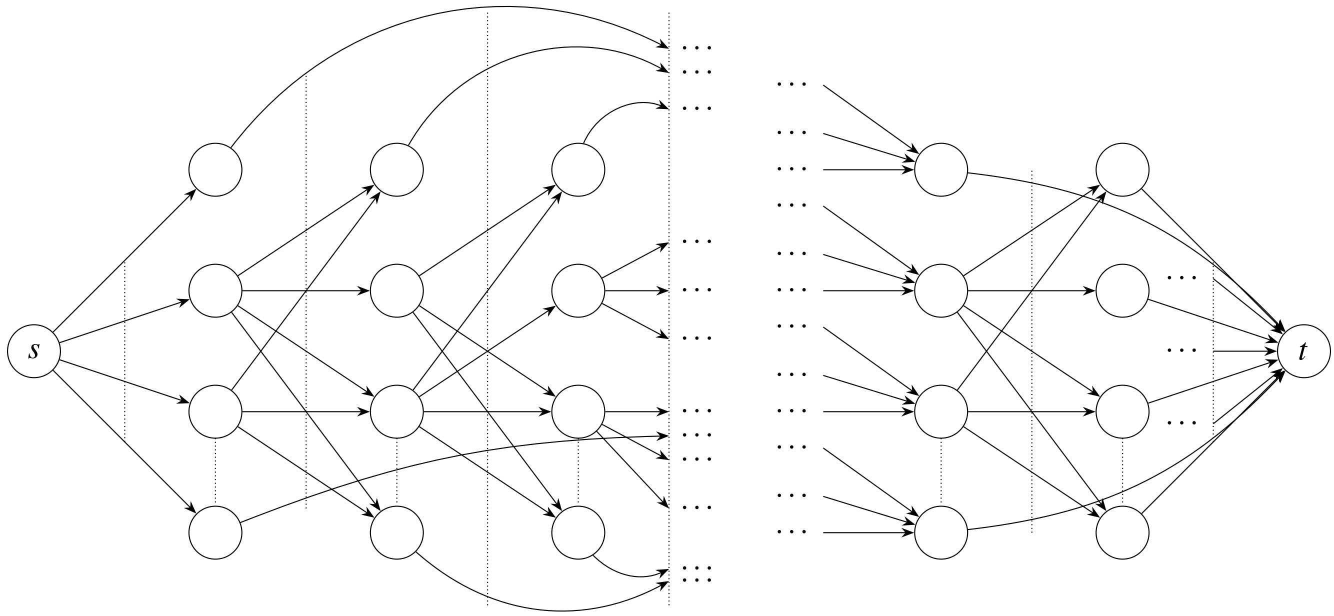

Figure 4.2: Present Biased Reward-Seeking Agents Layered Graph With Skips

4.2 Behavior of the Sophisticated Present Biased Agent

We now describe the behavior of a sophisticated present-biased agent on the Layered Graph With Skips. The agent is reward-seeking and traverses a path from s to t, collecting the true rewards on edges, but at each step it overweights the immediate reward of the next edge by a factor of b.

For each node u, we define $R_t(u)$ to be the true total reward obtained by the sophisticated agent when starting from u and following sophisticated behavior. The agent follows the present biased decision rule of Kleinberg et al [2] when standing at node u, it evaluates each outgoing edge (u, v) by

$$b \cdot r(u, v) + R_t(v),$$

and chooses a successor that maximizes this quantity. Formally, for every $u \neq t$ we set

$$S_s^{(R)}(u) \in \arg \max_{v:(u,v) \in E} \{b \cdot r(u, v) + R_t(v)\},$$

and then

$$R_t(u) = r\big(u, S_s^{(R)}(u)\big) + R_t\big(S_s^{(R)}(u)\big),$$

with boundary condition $R_t(t) = 0$.

By symmetry, all x-nodes in the same layer have the same value R_t, and all y-nodes in the same layer have the same value R_t. We write $R_t(X_i)$ for the common value $R_t(x_{i,k})$ and $R_t(Y_i)$ for the common value $R_t(y_{i,l})$.

