

Structural Approach to Guiding a Present-Biased Agent

Tatiana Belova¹, Yuriy Dementiev¹, Artur Ignatiev¹, Danil Sagunov¹

¹ITMO University

Abstract

Time-inconsistent behavior, such as procrastination or abandonment of long-term goals, arises when agents evaluate immediate outcomes disproportionately higher than future ones. This leads to globally suboptimal behavior, where plans are frequently revised or abandoned entirely. In the influential model of Kleinberg and Oren (2014) such behavior is modeled by a present-biased agent navigating a task graph toward a goal, making locally optimal decisions at each step based on discounted future costs. As a result, the agent may repeatedly deviate from initially intended plans.

Recent work by Belova et al. (2024) introduced a two-agent extension of this model, where a fully-aware principal attempts to guide the present-biased agent through a specific set of critical tasks without causing abandonment. This captures a rich class of principal–agent dynamics in behavioral settings.

In this paper, we provide a comprehensive algorithmic characterization of this problem. We analyze its computational complexity through the framework of parameterized algorithms, focusing on graph parameters that naturally emerge in this setting, such as treewidth, vertex cover, and feedback vertex set. Our main result is a fixed-parameter tractable algorithm when parameterized by the treewidth of the task graph and the number of distinct (v, t) -path costs. Our algorithm encaptures several input settings, such as bounded edge costs and restricted task graph structure. We demonstrate that our main result yields efficient algorithms for a number of such configurations.

We complement this with tight hardness results, that highlight the extreme difficulty of the problem even on simplest graphs with bounded number of nodes and constant parameter values, and motivate our choice of parameters. We delineate tractable and intractable regions of the problem landscape, which include answers to open questions of Belova et al. (2024).

Introduction

Present bias—the tendency to overvalue immediate outcomes relative to future ones—is a well-studied phenomenon in behavioral economics (Laibson 1997; Frederick, Loewenstein, and O’Donoghue 2002). It explains why

individuals often abandon long-term beneficial plans in favor of short-term ease, leading to suboptimal results in domains such as health, finance, education, and productivity.

(Kleinberg and Oren 2014) introduced a graph-theoretic model to formalize such behavior in sequential decision-making. In their model, a time-inconsistent agent with present bias navigates a directed acyclic graph (DAG) of tasks, choosing each step based on a discounted evaluation of future costs. The agent may abandon the plan before reaching the goal if their perceived cost exceeds the discounted reward. This model provides a powerful abstraction for reasoning about procrastination and dynamic decision-making in complex systems.

A natural extension of this framework, described in (Belova et al. 2024), involves a principal (e.g., a teacher, system designer, or platform) who can intervene by modifying the graph to guide the agent toward desired outcomes. Such interventions may involve removing distracting tasks or adding helpful shortcuts. Our work focuses on this setting, where the principal seeks to ensure that a present-biased agent completes a sequence of important tasks and reaches the final goal.

Problem Setting. We study intervention strategies in the Kleinberg–Oren model, where an agent operates on a time-inconsistent planning model $M = (G, w, s, t, \beta, r)$:

- $G = (V, E)$ is a directed acyclic graph with parallel arcs representing the task structure,
- $w : E \rightarrow \mathbb{N}_0$ assigns costs to arcs,
- s and t are the start and goal vertices,
- $\beta \in (0, 1]$ is the present-bias factor,
- r is the reward obtained upon reaching t .

When located at vertex v , the agent evaluates all v - t paths and selects the one $P = e_1 e_2 \dots e_k$ that minimizes the *perceived cost*

$$\zeta_M(P) = w(e_1) + \beta \cdot \sum_{i=2}^k w(e_i).$$

If $\zeta_M(P) > \beta \cdot r$, the agent abandons the task entirely. Otherwise, he commits to the first arc $e_1 = (v, u)$ and moves to vertex u , where the decision process is repeated recursively. If several paths have a minimum perceived cost, the agent

selects one of them deterministically, according to a given lexicographic order on the arcs.

Recent work by (Belova et al. 2024) initiated the study of *principal-agent* problems, where a principal aims to ensure that the agent follows a path from s to t that includes a designated set of critical arcs $T \subseteq E(G)$, using limited graph modifications. Two types of interventions were considered:

- *T*-PATH-DELETION: Can we delete at most k arcs from G so that the agent follows a path from s to t that includes all of T ?
- *T*-PATH-ADDITION: Given a set of auxiliary weighted arcs A , can we add at most k arcs from A (without creating cycles) so that the agent follows a path from s to t that includes all of T ?

In this work, we unify these two problems into a single general formulation:

***T*-PATH-EDITING**

Input: A time-inconsistent planning model $M = (G, w, s, t, \beta, r)$, a set of additional arcs A such that $G + A$ is acyclic and w maps w maps $E(G) \cup A$ into \mathbb{N}_0 , and a set of critical arcs $T \subseteq E(G)$.

Task: Compute the minimum number of arc edits (deletions from $E(G)$ or additions from A) needed so that the agent follows an s - t path (path in G from s to t) that traverses all arcs in T .

Throughout this paper, we say that the agent follows a path P if the agent, acting according to the Kleinberg–Oren model, traverses every arc of P without abandoning the task at any intermediate vertex.

Notably, unlike the previous formulations that fix a budget k , the optimization version allows a more global view of tractability.

Problem Motivation. In emerging computational systems involving human or autonomous agents, effective task delegation and coordination increasingly rely on explicit modeling of agent behavior. Modern multi-agent architectures often follow hierarchical paradigms, where higher-level entities (principals, mediators) interact with time- or resource-constrained agents to ensure alignment with overarching objectives.

