

Neither Anarchy nor Centralization: Optimal Routing Requires Competing Planners

Anonymous submission

Abstract

The inefficiency of selfish routing is classically measured against a centrally optimal flow, implicitly treating full centralization as the ideal. This view is incomplete once the incentives of both drivers and routing platforms are taken into account: recommendations must be individually rational for the underlying users, and in realistic markets traffic is controlled not by a single platform but by several competing planners (navigation providers, fleet operators, mobility services), each routing a fraction of the flow.

We study repeated weighted atomic splittable routing games in which planners act as mediators for their assigned users, and ask when their strategies can simultaneously guarantee individual rationality, resilience to competition, convergence to a socially optimal flow, and a mild admissibility condition ruling out self-harming recommendations. Our main results characterize when these desiderata are achievable, and yield a sharp dichotomy: the desiderata can be met whenever every planner controls a sufficiently small share of the traffic, and cannot be met when some planner's share is too large; in particular, a single central planner always fails. This shows that neither full anarchy nor full centralization achieves robust optimal outcomes under natural incentive constraints. Instead, a nontrivial level of competition among several planners is not merely compatible with optimal routing, but necessary to sustain it.

1 Introduction


The rapid growth of algorithmic routing services is reshaping the way we think about traffic management and network optimization. Traditional traffic flow models assume that drivers act selfishly, selecting routes that minimize their own travel time without regard for the impact on others. This decentralized behavior can lead to highly inefficient outcomes, particularly in congested networks. Increasingly, however, routing decisions are not made directly by individual drivers: they are delegated to autonomous planners such as navigation platforms, fleet operators, and mobility providers that compute routes on users' behalf. This delegation creates a new kind of *coordination potential* since, in principle, planners can incorporate network-wide information, anticipate congestion ef-

This is an anonymized submission for review purposes only. Distribution, citation, or public sharing of this manuscript is strictly prohibited. Copyright and publication details will appear in the final version if accepted.

fects, and recommend traffic assignments that improve overall performance. This raises a fundamental question at the intersection of computer science, game theory, and transportation systems: can autonomous planners be leveraged to eliminate the inefficiencies of selfish routing, and what assumptions are actually needed for such improvements to persist?

At first glance, the existence of a powerful planner with access to global traffic data seems to offer a straightforward solution: compute a socially optimal flow that minimizes total travel time and route everyone accordingly. However, this view overlooks two key challenges inherent in real-world planner-mediated routing. First, routing recommendations must be individually rational for the underlying users: when a suggested route is noticeably worse for a user than available alternatives, compliance cannot be taken for granted, regardless of how sophisticated the planner is. Second, traffic is rarely controlled by a single entity. Multiple autonomous planners (e.g., competing navigation providers or fleet operators) simultaneously route different populations, each optimizing its own objective and reacting strategically to the others. The resulting planner competition can reshape incentives and outcomes in ways that are not captured by either classical selfish-routing models or idealized full centralization.

For instance, consider a canonical Pigou network (Pigou 1920) as follows.



Here, traffic flows from point A to point B. The top path has a constant travel cost of 1, while the bottom path's cost is the fraction of total traffic using it. It is straightforward to verify that the socially optimal flow splits traffic evenly between the two paths. However, this allocation is not stable under individual incentives: drivers assigned to the top path would prefer to (manually) switch to the bottom path, where their individual travel time would be lower.

Now suppose that there are four independent planners,

each responsible for routing an equal portion of the total traffic. If each planner adheres to the socially optimal strategy (i.e., splitting traffic equally), the resulting flow remains globally optimal. Nevertheless, when each planner aims to minimize the total travel time for their own share of traffic, their best response becomes routing all of it through the bottom path, assuming the others maintain a balanced split. These examples highlight the limitations of simply enforcing optimal flow patterns without accounting for the underlying incentives of the agents involved. If these are considered, the flow proposed by the planners should satisfy the following two conditions:

- (1) **Individual rationality:** Following the proposed flow must be a best-response for the routed agents at all stages of the game.
- (2) **Resilience to competition:** The flow proposed by the planners must be optimal for their fraction of the traffic, given the flow proposed by the remaining planners.

Unfortunately, (1) and (2) are generally incompatible in isolation. In the Pigou network above, individual agents have a strict incentive to deviate from any assignment that places them on the top edge, and with multiple planners, each planner similarly benefits from unilaterally shifting additional mass to the bottom edge. The core issue is that, in a single period, there is no credible way to *discipline* such deviations: neither agents nor planners can be punished tomorrow for defecting today. By contrast, when routing decisions recur over time (e.g., trips from A to B each day), the repeated-game structure creates leverage. An analogue of the folk theorem allows compliance to be supported by strategies that condition on past behavior, using future routing decisions as a credible threat. In fact, the folk theorem provides strategies that satisfy (1) and (2) in addition to the following guarantee.

- (3) **Optimality:** The strategy profile must always converge to an optimum flow.

Indeed, in the canonical Pigou network, planners could try to deter deviations by congesting the bottom path: after observing a defection, they route all of their traffic through the bottom edge for several consecutive days, making defections unprofitable in the future. While this can support a Nash equilibrium, it relies on *empty threats*: once a deviation has occurred, carrying out the punishment is not in any planner's interest, since deliberately collapsing the bottom path also increases their own travel time. As established in the foundational literature on repeated games (most notably (Fudenberg and Maskin 1986)), credibility can be restored by requiring higher-order punishments. That is, planners must punish not only defections, but also *failures to punish*, so that if a planner refrains from punishing today, it triggers punishment tomorrow. While this refinement yields a subgame perfect Nash equilibrium, it does so through optimal penal codes (Abreu 1988) that rely on a fragile and counterintuitive *threat cascade*. Planners carry out inefficient punishments not because punishing is individually beneficial, but to avoid being punished themselves for refusing to punish others. The effect is particularly stark with two planners: planner A is supposed to punish planner B for defecting, lest B punish A for failing to punish B .

This leaves us between two unsatisfactory extremes. If we weaken the solution concept to plain Nash equilibrium, we reintroduce exactly the empty threats we set out to avoid: non-credible punishments that no planner would carry out. If we retain subgame perfection but place no discipline on what planners may propose, we are forced into the higher-order penal codes above, with their fragile threat cascade. We therefore take a middle path, keeping credibility but restricting punishments to their simplest natural form, in which a planner disciplines deviations by becoming more selfish itself rather than by taking actions that harm everyone, itself included. This is the same logic that sustains cooperation in the repeated Prisoner's Dilemma, where defection is met with defection and carrying out the punishment is never against the punisher's own interest. What must be excluded, then, is any recommendation a planner would never rationally make: one strictly worse both for its own users and for the system than an available alternative. This motivates the following requirement.

- (4) **Admissibility:** At each stage, a planner never proposes a flow that is strictly dominated, both in its own cost and in social cost, by another flow available to it¹.

The four desiderata are demanding, and it is not obvious that any strategy profile satisfies them. Our first result is that a fully centralized system never does: on any network whose price of stability exceeds one, a single planner controlling all traffic cannot satisfy individual rationality together with either resilience to competition or admissibility. Notably, with a single planner these two requirements coincide: the planner must propose a socially optimal flow at each time step, regardless of the game's history. The failure itself is then immediate. When the price of stability exceeds one, no optimal flow is a Wardrop equilibrium, so some route is always strictly cheaper than what its users are asked to take, and those users are better off ignoring the recommendation.

Perhaps surprisingly, competition repairs precisely what centralization cannot. We show that on every single-commodity network, provided each planner's share is sufficiently small, all four desiderata can be met simultaneously. The strategies play a socially optimal flow on the equilibrium path and react to deviations with a temporary reversion to planner-optimal behavior, so that a deviator's short-run gain is outweighed by its discounted long-run loss. What makes these punishments admissible, and what a monopolist structurally lacks, is that they are carried out by other planners, for whom routing more selfishly is not a sacrifice but simply their own interest. The message is conceptual as much as technical: a nontrivial level of decentralized competition among autonomous planners is not merely compatible with efficiency, but necessary to sustain it.

We summarize our contributions as follows.

- We introduce a model of *repeated weighted atomic splittable routing games* with multiple competing planners,

¹Although stated for every stage, admissibility has bite only during punishment: the cooperative flow is socially optimal and therefore undominated, so on the equilibrium path the condition holds automatically.

each controlling a fraction of the flow and mediating on behalf of its users. Within this framework, we formalize four desiderata for planner strategies: *individual rationality*, *resilience to competition*, *convergence to an optimal flow*, and *admissibility*.

- We characterize when these desiderata can be met on a given instance, reducing an infinite-horizon question about strategy profiles to a static one about one-shot flows. At a high level, they are achievable when there exists an admissible strategy profile whose *cheapest* path costs strictly more than the average cost of the optimal flow, and they are not achievable when all admissible strategy profiles have a path strictly cheaper than the average cost of the optimal flow².
- We prove an impossibility result: for every single-commodity network with price of stability strictly greater than one, if some planner controls a sufficiently large share of the traffic, then no strategy profile satisfies individual rationality and admissibility simultaneously. As a special case, a fully centralized (single-planner) system cannot satisfy all four desiderata.
- We establish a matching existence theorem: for any such network, if each planner controls a sufficiently small share of the total flow, then there exist strategy profiles satisfying all four desiderata simultaneously. The construction uses equilibrium behavior as a discipline device in the repeated game, implemented via karma-based strategies.
- We then turn to deciding this condition on a given instance, where a natural shortcut suggests itself: inspect only the *planner equilibrium*, in which every planner is fully selfish, on the intuition that selfish planners crowd onto the cheapest paths and drive up their cost. We show that the shortcut fails even for single-commodity networks with convex costs.
- On parallel networks with convex costs, however, the conjecture holds. This means that, in this case, our reduction to one-shot flows is polynomial, and the characterization is polynomial end-to-end whenever finding the planner equilibrium is polynomial as well (e.g., (affine costs, uniform shares) (Cominetti, Correa, and Stier-Moses 2009) or small number of planners (Bhaskar and Lolakapuri 2018)).
- Using the results above, we give tight bounds on the amount of competition required for the canonical Pigou network: we characterize the exact thresholds on planner sizes and show that at least three planners are necessary to satisfy the desiderata.

Our main results rely on minimal public monitoring: after each round, planners observe only the realized congestion on each edge, and our constructions never need to identify which users deviated or to observe other planners' recommendations. Under such coarse monitoring, however, unintentional user deviations are indistinguishable from strategic ones and

²The two conditions do not quite coincide, so a small gap remains between the instances we can certify and those we can rule out.

may repeatedly re-trigger punishment. Section H of the Supplementary Material therefore develops, under stronger observability assumptions, a “ticket” mechanism that preserves all incentives while absorbing low-rate random defections.

Related Work

Our planner model is closely related to *atomic splittable routing games*, in which a finite number of players each control a non-negligible, divisible amount of flow and may split it across multiple routes. This framework originates in work on competitive routing in multiuser communication networks by Orda, Rom, and Shimkin (1993), and has since been developed in a series of papers analyzing efficiency and equilibrium properties in network games with atomic players (Cominetti, Correa, and Stier-Moses 2006). Subsequent work has investigated existence and uniqueness of equilibria in weighted congestion and atomic splittable congestion games (Harks and Klimm 2012; Harks and Timmermans 2018), algorithmic aspects of equilibrium computation (Bhaskar and Lolakapuri 2018), and the effect of collusion among atomic players on social cost (Huang 2013). Our setting can be viewed as an atomic splittable routing game where each planner is an atomic player controlling a continuous flow, but with two crucial differences: planners act as mediators on behalf of underlying users rather than as final decision-makers, and their interaction unfolds in repeated play under additional behavioral constraints such as admissibility. These distinctions lead to qualitatively different questions about when planner competition can be leveraged to guarantee optimal outcomes.

Closest in spirit is Stackelberg routing, in which a single coordinator controls a fixed fraction of the flow and routes it so as to induce favorable behavior from the selfish remainder (Korilis, Lazar, and Orda 1997; Roughgarden 2004; Swamy 2012). Our planners differ in two essential ways: their recommendations must be individually rational for the users they route, and several planners interact strategically in repeated play, rather than a single leader committing once. The limiting behavior of atomic equilibria as players' shares vanish, which underlies our small-share analysis, goes back to Haurie and Marcotte (1985).

A central theme in mechanism design is the role of mediators in improving strategic outcomes. In fact, the inefficiency of selfish behavior in routing games is classically captured by the concept of the *Price of Anarchy*, introduced by Koutsoupias and Papadimitriou (1999), and further explored by Roughgarden and Tardos (2002) in the context of selfish vs. centrally optimized routing. Aside from selfish routing, mediators can be beneficial in various ways: through signaling (Arieli, Geffner, and Tennenholtz 2023a,b; Corrao and Dai 2023), messaging (Ben-Porath 2003; Gerardi 2004; Abraham et al. 2019; Geffner and Halpern 2024) promising payments (Monderer and Tennenholtz 2003; Jackson and Wilkie 2005; Geffner, Oesterheld, and Conitzer 2025), expanding the available strategy sets (Geffner and Tennenholtz 2024), or enabling commitment to strategies (Tennenholtz 2004; Kovařík, Oesterheld, and Conitzer 2023). However, in our setting, introducing a single central mediator is not sufficient to achieve optimal outcomes; rather, competition

between multiple mediators (i.e., planners) is essential. The idea of competition between mediators has been studied in other domains, such as auctions (Feldman et al. 2010; Leyton-Brown, Shoham, and Tennenholtz 2002; Ashlagi, Monderer, and Tennenholtz 2011; Monderer and Tennenholtz 2004) and facility location problems (Babaioff, Feldman, and Tennenholtz 2016).

One might expect that repeated interactions would trivialize coordination via the folk theorem, making efficient equilibria easy to sustain. However, this intuition can sometimes be misleading (Berker and Conitzer 2024; Borgs et al. 2008). Repeated games are also closely tied to learning dynamics, and notions such as learning equilibria have yielded positive results in certain congestion games (Tennenholtz and Zohar 2009).

A second line of work concerns collective deviations. Since our planners act on behalf of subsets of users, their strategic choices can be read as group deviations, a perspective developed in the literature on strong equilibria in congestion games (Holzman and Law-Yone 2003; Holzman and Monderer 2015), which asks when no coalition can jointly deviate profitably, and in congestion games with partitioned agents, where deviation capabilities are structured around fixed groups (Feldman and Tennenholtz 2009; Kuniavsky and Smorodinsky 2014). What distinguishes our setting from all of the above is that planners are mediators rather than players in their own right: they only recommend routes, and the users they route retain the option to ignore them.

2 Model

While user-level deviations are inherently discrete, we model traffic as a continuous (divisible) flow. This choice is both standard in the routing literature and technically convenient for our purposes. Besides being a natural approximation when demand is large, it allows planners to split their demand across paths with arbitrary precision, which makes the strategies and incentive arguments transparent and avoids inessential artifacts of discretization (e.g., indivisibilities when a single vehicle tips a threshold). Accordingly, we represent demand as a unit mass that can be divided continuously among available routes, and we let both planners and users control positive-measure portions of this flow. Formal definitions appear next.

Single-Commodity Flow Networks A single-commodity flow network G is a tuple (V, E, s, t, c) where

- (V, E) is a directed multigraph (V is the set of nodes and E is the set of edges.)
- s is the source node and t is the sink node.
- $c : E \times [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ is a cost function such that $c(e, x)$ describes the cost (or latency) of going through edge e given the amount of flow that traverses it. For each edge e , the restriction $c_e(\cdot)$ of c to e must be $\mathcal{C}^1$ (i.e., continuously differentiable) and non-decreasing.