We compute R_t using backward induction on the layers.

At the terminal node t we have

$$R_t(t) = 0.$$

For any $y_{i,l} \in Y_i$, the only outgoing edge is $y_{i,l} \to t$ with reward 1. Thus

$$R_t(y_{i,l}) = 1 \quad \text{for all } i,l,$$

and hence

$$R_t(Y_i) = 1 \quad \text{for all } i.$$

At layer n, each $x_{n,k}$ has a unique outgoing edge to t with reward c^n, so

$$R_t(X_n) = c^n.$$

At layer $n-1$, each $x_{n-1,k}$ has only edges to nodes in X_n with reward 0, and from X_n the agent goes to t with reward c^n. Thus

$$R_t(X_{n-1}) = R_t(X_n) = c^n.$$

At layer $n-2$, the agent at any $x_{n-2,k}$ has two types of choices: continuing to some $x_{n-1,\ell} \in X_{n-1}$, or skipping to one of the traps $y_{n-1,1}, y_{n-1,2}$.

- *Continuation.* Choosing any edge $x_{n-2,k} \to x_{n-1,\ell}$ gives perceived value

$$b \cdot 0 + R_t(X_{n-1}) = c^n.$$

- *Skip.* Choosing an edge $x_{n-2,k} \to y_{n-1,1}$ or $x_{n-2,k} \to y_{n-1,2}$ gives immediate reward c^{n-1} and then reward 1 when moving from the y-node to t. The perceived value is

$$b \cdot c^{n-1} + R_t(Y_{n-1}) = bc^{n-1} + 1.$$

Since $b > c$ and $c > 1$, we have

$$bc^{n-1} > c \cdot c^{n-1} = c^n,$$

and hence $bc^{n-1} + 1 > c^n$. Therefore, the sophisticated agent strictly prefers the skip edges at layer $n-2$, and its true reward from layer $n-2$ is

$$R_t(X_{n-2}) = c^{n-1} + 1.$$

The same reasoning applies inductively to all earlier layers.

Lemma 7. *For every $i = 0, 1, \ldots, n-2$ we have*

$$R_t(X_i) = c^{i+1} + 1.$$

In particular,

$$R_t(s) = c + 1.$$

Proof. We have already established that $R_t(X_{n-2}) = c^{n-1} + 1$, which matches the claimed formula for $i = n-2$.

Suppose that for some $i + 1 \leq n-2$ we have

$$R_t(X_{i+1}) = c^{i+2} + 1.$$

Consider any node $x_{i,k} \in X_i$. There are two types of outgoing edges.

- *Continuation.* For any child in X_{i+1} we have reward 0 and then continuation from layer $i+1$, so the perceived value is

$$b \cdot 0 + R_t(X_{i+1}) = c^{i+2} + 1.$$

- *Skip.* For a child in Y_{i+1} the immediate reward is c^{i+1}, and from Y_{i+1} the agent receives 1 when moving to t. Since $R_t(Y_{i+1}) = 1$, the perceived value is

$$b \cdot c^{i+1} + R_t(Y_{i+1}) = bc^{i+1} + 1.$$

Because $b > c$ and $c > 1$, we have

$$bc^{i+1} > c \cdot c^{i+1} = c^{i+2},$$

and therefore

$$bc^{i+1} + 1 > c^{i+2} + 1.$$

Thus the sophisticated agent strictly prefers the skip edges at layer i, and its total reward from X_i is

$$R_t(X_i) = c^{i+1} + 1.$$

This completes the induction. Setting $i = 0$ and using the fact that $s = x_{0,1} \in X_0$ yields

$$R_t(s) = R_t(X_0) = c^1 + 1 = c + 1.$$

□

Lemma 7 shows that, regardless of the number of x-nodes in each layer, the sophisticated present biased agent obtains a total reward of exactly $c+1$ when starting from s on the Layered Graph With Skips.