Agents in such systems may exhibit irrationalities, including time inconsistency or locally myopic behavior. Understanding how and when such agents can be steered toward desired behaviors—especially under minimal interventions—forms a key building block in the design of robust, interpretable, and strategically aligned agent systems. The *T*-PATH-EDITING problem formalizes one such scenario, capturing goal-aligned influence via constrained structural modifications.

Model Example. To illustrate our framework, consider a scenario in which Bob, an AI engineer at a large tech company, is tasked with developing a production-ready LLM-based agent for enterprise document analysis. His ultimate goal is to deploy the agent after thorough benchmarking on realistic tasks (vertex t). There are multiple possible strategies to reach this goal, each composed of subtasks such as

data preprocessing, model fine-tuning, prompt engineering, integration, and evaluation. These subtasks form a directed acyclic graph (see Figure 1), where each arc represents an action, and its weight corresponds to the estimated engineering effort or time cost.

However, Bob exhibits present-biased behavior: when evaluating a future plan, he gives full weight to the immediate next task, but discounts the remaining effort by a factor $\beta = 1/2$. Starting at node s , he considers several possible execution paths:

- $P_1 = sbet$, $P_2 = scft$: thorough and reliable pipelines involving data cleaning, fine-tuning, and full evaluation: $\zeta_M(P_1) = \zeta_M(P_2) = 10 + \frac{10+10}{2} = 20$;
- $P_3 = sadt$: a simplified path using default weights and partial testing: $\zeta_M(P_3) = 10 + \frac{8+8}{2} = 18$;
- $P_4 = sbft$: a shortcut based on prompt-only adaptation with minimal validation: $\zeta_M(P_4) = 10 + \frac{2+10}{2} = 16$.

Bob chooses the path P_4 that minimizes the perceived cost, and it is easy to see that at each next vertex he will continue to choose this path, and will eventually follow *sbft*.

Now consider the role of Alice, Bob’s team lead and project owner. Her priority is to ensure that the critical evaluation step (arc (b, e)) is completed, as it directly affects the model’s reliability and regulatory compliance. Alice provides Bob with a promised reward of $r = 36$ upon successful deployment. However, due to Bob’s present bias, he perceives this as only $\beta \cdot r = 18$ at each decision point.

Suppose Alice simply disables the prompt-only shortcut by removing arc (b, f) , hoping to block the path P_4 , the perceived cost at s becomes $10 + \frac{8+8}{2} = 18 = \beta \cdot r$, which achieve only on path P_3 . Bob would follow the arc (s, a) . In contrast, the path $P_1 = sbet$, which only includes the crucial arc (b, e) , has a perceived cost of 20—too high for Bob to choose under present bias.

To resolve this, Alice applies a strategic edit: she introduces a new auxiliary arc (a, b) , allowing Bob to access path $P = sabet$. This new path includes the critical task (b, e) , and thanks to the adjusted structure, it has a lower perceived cost $1 + \frac{10+10}{2} = 11$ at key decision point a . As a result, Bob follows $P = sabet$, completing all required stages and aligning his behavior with Alice’s objectives.




Figure 1: Example: $T = \{(b, e)\}$, $A = \{(a, b)\}$. For initial graph G , the agent follows the path *sbft*. After deletion (b, f) —path *sadt*. And after adding (a, b) —*T*-path *sabet*.

Why Graph Parameters Matter. In real-world planning graphs modeling behavioral agents—such as workflows in education, research, digital assistants, or financial

planning—the underlying structure is rarely arbitrary. Instead, it tends to reflect the limited complexity of human-designed or human-operated systems. This makes structural graph parameters highly relevant for understanding algorithmic tractability in practice (Maniu, Senellart, and Jog 2019).

Below, we summarize the key graph parameters considered in our study. All of these parameters are undirected graph parameters. We consider them with respect to the underlying undirected graph of $G + A$.

The *treewidth* (tw) of a graph G is the minimum width among all possible tree decompositions of G , where the width of a tree decomposition is defined as the size of its largest bag minus one. Intuitively, treewidth quantifies how “tree-like” a graph is. Many real-world planning graphs exhibit low treewidth due to hierarchical structure, and dynamic programming algorithms often become efficient on graphs of bounded treewidth (Cygan et al. 2015).

The *feedback vertex set* number (fvs) is the size of the smallest set of vertices whose removal renders the undirected graph skeleton acyclic. This parameter captures the extent to which cyclic dependencies exist and is often small in acyclic or near-acyclic planning systems.

The *vertex cover* number (vc) is the size of the smallest set $C \subseteq V(G)$ such that every edge $(u, v) \in E(G)$ is incident to at least one vertex in C . In task graphs with a few highly connected control points or verification steps, the vertex cover tends to be small.

The *path length* (p) refers to the number of arcs on the longest directed path from source s to target t in the DAG. It reflects the maximum depth of planning or reasoning required by the agent.

The *tree-depth* (td) of a graph G is the minimum height of a rooted tree T such that G is a subgraph of the closure of T ; that is, for every edge $(u, v) \in E(G)$, one of u or v must be an ancestor of the other in T . Equivalently, tree-depth measures how “deeply nested” a graph is in terms of ancestor-descendant relations and reflects hierarchical complexity.

These parameters are not only natural from an algorithmic perspective, but also arise organically in structured task graphs. Treewidth and feedback vertex set are small in many systems where tasks are layered or hierarchically decomposed—such as curriculum maps (Stavrinos and Zuev 2023; Yang et al. 2024), multi-stage workflows, or modular design pipelines. Even when cycles are present, they are typically localized and limited in scope. Vertex cover tends to be small in graphs where dependencies are funneled through a small number of coordination or bottleneck nodes—common in multi-agent or supervisory planning setups (Cygan et al. 2015). Path length is inherently bounded in most practical systems: no realistic agent executes hundreds of sequential steps before reaching a goal. In human-facing systems, this is constrained by attention spans, task fatigue, or real-time processing limits.