Given a single-commodity flow network $G = (V, E, s, t, c)$, let $\mathcal{P} := \{P_1, P_2, \dots, P_M\}$ denote the set of all paths from source s to sink t . A flow is a function $f : \mathcal{P} \rightarrow [0, 1]$ such that $\sum_{P \in \mathcal{P}} f(P) = 1$. If we view the

network as a system of one-way roads, a flow describes how traffic is split among the several paths from the source to the sink. Denote by $\mathcal{P}_e$ the set of all paths containing edge e . Given a flow f , we also denote $f_e := \sum_{P \in \mathcal{P}_e} f(P)$. Intuitively, f_e denotes the amount of traffic that goes through some edge $e \in E$. With these definitions, we can define the cost of traversing a path $P \in \mathcal{P}$ as $c_P(f) := \sum_{e \in P} c_e(f_e)$, the total flow cost as $c(f) := \sum_{P \in \mathcal{P}} f(P) \cdot c_P(f)$, and the cheapest path cost as $c_{\min}(f) := \min_{P \in \mathcal{P}} c_P(f)$.

If f is a flow on a single-commodity flow network $G = (V, E, s, t, c)$, we say that f is a **Wardrop equilibrium** if no subset of the flow can decrease their cost by switching to a different path. Since the cost functions are continuous, we have that f is a Wardrop equilibrium if and only if, for all paths P such that $f(P) > 0$, $c_{P'}(f) \geq c_P(f)$ for all paths P' . In particular, in a Wardrop equilibrium, all paths used by f must have the same cost.

Two relevant concepts associated with flow networks are the **price of anarchy** (PoA) and the **price of stability** (PoS). The price of anarchy is defined as the ratio between the (social) cost of the worst Wardrop equilibrium and the cost of an optimal solution. The price of stability is defined analogously, except that we take the cost of the best Wardrop equilibrium instead. Since Wardrop equilibria are defined at the level of the aggregate (non-atomic) flow and the edge costs are continuous and non-decreasing, all Wardrop equilibria of G induce the same edge costs, and hence the same social cost (Roughgarden and Tardos 2002); the two notions therefore coincide here. In particular, $\text{PoA}(G) > 1$ holds if and only if no optimal flow of G is a Wardrop equilibrium, which is the form in which we use the hypothesis throughout.

Routing Games with Multiple Planners A **weighted atomic splittable routing game** (WASRG) is a tuple (G, C, α) where $G = (V, E, s, t, c)$ is a single-commodity flow network, $C = \{1, \dots, N\}$ is the set of planners, and $\alpha = (\alpha_1, \dots, \alpha_N)$ is a partition of unity where α_i represents the fraction of flow (traffic) assigned to planner i .

Given a WASRG (G, C, α) , we can identify the strategy of each planner $i \in C$ with a point $(x_1^i, x_2^i, \dots, x_M^i)$ on the $(M-1)$ -dimensional standard simplex (i.e., each coordinate must be non-negative and they must all add up to 1), in which x_j^i denotes the portion of i 's traffic that i assigns to P_j . Given a deterministic strategy profile $\vec{\sigma} = (\vec{x}^1, \dots, \vec{x}^N)$, the (average) cost $c^i(\vec{\sigma})$ of planner i can be expressed as $\sum_{j=1}^M x_j^i \cdot c_{P_j}(f)$, where f is the flow induced by $\vec{\sigma}$. If the strategy includes randomization, the cost is calculated by taking the expected value of the previous expression.

We say that a strategy profile $\vec{\sigma}$ is a Nash equilibrium if no planner can decrease its own cost by switching to a different strategy. Equivalently, $\vec{\sigma}$ is a Nash equilibrium if $c^i(\vec{\sigma}) \leq c^i(\vec{\sigma}_{-i}, \sigma'_i)$ for all strategies σ'_i for planner i . To distinguish this notion from Nash equilibria in other settings, we also use the term **planner equilibrium**.

From a computational point of view, we treat both socially optimal flows and planner equilibria as primitives, but in general they are hard to compute. For atomic congestion games with general latency functions, minimizing the

social cost is NP-hard to approximate within a constant factor (Paccagnan and Gairing 2021), and computing planner equilibria in atomic splittable congestion games is PPAD-complete even with player-specific affine costs (Klimm and Warde 2020). For atomic splittable routing games with general convex costs, Bhaskar and Lolakapuri (Bhaskar and Lolakapuri 2018) give polynomial-time algorithms only in restricted cases (e.g., parallel links with a small number of planners or edges) and show hardness results for general networks. On the positive side, in many classical single-commodity routing models the system-optimal flow can be computed efficiently: when the edge cost functions are non-decreasing and the functions $x \mapsto x \cdot c_e(x)$ are convex, the optimal flow is the solution of a convex program and can be found in polynomial time, up to arbitrary precision (Swamy 2012).

A **repeated weighted atomic splittable routing game** (RWASRG) is defined exactly in the same way as a standard WASRG, except that the game is played repeatedly in discrete stages. Here, planners must decide not only how much flow they are sending through each path on each stage, but also which section of the flow they are sending. As a consequence, the strategy space of each planner becomes considerably more convoluted than in the one-shot game (although, for most of our results, we will focus on the much simpler karma-based strategies described below). The strategy of each planner i depends on its local history, which is the sequence of all the decisions made by i , and the flow distribution in all previous stages of the game. We model this as follows: the flow of each planner is modeled by $[0, 1]$. Planners assign measurable portions of flow to paths, and a defection is a positive-measure subset of a planner’s users following an alternative assignment. Stage and discounted costs are defined in the natural way, with formal definitions in the supplementary material (Section A). For future reference, given a discount factor λ , a strategy profile $\vec{\sigma}$, and a planner i , we denote by $c_i^\lambda(\vec{\sigma})$ the total discounted cost of planner i when playing $\vec{\sigma}$ with discount factor λ , and we denote by $c_{(i,T)}^\lambda(\vec{\sigma})$ the total discounted cost of a subset $T \subseteq [0, 1]$ of i ’s flow. If h is a global history, we denote by $c^\lambda(\vec{\sigma} \mid h)$ the total discounted cost conditioned on reaching history h ; $c_i^k(\vec{\sigma} \mid h)$, $c_i^\lambda(\vec{\sigma} \mid h)$, $c_{(i,T)}^k(\vec{\sigma} \mid h)$, and $c_{(i,T)}^\lambda(\vec{\sigma} \mid h)$ are defined analogously.

Desiderata

The notation and definitions presented in the previous sections allow us to rigorously define the properties (1)-(4) shown in the introduction. Given a routing game (G, C, α) and a strategy profile $\vec{\sigma}$ for the planners, these are the desired properties that we want $\vec{\sigma}$ to satisfy.

- (1) **Individual rationality:** Regardless of the state (history) of the game, for all planners i and all positive-measure subsets S of i ’s assigned agents, there is no defection $\mathcal{D} = (i, S, D)$ that strictly decreases the average cost of i ’s agents in S . In other words, for all global histories h and all defections $\mathcal{D} = (i, S, D)$, there exists $\lambda_0 \in (0, 1)$ such that, for all $\lambda > \lambda_0$, $c_{(i,S)}^\lambda(\vec{\sigma} \mid h) \leq c_{(i,S)}^\lambda((\vec{\sigma}, \mathcal{D}) \mid h)$.

- (2) **Resilience to competition:** For all global histories h , all planners i , and all strategies σ'_i for i , there exists $\lambda_0 \in (0, 1)$ such that, for all $\lambda > \lambda_0$, $c_i^\lambda(\vec{\sigma} \mid h) \leq c_i^\lambda((\sigma'_i, \vec{\sigma}_{-i}) \mid h)$.
- (3) **Optimality:** Let $c_{\text{opt}}(G)$ be the minimum total cost of a unit flow in G . Then, for all global histories h , $\lim_{k \rightarrow \infty} c^k(\vec{\sigma} \mid h) = c_{\text{opt}}(G)$.
- (4) **Admissibility:** For all histories h , all stages k , all planners i , and all strategies σ'_i for i , we have that $c_i^k((\sigma'_i, \vec{\sigma}_{-i}) \mid h) < c_i^k(\vec{\sigma} \mid h) \implies c^k((\sigma'_i, \vec{\sigma}_{-i}) \mid h) \geq c^k(\vec{\sigma} \mid h)$.

Why the discount threshold depends on history. Our incentive guarantees are necessarily history-dependent: for each history h we require the existence of a threshold $\lambda_0(h)$ such that deviations are unprofitable for all $\lambda > \lambda_0(h)$. We do *not* fix a single global discount factor that works uniformly over all histories. The reason is structural in our setting: effective discipline requires punishments whose duration may increase after successive deviations (punishment phases “stack”). With R rounds of punishment still outstanding, the extra loss from a fresh deviation arrives only after those R rounds and is scaled by roughly λ^R , while the gain is immediate; since R grows without bound as punishments stack, no fixed $\lambda < 1$ works for every history.

Karma-based strategies

For future reference, some of our proofs will focus on a special class of strategy profiles that allows us to translate one-shot flow properties into repeated-game guarantees for all subsets of a planner’s flow. At a high level, in a karma-based strategy a planner always assigns the currently shortest routes to those parts of its flow that have accumulated the most travel time so far. One can think of this as a “karma system”: driving longer earns more karma, and at each stage the routes with the lowest current cost are allocated to the drivers with the highest karma.

Since we are dealing with continuous flows, a more precise definition is required. Given a weighted atomic splittable game $\Gamma = (G, C, \alpha)$ and a deterministic strategy $\vec{x}$ for planner $i \in C$, we define the **karma-implementation** of $\vec{x}$ as the strategy $K(\vec{x})$ in the repeated version of Γ in which, in stage k , the routing is assigned as follows. In the first stage, all routes are assigned lexicographically (i.e., P_1 is assigned to $[0, x_1]$, P_2 is assigned to $[x_1, x_1 + x_2]$, etc.). If $k > 1$, it can be easily shown inductively that $[0, 1]$ can be partitioned into a finite number of intervals $I_1 = [a_1, b_1]$, $I_2 = [a_2, b_2], \dots, I_m = [a_m, b_m]$ such that the total cost taken by i ’s flow in each of these intervals is constant. Suppose that these intervals are sorted according to their total cost and that paths $P_1, \dots, P_M$ are sorted by their cost when playing $\vec{x}$. Then, we assign the routes following Algorithm 1.

For instance, the karma-implementation of a strategy profile in the Pigou network that splits the flow evenly between the two paths is the strategy profile in which, in stage 1, it assigns the first half of the flow to the top path and the second half of the flow to the bottom path; in stage 2 it assigns the bottom path to the first half of the flow and the top path to the second half, and so on.

Algorithm 1: Karma-based assignment

```

1: Initialize  $\ell = 1, j = 1$ 
2: while  $\ell \leq M$  and  $j \leq m$  do
3:   if  $x_\ell < b_j - a_j$  then
4:     Assign route  $P_\ell$  to interval  $[a_j, a_j + x_\ell]$ 
5:      $a_j \leftarrow a_j + x_\ell$ 
6:      $\ell \leftarrow \ell + 1$ 
7:   else
8:     Assign route  $P_\ell$  to interval  $[a_j, b_j]$ 
9:      $x_\ell \leftarrow x_\ell - (b_j - a_j)$ 
10:     $j \leftarrow j + 1$ 
11:   end if
12: end while

```

At a high level, the most important property of these karma-implementations is that, on average, all users' discounted cost is roughly equal. This is captured by the following proposition, which states that, for a sufficiently large discount factor, the normalized discounted cost of all positive-measure subset of flow gets arbitrarily close to its associated planner's cost.

Proposition 1. *Let $\vec{\sigma} = (\vec{x}^1, \dots, \vec{x}^N)$ be a deterministic strategy profile. Then, for all planners i and all $\varepsilon > 0$, there exists $\delta > 0$ such that, for all $1 - \delta < \lambda < 1$,*

$$\left| c_i(\vec{\sigma}) - (1 - \lambda) \sum_{k \geq 0} c_{(i,T)}^k(K(\vec{x}^i), \vec{\sigma}_{-i}) \cdot \lambda^k \right| < \varepsilon$$

for all positive-measure subsets $T \subseteq [0, 1]$.

Although the result may seem intuitive by construction (karma-based strategies try to equalize the accumulated cost of each flow segment), the proof is quite technical and can be found in Section B of the supplementary material.

3 Main Results

Our main results characterize when we can satisfy the desiderata, both at a theoretical/existential level, and on a practical level. We begin with the following characterization as a function of sufficiently inefficient admissible strategy profiles.

Proposition 2. *Let G be a single commodity flow network with $\text{PoA}(G) > 1$, let $\Gamma = (G, C, \alpha)$ be a (non-repeated) WASRG, and let Γ' be its repeated extension.*

(Sufficiency:) *There exists a strategy profile $\vec{\sigma}'$ in Γ' satisfying all four desiderata whenever there exists an admissible strategy profile $\vec{\sigma}$ in Γ such that its induced flow f satisfies*

- $c_{\min}(f) \geq c_{\text{opt}}(G)$, and
- $c^i(\vec{y}^i, \vec{\sigma}_{-i}) > c_{\text{opt}}(G)$ for all planners i and (selfish) best-responses $\vec{y}^i$ of i to $\vec{\sigma}_{-i}$.

Moreover, given $\vec{\sigma}$ and the social optimum, each planner can compute its stage- k recommendation in $\vec{\sigma}'$ in polynomial time from the history of the game.

(Necessity:) *Suppose that there exists a strategy profile $\vec{\sigma}'$ in Γ' satisfying all four desiderata. Then, for each planner i , there exists an admissible strategy profile $\vec{\sigma}$ in Γ such that its induced flow f satisfies*

- $c_{\min}(f) \geq c_{\text{opt}}(G)$, and
- $c^i(\vec{\sigma}) > c_{\text{opt}}(G)$.

Note that a subtle gap remains between sufficiency and necessity. Both directions revolve around admissible strategy profiles $\vec{\sigma}$ whose induced flow f satisfies $c_{\min}(f) \geq c_{\text{opt}}(G)$, and any such profile automatically satisfies $c^i(\vec{\sigma}) \geq c_{\text{opt}}(G)$ for every planner i , since each planner's cost is an average over paths that cost at least $c_{\min}(f)$. The gap concerns strictness: sufficiency requires a single profile at which the inequality is strict for every planner simultaneously, even when planners best-respond, whereas necessity guarantees strictness only one planner at a time, with a possibly different profile for each. At a high level, the main idea for proving Proposition 2 is the fact that planners can punish deviating users and planners by shifting to an inefficient admissible flow: whenever a suitable such flow exists, it deters all deviations, and conversely, when none exists, we show that deviators can always profit by repeatedly switching to the cheapest available paths. The full proof can be found in the Supplementary Material (Section C).

Proposition 2 yields our main dichotomy: a small-share regime in which competition makes all four desiderata achievable, and a large-share regime in which individual rationality and admissibility become incompatible.