4.3 Behavior of the Optimal Reward-Seeking Agent

We now analyze the behavior of an optimal, unbiased reward-seeking agent on the same Layered Graph With Skips. This agent does not suffer from present bias and simply chooses a path that maximizes the true total reward.

For each node u, we let $R_o(u)$ denote the maximum total reward obtainable when starting from u and following an optimal strategy. As the graph is acyclic, $R_o(u)$ satisfies

$$R_o(u) = \begin{cases} 0, & u = t, \\ \max\limits_{v:(u,v)\in E} \{r(u,v) + R_o(v)\}, & u \neq t. \end{cases}$$

Again, by symmetry all nodes in X_i share the same value $R_o(X_i)$ and all nodes in Y_i share the same value $R_o(Y_i)$.

We distinguish two natural types of s–t paths.

The backbone path. We fix an arbitrary backbone path

$$s = x_{0,1} \to x_{1,1} \to \cdots \to x_{n,1} \to t.$$

All edges $x_{i-1,1} \to x_{i,1}$ have reward 0, and the last edge $x_{n,1} \to t$ has reward c^n. Therefore the total reward of the backbone path is

$$R_{\text{backbone}} = c^n.$$

A path using a trap. Alternatively, the agent may follow the backbone up to some layer i and then take a skip edge to a trap in layer $i+1$, followed by the edge to t:

$$s \to \cdots \to x_{i,1} \to y_{i,l} \to t, \quad 1 \le i+1 \le n-1.$$

On such a path, all rewards are zero except:

- the skip edge $x_{i,1} \to y_{i,l}$ with reward c^{i+1},
- the edge $y_{i,l} \to t$ with reward 1.

Hence the total reward of this path is

$$R_{\text{skip}(i+1)} = c^{i+1} + 1, \quad 1 \le i+1 \le n-1.$$

We now show that, for large enough n, the backbone path is optimal.

Lemma 8. *Fix $c > 1$. For all integers n such that $c^{n-1}(c-1) > 1$, the optimal agent on Layered Graph With Skips chooses a backbone path and $R_o(s) = c^n$.*

Proof. For any $i+1$ with $1 \le i+1 \le n-1$ we have

$$R_{\text{skip}(i+1)} = c^{i+1} + 1 \le c^{n-1} + 1.$$

If $c^{n-1}(c-1) > 1$, then

$$c^{n-1}(c-1) - 1 > 0 = c^n - (c^{n-1} + 1) > 0,$$

so $c^n > c^{n-1} + 1$. Hence

$$c^n > c^{i+1} + 1$$

for all $1 \le i+1 \le n-1$, and every path that uses a trap yields less reward than the backbone path. Therefore the optimal agent selects a backbone path and obtains reward $R_o(s) = c^n$. □

For any fixed $c > 1$, the condition $c^{n-1}(c-1) > 1$ holds for all sufficiently large n, so asymptotically we have $R_o(s) = c^n$.

4.4 Comparison Between Sophisticated and Optimal Agents

We now compare the rewards of the sophisticated and optimal agents on the present biased layered graph.

By Lemma 7 we have

$$R_t(s) = c + 1,$$

while by Lemma 8 we have

$$R_o(s) = c^n$$

for all sufficiently large n. The reward ratio is therefore

$$\frac{R_o(s)}{R_t(s)} = \frac{c^n}{c+1}.$$

For any fixed $c > 1$, this ratio grows exponentially in n:

$$\frac{c^n}{c+1} = \Theta\big(c^{n-1}\big) \quad \text{as } n \to \infty.$$

Thus, on the Layered Graph With Skips, a sophisticated present biased agent can lose an exponential factor in n relative to the optimal reward.