Taken together, these structural properties motivate a parameterized approach to the T -PATH-EDITING problem. By isolating parameters that remain small in realistic deployments, we identify meaningful tractable regimes where efficient intervention strategies for present-biased agents can be

computed. This connection between structure and tractability bridges theoretical models with real-world applicability.

Parameterized complexity. To rigorously analyze the computational complexity of the T -PATH-EDITING problem, we rely on the framework of parameterized complexity. This framework allows us to distinguish between sources of hardness by isolating specific aspects of the input—such as structural graph parameters—as formal *parameters*.

Formally, a parameterized problem is a language $Q \subseteq \Sigma^* \times \mathbb{N}$, where an instance is given as a pair (I, k) consisting of the main input I and a parameter k . A problem is said to be *fixed-parameter tractable* (FPT) if it can be solved in time $f(k) \cdot |I|^{\mathcal{O}(1)}$ for some computable function f depending only on k . This contrasts with the broader class XP, where the running time is allowed to be $|I|^{g(k)}$ for some computable function g .

To classify intractable problems, the W-hierarchy provides a series of hardness classes: $\text{FPT} = \text{W}[0] \subseteq \text{W}[1] \subseteq \text{W}[2] \subseteq \dots \subseteq \text{W}[P]$. It is widely believed that $\text{FPT} \neq \text{W}[1]$, so proving $\text{W}[1]$ -hardness is strong evidence that no fixed-parameter tractable algorithm exists. Such hardness proofs are typically obtained via parameterized reductions from known $\text{W}[1]$ -complete problems.

In addition, we use the notion of Para-NP-hardness, which refers to parameterized problems that remain NP-hard even when the parameter is fixed to a constant. This notion is useful for ruling out tractability in low-parameter regimes.

For a general reference on the theory and techniques of parameterized complexity, we refer the reader to (Cygan et al. 2015).

Our Contribution. We present algorithmic results for T -PATH-EDITING, highlighting structural and numerical conditions under which efficient interventions are possible. We comprehensively describe a complexity landscape of these problems. Our main result is an algorithm parameterized by the treewidth (tw) of the graph $G + A$ and the diversity of path costs:

Theorem 1. *T -PATH-EDITING admits an algorithm running in $|L|^{\mathcal{O}(\text{tw})} \cdot m^{\mathcal{O}(1)}$ time, where tw equals the treewidth of $G + A$ and*

$$L = \left| \bigcup_{v \in V(G)} \{w(P) \mid P \text{ is a } (v, t)\text{-path in } G + A\} \right|.$$

This shows that T -PATH-EDITING is fixed-parameter tractable when both the graph structure (via treewidth tw) and weight diversity are bounded. We also show:

- NP-hardness on graphs of treewidth 2,
- $\text{W}[1]$ -hardness w.r.t. the number of vertices n and the combined parameter (p, vc) ,
- XP-algorithms with respect to feedback vertex set, vertex cover, and tree-depth,
- FPT-algorithms for cost-limited instances with bounded cost diversity $|\{w(e)\}|$.

Related Work. The study of present-biased behavior originates from classical economic theory. Samuelson’s discounted utility model (Samuelson 1937) laid the foundation for formal models of intertemporal choice. This was later extended into the hyperbolic and quasi-hyperbolic discounting frameworks (Laibson 1994; McClure et al. 2004), which more accurately capture behavioral inconsistencies over time. The model we adopt follows the framework introduced by Kleinberg and Oren (Kleinberg and Oren 2014, 2018), where a present-biased agent navigates a task graph by selecting paths according to discounted cost evaluations. Their model can be viewed as a discrete analogue of quasi-hyperbolic discounting and incorporates Akerlof’s notion of salience (Akerlof 1991).

While the Kleinberg–Oren model has strong empirical motivation, it does not capture all known psychological aspects of time-inconsistent behavior (Frederick, Loewenstein, and O’Donoghue 2002). Nevertheless, it has spurred extensive algorithmic research, including complexity and approximation analyses (Gravin et al. 2016; Kleinberg, Oren, and Raghavan 2016, 2017; Halpern and Saraf 2023; Meyer, Pomplun, and Schill 2022; Dementiev, Fomin, and Ignatiev 2022; Fomin and Strømme 2020).

Several recent works have extended the Kleinberg–Oren model in complementary directions. (Akagi, Marumo, and Kurashima 2024; Akagi, Kim, and Kurashima 2025) propose a continuous-time variant with closed-form solutions for optimal reward placement and abandonment conditions. (Belova et al. 2024) introduce the T -PATH-EDITING framework and analyze structural interventions in two-agent settings.

The principal–agent formulation we study can be seen as a natural and timely generalization of the classical motivating subgraph and P -motivating subgraph problems, where the goal is to identify a subgraph that leads a present-biased agent to the goal (through a specific path P) without abandonment. (Tang et al. 2017; Albers and Kraft 2017, 2019; Oren and Soker 2019; Fomin and Strømme 2020).

Finally, our work is also situated in the broader context of graph modification problems, a well-established domain in algorithmic graph theory. These problems aim to alter a graph’s structure—via edge or vertex deletions/additions—to satisfy desired properties. We refer the reader to surveys such as (Burzyn, Bonomo, and Durán 2006; Natanzon, Shamir, and Sharan 2001; Crespelle et al. 2023) for comprehensive overviews.