Theorem 1. *Given a single-commodity flow network G with $\text{PoA}(G) > 1$, there exist $\alpha_{\min}, \alpha_{\max} \in (0, 1)$ such that, for all repeated weighted atomic splittable routing games $\Gamma = (G, C, \alpha)$,*

- If $\alpha_i < \alpha_{\min}$ for all $i \in C$, there exists a strategy profile $\vec{\sigma}$ in Γ that satisfies individual rationality, resilience to competition, optimality, and admissibility.*
- If $\alpha_i > \alpha_{\max}$ for some $i \in C$, there are no strategy profiles that satisfy individual rationality and admissibility.*

Note that Theorem 1 immediately implies that, in order to satisfy all the desiderata, we need at least a certain number of distinct participating planners.

Corollary 1. *Given a single-commodity flow network G with $\text{PoA}(G) > 1$, there exists $N_{\min} \geq 2$ such that, for all repeated weighted atomic splittable routing games $\Gamma = (G, C, \alpha)$ with $|C| < N_{\min}$, there are no strategy profiles that satisfy individual rationality and admissibility.*

Theorem 1 follows from two limiting arguments. When every share is small, planner equilibria approach a Wardrop equilibrium; since $\text{PoA}(G) > 1$, its common path cost strictly exceeds $c_{\text{opt}}(G)$, so nearby planner equilibria satisfy the sufficiency condition of Proposition 2. When some share approaches one, admissibility forces play toward socially optimal flows, which contain a path strictly cheaper than $c_{\text{opt}}(G)$ and therefore fail the necessity condition. The proof is quite technical and can be found in the Supplementary Material (Section D).

Although Proposition 2 is useful for existence theorems like the one above, it may not seem as useful on a practical level. Indeed, in order to determine if the desiderata can be satisfied in a given instance, it requires us to search over all admissible strategy profiles. We may conjecture that, if

there is any admissible strategy profile that satisfies the sufficiency or necessity conditions of Proposition 2, it must be the *planner equilibrium*. After all, it would seem that the more selfish the planners are, the more flow they will concentrate through the cheapest paths, and the more likely that the sufficiency conditions are satisfied (note that, additionally, all planners would already be best-responding). Unfortunately, this conjecture is not true in general, as shown in the Supplementary Material (Section E). There, we provide an instance where a non-planner equilibrium satisfies both conditions of Proposition 2 but the planner equilibrium fails. However, the conjecture is true for a large class of networks: parallel networks with convex costs. It follows from the following result, shown in the Supplementary Material (Section F).

Proposition 3. *Consider a WASRG $\Gamma = (G, C, \alpha)$ where G is a parallel network, $|C| \geq 2$, and all cost functions are convex. Let f be the flow induced by planner equilibrium $(\vec{x}^1, \dots, \vec{x}^N)$, and let f' be the flow induced by an admissible strategy profile $(\vec{y}^1, \dots, \vec{y}^N)$. Then, $c_{\min}(f) \geq c_{\min}(f')$.*

Proposition 3 implies that, in parallel networks with convex costs, we can restrict our search of admissible profiles satisfying the sufficiency and necessity conditions of Proposition 2 to simply checking the planner equilibrium. Note that, if a planner equilibrium induces a flow f with $c_{\min}(f) > c_{\text{opt}}(G)$, the second sufficiency condition of Proposition 2 is automatically satisfied since every planner is already best-responding. This implies the following.

Theorem 2. *Let $\Gamma = (G, C, \alpha)$ be a RWASRG where G is a parallel network with convex costs, and let f be the flow induced by the (unique (Orda, Rom, and Shimkin 1993)) planner equilibrium of the associated one-shot game. If $c_{\min}(f) \neq c_{\text{opt}}(G)$, then a strategy profile satisfying all four desiderata exists if and only if $c_{\min}(f) > c_{\text{opt}}(G)$. Moreover, given f and an optimal flow, this condition can be checked with $O(|E|)$ evaluations of the edge cost functions, and whenever it holds, each planner’s stage- k recommendation can be computed in polynomial time from the history of the game.*

Corollary 2. *The test of Theorem 2 runs end to end in polynomial time whenever the planner equilibrium and an optimal flow are polynomial-time computable. This holds in particular (i) for affine costs, where both are solutions of convex quadratic programs (Cominetti, Correa, and Stier-Moses 2009); (ii) for uniform shares $\alpha_i = 1/N$, where the planner equilibrium minimizes the convex potential $\sum_e [(1 - \frac{1}{N}) \int_0^{f_e} c_e(t) dt + \frac{1}{N} f_e c_e(f_e)]$ (Cominetti, Correa, and Stier-Moses 2009); and (iii) when the number of planners or edges is bounded (Bhaskar and Lolakapuri 2018).*

Remark 1: For non-affine convex costs the equilibrium may be irrational (Bhaskar and Lolakapuri 2018), so in cases (ii) and (iii) the flow is computed to any desired accuracy ε ; the test is then correct whenever the gap $|c_{\min}(f) - c_{\text{opt}}(G)|$ exceeds the resulting error bound, which the hypothesis $c_{\min}(f) \neq c_{\text{opt}}(G)$ guarantees for any fixed instance. For affine costs all comparisons are exact.

Remark 2: Theorem 2 leaves open the case when $c_{\min}(f) = c_{\text{opt}}(G)$, which corresponds to the open gap of

Proposition 2. In this case, we know that, if $\vec{\sigma}$ is the planner equilibrium and $c^i(\vec{\sigma}) > c_{\text{opt}}(G)$ for all planners i , then the desiderata can be satisfied (and efficiently constructed). However, if a planner i satisfies that $c^i(\vec{\sigma}) = c_{\text{opt}}(G)$, the necessity clause may still be satisfied, since there could potentially exist other admissible strategy profiles $\vec{\sigma}'$ inducing a flow f' with $c_{\min}(f') = c_{\text{opt}}(G)$ and $c^i(\vec{\sigma}') > c_{\text{opt}}(G)$. Closing the gap, both in Proposition 2 and in Theorem 2, is an **open problem**.

The last of our results uses the results above to compute tight bounds for $\alpha_{\min}$, $\alpha_{\max}$ and $N_{\min}$ in the canonical Pigou network (where the cost function in the bottom path is the identity function). These bounds are illustrated by the following Proposition, proved in the Supplementary Material (Section G).

Proposition 4. *Let G be the classical Pigou network. Then, the following holds.*

- (i) *The supremum of all $\alpha_{\min}$ that satisfy Theorem 1 is $\frac{3}{8}$.*
- (ii) *The infimum of all $\alpha_{\max}$ that satisfy Theorem 1 is $\frac{1}{2}$.*
- (iii) *The maximum of all $N_{\min}$ that satisfy Corollary 1 is 3.*

4 Discussion

We introduced repeated weighted atomic splittable routing games, in which competing planners mediate route choices on behalf of their users, and identified four desiderata for planner strategies: individual rationality, resilience to competition, convergence to optimality, and admissibility. Our results yield a sharp dichotomy on any network whose price of anarchy exceeds one: the desiderata are jointly achievable whenever every planner’s share of the traffic is sufficiently small, and individual rationality and admissibility become incompatible whenever some share is too large (in particular, under full centralization). Crucially, the punishments that sustain cooperation are carried out by competitors for whom routing more selfishly is in their own interest, so credibility requires no counterintuitive threat cascades. In this precise sense, competition among planners is not an obstacle to optimal routing but the resource that sustains it.

Several directions remain open. Closing the strictness gap between the sufficient and necessary conditions of Proposition 2 is one. Another is extending our polynomial-time test beyond parallel networks: for affine costs, equilibrium computation is not the bottleneck (planner equilibria are solutions of convex quadratic programs (Cominetti, Correa, and Stier-Moses 2009)), so the question is whether the planner equilibrium still maximizes the cheapest-path cost over admissible profiles; our counterexample for the general case uses nonlinear costs and leaves this open. Finally, our karma-based strategies suggest a welfare agenda beyond total travel time: treating priority as a non-monetary currency with equal initial endowments would let users with heterogeneous, time-varying urgency trade priority across days, without the equity concerns of willingness-to-pay-based congestion pricing. Developing urgency-aware welfare benchmarks and principled karma mechanisms is a promising avenue for future work.

References

- Abraham, I.; Dolev, D.; Geffner, I.; and Halpern, J. Y. 2019. Implementing mediators with asynchronous cheap talk. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, 501–510.
- Abreu, D. 1988. On the theory of infinitely repeated games with discounting. *Econometrica: Journal of the Econometric Society*, 383–396.
- Arieli, I.; Geffner, I.; and Tennenholtz, M. 2023a. Mediated cheap talk design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 5456–5463.
- Arieli, I.; Geffner, I.; and Tennenholtz, M. 2023b. Resilient Information Aggregation. In *Proceedings Nineteenth conference on Theoretical Aspects of Rationality and Knowledge*, volume 379 of *EPTCS*, 31–45.
- Ashlagi, I.; Monderer, D.; and Tennenholtz, M. 2011. Simultaneous ad auctions. *Mathematics of Operations Research*, 36(1): 1–13.
- Babaioff, M.; Feldman, M.; and Tennenholtz, M. 2016. Mechanism Design with Strategic Mediators. *ACM Trans. Econ. Comput.*, 4(2).
- Ben-Porath, E. 2003. Cheap talk in games with incomplete information. *Journal of Economic Theory*, 108(1): 45–71.
- Berker, R. E.; and Conitzer, V. 2024. Computing optimal equilibria in repeated games with restarts. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI '24*. ISBN 978-1-956792-04-1.
- Bhaskar, U.; and Lolakapuri, P. R. 2018. Equilibrium computation in atomic splittable routing games. In *26th Annual European Symposium on Algorithms (ESA 2018)*, 58–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- Borgs, C.; Chayes, J.; Immorlica, N.; Kalai, A. T.; Mirrokni, V.; and Papadimitriou, C. 2008. The myth of the folk theorem. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, 365–372.
- Cominetti, R.; Correa, J. R.; and Stier-Moses, N. E. 2006. Network games with atomic players. In *International colloquium on automata, languages, and programming*, 525–536. Springer.
- Cominetti, R.; Correa, J. R.; and Stier-Moses, N. E. 2009. The impact of oligopolistic competition in networks. *Operations Research*, 57(6): 1421–1437.
- Corrao, R.; and Dai, Y. 2023. Mediated communication with transparent motives. In *Proceedings of the 24th ACM Conference on Economics and Computation*, 489–489.
- Ehrgott, M. 2005. *Multicriteria Optimization*. Berlin ; New York: Springer, 2nd ed edition. ISBN 978-3-540-21398-7.
- Feldman, J.; Mirrokni, V.; Muthukrishnan, S.; and Pai, M. M. 2010. Auctions with intermediaries: extended abstract. In *Proceedings of the 11th ACM Conference on Electronic Commerce, EC '10*, 23–32. New York, NY, USA: Association for Computing Machinery. ISBN 9781605588223.
- Feldman, M.; and Tennenholtz, M. 2009. Partition equilibrium. In *International Symposium on Algorithmic Game Theory*, 48–59. Springer.
- Fudenberg, D.; and Maskin, E. 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3): 533–554.
- Geffner, I.; and Halpern, J. Y. 2024. Communication games, sequential equilibrium, and mediators. *Journal of Economic Theory*, 221: 105890.
- Geffner, I.; Oesterheld, C.; and Conitzer, V. 2025. Maximizing Social Welfare with Side Payments. *arXiv preprint arXiv:2508.07147*.
- Geffner, I.; and Tennenholtz, M. 2024. Making a Nash Equilibrium Resilient to Coalitions. In *Proceedings of the 25th ACM Conference on Economics and Computation*, 213–238.
- Gerardi, D. 2004. Unmediated communication in games with complete and incomplete information. *Journal of Economic Theory*, 114(1): 104–131.
- Harks, T.; and Klimm, M. 2012. On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research*, 37(3): 419–436.
- Harks, T.; and Timmermans, V. 2018. Uniqueness of equilibria in atomic splittable polymatroid congestion games. *Journal of Combinatorial Optimization*, 36(3): 812–830.
- Haurie, A.; and Marcotte, P. 1985. On the relationship between Nash—Cournot and Wardrop equilibria. *Networks*, 15(3): 295–308.
- Holzman, R.; and Law-Yone, N. 2003. Network structure and strong equilibrium in route selection games. *Mathematical social sciences*, 46(2): 193–205.
- Holzman, R.; and Monderer, D. 2015. Strong equilibrium in network congestion games: increasing versus decreasing costs. *International Journal of Game Theory*, 44(3): 647–666.
- Huang, C.-C. 2013. Collusion in atomic splittable routing games. *Theory of Computing Systems*, 52(4): 763–801.
- Jackson, M. O.; and Wilkie, S. 2005. Endogenous games and mechanisms: Side payments among players. *The Review of Economic Studies*, 72(2): 543–566.
- Klimm, M.; and Warode, P. 2020. Complexity and parametric computation of equilibria in atomic splittable congestion games via weighted block Laplacians. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, 2728–2747. SIAM.
- Korilis, Y. A.; Lazar, A. A.; and Orda, A. 1997. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM transactions on networking*, 5(1): 161–173.
- Koutsoupias, E.; and Papadimitriou, C. 1999. Worst-case equilibria. In *Annual symposium on theoretical aspects of computer science*, 404–413. Springer.
- Kovářík, V.; Oesterheld, C.; and Conitzer, V. 2023. Game Theory with Simulation of Other Players. In Elkind, E., ed., *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, 2800–2807. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Kuniavsky, S.; and Smorodinsky, R. 2014. Equilibrium and potential in coalitional congestion games. *Theory and decision*, 76(1): 69–79.

- Leyton-Brown, K.; Shoham, Y.; and Tennenholtz, M. 2002. Bidding clubs in first-price auctions. In *AAAI/IAAI*, 373–378.
- Monderer, D.; and Tennenholtz, M. 2003. k-Implementation. In *Proceedings of the 4th ACM conference on Electronic Commerce*, 19–28.
- Monderer, D.; and Tennenholtz, M. 2004. K-price auctions: Revenue inequalities, utility equivalence, and competition in auction design. *Economic Theory*, 24(2): 255–270.
- Nocedal, J.; and Wright, S. J. 2006. *Numerical Optimization*. Springer Series in Operations Research. New York: Springer, 2nd ed edition. ISBN 978-0-387-30303-1.
- Orda, A.; Rom, R.; and Shimkin, N. 1993. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking*, 1(5): 510–521.
- Paccagnan, D.; and Gairing, M. 2021. In congestion games, taxes achieve optimal approximation. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, 743–744.
- Pigou, A. C. 1920. *The economics of welfare*. Macmillan and Co.
- Prokhorov, Y. V. 1956. Convergence of random processes and limit theorems in probability theory. *Theory of Probability & Its Applications*, 1(2): 157–214.
- Roughgarden, T. 2004. Stackelberg scheduling strategies. *SIAM journal on computing*, 33(2): 332–350.
- Roughgarden, T. 2007. Routing Games. In *Algorithmic Game Theory*. New York: Cambridge university press. ISBN 978-0-521-87282-9.
- Roughgarden, T.; and Tardos, É. 2002. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2): 236–259.
- Swamy, C. 2012. The effectiveness of Stackelberg strategies and tolls for network congestion games. *ACM Transactions on Algorithms (TALG)*, 8(4): 1–19.
- Tennenholtz, M. 2004. Program equilibrium. *Games and Economic Behavior*, 49(2): 363–373.
- Tennenholtz, M.; and Zohar, A. 2009. Learning equilibria in repeated congestion games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 233–240.