In the fan-graph examples of Kleinberg et al. [2], the behavior of a present biased sophisticated agent also leads to an exponential gap between the optimal reward and the reward achieved by the sophisticated agent for suitable parameters, one has

$$R_o(s) = c^n, \qquad R_t(s) = 1, \qquad \frac{R_o(s)}{R_t(s)} = c^n.$$

In our Layered Graph With Skips construction, we obtain

$$R_o(s) = c^n, \qquad R_t(s) = c+1, \qquad \frac{R_o(s)}{R_t(s)} = \frac{c^n}{c+1} = \frac{1}{c}c^n = c^{n-1}.$$

The gap is therefore exponential in n with the same base c it differs from the fan graph only by a constant multiplicative factor depending on c.

From an asymptotic point of view, the Layered Graph With Skips is therefore as bad as the fan graph for sophisticated present biased agents. Combined with the sunk cost analysis of Chapters 3, this shows that Layered Graph With Skips form a flexible framework within which one can obtain worst-case bounds for different behavioral biases using a common structural template.

Chapter 5

Conclusion

5.1 Summary of Contributions

In this thesis we studied planning problems on directed acyclic graphs under behavioral deviations from rationality, focusing primarily on sunk cost bias and sophisticated agents as introduced in the framework of Kleinberg et al. [1]. We introduced a new structured family of graphs, the *Layered Graphs With Skips*, which generalizes layered DAGs by allowing shortcut edges that reshape both the accumulation of sunk costs and the evolution of expected rewards.

Within this framework we identified a set of structural parameters that describe branching behavior, reward normalization, and layered depth. Using these parameters, we established a general analytical upper bound on the performance loss of sophisticated sunk cost agents relative to an optimal unbiased agent. We then constructed an explicit worst-case family of instances showing that this bound is asymptotically tight. A key implication of our analysis is that the resulting loss can exceed the classical $\frac{\lambda R}{e}$ bound known from fan graphs, proving that fan graphs are not necessarily worst-case structures in richer stochastic environments. Finally, we demonstrated that the same layered-with-skips template can also generate severe inefficiencies for present-biased agents in reward-based planning, revealing the versatility of the model.

5.2 Conclusion and Future Work

Overall, the results of this thesis show that structured stochastic environments can significantly amplify the inefficiency of sophisticated biased agents. The Layered Graph With Skips framework provides a powerful and unifying way to analyze worst-case behavior, to isolate the structural sources of inefficiency, and to demonstrate separations between biased and unbiased planning performance. The thesis therefore contributes both conceptual insight and precise quantitative bounds to the study of behavioral planning in graph-theoretic settings, and highlights that richer structural models are essential for fully understanding the impact of cognitive biases in sequential decision-making.

Building on these insights, a natural direction for future work is to investigate how the identified worst-case phenomena extend to more general and heterogeneous structural settings. In particular, it would be interesting to relax the homogeneity assumptions imposed on Layered Graphs With Skips by allowing paths of non-uniform lengths, as well as instances in which the structural parameters a and b, the transition probabilities, and the incurred costs vary across layers or nodes. Such heterogeneity arises naturally in stochastic planning environments and may lead to qualitatively different worst-case behavior. Beyond structural considerations, another promising direction is to examine how the framework developed in this thesis interacts with other behavioral distortions studied in graph-theoretic planning models, such as projection bias [6] or combinations of multiple biases, including the interaction of sunk-cost bias with other forms of behavioral deviation [4], and to assess how these additional distortions affect the structure and severity of worst-case instances.

Bibliography

[1] J. Kleinberg, S. Oren, M. Raghavan, and N. Sklar. A stochastic model for sunk-cost bias. In *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2021.

[2] J. Kleinberg, S. Oren, and M. Raghavan. Planning problems for sophisticated agents with present bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, 2016.

[3] S. Oren and D. Soker. Principal-agent problems with present-biased agents. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, 2019.

[4] J. Kleinberg, S. Oren, and M. Raghavan. Planning with multiple biases. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, 2017. 2.

[5] A. Cristi and S. Oren. Planning against a prophet: A graph-theoretic framework for making sequential decisions.

[6] S. Oren and N. Sklar. Planning on an empty stomach: On agents with projection bias. In *Proceedings of the 17th International Conference on Web and Internet Economics (WINE)*, 2021.