Polynomial Algorithm on Bounded Vertex Cover Graphs

In this section, we prove that T -PATH-EDITING can be solved in polynomial time when vertex cover of $G + A$ has constant size.

The first observation about graphs with bounded vertex cover is that such graphs have bounded path lengths.

Observation 1. *In a DAG G , every path consists of at most $2 \cdot \text{vc}(G) - 1$ arcs.*

Before moving on to the main result of this section, we

show an auxiliary algorithmic result. We will use it as a subroutine in our main algorithm.

Proposition 1. *There is a polynomial-time algorithm that, given two disjoint sets A, B and a sequence of m pairs $(a_1, b_1), (a_2, b_2), \dots, (a_m, b_m) \in A \times B$, and a set $R \subset (A \cup B)$ finds minimum possible $X \subset (A \cup B)$ such that*

- for each $i \in [m]$, $\{a_i, b_i\} \cap X \neq \emptyset$ and
- $X \cap R = \emptyset$,

or reports that X does not exist.

We note that Proposition 1 is folklore and is essentially a problem of finding a vertex cover of minimum cardinality in a bipartite graph. For completeness, we provide the proof with necessary references in the appendix.

We now show the XP-algorithm for T -PATH-EDITING parameterized by vc . In the following proof, our algorithm makes comparisons between perceived costs.

Theorem 2. *T -PATH-EDITING admits an algorithm with running time $m^{\mathcal{O}(\text{vc}^2)}$, where vc is the vertex cover number of $G + A$.*

Proof. We present an algorithm that finds an optimal solution to the given instance $((G, w, s, t, \beta, r), A, T)$ of T -PATH-EDITING. If no solution exists, the algorithm correctly reports it.

We assume without loss of generality that the algorithm is additionally given a vertex cover C of $G + A$ with $|C| \leq 2 \cdot \text{vc}(G + A)$, as such vertex cover can be found in polynomial time (see, e.g., (Papadimitriou and Steiglitz 1982) for 2-approximation of VERTEX COVER).

Let G^* be the graph obtained via optimal sequence of arc edits. Speaking informally, our algorithm does not know G^* in advance, so its goal is to guess only the *relevant* structure of G^* : the agent’s s - t path and *important* shortest paths in G^* . Since all path length are bounded in $G + A$, (therefore in G^*), the search space for these guesses is bounded with $m^{\mathcal{O}(\text{vc}^2)}$. The high-level sketch of our algorithm is the following.

1. Iterate over all s - t paths P in $G + A$ (the guess for the agent’s path);
2. Construct a set S consisting of all vertices of C and all vertices of the path P . That is, each arc of P has both endpoints in S and each arc of G starts in S or ends in S or both;
3. For each $v \in S$, iterate *one* path R_v from v to t in $G + A$ (the guess for the shortest path);
4. Find a smallest possible set of arc deletions that
 - forces the agent to follow P , and
 - for each $v \in S$, breaks every v - t path shorter than R_v .

We move on to the formal part of the proof. We describe the steps of the algorithm formally and supply it with correctness claims when needed.

Iterating the agent’s s - t path and critical vertices. In its outer loop, the algorithm iterates an arc set $P \subset (E(G) \cup A)$ such that $P \supset T$ and P forms a simple s - t path in $G + A$. In particular, $|P| \leq 2|C| - 1$.

Slightly abusing the notation, we refer to P both as an arc set and a path in $G + A$. Let d be the length of P , and let u_0, u_1, \dots, u_d be the vertices on P in the natural order, so $u_0 = s$ and $u_d = t$. For each $i \in [d]$, the algorithm denotes by e_i the edge between u_{i-1} and u_i in P .

The algorithm then constructs a set $S := \{u_0, u_1, \dots, u_d\} \cup C$ of *critical vertices*.

Iterating the critical shortest paths edges. The algorithm then iterates $R \subset (E(G) \cup A)$ such that $|R| \leq (|S| + 1) \cdot (2|C| - 1)$, and $R \supset P$ and

$$(R1) \text{ for each } i \in [d], w(e_i) + \beta \cdot \text{dist}_R(u_i, t) \leq \beta \cdot r.$$

Here and further in this proof, we slightly abuse the notation and denote by $\text{dist}_R(x, y)$ the distance between x and y in the graph $(V(G), R)$.

An intuition behind R is that it is (when guessed correctly) a union of all shortest paths that are required for the agent to follow P . (R1) specifically ensures that the agent does not abandon P .

Highlighting local obstructions. Having P and R both fixed, the algorithm now aims to find the minimum size subset $X \subset E(G)$ (the edges one has to delete) such that, for $G' = (V(G), R \cup (E(G) \setminus X))$: (a) we do not delete any edge in R ; (b) distances in G' agree with distances in R ; (c) the agent follows every edge $e_i \in P$ in G' . Expressed formally, this becomes

$$(X1) X \cap R = \emptyset;$$

$$(X2) \text{ For each } v \in S, \text{dist}_{G'}(v, t) = \text{dist}_R(v, t);$$

$$(X3) \text{ For each } i \in [d], \text{ for each vertex } v \in V(G) \text{ and for each edge } e' \in E(G') \text{ that goes from } u_{i-1} \text{ to } v \text{ and } e' \neq e_i, \text{ holds}$$

$$w(e_i) + \beta \cdot \text{dist}_{G'}(u_i, t) \prec w(e') + \beta \cdot \text{dist}_{G'}(v, t).$$

Here and forth, we use \prec to indicate the comparison between perceived costs according to the agent's behavior. That is, if the parts to the left and to the right of \prec are equal, the result of comparison is equivalent to the comparison of e_i and e' according to the ordering of edges.