A Formal Model

Following the notation introduced in Section 2, we can represent the route recommendations made by planner i at stage k as a measurable function³ $g_{i,k} : [0, 1] \rightarrow \mathcal{P}$. Intuitively, $[0, 1]$ represents i 's share of the total traffic, and $g_{i,k}$ denotes the path that i assigns to each of the (infinitesimal) cars. We can represent car defections as tuples (i, S, D) , where i is a given planner, S is a subset of $[0, 1]$ and D is a function that takes as input their local history (i.e., the sequence of all flows and recommendations from planner i in earlier stages) and outputs a function $d : S \rightarrow \mathcal{P}$. At a high level, we can view (i, S, D) as a deviation performed by the subset S of cars assigned to planner i in which, given all the information available, they decide to take the paths according to function d instead of the paths recommended by the planner.

Given a deterministic strategy profile $\vec{\sigma} = (\sigma_1, \dots, \sigma_N)$, denote by f_e^k the total flow that traverses edge e at stage k . If there are no defections, its actual value can be computed by aggregating the amount of flow that each planner i sends through e , scaling it with the portion of the traffic managed by planner i . This can be expressed as

$$f_e^k = \sum_{i=1}^N \alpha_i \cdot \mu \left(g_{i,k}^{-1}(\mathcal{P}_e) \right),$$

where $\mu(S)$ represents the Lebesgue measure of $S \subseteq [0, 1]$. If there is a defection $\mathcal{D} = (i, S, D)$, the expression above looks the same, except that the contribution of planner i 's traffic is $\alpha_i \cdot \mu \left(\left(g_{i,k}^{-1}(\mathcal{P}_e) \cap \bar{S} \right) \cup d_k^{-1}(\mathcal{P}_e) \right)$ instead, where $d_k : S \rightarrow \mathcal{P}$ encodes the defection of agents in S at stage k . The total cost of the flow at stage k is given by $c^k(\vec{\sigma}) := \sum_{e \in E(G)} f_e^k \cdot c_e(f_e^k)$. In a similar fashion, when there are no defections, the average cost $c_{(i,T)}^k(\vec{\sigma})$ of a positive-measure subset T of flow assigned to planner i is given by

$$c_{(i,T)}^k = \frac{1}{\mu(T)} \sum_{e \in E(G)} \mu_T \left(g_{i,k}^{-1}(\mathcal{P}_e) \right) \cdot c_e(f_e^k),$$

where $\mu_T(A) := \mu(A \cap T)$. If there is a defection $\mathcal{D} = (i, S, D)$, the average cost becomes

$$\frac{1}{\mu(T)} \sum_{e \in E(G)} \mu_T \left(\left(g_{i,k}^{-1}(\mathcal{P}_e) \cap \bar{S} \right) \cup d_k^{-1}(\mathcal{P}_e) \right) \cdot c_e(f_e^k).$$

Note that, if $T = S$, the expression above is equal to $\frac{1}{\mu(S)} \sum_{e \in E(G)} \mu(d_k^{-1}(\mathcal{P}_e)) \cdot c_e(f_e^k)$.

We also denote average cost of planner i by $c_i^k(\vec{\sigma})$. This can be computed by setting $T = [0, 1]$ in $c_{(i,T)}^k(\vec{\sigma})$, resulting in $c_i^k(\vec{\sigma}) = \sum_{e \in E(G)} \mu \left(g_{i,k}^{-1}(\mathcal{P}_e) \right) \cdot c_e(f_e^k)$. If there is a defection $\mathcal{D}$, $c_i^k(\vec{\sigma}, \mathcal{D})$ is computed analogously.

Given a discount factor $\lambda \in (0, 1)$, the total cost of the flow is given by

$$c^\lambda(\vec{\sigma}) := \sum_{k \geq 0} \lambda^k c^k(\vec{\sigma}).$$

³Note that the definition is also valid if the domain is $[0, 1]$ instead of $[0, 1]$. However, some definitions are simpler this way.

The total cost $c_i^\lambda(\vec{\sigma})$ of planner i and $c_{(i,T)}^\lambda(\vec{\sigma})$ of a subset T of i 's agents are defined analogously. Even though all these values have been defined for pure strategies, we simply take their expectation when considering mixed strategies⁴. In some scenarios we also want to analyze the expected costs after a concrete sequence of decisions made by planners and cars. If h is the global history at that point (i.e., the sequence of flows and recommendations by planners), we denote by $c^\lambda(\vec{\sigma} \mid h)$ the total cost conditioned on reaching history h . All other values $c_i^k(\vec{\sigma} \mid h)$, $c_{(i,T)}^k(\vec{\sigma} \mid h)$, and $c_{(i,T)}^\lambda(\vec{\sigma} \mid h)$ are defined analogously.

B Proof of Proposition 1

Throughout, fix a planner i and write $\hat{\sigma} := (K(\vec{x}^i), \vec{\sigma}_{-i})$ for the profile in which i plays the karma-implementation of $\vec{x}^i$ and every other planner plays its (stage-independent) strategy from $\vec{\sigma}$.

We first record that the karma-implementation does not change the flow. By construction, $K(\vec{x}^i)$ only reassigns which of i 's users are routed along which path: at every stage k the measure of i 's flow on each path P_j is x_j^i , exactly as under $\vec{x}^i$. Since the remaining planners play the same strategies as in $\vec{\sigma}$ at every stage, the induced aggregate flow at every stage k is the flow f induced by $\vec{\sigma}$. In particular the edge costs $c_e(f_e)$, and hence the path costs $c_{P_j}(f)$, are the same at every stage, and

$$c_{(i,[0,1])}^k(\hat{\sigma}) = \sum_{j=1}^M x_j^i \cdot c_{P_j}(f) = c^i(\vec{\sigma}) \quad \text{for all } k \geq 0.$$

The proof of Proposition 1 relies on the following key fact: using a karma-based strategy, the difference between the accumulated (undiscounted) cost of two users from the same planner is uniformly bounded.

Lemma 1. Fix planner i . Let $A_i^k(u)$ be the cost of user u of planner i at stage k when players play $\hat{\sigma}$. Let $S_i^k(u) := \sum_{\ell=0}^k A_i^\ell(u)$, and let $c_{\max} := \max_j c_{P_j}(f)$ and $c_{\min} := \min_j c_{P_j}(f)$ denote the largest and smallest path costs under f . Then, we have that

$$|S_i^k(u) - S_i^k(v)| \leq c_{\max} - c_{\min}$$

for all $u, v \in [0, 1]$.

Proof. We show this by induction on k . By induction hypothesis, the absolute difference between the accumulated cost of two users u and v is less or equal than $c_{\max} - c_{\min}$. We also have, by construction, that users with the most accumulated cost are given the fastest paths. In particular, if the accumulated cost of u is larger than the accumulated cost of v , then u is guaranteed to get a route at least as fast as v . This means that $S_i^k(u) - S_i^k(v)$ and $A_i^{k+1}(u) - A_i^{k+1}(v)$ have opposite signs, which implies that

$$\begin{aligned} & |S_i^{k+1}(u) - S_i^{k+1}(v)| \\ & \leq \max(|S_i^k(u) - S_i^k(v)|, |A_i^{k+1}(u) - A_i^{k+1}(v)|) \\ & \leq c_{\max} - c_{\min}. \end{aligned}$$

⁴Note that only planners can randomize. We assume that traffic defections are pure without loss of generality.

□

Let $T \subseteq [0, 1)$ be any positive-measure subset of $[0, 1)$ and let $\Delta := c_{\max} - c_{\min}$. Averaging S_i^k over T and using Lemma 1, we get that

$$\left| \frac{1}{\mu(T)} \int_T S_i^k d\mu - \int_{[0,1)} S_i^k d\mu \right| \leq \Delta. \quad (1)$$

Writing $c_i(\vec{\sigma}) = (1 - \lambda) \sum_{k \geq 0} \lambda^k \cdot c_i(\vec{\sigma})$, we get that Proposition 1 is equivalent to showing that

$$\lim_{\lambda \rightarrow 1^-} (1 - \lambda) \sum_{k \geq 0} \left(c_i(\vec{\sigma}) - c_{(i,T)}^k(\hat{\sigma}) \right) \lambda^k = 0.$$

By Equation 1, we have that

$$\left| \sum_{\ell=0}^k \left(c_i(\vec{\sigma}) - c_{(i,T)}^\ell(\hat{\sigma}) \right) \right| \leq \Delta.$$

Therefore, Proposition 1 follows from the fact that $\sum_{k \geq 0} \left(c_i(\vec{\sigma}) - c_{(i,T)}^k(\hat{\sigma}) \right) \lambda^k$ is uniformly bounded, which is a consequence of the following result.

Lemma 2. Let $\{a_k\}_{k \in \mathbb{N}_0}$ be a sequence of real numbers such that there exists a constant C such that, for all k ,

$$\left| \sum_{i=0}^k a_i \right| \leq C.$$

Then,

$$\left| \sum_{k \geq 0} a_k \lambda^k \right| \leq C$$

for all $0 < \lambda < 1$.

Proof. Let $R_k := \sum_{i=0}^k a_i$. By Abel's lemma, we have that

$$\sum_{k \geq 0} a_k \lambda^k = (1 - \lambda) \sum_{k \geq 0} R_k \lambda^k.$$

Therefore,

$$\sum_{k \geq 0} a_k \lambda^k \leq (1 - \lambda) \sum_{k \geq 0} |R_k| \lambda^k \leq C.$$

□

C Proof of Proposition 2

Consider the perspective of a potential deviator, which can be both a planner or a set of vehicles deviating on their own. After deviating, the deviator(s) gain some utility in that round, which we call the **one-shot gain**. Upon detection, all planners (other than a possible deviating planner) switch to the punishment strategy for R rounds, during which the deviator can at best play their best response against the punishment. The deviator therefore trades the one-shot gain for R rounds of suboptimal play. They will choose not to deviate provided that the one-shot gain is lower than R times the difference between the cost of the best-response to punishment

and cost of the optimal flow. The value of R in our strategy is the minimum integer satisfying this inequality. We will end up developing several inequalities of this type, each for a different type of deviator. We prove sufficiency first, then necessity.

For the sufficiency direction, the cooperative phase has every planner route its own share in the same proportions as the optimal aggregate flow f_{opt} , that is, $x_j^i = f_{\text{opt}}(P_j)$ for every planner i and path P_j , and play the karma-based assignment of this strategy. The induced aggregate flow is f_{opt} and $c^i(\vec{\sigma}_{\text{opt}}) = c_{\text{opt}}(G)$ for every planner i . By Proposition 1, the normalized discounted cost of every positive-measure subset of a planner's flow gets arbitrarily close to that planner's cost as $\lambda \rightarrow 1$, hence arbitrarily close to $c_{\text{opt}}(G)$. Note that Proposition 1 still applies when every planner switches to its karma-based assignment simultaneously: a planner's karma-implementation only permutes which of its users are routed along which path, leaving its path-flow vector, and therefore the aggregate flow at every stage, unchanged. Since all the comparisons we make below are strict (the second hypothesis of the proposition is a strict inequality, as is $c_{\min}(f) > c_{\text{opt}}(G)$ whenever it holds), these approximations are immaterial, and for all our purposes we may assume that the cost of each user on the cooperative path is exactly $c_{\text{opt}}(G)$.

Write $\vec{\sigma}$ for the admissible strategy profile of the hypothesis and f for its induced flow, and let $\vec{y}^i$ denote the selfish best response of planner i against $\vec{\sigma}_{-i}$ (and therefore, against its karma implementation).

Sufficiency: case $c_{\min}(f) > c_{\text{opt}}(G)$. Write $K(\vec{\sigma}_{\text{opt}})$ and $K(\vec{\sigma})$ for the profiles in which every planner plays the karma-based assignment of its component of $\vec{\sigma}_{\text{opt}}$ and $\vec{\sigma}$ respectively. The strategy starts with all planners playing $K(\vec{\sigma}_{\text{opt}})$. If a deviation is detected, the planners switch to $K(\vec{\sigma})$ for $R + 1$ rounds, where R is the smallest integer such that

$$R \cdot (c_{\min}(f) - c_{\text{opt}}(G)) > c_{\max}(f_{\text{opt}}) - c_{\min}(f_{\text{opt}}) \quad (2)$$

and

$$R \cdot (c^i(\vec{y}^i, \vec{\sigma}_{-i}) - c_{\text{opt}}(G)) > c_{\text{opt}}(G) - c^i(\vec{y}_{\text{opt}}^i, (\vec{\sigma}_{\text{opt}})_{-i}) \quad (3)$$

for every planner i , where $\vec{y}_{\text{opt}}^i$ denotes the selfish best response of planner i against $(\vec{\sigma}_{\text{opt}})_{-i}$ (and, therefore, also against its karma implementation). Both right-hand sides are finite and both left-hand coefficients are strictly positive, by hypothesis, so such an R exists. Punishment phases stack: every further deviation adds another $R + 1$ rounds.

Equation (2) ensures individual rationality and Equation (3) ensures resilience to competition. We check the four desiderata in turn.

Optimality. After the last deviation only finitely many rounds are non-optimal, so $c^k \rightarrow c_{\text{opt}}(G)$ from every history.

Admissibility. Planners only ever play $\vec{\sigma}_{\text{opt}}$ or $\vec{\sigma}$. The former induces a socially optimal flow, so no alternative can be strictly better in social cost; the latter is admissible by hypothesis.

Individual rationality. A deviating subset of users gains at most $c_{\max}(f_{\text{opt}}) - c_{\min}(f_{\text{opt}})$ in the round in which it deviates. During the $R + 1$ punishment rounds that follow,

even the cheapest available path costs $c_{\min}(f) > c_{\text{opt}}(G)$, so the subset loses at least $c_{\min}(f) - c_{\text{opt}}(G)$ per round relative to complying. By Equation (2) the loss exceeds the gain, so for λ close enough to 1 the deviation is unprofitable. The same holds for a deviation during a punishment phase, which merely adds $R + 1$ further rounds at a cheapest path costing more than $c_{\text{opt}}(G)$.

Resilience to competition. A planner deviating from $K(\vec{\sigma}_{\text{opt}})$ gains at most $c_{\text{opt}}(G) - c^i(\vec{y}_{\text{opt}}^i, (\vec{\sigma}_{\text{opt}})_{-i})$ in that round, and pays at least $c^i(\vec{y}^i, \vec{\sigma}_{-i}) - c_{\text{opt}}(G)$ in each of the $R + 1$ punishment rounds, since $\vec{y}^i$ is its best response against $\vec{\sigma}_{-i}$. Equation (3) and the same discounting argument make the deviation unprofitable.

Sufficiency: case $c_{\min}(f) = c_{\text{opt}}(G)$. We first record two consequences of the hypothesis. Call a path *cheapest* if its cost under f equals $c_{\text{opt}}(G)$.

First, no planner routes all of its flow onto cheapest paths: such a planner would have cost exactly $c_{\text{opt}}(G)$ at $\vec{\sigma}$, and hence best-response cost at most $c_{\text{opt}}(G)$, contradicting $c^i(\vec{y}^i, \vec{\sigma}_{-i}) > c_{\text{opt}}(G)$.

Second, no cheapest path has constant cost. Suppose P had $c_P(f) = c_{\text{opt}}(G)$ with c_P constant. Then any planner using a path of cost above $c_{\text{opt}}(G)$ could shift flow onto P , strictly lowering both its own cost and the social cost, contradicting admissibility of $\vec{\sigma}$; hence every planner uses only cheapest paths and has cost $c_{\text{opt}}(G)$, contradicting the previous paragraph.