We claim that these constraints are essentially equivalent to the agent following P in G' .

Claim 1. *Let $X \subset E(G)$ satisfy (X2). Then R satisfies (R1) and X satisfies (X3) if and only the agent follows P in G' .*

Proof of Claim 1. Since (X2) holds, $\text{dist}_R(v, t)$ and $\text{dist}_{G'}(v, t)$ are equivalent if $v \in S$. The agent follows P in G' if and only if:

- The agent does not abandon the task in neither of u_0, u_1, \dots, u_{d-1} , so for each $i \in [d]$ we have $w(e_i) + \beta \cdot \text{dist}_{G'}(u_i, t) \leq \beta \cdot r$. This is equivalent to (R1).
- For each $i \in [d]$, the agent always prefers the edge e_i when goes from u_{i-1} , meaning taking any other edge e' outgoing from u_{i-1} gives greater cost estimations (or estimations for e' and e_i are equal but e' goes after e_i in the ordering of edges). This is equivalent to (X3).

The equivalences prove the claim. \lrcorner

We now explain how the algorithm finds suitable X algorithmically.

Obstructions between critical vertices. The algorithm first constructs X_1 by finding the arcs with critical endpoints that have to be deleted necessarily to fulfill the properties. X_1 consists of all arcs e' of G such that e' starts in some $x \in S$ and ends in some $y \in S$, and either

- (i) $\text{dist}_R(x, t) > w(e') + \text{dist}_R(y, t)$, or
- (ii) $x = u_{i-1}$ for some $i \in [d]$, $e' \neq e_i$ and

$$w(e_i) + \beta \cdot \text{dist}_R(u_i, t) \succ w(e') + \beta \cdot \text{dist}_R(y, t).$$

We claim that all arcs in X_1 are necessary to delete.

Claim 2. *If X satisfies (X2), (X3) then $X_1 \subset X$.*

Proof of Claim 2. Let $X \subset E(G)$ satisfy (X2), (X3). Targeting towards a contradiction, suppose $e' \in X_1 \setminus X$. Denote by $x, y \in S$ the start and end points of e' , i.e. e' goes from $x \in S$ to $y \in S$.

If (i) is true, then $\text{dist}_R(x, t) > w(e') + \text{dist}_R(y, t)$, then by (X2) we have $\text{dist}_{G'}(x, t) > w(e') + \text{dist}_{G'}(y, t)$, which contradicts the definition of shortest distances, as e' goes from x to y .

If (ii) is true, then by (X2) we have

$$w(e_i) + \beta \cdot \text{dist}_{G'}(u_i, t) \succ w(e') + \beta \cdot \text{dist}_{G'}(y, t),$$

and this clearly violates (X3), since both e_i and e' start in u_{i-1} . The obtained contradiction finishes the proof of the claim. \lrcorner

We will show a little later that in any optimal solution, there is no arc between vertices in S that should be deleted but does not belong to X_1 .

Obstructions incident to non-critical vertices. Then, for each non-critical vertex $v \in V(G) \setminus S$, the algorithm constructs an arc pair set \mathcal{X}_v . This is done in the following way. Consider an arc e' that ends in v and an arc e'' that starts in v . Let x and y be the start point of e' and the end point of e'' respectively. Note that $x, y \in S$, since we cannot have arcs between two non-critical vertices. The pair $\{e', e''\}$ is added to \mathcal{X}_v if

- (i) $\text{dist}_R(x, t) > w(e') + w(e'') + \text{dist}_R(y, t)$, or
- (ii) $x = u_{i-1}$ for some $i \in [d]$ and

$$w(e_i) + \beta \text{dist}_R(u_i, t) \succ w(e') + \beta(w(e'') + \text{dist}_R(y, t)).$$

We show that its necessary for X to hit every pair in \mathcal{X}_v .

Claim 3. *Let $X \subset E(G)$ satisfy (X2), (X3). Then, for each $v \in V(G) \setminus S$ and for each $Y \in \mathcal{X}_v$, we have $X \cap Y \neq \emptyset$.*

Proof of Claim 3. The proof resembles the proof of Claim 2 with a pair of edges e', e'' instead of a single edge e' . Targeting towards a contradiction, assume there exists $v \in V(G) \setminus S$ and $\{e', e''\} \in \mathcal{X}_v$ such that $e', e'' \notin X$.

If $\{e', e''\}$ satisfies (i), then by (X2)

$$\text{dist}_{G'}(x, t) \succ w(e') + w(e'') + \text{dist}_{G'}(y, t),$$

which contradicts the definition shortest distances identically to the proof of Claim 2.

The remaining case is when $\{e', e''\}$ satisfies (ii). By (X2) and $u_i, y \in S$ we have $\text{dist}_R(y, t) = \text{dist}_{G'}(y, t)$ and $\text{dist}_R(u_i, t) = \text{dist}_{G'}(u_i, t)$. Since e'' goes from v to y , we have $\text{dist}_{G'}(v, t) \leq w(e'') + \text{dist}_{G'}(y, t)$ by definition of dist . Combining all of these with (ii) yields

$$w(e_i) + \beta \cdot \text{dist}_{G'}(u_i, t) \succ w(e') + \beta \cdot \text{dist}_{G'}(v, t),$$

contradicting (X3). The proof of the claim is complete. \square

Inner step of the algorithm. Having P and R fixed by the two outer loops, and X_1 and \mathcal{X}_v constructed for each $v \notin S$, the algorithm finds the set X of minimum cardinality that satisfies (X1) and the right parts Claim 2 and Claim 3, that is, X does not intersect R , X contains X_1 , and X intersects every pair in $\bigcup_{v \notin S} \mathcal{X}_v$.

The algorithm might also determine that X does not exist (this would mean basically that the choice of P and R is wrong, or no solution exists at all). This happens, for instance, when $X_1 \cap R \neq \emptyset$.

Since \mathcal{X}_v and $\mathcal{X}_{v'}$ contains pairwise distinct arcs, the minimum possible set intersecting every pair of \mathcal{X}_v can be done independently for each vertex $v \notin S$. For each $v \notin S$, the algorithm uses polynomial-time subroutine Proposition 1 to find minimum-size set X_v that hits every pair in \mathcal{X}_v and does not intersect R . If the subroutine of Proposition 1 reports that X_v does not exist, the algorithm correctly determines that X does not exist and moves to the next iteration of the outer loops.

If X_v was successfully constructed for each $v \notin S$, the algorithm constructs X via $X := X_1 \cup \bigcup_{v \in V(G) \setminus S} X_v$. The algorithm obtains a solution of size $|R \cap A| + |X|$: edges in $(R \cap A)$ should be added to G , edges in X should be removed from G . The algorithm updates its best solution with this one and moves to the next choice of P and R .

The description of the algorithm is finished. We move on to proving its correctness.

The algorithm never outputs incorrect solutions. We first show that $(R \cap A) \cup X$ is always a correct solution. Note that if we add $R \cap A$ to G and remove X from G , we obtain exactly $G' = (V(G), R \cup (E(G) \setminus X))$. Therefore, it is enough to prove that the agent always follows P in G' .

We show that X constructed in the inner step of the algorithm satisfies (X2) and (X3).

Claim 4. *If X satisfies (X1), X contains X_1 and X intersects every pair of $\bigcup_{v \notin S} \mathcal{X}_v$, then X satisfies (X2) and (X3).*

Combining Claim 4 and Claim 1, we obtain that the agent follows P in G' , and P traverses all arcs T by construction. Therefore, the algorithm outputs only correct solutions.

The algorithm finds a minimum-size solution. To see that the algorithm outputs a solution of minimum size, we show a valid choice of P and R that leads to the optimal answer.

To see this choice, let G^* be the graph obtained from G with the optimal edit sequence. Let P^* be the agent's s - t path in G^* . We have that $T \subset P^*$, and $|P^*| \leq 2 \text{vc}(G + A) - 1 \leq 2|C| - 1$. Let S be the union of all vertices of P^*

and C . For each $v \in S$, let R_v^* be the edge set of the shortest path between v and t in G^* . Construct R^* as a union of P^* and $\bigcup_{v \in S} R_v^*$, clearly

$$|R^*| \leq (2|C| - 1) + |S| \cdot (2|C| - 1) \leq (|S| + 1) \cdot (2|C| - 1),$$

and R^* satisfies (R1) since the agent does not abandon P^* in G^* .

Therefore, P^* and R^* are valid choices of P and R that the algorithm will consider in one of its iterations. Note that X^* (a sequence of arc deletions for G^*) satisfies (X1),(X2), (X3) for this choice of P and R . Hence, the algorithm will find edge deletion set X with $|X| = |X^*|$ and update the answer with an arc edit set of size $|R^* \cap A| + |X^*|$, which is minimum possible by definition of G^* .

Running time analysis. The inner step of the algorithm is polynomial in m , so we have to analyze the number of iteration given by its two loops. In its first loop, the algorithm iterates arc sets of size at most $\mathcal{O}(\text{vc})$, and in its second loop it iterates arc sets of size $\mathcal{O}(\text{vc}^2)$. Multiplied, this gives a total of $m^{\mathcal{O}(\text{vc}^2)}$ iterations. \square

Parameter Landscape

The proof of our main algorithmic result, Theorem 1, shares its basic idea with the proof of Theorem 2: guess each distance from v to t , ensure that necessary arcs are added, and delete necessary arcs so that the resulting distances agree with the guess and the resulting agent's path traverses each arc in T . To achieve the running time of Theorem 1, we avoid guessing the agent's path P , as we did in Theorem 2, because this approach requires comparing arbitrary arcs of $G + A$ lexicographically. In the proof of Theorem 1, we highly rely on that it's enough to compare arcs of $G + A$ only to arcs from T .

Other than that, Theorem 1 is a technical dynamic programming algorithm over a tree decomposition, common to the field of parameterized complexity (see, e.g. (Cygan et al. 2015)). The complete proof of Theorem 1 can be found in the appendix.

We now demonstrate that Theorem 1 captures several structural settings for T -PATH-EDITING, including the bounded vertex cover scenario and constant arc costs.

Lemma 3. *Let (G, w) be a DAG with non-negative integer arc costs $w : E(G) \rightarrow W$. Let t be a vertex in G . Let p be a maximum number of arcs in a path in G . Then*

$$|L| \leq \min\{(p + 1)^{|W|}, 1 + p \cdot \max W\} \text{ and}$$

$$|L| \leq \min\{m^{p+1}, m^{2 \cdot \text{vc}}, m^{2^{\text{td}+1}}, m^{2 \cdot \text{fvs} + 1}\}.$$

Theorem 3 has interesting consequences. Because each of the three parameters vc , fvs , td is at least tw , the running time of the algorithm of Theorem 1 is upper-bounded by $m^{\mathcal{O}(\text{vc}^2)}, m^{\mathcal{O}(\text{fvs}^2)}, m^{\mathcal{O}(\text{td} \cdot 2^{\text{td}})}$. That is, in particular, Theorem 1 generalizes Theorem 2.