For each planner i , let $x_{\min}^i$ be the fraction of its flow that $\vec{\sigma}$ places on cheapest paths and γ^i the per-round cost of the remainder, so that $c^i(\vec{\sigma}) = \gamma^i + x_{\min}^i c_{\text{opt}}(G)$. By the first observation, $x_{\min}^i < 1$, and since $c^i(\vec{\sigma}) \geq c^i(\vec{y}^i, \vec{\sigma}_{-i}) > c_{\text{opt}}(G)$, we have $\gamma^i + x_{\min}^i c_{\text{opt}}(G) > c_{\text{opt}}(G)$.

The strategy is as in the strict case; only the choice of R is refined. When a deviation is detected, let f' be the realized flow, which places strictly more flow on some cheapest path than f does. Let R be the smallest integer satisfying

$$R \cdot (c_P(f') - c_{\text{opt}}(G)) > c_{\max}(f_{\text{opt}}) - c_{\min}(f_{\text{opt}}) \quad (4)$$

for every P with $c_P(f') > c_P(f)$,

$$R \cdot (c^i(\vec{y}^i, \vec{\sigma}_{-i}) - c_{\text{opt}}(G)) > c_{\text{opt}}(G) - c^i(\vec{y}_{\text{opt}}^i, (\vec{\sigma}_{\text{opt}})_{-i}) \quad (5)$$

for every planner i , and

$$\begin{aligned} & R \cdot (\gamma^i + x_{\min}^i c_{\text{opt}}(G) - c_{\text{opt}}(G)) \\ & > (c_{\max}(f_{\text{opt}}) - c_{\min}(f_{\text{opt}})) + \\ & \quad (c_{\max}(f) - c_{\min}(f)) \end{aligned} \quad (6)$$

for every planner i .

All three left-hand coefficients are strictly positive: the first because a deviation onto a cheapest path strictly raises its cost, by the second observation above; the second by hypothesis; the third by the display above.

Optimality and admissibility hold exactly as in the strict case.

Individual rationality. A deviating subset gains at most $c_{\max}(f_{\text{opt}}) - c_{\min}(f_{\text{opt}})$ in the round in which it deviates. If it complies thereafter, the planner pays at least $(c_{\max}(f_{\text{opt}}) - c_{\min}(f_{\text{opt}})) + (c_{\max}(f) - c_{\min}(f))$ additional

cost distributed among the following rounds, for a sufficiently large discount factor λ . Since the planner follows a karma-based implementation, by Lemma 1, each subset of its users pays at least $c_{\max}(f_{\text{opt}}) - c_{\min}(f_{\text{opt}})$ additional cost, making the deviation unprofitable. If instead it keeps deviating onto a cheapest path, that path's cost rises above $c_{\text{opt}}(G)$, and Equation (4) makes this unprofitable as well for a sufficiently large λ . Since these are the only two options available to the subset, individual rationality holds.

Resilience to competition. As in the strict case: a planner's one-shot gain is outweighed by R rounds at $c^i(\vec{y}^i, \vec{\sigma}_{-i}) > c_{\text{opt}}(G)$, by Equation (5). This concludes the sufficiency part of the proof.

Before we go with the proof for the necessity condition, we introduce the following lemma.

Lemma 3. *The set $\mathcal{A}$ of admissible strategy profiles of Γ is nonempty and compact, and $c_{\min}$ attains a maximum on it.*

Proof. $\mathcal{A}$ is nonempty: any profile inducing a socially optimal flow is admissible, since no alternative can strictly lower the social cost, so the consequent of the admissibility implication holds trivially.

$\mathcal{A}$ is closed. Its complement is open: if $\vec{\sigma}$ is not admissible, some planner i and strategy $\vec{\tau}^i$ satisfy $c^i(\vec{\tau}^i, \vec{\sigma}_{-i}) < c^i(\vec{\sigma})$ and $c(\vec{\tau}^i, \vec{\sigma}_{-i}) < c(\vec{\sigma})$. Both c^i and c are continuous in the profile, and both inequalities are strict, so they persist for all profiles in a neighbourhood of $\vec{\sigma}$, every one of which is therefore inadmissible.

Since $\mathcal{A}$ is a closed subset of the product of the planners' simplices, it is compact. Finally $c_{\min}$ is continuous in the profile, as a minimum of finitely many continuous path costs composed with the (linear) map from profiles to induced flows, so it attains a maximum on $\mathcal{A}$. $\square$

With Lemma 3 we can continue with the proof. Write $\mathcal{A}$ for the set of admissible strategy profiles of Γ and, by Lemma 3, let

$$M := \max_{\vec{\tau} \in \mathcal{A}} c_{\min}(f_{\vec{\tau}}),$$

where $f_{\vec{\tau}}$ denotes the flow induced by $\vec{\tau}$. Throughout we treat deterministic planner strategies; the probabilistic case is identical with all quantities replaced by their expectations. We first show $M \geq c_{\text{opt}}(G)$, which gives the first condition, and then produce, for each planner i , a profile satisfying both.

First condition. We prove the contrapositive: if $M < c_{\text{opt}}(G)$, no strategy profile of Γ' satisfies the four desiderata.

Since the edge costs are continuous and the space of flows is compact, we may fix a positive integer R large enough that adding $\frac{1}{R}$ units of flow to any single path raises its cost by less than $c_{\text{opt}}(G) - M$. Suppose for contradiction that $\vec{\sigma}'$ satisfies all four desiderata. By admissibility, the profile $\vec{\tau}^k$ played at every stage k of every history lies in $\mathcal{A}$, so its induced flow f^k satisfies $c_{\min}(f^k) \leq M$.

Partition each planner's users into R pieces $S_j = [\frac{j-1}{R}, \frac{j}{R})$. Fix a planner i and an index j and consider the defection $\mathcal{D}_i^j = (i, S_j, D_i^j)$ that routes S_j onto the cheapest available path at every stage. At each stage this piece adds at most $\frac{1}{R}$ units to that path, so it pays less than

$c_{\min}(f^k) + (c_{\text{opt}}(G) - M) \leq c_{\text{opt}}(G)$. Individual rationality forces the defection to be unprofitable for λ near 1, so

$$c_{(i,S_j)}^\lambda(\vec{\sigma}') \leq c_{(i,S_j)}^\lambda((\vec{\sigma}', \mathcal{D}_i^j)) < \sum_{k \geq 0} \lambda^k c_{\text{opt}}(G)$$

for all i, j . Since the S_j partition $[0, 1]$ into pieces of measure $\frac{1}{R}$, we have $c_i^k = \frac{1}{R} \sum_j c_{(i,S_j)}^k$, and hence

$$c^\lambda(\vec{\sigma}') = \sum_i \alpha_i \frac{1}{R} \sum_j c_{(i,S_j)}^\lambda(\vec{\sigma}') < \sum_{k \geq 0} \lambda^k c_{\text{opt}}(G).$$

But $c_{\text{opt}}(G)$ is the minimum social cost of a unit flow, so $c^k(\vec{\sigma}') \geq c_{\text{opt}}(G)$ at every stage and $c^\lambda(\vec{\sigma}') \geq \sum_{k \geq 0} \lambda^k c_{\text{opt}}(G)$, a contradiction. Hence $M \geq c_{\text{opt}}(G)$, and any maximizer $\vec{\tau}^*$ satisfies the first condition.

Second condition, case $M > c_{\text{opt}}(G)$. Let $\vec{\tau}^*$ attain M . For every planner i , the cost $c^i(\vec{\tau}^*) = \sum_j x_j^i c_{P_j}(f_{\tau^*})$ is a convex combination of the costs of the paths i uses, each of which is at least $c_{\min}(f_{\tau^*}) = M$. Hence $c^i(\vec{\tau}^*) \geq M > c_{\text{opt}}(G)$ for every i , and $\vec{\tau}^*$ witnesses both conditions simultaneously for all planners.

Second condition, case $M = c_{\text{opt}}(G)$. Fix a planner i and suppose for contradiction that every $\vec{\tau} \in \mathcal{A}$ with $c_{\min}(f_{\tau}) \geq c_{\text{opt}}(G)$ has $c^i(\vec{\tau}) \leq c_{\text{opt}}(G)$. By the convex-combination bound above, any such $\vec{\tau}$ has $c^i(\vec{\tau}) \geq c_{\min}(f_{\tau}) = c_{\text{opt}}(G)$, so in fact $c^i(\vec{\tau}) = c_{\text{opt}}(G)$ and planner i uses only paths of cost exactly $c_{\text{opt}}(G)$ at $\vec{\tau}$.

Every admissible profile has $c_{\min} \leq M = c_{\text{opt}}(G)$, so a sufficiently small piece routed onto the cheapest available path pays at most $c_{\text{opt}}(G) + \eta$ in every round, where η can be made arbitrarily small by shrinking the piece. By optimality the cooperative flow is f_{opt} , and since $\text{PoA}(G) > 1$ it is not a Wardrop equilibrium, so some path P has $c_P(f_{\text{opt}}) = c_{\text{opt}}(G) - 2\theta$ with $\theta > 0$; fix the piece small enough that $\eta < \theta$.

First, $c^i(\vec{\sigma}_{\text{opt}}) < c_{\text{opt}}(G)$. Otherwise planner i uses a path of cost at least $c_{\text{opt}}(G)$ in cooperative rounds; a small piece of that flow routed onto the cheapest available path at every stage pays at most $c_{\text{opt}}(G) - \theta$ in cooperative rounds and at most $c_{\text{opt}}(G) + \eta$ in all others, while under $\vec{\sigma}'$ it pays at least $c_{\text{opt}}(G)$ in cooperative rounds and, by the assumption on planner i , exactly $c_{\text{opt}}(G)$ in any round whose profile has $c_{\min} = c_{\text{opt}}(G)$. For λ near 1 the defection is strictly profitable, contradicting individual rationality.

Since the planners' costs at f_{opt} are an α -weighted average equal to $c_{\text{opt}}(G)$, some planner j has $c^j(\vec{\sigma}_{\text{opt}}) = c_{\text{opt}}(G) + \delta$ with $\delta > 0$, and therefore uses a path Q with $c_Q(f_{\text{opt}}) \geq c_{\text{opt}}(G) + \delta$. Take a piece S of j 's flow on Q , small enough that $\eta < \delta$, and route it onto the cheapest available path at every stage. It pays less than $c_{\text{opt}}(G) + \delta$ in every round, while on the equilibrium path its users pay $c_Q(f_{\text{opt}}) \geq c_{\text{opt}}(G) + \delta$ in every cooperative round. For λ near 1 this defection is strictly profitable, again contradicting individual rationality. Hence no such planner i exists.

D Proof of Theorem 1

At a high level, the proof of Theorem 1 goes as follows. For part (a), we show that if α_i approaches 0 for all i , then some of

the planner equilibria approach to the Wardrop equilibria of the system. In particular, for a sufficiently small $\alpha_{\min}$, we can always select a planner equilibrium such that the flow cost for every path is strictly larger than the optimal expected utility. This way, the planner equilibrium can be used as punishment in the same way as in the Pigou case. For part (b), we show that if α_i approaches 1 for some i , then all of i 's strategies in strategy profiles that satisfy the desiderata approach to an optimal flow. Since the price of stability is strictly greater than 1, no optimal flow is a Wardrop equilibrium. Therefore, for a sufficiently large $\alpha_{\max}$, there is always a path with strictly lower cost than the expected cost of an optimal flow. Thus, for sufficiently small fractions of the total flow, it is always beneficial to defect and take these lower cost routes instead.

To prove this result in its full generality in the case of single-commodity flow networks we need the following additional notation. Given a single-commodity flow network G , we denote by $\mathcal{F}_{\text{opt}}^G$ the set of (socially) optimal flows on G . Moreover, given a flow f in G , we define

$$P_{\min}^f := \operatorname{argmin}_{P \in \mathcal{P}} \{c_P(f)\} \\ P_{\max}^f := \operatorname{argmax}_{P \in \mathcal{P} \mid f(P) > 0} \{c_P(f)\}.$$

Intuitively, $P_{\max}^f$ is the path used by f that maximizes the total cost from source to sink, while $P_{\min}$ is the path that minimizes it (not necessarily used by f). With this, we can define

$$c_{\min}(f) := c_{P_{\min}^f}(f) \\ c_{\max}(f) := c_{P_{\max}^f}(f).$$

We also define $\Delta c(f) := c_{\max}(f) - c_{\min}(f)$ and $c_{\text{opt}}(G)$ as the total cost of any optimal flow of G . Using this notation, a full proof for the single-commodity network case is given next. The generalization to multiple commodities is straightforward.

Part (a)

Proposition 2 shows that we can reduce part (a) to showing that there exists a sufficiently small value $\alpha_{\min}$ such that, if $\alpha_i < \alpha_{\min}$ for all $i \in C$, there exists a planner equilibrium for the one-shot version of Γ in which the resulting flow f satisfies $c_{\min}(f) > c_{\text{opt}}(G)$.

Suppose that such $\alpha_{\min}$ does not exist. This means that there exists a sequence of strictly positive real values $\alpha_{\min}^1, \alpha_{\min}^2, \dots$ with $\lim_{i \rightarrow \infty} \alpha_{\min}^i = 0$ such that, for all i , there exists $N_i \in \mathbb{N}$ and a partition $\alpha^i = (\alpha_1^i, \dots, \alpha_N^i)$ with $\alpha_j^i < \alpha_{\min}^i$ for all $j \in [N_i]$, such that every flow f_i induced by a deterministic planner equilibrium in $(G, [N_i], \alpha^i)$ satisfies $c_{\min}(f_i) \leq c_{\text{opt}}(G)$.

Let $\vec{\sigma}_1, \vec{\sigma}_2, \dots$ be any deterministic planner equilibria of $(G, [N_1], \alpha^1), (G, [N_2], \alpha^2), \dots$, respectively, and let $f_1, f_2, \dots$ be their induced flows. Since the space of possible flows in G is compact, the Bolzano-Weierstrass theorem implies that there is a sequence of indices $i_1 < i_2 < \dots$ such that the subsequence of flows $f_{i_1}, f_{i_2}, \dots$ converges to some flow f . The remainder of the proof consists of proving the following proposition.

Proposition 5.

$$\Delta c(f) = 0.$$

Note that Proposition 5 implies that f is a Wardrop equilibrium of G . By construction, f also satisfies $c_{\min}(f) \leq c_{\text{opt}}(G)$. The combination of these properties would imply that f is simultaneously a Wardrop equilibrium and an optimal flow, which contradicts the assumption that the price of stability of G is strictly greater than 1.