For constant number of distinct arc costs, the algorithm of Theorem 1 is more efficient with running time $p^{\mathcal{O}(\text{tw} \cdot |W|)} \cdot m^{\mathcal{O}(1)}$. In particular, T -PATH-EDITING is fixed-parameter tractable with respect to both $\text{fvs} + |W|$ and $\text{td} + |W|$.

Hardness of T -PATH-EDITING

In this section we prove hardness of T -PATH-EDITING, showing that our XP-algorithms cannot be turned into FPT-algorithms without breaking $\text{FPT} \neq \text{W}[1]$.

Theorem 4. T -PATH-EDITING is $\text{W}[1]$ -hard with respect to:

- n , when parallel arcs in G are allowed,
- $p + \text{vc}$, when parallel arcs in G are forbidden.

Proof. We give a parameterized gap reduction from MODIFIED k -SUM (MkS), which is $\text{W}[1]$ -hard (Dementiev, Fomin, and Ignatiev 2022). In MkS , we are given k integer sets X_1, \dots, X_k and the target sum S . The goal is to determine whether there exists a choice of k integers, one from each of X_1, X_2, \dots, X_k , with a total sum S .

Given an instance of (X_1, \dots, X_k, S) of MkS , we construct a corresponding T -PATH-EDITING instance as illustrated in Figure 2. For each set X_i , we create a gadget consisting of two vertices v_i and v_{i+1} connected by multiple parallel arcs, each representing an element $x \in X_i$ with weight $w = x$. The goal is to ensure that the agent reaches vertex t via the unique T -path $sabt$ including the critical arc (b, t) .

The construction uses the following parameters: $\beta = \frac{1}{4} + \varepsilon$ for arbitrary $0 < \varepsilon < 1$; $r = 10 \cdot S$; $A = \emptyset$; $T = \{(b, t)\}$; $x = S + 2$. We duplicate each of the arcs of (s, c) , (c, t) into $z := |X_1| + \dots + |X_k|$ identical copies with the same arc cost. Breaking the path $s \rightarrow c \rightarrow t$ in G uses at least z edits.

We will show that the answer to the MODIFIED k -SUM problem is “Yes” if and only if there is a solution for the T -PATH-EDITING problem of size no greater than $z - k$.

Suppose that the T -PATH-EDITING instance has a feasible solution of size at most $z - k$. Then, the only arcs removed are those between vertices v_1, \dots, v_k, t in gadgets; removing any arc from the lower part of the graph would be too costly or would block the required edge (b, t) . Let G' denote the modified graph, and let S' be the cost of the shortest v_1 - t path in G' .

We now analyze possible values of S' . If $S' > x - 2$, then $S' \geq x - 1$, and the perceived cost of any path starting with $s \rightarrow a \rightarrow \dots$ is at least $\beta \cdot (1 + x - 1)$, which is worse than the perceived cost of the lower path $s \rightarrow c \rightarrow t$, that is equal to $\beta \cdot (x - \varepsilon)$. Hence, the agent would not take the $s \rightarrow a \rightarrow b \rightarrow t$ path. If $S' < x - 2$, then $S' \leq x - 3$, and, when he makes decision in the vertex a , the agent compares

$$1 + \beta \cdot S' < 0 + \beta \cdot (x + 1),$$

and prefers to go through v_1 rather than directly to b .

Therefore, the only way for the agent to traverse the desired arc (b, t) is when $S' = S = x - 2$. This implies the existence of a sequence of arc weights corresponding to a valid solution of the original MODIFIED k -SUM instance. Conversely, any such solution induces a deletion pattern in T -PATH-EDITING with no more than $z - k$ edits, that forces the agent to pass through the arc (b, t) as desired.

At the same time, as is evident from the proof, the only solution of another type for such T -PATH-EDITING instance—to break the path $s \rightarrow c \rightarrow t$ —will have a size at least z . Which completes our reduction.




Figure 2: The construction of graph G from Theorem 4.

Note that the construction of the graph G involves only $k + 5$ vertices. In the same time, this example can be easily modified into an example without parallel arcs by adding an additional vertex in the middle of each arc. In such a graph, the vertex cover number will not exceed $k + 5$, and the parameter p , the maximal arc length of a path in G , is equal to $2k + 2$. Thus, we have obtained a parameterized reduction to T -PATH-EDITING with parameter $n = k + 5$ in one case, and $p + \text{vc} \leq 3k + 7$ in the other.

This completes the reduction, establishing $\text{W}[1]$ -hardness for both settings. \square

The graph in the reduction of Theorem 4 is a series-parallel graph, and such graphs are known to have treewidth at most two (Brandstädt, Le, and Spinrad 1999).

Corollary 5. T -PATH-EDITING is NP-hard for $\text{tw} = 2$.

Note that this result is tight: when $\text{tw}(G + A) = 1$, the graph is a tree, and the T -PATH-EDITING problem becomes trivial. In this case, the agent has a unique s - t path, and the decision reduces to checking whether agent follow this path and whether it contains all arcs in T , with no modifications required.

On the other hand, (Belova et al. 2024) proved that the T -PATH-DELETION problem—a special case of our T -PATH-EDITING formulation—remains NP-hard even when the maximum length p of any s - t path in G is bounded by 8. This result immediately implies that T -PATH-EDITING is Para-NP-hard with respect to the parameter p .

Conclusion

We have introduced and studied the T -PATH-EDITING problem, a general framework for designing structural interventions that guide time-inconsistent agents toward completing critical tasks. Our model extends the Kleinberg–Oren framework by formalizing the principal’s role in influencing agent behavior through minimal graph edits.