Proof of Proposition 5. Since the cost functions are continuous and the space of possible flows is compact, the Heine-Cantor theorem implies that, for all $\epsilon > 0$ there exists $\delta > 0$ such that, for all paths $P \in \mathcal{P}$ and pairs of flows f' and f'' ,

$$\|f' - f''\|_1 < \delta \implies |c_P(f') - c_P(f'')| < \epsilon,$$

where $\|f' - f''\|_1 := \sum_{P \in \mathcal{P}} |f'(P) - f''(P)|$. In particular, there exists $\delta_0 > 0$ such that, for all paths $P \in \mathcal{P}$ and pairs of flows f' and f'' ,

$$\|f' - f''\|_1 < \delta_0 \implies |c_P(f') - c_P(f'')| < \frac{\Delta c(f)}{4}.$$

Moreover, for each $m > 0$, there exists $\delta_m > 0$ such that, for all paths $P \in \mathcal{P}$ and pairs of flows f' and f'' ,

$$\|f' - f''\|_1 < \delta_m \implies |c_P(f') - c_P(f'')| < \frac{1}{m}.$$

Since f_{i_n} converges to f , for each $m > 0$ there exists an integer n_m such that, if $n \geq n_m$, then both $\|f - f_{i_n}\|_1 < \min\{\delta_0, \delta_1, \dots, \delta_m\}$ and $\alpha_{\min}^{i_n} < \min\{\delta_0, \delta_1, \dots, \delta_m\}$. Consider an n strictly larger than n_m , and let k be any of the planners in $[N_{i_n}]$. Suppose that k moves all its flow going through $P_{\max}^f$ to $P_{\min}^f$. Then, since the total flow moved cannot exceed $\alpha_{\min}^{i_n}$ (which is smaller than δ_m), the total cost increase in $P_{\max}^f$ is at most $\frac{1}{m}$. We compute next the additional cost for the portion of k 's flow that does the transition. Since $\|f - f_{i_n}\|_1 < \delta_0$, we have that

$$\begin{aligned} c_{P_{\max}^f}(f_{i_n}) &> c_{P_{\max}^f}(f) - \frac{\Delta c(f)}{4} \\ c_{P_{\min}^f}(f_{i_n}) &< c_{P_{\min}^f}(f) + \frac{\Delta c(f)}{4} \end{aligned}$$

and, since $\alpha_{\min}^{i_n} < \delta_0$, we also have that the total cost increase in $P_{\min}^f$ after the transition cannot exceed $\frac{\Delta c(f)}{4}$. Therefore, the additional cost for the portion of k 's flow originally in $P_{\max}^f$ is bounded by

$$\left(c_{P_{\min}^f}(f) + \frac{\Delta c(f)}{4} + \frac{\Delta c(f)}{4} \right) - \left(c_{P_{\max}^f}(f) - \frac{\Delta c(f)}{4} \right),$$

which is at most $\frac{-\Delta c(f)}{4}$. Altogether, if x_k^n is k 's fraction of the traffic in f_{i_n} that goes through $P_{\max}^f$, we have that the total cost difference for k when performing this transition is bounded by

$$\frac{1}{m} - x_k^n \cdot \frac{\Delta c(f)}{4}.$$

Since $\vec{\sigma}_{i_n}$ is a planner equilibrium, k should not be able to decrease its total cost with this operation. Therefore, we have that $\frac{1}{m} - x_k^n \cdot \frac{\Delta c(f)}{4} \geq 0$, which is equivalent to

$$x_k^n \leq \frac{4}{m \cdot \Delta c(f)}.$$

This argument implies that, if $n \geq n_m$, the total fraction of each planner's flow that is assigned to path $P_{\max}^f$ is at most $\frac{4}{m \cdot \Delta c(f)}$, which means that $f_{i_n}(P_{\max}^f) \leq \frac{4}{m \cdot \Delta c(f)}$ and therefore that, in the limit, $f(P_{\max}^f) = 0$. This contradicts the fact that $P_{\max}^f$ must be used by f . $\square$

Part (b)

Again, Proposition 2 implies that part (b) reduces to show that there exists $\alpha_{\max} < 1$ such that, for all routing games $\Gamma = (G, C, \alpha)$ such that $\alpha_1 > \alpha_{\max}$, in all strategy profiles $\vec{\sigma}$ that satisfy admissibility, there exists a path P such that the expected cost of P after playing $\vec{\sigma}$ is strictly less than c_{opt} .

Suppose that such $\alpha_{\max}$ does not exist. This means that there exists a sequence $\alpha_{\max}^1, \alpha_{\max}^2, \dots$ with $\lim_{i \rightarrow \infty} \alpha_{\max}^i = 1$ in which, for all i , there exists $N_i > 0$ and a partition α^i with $\alpha_1^i > \alpha_{\max}^i$ such that, for all strategy profiles $\vec{\sigma}_i$ for $(G, [N_i], \alpha^i)$ that satisfy admissibility, the expected cost of each path after playing $\vec{\sigma}_i$ is at least as large as c_{opt} .

Let $\vec{\sigma}_1, \vec{\sigma}_2, \dots$ be a sequence of strategy profiles for $(G, [N_1], \alpha_1), (G, [N_2], \alpha_2), \dots$, respectively, such that the expected cost of each path after playing $\vec{\sigma}_i$ is at least as large as c_{opt} . Then, we have the following proposition, whose proof is a direct implication of Prokhorov's theorem (Prokhorov 1956).

Proposition 6. *There exists a sequence of indices $i_1, i_2, \dots$ such that the strategy profiles $\vec{\sigma}_{i_1}, \vec{\sigma}_{i_2}, \dots$ converge to some strategy profile $\vec{\sigma}$ in the weak topology. In particular, the expected cost of each path P after playing $\vec{\sigma}_{i_n}$ converges to the expected cost of P after playing $\vec{\sigma}$.*

Let $\vec{\sigma}$ be the sequence obtained in Proposition 6. Because of the convergence of expectations, we have that

$$\lim_{n \rightarrow \infty} c(\vec{\sigma}_{i_n}) = \lim_{n \rightarrow \infty} c^1(\vec{\sigma}_{i_n}) = c(\vec{\sigma}).$$

Therefore, $c(\vec{\sigma})$ must be equal to c_{opt} . Otherwise, near the limit, both the total cost and planner 1's cost decrease if planner 1 switches to an optimal flow, contradicting the fact that all strategy profiles in the sequence satisfy admissibility. This means that $\vec{\sigma}$ is an optimal flow in which all paths of G cost at least c_{opt} . This would imply that, after playing $\vec{\sigma}$, all paths in G cost exactly c_{opt} , which further implies that $\vec{\sigma}$ is both an optimal flow and a Wardrop equilibrium. This contradicts the assumption that the price of stability of G is strictly greater than 1.

E Planner Equilibrium Fails to Maximize Cheapest Path

We will construct a WASRG instance which satisfies the property that the planner equilibrium does not maximize the cost of the cheapest path. The counterexample will use a Braess network topology (Roughgarden 2007). The underlying single-commodity flow network we play on is given by fig. 1. The WASRG we play has $n = 6$ planners with uniform weights, that is $\alpha = \left(\frac{1}{6}, \dots, \frac{1}{6}\right)$. The three paths through

the network are given by

$$P_1 : s \rightarrow a \rightarrow t, \quad P_2 : s \rightarrow b \rightarrow t, \quad P_3 : s \rightarrow a \rightarrow b \rightarrow t.$$

For this instance, the cost of the socially optimal flow is $c_{\text{opt}} = 1.41337$.




Figure 1: A single commodity flow network based on Braess network. This network has four nodes, $\{s, a, b, t\}$ where s and t are the source and sink nodes respectively. The total flow to be routed through the network is $F = 1$.

Due to all cost functions being convex in this example, we can find all admissible strategy profiles by using weighted sum optimization (Ehrgott 2005). We assign each planner i a value θ^i and an objective function $\theta^i c^i + (1 - \theta^i)c$, where c^i is their own cost and c is the cost over all users. This allows us to easily find and verify that the resulting profiles are admissible.

We sweep a subset of the admissible strategy profiles where we fix the selfishness of every planner to be uniform. That is, $\theta^i = \theta$ for all planners i . The costs of paths P_1 and P_3 across this sweep are plotted in fig. 2. Path P_2 , which is more expensive than P_1 and P_3 at all θ , is left out since the gap between the cost of P_1 at $\theta = 1$ would not be visible if P_2 was included.




Figure 2: The path costs of two paths in an WASRG played on Braess network. Shown are the costs of paths P_1 and P_3 versus the selfishness of all planners. The dotted line shows the cost of the optimal flow of the network. Path P_2 was left out and the range of θ was limited to $[0.70, 1.00]$ to ensure the difference between c_{opt} and the cost of c_{P_1} is visible at $\theta = 1.00$.

fig. 2 shows that the planner equilibrium does not attain the largest $c_{\min}$ across the sweep. Instead an admissible strategy profile attains it at $\theta = 0.916$. Let f^θ represent the aggregate flow to the equilibrium achieved at the value θ . Numerically, we can then verify that:

$$\begin{aligned} c_{\min}(f^{0.916}) &= 1.42613 \\ &> c_{\text{opt}} &= 1.41337 \\ &> c_{\min}(f^1) &= 1.41192, \end{aligned}$$

so the cheapest path at the planner equilibrium drops below c_{opt} , while at the equilibrium at $\theta = 0.916$ it stays above this value. The planner equilibrium therefore fails the condition $c_{\min} \geq c_{\text{opt}}$ of Proposition 2, whereas the equilibrium at $\theta = 0.916$ satisfies it.

Checking the second sufficiency condition of Proposition 2, we see that the best selfish response against $f^{0.916}$ is given by

$$x^{*i} = \begin{pmatrix} P_1 & P_2 & P_3 \\ 0.94418 & 0.05581 & 0 \end{pmatrix}.$$

After this deviation, the deviating planner pays a cost of $c^i(f^{-i}, x^{*i}) = 1.45543$, above $c_{\text{opt}} = 1.41337$ by 0.042. The equilibrium at $\theta = 0.916$ therefore meets both sufficiency conditions of Proposition 2: it is an admissible strategy profile with $c_{\min} \geq c_{\text{opt}}$, and every planner's selfish best response costs strictly more than c_{opt} . By that proposition a strategy satisfying all four desiderata exists for this instance.

F Proof of Proposition 3

The main tool we will use to prove the result are **marginal costs**.

Definition 1. The individual marginal cost of a player i on a path P is given by

$$l_P^i(f^{-i}, x^i) = c_P(f) + \alpha_i \sum_{e \in P} x_e^i c'_e(f_e). \quad (7)$$

The social marginal cost on a path P is given by

$$l_P(f) = c_P(f) + \sum_{e \in P} f_e c'_e(f_e). \quad (8)$$

These marginal costs take on the role of path costs for the planners and for the social cost respectively. Where in a regular Wardrop equilibrium the players seek to equalize the path costs, in a planner equilibrium the agents seek to equalize their individual marginal costs. This result comes from an application of the *Karush-Kuhn-Tucker* conditions to the minimization problem a single planner solves (Nocedal and Wright 2006)

Lemma 4. *A flow distribution $(x^1, x^2, \dots, x^N)$ is a Nash equilibrium for a WASRG if and only if, for every planner i , every path P with $x_P^i > 0$, and every other path $\tilde{P}$,*

$$l_P^i(f^{-i}, x^i) \leq l_{\tilde{P}}^i(f^{-i}, x^i).$$

By considering the perspective of a imaginary planner which controls all flow, we can extend the above result to also apply to the socially optimal flow.

Corollary 3. *An aggregate flow f is optimal if and only if, for every path P with $f(P) > 0$ and every other path $\tilde{P}$,*

$$l_P(f) \leq l_{\tilde{P}}(f).$$

As mentioned, individual marginal costs take the place of path costs for the planners. However, marginal costs and path costs do differ in some ways. Path costs have the property that inequalities between path costs immediately tells us which deviations are profitable and profitable deviations tells us which paths are more expensive and which are cheaper. This only partially translates to marginal costs. The following result describes the relation between inequalities of marginal costs and profitable deviations.

Lemma 5. *Let $(x^1, \dots, x^N)$ be any feasible flow. Then, for any two paths $P, \tilde{P}$ and for any planner i with $x_P^i > 0$ we have that:*

($\Rightarrow$) *If $l_P^i(f^{-i}, x^i) < l_{\tilde{P}}^i(f^{-i}, x^i)$, then there exists $\varepsilon_0 \in (0, x_P^i]$ such that shifting ε_0 flow from P to $\tilde{P}$ strictly decreases planner i 's cost.*

($\Leftarrow$) *Assume in addition that every edge $e \in P \cup \tilde{P}$ has a convex cost function. Then, if there is a profitable deviation from P to $\tilde{P}$, then $l_P^i(f^{-i}, x^i) < l_{\tilde{P}}^i(f^{-i}, x^i)$.*

Proof of Lemma 5. Let planner i play the strategy $\tilde{x}^i$ and fix the other planners to any strategy with resulting background flow f^{-i} . Any deviation can be captured by a vector of values $\varepsilon \in \mathbb{R}^{|\mathcal{P}|}$ where each path gains or loses $\varepsilon_P \in [-\tilde{x}_P^i, 1 - \tilde{x}_P^i]$ units of flow. Specifically, we focus on the deviation $\varepsilon = \varepsilon_{\tilde{P}} - \varepsilon_P$. After this deviation, the contribution of an edge e to the objective function of planner i is given by

$$H_e(\varepsilon) := (\tilde{x}_e^i + \varepsilon) c_e(f_e^{-i} + \alpha_i(\tilde{x}_e^i + \varepsilon)). \quad (9)$$

Since each edge cost function is convex and non-decreasing, the function (9) is convex in the region $\varepsilon \in$

$[-\tilde{x}_P^i, 1 - \tilde{x}_P^i]$. The change in player i 's total cost after deviating ε flow from P to $\tilde{P}$ is given by

$$\Delta c^i(\varepsilon) := \sum_{e \in P \setminus \tilde{P}} (H_e(-\varepsilon) - H_e(0)) + \sum_{e \in \tilde{P} \setminus P} (H_e(\varepsilon) - H_e(0)), \quad (10)$$

where we use the fact that shared edges $e \in P \cap \tilde{P}$ do not change in flow. A unilaterally profitable deviation is exactly an $\varepsilon_0 \in (0, x_P^i]$ with $\Delta c^i(\varepsilon_0) < 0$. Differentiating H_e with regards to ε at $\varepsilon = 0$ gives

$$H'_e(0) = \alpha_i c_e(f_e) + (\alpha_i)^2 x_e^i c'_e(f_e) = \alpha_i l_e^i(f^{-i}, x^i), \quad (11)$$

Taking the derivative of (10) and filling in (11) gives us

$$\frac{\partial \Delta c^i}{\partial \varepsilon} = \alpha_i (l_{\tilde{P}}^i(f^{-i}, x^i) - l_P^i(f^{-i}, x^i)). \quad (12)$$

($\Rightarrow$) The hypothesis gives $\left. \frac{\partial \Delta c^i}{\partial \varepsilon} \right|_{\varepsilon=0} < 0$ and per (10)

we have that $\Delta c^i(0) = 0$. So by first-order Taylor expansion there exists an $\varepsilon > 0$ such that $\Delta c^i(\varepsilon) < 0$.

($\Leftarrow$) We have that there is a value $\varepsilon_0 > 0$ such that $\Delta c^i(\varepsilon_0) < 0$. The function Δc^i is convex on $[0, x_P^i]$ as it is a sum of convex functions. For a convex differentiable function g and points u, v in its domain, the **tangent line inequality** states

$$g'(u)(v - u) \leq g(v) - g(u). \quad (13)$$

We can apply this on the points 0 and ε_0 to find that

$$\frac{\partial \Delta c^i}{\partial \varepsilon} \leq \frac{\Delta c^i(\varepsilon_0)}{\varepsilon_0} < 0. \quad (14)$$

The equations (12) and (14) combine to give the desired result of $l_P^i(f^{-i}, x^i) < l_{\tilde{P}}^i(f^{-i}, x^i)$. $\square$

We now have sufficient background knowledge on marginal costs to start working towards the main proof. First we start by stating several supporting results on admissible strategy profiles and how marginal costs are ordered in those equilibria. Then, we will examine how planners distribute flow in the planner equilibrium. After that we will seek to proof Proposition 3.

The first supporting result captures the idea that a cheap path must have a high social cost. We prove that if a path has a lower cost than another path which has positive flow, then the former path has a higher social marginal cost than the latter. This formalizes the idea that to improve social cost one has to move flow away from the cheapest paths towards more expensive ones.

Lemma 6. *At any admissible strategy profile, if two edges G, B satisfy*

$$c_G(f_G) \leq c_B(f_B)$$

and $f_B > 0$, then

$$l_G(f) \geq l_B(f).$$

Proof of Lemma 6. Let G and B be two edges satisfying the stated conditions. Assume that $l_G(f) \neq l_B(f)$, if this assumption does not hold we are automatically done. Per this assumption and Lemma 5 we can strictly improve social cost by deviating from G to B or vice-versa. This means that among users with $x_B^i > 0$ the quantities $l_G^i(f^{-i}, x^i) - l_B^i(f^{-i}, x^i)$ are either all non-negative or all non-positive. Else, if there are two players i, j such that

$$l_G^i(f^{-i}, x^i) - l_B^i(f^{-i}, x^i) < 0 < l_G^j(f^{-j}, x^j) - l_B^j(f^{-j}, x^j),$$

then regardless of the sign of $l_G(f) - l_B(f)$, one of these players will violate admissibility.

Without loss of generality, order the players such that the first k are all players with $x_B^i > 0$. As $f_B > 0$ we have that $k \geq 1$. Summing the individual marginal cost over these players gives

$$\begin{aligned} & \sum_{i=1}^k l_G^i(f^{-i}, x^i) - l_B^i(f^{-i}, x^i) \\ &= k(c_G(f_G) - c_B(f_B)) + \\ & \sum_{i=1}^k \alpha_i x_G^i c'_G(f_G) - f_B c'_B(f_B) \\ &\leq k(c_G(f_G) - c_B(f_B)) + f_G c'_G(f_G) - f_B c'_B(f_B). \end{aligned} \quad (15)$$

Since $c_G(f_G) - c_B(f_B) \leq 0$, we can write that

$$\begin{aligned} & k(c_G(f_G) - c_B(f_B)) + f_G c'_G(f_G) - f_B c'_B(f_B) \\ &\leq c_G(f_G) - c_B(f_B) + f_G c'_G(f_G) - f_B c'_B(f_B) \\ &= l_G(f) - l_B(f). \end{aligned} \quad (16)$$

Combining (15) and (16) gives us that

$$\sum_{i=1}^k l_G^i(f^{-i}, x^i) - l_B^i(f^{-i}, x^i) \leq l_G(f) - l_B(f). \quad (17)$$

Now, there are two possible cases remaining:

1. For all $i \leq k$ we have that $l_G^i(f^{-i}, x^i) - l_B^i(f^{-i}, x^i) \geq 0$
 2. Or, for all $i \leq k$ we have that $l_G^i(f^{-i}, x^i) - l_B^i(f^{-i}, x^i) \leq 0$ with the inequality being strict for at least one player.
- This last conditions is there since if all players have equality we revert back to the first case.

In the first case, we can use (17) to denote that

$$0 \leq \sum_{i=1}^k l_G^i(f^{-i}, x^i) - l_B^i(f^{-i}, x^i) \leq l_G(f) - l_B(f),$$

and we have shown that $l_G(f) \geq l_B(f)$.

In the latter case, there is a player for who deviating from B to G reduces their own cost function. This means that this deviation must increase the social cost, hence by Lemma 5 we have that

$$l_G(f) \geq l_B(f). \quad \square$$

Lemma 6 gives us the conditions for when an inequality between path costs leads to the reverse inequality between the social marginal costs. We will also require the opposite direction. We state here what requirements are necessary to ensure that the implication $l_G(f) > l_B(f) \implies c_G(f_G) < c_B(f_B)$ holds.

Corollary 4. *At any admissible strategy profile, if two edges G, B satisfy*

$$l_G(f) > l_B(f)$$

and $f_B, f_G > 0$, then

$$c_G(f_G) < c_B(f_B).$$

Proof of Corollary 4. We prove the contrapositive: assume $c_G(f_G) \geq c_B(f_B)$ and $f_G, f_B > 0$, and show $l_G(f) \leq l_B(f)$.

Since $c_B(f_B) \leq c_G(f_G)$ and $f_G > 0$, we can apply Lemma 6 with the roles of G and B swapped, which gives $l_B(f) \geq l_G(f)$. $\square$

Results Lemma 6 and Corollary 4 give us a correspondence between the ordering of the path costs and the social marginal costs. If we combine this with Lemma 5 and the admissibility condition, we can get a sense of how the individual marginal costs are ordered as well. This gives us a full picture of how the marginal costs look based on the costs of the paths.

Corollary 5. *At any admissible strategy profile, if edges G, B satisfy*

$$c_G(f_G) < c_B(f_B) \quad l_G(f) > l_B(f),$$

then for every player i ,

$$l_G^i(f^{-i}, x^i) \leq l_B^i(f^{-i}, x^i).$$

Proof of Corollary 5. The proof splits on whether player i uses edge G .

Case 1: player i does not use G . Since we have that $x_G^i = 0$, we immediately have that $l_G^i(f^{-i}, x^i) = c_G(f_G)$. Furthermore, $l_B^i(f^{-i}, x^i) = c_B(f_B) + \alpha_i x_B^i c'_B(f_B) \geq c_B(f_B)$ since $c'_B(f_B) \geq 0$ and $x_B^i \geq 0$. So $l_G^i(f^{-i}, x^i) = c_G(f_G) < c_B(f_B) \leq l_B^i(f^{-i}, x^i)$.

Case 2: player i uses G . Since we have that $x_G^i > 0$, we can consider the deviation of player i sending flow from G to B . By the assumption that $l_G(f) > l_B(f)$ this deviation decreases the social cost. To retain admissibility, this must then increase the cost for player i . By Lemma 5 this immediately implies that $l_G^i(f^{-i}, x^i) \leq l_B^i(f^{-i}, x^i)$. $\square$

Two more results we will need concern how the planner equilibrium behaves in these networks. The first result captures the following intuition: if a planner uses an edge, then they use all edges that are strictly cheaper than said edge.

Lemma 7. *Let $A^i = \{e : x_e^i > 0\}$ be the **active set** of player i , and let μ^i denote the individual marginal cost of i 's active edges at the planner equilibrium. At the planner equilibrium, for every player i and edge e :*

$$c_e(f_e) < \mu^i \implies e \in A^i$$

Proof of Lemma 7. At the planner equilibrium player i plays a selfish best response, so by the Nash characterization $l_e^i(f^{-i}, x^i) \geq \mu^i$ for every edge e , with equality on the active edges A^i .

Suppose $c_e(f_e) < \mu^i$. Were e inactive we would have $x_e^i = 0$ and hence $l_e^i(f^{-i}, x^i) = c_e(f_e) < \mu^i$, contradicting $l_e^i(f^{-i}, x^i) \geq \mu^i$. So $x_e^i > 0$, that is $e \in A^i$. $\square$

The second result we give on the planner equilibrium is about constant cost edges. We show that whenever edges with zero derivative they are used at any admissible strategy profile, then they are the most expensive edge. While this last result holds for all admissible strategy profiles, we will only apply it to the planner equilibrium.

Lemma 8. *At any admissible strategy profile, if some edge b_0 has $c'_{b_0}(f_{b_0}) = 0$ with cost $c_{b_0}(f_{b_0}) = v$, then every used edge a has $c_a(f_a) \leq v$, and any used edge with $c_a(f_a) = v$ also satisfies $c'_a(f_a) = 0$.*

Proof of Lemma 8. As $c'_{b_0}(f_{b_0}) = 0$ we have that $l_{b_0}^k(f^{-k}, x^k) = c_{b_0}(f_{b_0}) = v$ for every player k and $l_{b_0}(f) = v$. Let a be a used edge with a user k , such that $x_a^k > 0$. Suppose either $c_a(f_a) > v$ or $c_a(f_a) = v$ and a satisfies $c'_a(f_a) > 0$. In the first case:

$$l_a^k(f^{-k}, x^k) = c_a(f_a) + \alpha_k x_a^k c'_a(f_a) > v = l_{b_0}^k(f^{-k}, x^k)$$

and in the second case:

$$l_a(f) = c_a(f_a) + f_a c'_a(f_a) > v = l_{b_0}(f).$$

so shifting flow from a to b_0 strictly lowers both player k 's cost and the social cost, contradicting admissibility. Hence $c_a(f_a) \leq v$, with equality only when a is flat. $\square$

Before we prove Proposition 3 we first introduce some additional notation. For any player set $X \subseteq C$ and any edge e , write

$$f_e^X = \sum_{k \in X} \alpha_k x_e^k$$

for the aggregate flow from players in X on edge e . Throughout, μ^i denotes the individual marginal cost of player i 's active edges at the planner equilibrium and A^i its active set.

Proof of Proposition 3. We order the edges according to the cost at the planner equilibrium, so $c_1(f_1) \leq c_2(f_2) \leq \dots \leq c_m(f_m)$. Ties can use an arbitrary ordering.

Assume for contradiction that $c_{\min}(f') > c_{\min}(f)$. To accomplish this, every edge e with $c_e(f_e) = c_{\min}(f)$ must have that $f'_e > f_e$. Then, let b be the first edge where $f'_b < f_b$ and let $B = \{e : c_e(f_e) = c_b(f_b), f_e > 0\}$. B must be nonempty and we must have that $c_b(f_b) > c_{\min}(f)$ for all $b \in B$.

Let $L = \{e : c_e(f_e) < c_b(f_b)\}$ be the set of edges strictly cheaper than the edges in B . Since $b \in B$ is the cheapest edge that loses flow, we must have that every edge $e \in L$ gains flow, so

$$f'_e \geq f_e \quad \forall e \in L. \quad (18)$$

For edges with $c_e(f_e) = c_{\min}(f)$ this inequality is strict. This means that if we sum this inequality over all edges $e \in L$ we get

$$\sum_{e \in L} f'_e > \sum_{e \in L} f_e \quad (19)$$

Let $S_B = \{i : x_e^i > 0 \text{ for some } e \in B\}$ and let $T_B = C \setminus S_B$. By Lemma 7, every planner $i \in T_B$ has that $\mu^i \leq c_b(f_b)$ and hence all of their flow is in L at the planner equilibrium. These planners can only lose total flow in L when deviating towards f' :

$$\sum_{e \in L} x_e^i \leq \sum_{e \in L} x_e^i \quad \forall i \in T_B. \quad (20)$$

Multiplying (20) by α_i and summing over all $i \in T_B$ gives

$$\sum_{e \in L} f_e^{T_B} \leq \sum_{e \in L} f_e^{T_B}. \quad (21)$$

Combining (19) and (21) we can derive that

$$\sum_{e \in L} f_e^{S_B} = \sum_{e \in L} (f'_e - f_e^{T_B}) > \sum_{e \in L} (f_e - f_e^{T_B}) = \sum_{e \in L} f_e^{S_B}.$$

By pigeonhole, there must be an edge $G \in L$ such that

$$f'_G{}^{S_B} > f_G^{S_B}.$$

Summarizing, we have found edges G and b such that the flows satisfy

$$f'_G \geq f_G \quad f'_b < f_b \quad f'_G{}^{S_B} > f_G^{S_B}. \quad (22)$$

Note that $f_b > 0$, so together with $c_G(f_G) < c_b(f_b)$ we can apply Lemma 6 to conclude that

$$l_G(f) \geq l_b(f). \quad (23)$$

Moreover, G has nonzero derivative at f . Every $i \in S_B$ uses an edge in B , so $\mu^i \geq c_b(f_b) > c_G(f_G)$; by Lemma 7 this gives $G \in A^i$, so i uses G and $l_G^i(f^{-i}, x^i) = \mu^i$. Since $c_G(f_G) < \mu^i$, the indifference $c_G(f_G) + \alpha_i x_G^i c'_G(f_G) = \mu^i$ forces $c'_G(f_G) > 0$. By convexity $c'_G(f'_G) \geq c'_G(f_G) > 0$, so G has nonzero derivative at f' as well.

Combining (22) and (23) we also have

$$l_G(f') \geq l_G(f) \geq l_b(f) \geq l_b(f').$$

We will prove that at least one inequality is strict in this chain. Assume for contradiction that this is not the case, this is only possible if we have that

1. $f'_G = f_G$ (as $c'_G(f_G) > 0$),
2. $l_G(f) = l_b(f)$,
3. $c'_b(f_b) = 0$.

The third property together with Lemma 8 implies that $c'_e(f_e) = 0$ for all $e \in B$, hence all edges $e \in B$ are all constant cost edges with the same cost. This means that we can conclude

$$l_G^i(f^{-i}, x^i) = l_b^i(f^{-i}, x^i) \quad \forall i \in S_B.$$

Summing the above inequality over all $i \in S_B$ gives

$$f_G^{S_B} c'_G(f_G) = |S_B| (c_b(f_b) - c_G(f_G)). \quad (24)$$

Next, we write out the expression for the social marginal cost

$$l_G(f) = c_G(f_G) + (f_G^{S_B} + f_G^{T_B}) c'_G(f_G) \quad (25)$$

where we have used that $f_G = f_G^{S_B} + f_G^{T_B}$. Combining (24) and (25) gives us that

$$l_G(f) = c_b(f_b) + (|S_B| - 1)(c_b(f_b) - c_G(f_G)) + f_G^{T_B} c'_G(f_G). \quad (26)$$

Since $|S_B| \geq 1$ and $c_b(f_b) > c_G(f_G)$, the second term on the RHS is weakly positive. As $f'_G = f_G$ and $f'^{S_B}_G > f^{S_B}_G$, we must have that $f'^{T_B}_G < f^{T_B}_G$ so $f'^{T_B}_G > 0$. This means that the last term in (26) is strictly positive. We can conclude that

$$l_G(f) > c_b(f_b) = l_b(f)$$

which contradicts the second property. Concluding this part, we must have that

$$l_G(f') > l_b(f'). \quad (27)$$

We now wish to show that

$$l^i_G(f'^{-i}, x'^i) \leq l^i_b(f'^{-i}, x'^i) \quad \forall i \leq N.$$

If $f'_b > 0$, then by chaining Corollary 4 and Corollary 5 we have our result.