Our results delineate the tractability frontier of T -PATH-EDITING across a range of structural parameters, including treewidth, vertex cover, feedback vertex set, and path length. In particular, we demonstrate fixed-parameter tractability for combined parameters capturing graph structure and cost diversity, while also establishing lower bounds via Para-NP- and $\text{W}[1]$ -hardness.

An intriguing open question is to determine whether T -PATH-EDITING admits an $f(\text{tw}) \cdot m^{g(p)}$ -time algorithm for computable functions f and g .

Acknowledgments

This work was supported by the Ministry of Economic Development of the Russian Federation (IGK 000000C313925P4C0002), agreement №139-15-2025-010.

References

- Akagi, Y.; Kim, H.; and Kurashima, T. 2025. A Continuous-time Tractable Model for Present-biased Agents. In *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, 13510–13519. AAAI Press.
- Akagi, Y.; Marumo, N.; and Kurashima, T. 2024. Analytically Tractable Models for Decision Making under Present Bias. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20-27, 2024, Vancouver, Canada*, 9441–9450. AAAI Press.
- Akerlof, G. A. 1991. Procrastination and obedience. *American Economic Review: Papers and Proceedings*, 81(2): 1–19.
- Albers, S.; and Kraft, D. 2017. 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. 2019. Motivating Time-Inconsistent Agents: A Computational Approach. *Theory Comput. Syst.*, 63(3): 466–487.
- Belova, T.; Dementiev, Y.; Fomin, F. V.; Golovach, P. A.; and Ignatiev, A. 2024. How to Guide a Present-Biased Agent Through Prescribed Tasks? In *ECAI 2024 - 27th European Conference on Artificial Intelligence, 19-24 October 2024, Santiago de Compostela, Spain - Including 13th Conference on Prestigious Applications of Intelligent Systems (PAIS 2024)*, volume 392 of *Frontiers in Artificial Intelligence and Applications*, 3461–3468. IOS Press.
- Brandstädt, A.; Le, V. B.; and Spinrad, J. P. 1999. *Graph Classes: A Survey*. Society for Industrial and Applied Mathematics.
- Burzyn, P.; Bonomo, F.; and Durán, G. 2006. NP-completeness results for edge modification problems. *Discrete Applied Mathematics*, 154(13): 1824–1844.
- Crespelle, C.; Drange, P. G.; Fomin, F. V.; and Golovach, P. A. 2023. A survey of parameterized algorithms and the complexity of edge modification. *Comput. Sci. Rev.*, 48: 100556.
- Cygan, M.; Fomin, F. V.; Kowalik, L.; Lokshtanov, D.; Marx, D.; Pilipczuk, M.; Pilipczuk, M.; and Saurabh, S. 2015. *Parameterized Algorithms*. Springer.
- Dementiev, Y.; Fomin, F.; and Ignatiev, A. 2022. Inconsistent Planning: When in Doubt, Toss a Coin! *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 36(9): 9724–9731.
- Fomin, F. V.; and Strømme, T. J. F. 2020. Time-Inconsistent Planning: Simple Motivation Is Hard to Find. In *Proceeding of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 9843–9850. AAAI Press.
- Frederick, S.; Loewenstein, G.; and O’Donoghue, T. 2002. Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2): 351–401.
- Gravin, N.; Immorlica, N.; Lucier, B.; and Pountourakis, E. 2016. Procrastination with Variable Present Bias. In *ACM Conference on Economics and Computation (EC)*, 361.
- Halpern, J. Y.; and Saraf, A. 2023. Chunking Tasks for Present-Biased Agents. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 853–884.
- 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.
- Laibson, D. 1997. Golden Eggs and Hyperbolic Discounting*. *The Quarterly Journal of Economics*, 112(2): 443–478.
- Laibson, D. I. 1994. *Hyperbolic Discounting and Consumption*. Ph.D. thesis, Massachusetts Institute of Technology, Department of Economics.
- Maniu, S.; Senellart, P.; and Jog, S. 2019. An Experimental Study of the Treewidth of Real-World Graph Data. In *22nd International Conference on Database Theory, ICDT 2019, March 26-28, 2019, Lisbon, Portugal*, volume 127 of *LIPICs*, 12:1–12:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- McClure, S. M.; Laibson, D. I.; Loewenstein, G.; and Cohen, J. D. 2004. Separate Neural Systems Value Immediate and Delayed Monetary Rewards. *Science*, 306(5695): 503–507.
- Meyer, S. A.; Pomplun, J.; and Schill, J. 2022. Present bias in partially sophisticated and assisted agents. *Mathematical Social Sciences*, 118: 36–47.
- Natanzon, A.; Shamir, R.; and Sharan, R. 2001. Complexity classification of some edge modification problems. *Discrete Applied Mathematics*, 113(1): 109–128.
- Oren, S.; and Soker, D. 2019. Principal-agent problems with present-biased agents. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, 237–251. Springer.
- Papadimitriou, C. H.; and Steiglitz, K. 1982. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall. ISBN 0-13-152462-3.

- Samuelson, P. A. 1937. A Note on Measurement of Utility. *The Review of Economic Studies*, 4(2): 155–161.
- Stavrinides, P.; and Zuev, K. 2023. Course-prerequisite networks for analyzing and understanding academic curricula. *Applied Network Science*, 8.
- 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)*. AAAI Press.
- Yang, B.; Gharebhaygloo, M.; Rondi, H.; Hortis, E.; Lostalo, E.; Huang, X.; and Ercal, G. 2024. Comparative analysis of course prerequisite networks for five Midwestern public institutions. *Applied Network Science*, 9.