We now assume that $f'_b = 0$. Let planner j be a planner such that $x'^j_G > 0$. For this planner, by (27) and by admissibility

$$l^j_G(f'^{-j}, x'^j) \leq l^j_b(f'^{-j}, x'^j).$$

For planners with $x'^i_G = 0$ we have that

$$\begin{aligned} l^i_G(f'^{-i}, x'^i) &= c_G(f'_G) \leq l^j_G(f'^{-j}, x'^j) \\ &\leq l^j_b(f'^{-j}, x'^j) = c_b(f'_b) = l^i_b(f'^{-i}, x'^i) \end{aligned}$$

where the last step uses that $f'_B = 0$, so we must have that $x'^i_B = 0$ for all planners. So now we have that

$$l^i_G(f'^{-i}, x'^i) \leq l^i_b(f'^{-i}, x'^i) \quad \forall i \leq N. \quad (28)$$

This will finally allow us to set up the last part of the proof. For all planners i and edges e , define the quantity Δl^i_e as

$$\Delta l^i_e = l^i_e(f'^{-i}, x'^i) - l^i_e(f^{-i}, x^i). \quad (29)$$

Combining (28) and (29) gives us

$$l^i_G(f'^{-i}, x'^i) = l^i_G(f^{-i}, x^i) + \Delta l^i_G \leq l^i_b(f^{-i}, x^i) + \Delta l^i_b \quad (30)$$

for all $i \leq N$. Now, we will focus on the planners $i \in S_B$. We will show that

$$l^i_b(f^{-i}, x^i) \leq l^i_G(f^{-i}, x^i) \quad \forall i \in S_B. \quad (31)$$

This is obviously true if a planner has that $x^i_b > 0$. However, if the planner has that $x^i_b = 0$ there must be another edge $d \in B$ such that $x^i_d > 0$. Since each planner $i \in S_B$ has that $x^i_G > 0$ we have that $l^i_d(f^{-i}, x^i) = l^i_G(f^{-i}, x^i)$. That is why we can state that for these planners we have that

$$l^i_b(f^{-i}, x^i) = c_b(f_b) \leq l^i_d(f^{-i}, x^i) = l^i_G(f^{-i}, x^i).$$

We can combine (30) and (31) to write

$$\Delta l^i_G \leq \Delta l^i_b \quad \forall i \in S_B.$$

Sum the above expression over all $i \in S_B$ and write it out to get

$$\begin{aligned} &|S_B|(c_G(f'_G) - c_G(f_G)) + \\ &(f'^{S_B}_G c'_G(f'_G) - f^{S_B}_G c'_G(f_G)) \\ &\leq |S_B|(c_b(f'_b) - c_b(f_b)) + (f'^{S_B}_b c'_b(f'_b) - f^{S_B}_b c'_b(f_b)). \end{aligned}$$

We now show that the left hand side evaluates to ≥ 0 and the right hand side to ≤ 0 . To this end, recall the inequalities as given in (22). Using those inequalities we can state that

$$f'^{S_B}_b \leq f'_b < f_b = f^{S_B}_b. \quad (32)$$

The inequalities (22) and (32) coupled with the relevant monotonicity properties of the cost functions and their derivatives allow us to conclude that

$$\begin{aligned} &|S_B|(c_G(f'_G) - c_G(f_G)) \geq 0 \\ &(f'^{S_B}_G c'_G(f'_G) - f^{S_B}_G c'_G(f_G)) > 0 \\ &|S_B|(c_b(f'_b) - c_b(f_b)) \leq 0 \\ &(f'^{S_B}_b c'_b(f'_b) - f^{S_B}_b c'_b(f_b)) \leq 0 \end{aligned}$$

The left hand side is strictly positive while the right hand side is at most zero, contradicting the summed inequality. Hence $c_{\min}(f') \leq c_{\min}(f)$. $\square$

G Proof of Proposition 4

In the canonical Pigou case, the repeated weighted atomic splittable routing games in which the desiderata can be achieved by a strategy profile are characterized by the planner equilibria in the non-repeated version of their respective games. The following proposition characterizes such equilibria as a solution of a piecewise linear system of equations.

Proposition 7. *Let G be the canonical Pigou network and let $\Gamma = (G, C, \alpha)$ be a weighted atomic splittable routing game. Then, there exists a unique planner equilibrium $\bar{x}$ that is deterministic and in which, if x_i is the portion of planner i 's share of the traffic that is routed through the bottom path, then*

$$x_i = \begin{cases} 1 & \text{if } \frac{1 - \sum_{j \geq 1} \alpha_j x_j}{\alpha_i} > 1 \\ \frac{1 - \sum_{j \geq 1} \alpha_j x_j}{\alpha_i} & \text{otherwise.} \end{cases} \quad (33)$$

The solution of this system can be found in $O(n \log n)$ time, where n is the number of planners. If the partition is sorted by size beforehand (e.g., $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$), the solution can be found in $O(n)$ time.

The proof of Proposition 7 is given below. Since canonical Pigou has a unique planner equilibrium, we can characterize the games in which there is a strategy profile that satisfies the desiderata as follows.

Proposition 8. *Let G be the canonical Pigou network, C be a set of planners, and α be a unit partition. Let F be the total flow through the bottom path in the planner equilibrium of $\Gamma = (G, C, \alpha)$. If $F > \frac{3}{4}$, there exists a strategy profile $\bar{\sigma}$ in the repeated version of Γ that satisfies individual*

rationality, resilience to competition, optimality and admissibility. If $F < \frac{3}{4}$, there exists no strategy profile for the repeated version of Γ that satisfies individual rationality and admissibility. If $F = \frac{3}{4}$, if no planner sends all their flow to the bottom path in the planner equilibrium, there exists a strategy profile $\vec{\sigma}$ in the repeated version of Γ that satisfies individual rationality, resilience to competition, optimality and admissibility. Otherwise, there is no strategy profile $\vec{\sigma}$ in the repeated version of Γ that satisfies individual rationality and admissibility.

The proof of Proposition 8 combines Propositions 2 and 3 with the following characterization of admissibility in the Pigou network.

Lemma 9. *Let $\vec{\sigma}$ be a strategy profile in the canonical Pigou network with bottom-path flow $F > \frac{1}{2}$, and let x_j be the fraction of planner j 's flow on the bottom path. Then $\vec{\sigma}$ is admissible if and only if $\alpha_j x_j \leq 1 - F$ for every planner j .*

Proof. Planner j 's cost as a function of its bottom fraction is the convex parabola $c^j(x_j) = \alpha_j x_j^2 + (F_{-j} - 1)x_j + 1$ with $F_{-j} = F - \alpha_j x_j$ and derivative $F + \alpha_j x_j - 1$, while the social cost $F^2 + (1 - F)$ is strictly increasing in F on $(\frac{1}{2}, 1]$. Increasing x_j strictly raises the social cost, so no upward move is both-improving. If $\alpha_j x_j > 1 - F$, the derivative of c^j is strictly positive, so a small decrease strictly lowers both j 's cost and the social cost, and $\vec{\sigma}$ is inadmissible. If $\alpha_j x_j \leq 1 - F$, the derivative is nonpositive at x_j and strictly negative below it, so every decrease strictly increases j 's cost, and no alternative dominates. $\square$

By Proposition 3 the planner equilibrium maximizes F (equivalently $c_{\min}$) over admissible profiles. If $F^* > \frac{3}{4}$ there, the sufficiency conditions of Proposition 2 hold at the planner equilibrium. If $F^* < \frac{3}{4}$, no admissible profile attains $c_{\min} \geq \frac{3}{4}$ and the necessity condition fails. At the boundary $F^* = \frac{3}{4}$, Lemma 9 gives $\alpha_j x_j \leq \frac{1}{4}$ for all j in every admissible profile with bottom flow $\frac{3}{4}$; since the planner equilibrium has $\alpha_j x_j = \min\{\alpha_j, \frac{1}{4}\}$ and these caps sum to exactly $\frac{3}{4}$, all caps bind in every such profile. If some planner i has $\alpha_i \leq \frac{1}{4}$, then $x_i = 1$ and $c^i = \frac{3}{4} = c_{\text{opt}}$ in every candidate certificate, so the necessity condition fails for i . If instead $\alpha_j > \frac{1}{4}$ for all j , the planner equilibrium satisfies the equal-case sufficiency conditions, since $c^i = 1 - \frac{x_i}{4} > \frac{3}{4}$ for every planner.

Proposition 9. *Let G be the classical Pigou network and let $\Gamma = (G, C, \alpha)$ be a RWASRG with at most two planners. Then, there are no strategy profiles that satisfy individual rationality and admissibility.*

Proof. This follows immediately from the fact that, in the non-repeated version of Γ , the flow F that goes through the bottom path in the planner equilibrium satisfies $F \leq \frac{n}{n+1}$, where n is the number of planners. To see this, note that Proposition 7 implies that, in this planner equilibrium, $x_i \leq \frac{1-F}{\alpha_i}$ for all i . Therefore, $\sum_{i=1}^n \alpha_i x_i \leq n(1-F)$, which is equivalent to $F \leq n(1-F)$ and therefore to $F \leq \frac{n}{n+1}$ as desired. $\square$

Proposition 10. *Let G be the classical Pigou network and let $\Gamma = (G, C, \alpha)$ be a RWASRG. If a planner controls more than half of the flow, there are no strategy profiles that satisfy individual rationality and admissibility.*

Proof. Suppose that $\alpha_1 > \frac{1}{2}$ and that there exists a strategy profile $\vec{\sigma}$ that satisfies individual rationality and admissibility. Then, by Proposition 7 and 8, there exists a planner equilibrium in the non-repeated version of Γ that satisfies Equation 33 and in which $F := \sum_{i \geq 1} \alpha_i x_i \geq \frac{3}{4}$. This means that $x_1 = \frac{1-F}{\alpha_1} < \frac{1}{2}$, and therefore $F \leq \frac{\alpha_1}{2} + (1 - \alpha_1) < \frac{3}{4}$. This contradicts the assumption that $F \geq \frac{3}{4}$. $\square$

Proposition 11. *Let G be Pigou's network and let $\Gamma = (G, C, \alpha)$ be a RWASRG. If no planner controls at least $\frac{3}{8}$ of the total flow, there exists a strategy profile $\vec{\sigma}$ that satisfies individual rationality, resilience to competition, optimality, and admissibility.*

Proof. Suppose that there exists no such strategy profile. We show next that at least one planner controls more than $\frac{3}{8}$ of the total flow. By Proposition 8, there is no such strategy profile if and only if the flow F through the bottom path in a planner equilibrium of the non-repeated version of Γ satisfies $F < \frac{3}{4}$ or $F = \frac{3}{4}$ and at least one planner sends all its flow through the bottom path. Let $S \subseteq C$ be the subset of planners that split their flow in the planner equilibrium of the non-repeated version of Γ (i.e., S consists of planners that do not send all their flow through the bottom path). Then, we have that $\alpha_i x_i = 1 - F$ for all $i \in S$ and $\alpha_i x_i = \alpha_i$ for all $i \notin S$. Adding all these equations together, we get that $F = |S| \cdot (1 - F) + \sum_{i \notin S} \alpha_i$, which is equivalent to

$$F = \frac{|S| + \sum_{i \notin S} \alpha_i}{|S| + 1}.$$

Since we are assuming that $F \leq \frac{3}{4}$, $|S|$ can be at most 3. Thus, we have the following cases.

- $|S| = 0$ implies that all the flow goes through the bottom path, contradicting our initial assumption.
- $|S| = 1$ implies that $\sum_{i \notin S} \alpha_i \leq \frac{1}{2}$, and therefore that the only planner in S controls at least half of the flow.
- $|S| = 2$ implies that $\sum_{i \notin S} \alpha_i \leq \frac{1}{4}$, and therefore that at least one of the planners in S controls $\frac{3}{8}$ of the total flow.
- $|S| = 3$ implies that $\sum_{i \notin S} \alpha_i = 0$, and therefore, by Proposition 8, that there exists $\vec{\sigma}$ that satisfies the desiderata. This contradicts the original assumption. $\square$

It remains to show that the values given in Propositions 9, 10 and 11 are tight. To see this, note that

- If $\alpha = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, each planner sends $\frac{3}{4}$ of the flow in the planner equilibrium of Γ . Thus, by Proposition 8, there exists a strategy profile that satisfies the desiderata.

- If $\alpha = (\frac{1-\epsilon}{2}, \frac{1+\epsilon}{6}, \frac{1+\epsilon}{6}, \frac{1+\epsilon}{6})$, there also exists a strategy profile that satisfies the desiderata. To see this, suppose that the total flow F that goes through the bottom path in the planner equilibrium of Γ satisfies $F \leq \frac{3}{4}$. If this was the case, by Proposition 7, planners 2, 3 and 4 would send all their flow through the bottom path, and 1 would send half of its flow as well. This would imply that the total flow through the bottom path is $\frac{3+\epsilon}{4}$, contradicting our original assumption.
- If $\alpha = (\frac{1}{4}, \frac{3}{8}, \frac{3}{8})$, planner 1 sends all its flow through the bottom path in the planner equilibrium of Γ , while planners 2 and 3 send $\frac{2}{3}$ of their flow. This means that the total flow through the bottom path is exactly $\frac{3}{4}$. By Proposition 8, there is no strategy profile that satisfies individual rationality and admissibility.

This proves Proposition 4.

Proof of Proposition 7

Suppose that $\vec{x}$ is a planner equilibrium of Γ in which each planner i sends, in expectation, a x_i fraction of their assigned agents to the bottom path. Then, the total cost of planner i is given by

$$c_i(\vec{\sigma}) = (1 - x_i) + x_i \left(\alpha_i x_i + \sum_{j \neq i} \alpha_j x_j \right).$$

If we fix x_j for all $j \neq i$, we get that the expression is a polynomial of degree 2 on x_i , which can be written as

$$c_i(\vec{\sigma}) = \alpha_i x_i^2 + (F_i - 1)x_i + 1,$$

where $F_i = \sum_{j \neq i} \alpha_j x_j$.

Since $\vec{\sigma}$ is a planner equilibrium, x_i must be the value that minimizes this expression in $[0, 1]$, which is either the vertex v_i of the parabola, or 1 if $v_i > 1$ (note that this implies that the resulting strategy profile is deterministic). The vertex v_i can be expressed as

$$v_i = \frac{1 - F_i}{2\alpha_i} = \frac{1 - F + \alpha_i x_i}{2\alpha_i}.$$

Setting this expression equal to x_i , we get that $x_i = \frac{1-F}{\alpha_i}$. Thus, all planner equilibria of Γ satisfy the following.

$$x_i = \begin{cases} 1 & \text{if } \frac{1-F}{\alpha_i} > 1 \\ \frac{1-F}{\alpha_i} & \text{otherwise.} \end{cases},$$

To check that the solution of the resulting system exists and is unique, note that solutions must satisfy $\alpha_i x_i = \min\{\alpha_i, 1 - F\}$. Consider the function $g : [0, 1] \rightarrow [0, 1]$ defined by

$$g(F) := \sum_{i=1}^N \min\{\alpha_i, 1 - F\}.$$

F is the resulting flow through the bottom path of a solution if and only if $g(F) = F$. However, $g(F)$ is (weakly) decreasing w.r.t. F . Therefore, there is exactly one solution.

Solving Equation 33 Efficiently

The key to solving Equation 33 is the following observation.

Lemma 10. *Suppose that $\alpha_1 \leq \alpha_2 \leq \dots \alpha_n$. Then the unique solution of Equation 33 satisfies that there exists a $k \leq n$ such that*

$$x_i = \begin{cases} 1 & \text{if } i \leq k \\ \frac{1-F}{\alpha_i} & \text{otherwise.} \end{cases}$$

Intuitively, Lemma 10 says that the planners that send all their traffic to the bottom path in equilibrium are precisely those that have the least portion of the traffic, and the proof follows immediately from the fact that $x_i = 1$ if and only if $\alpha_i < 1 - F$. Once k is fixed, we can do the following to check if we get a correct solution.

From Lemma 10, we have that

$$\alpha_i x_i = \begin{cases} \alpha_i & \text{if } i \leq k \\ 1 - F & \text{otherwise.} \end{cases}$$

Adding all these equations together, we get that $F = \sum_{i=1}^k \alpha_i + (n - k) \cdot F$. Setting $S_k = \sum_{i=1}^k \alpha_i$, we get that

$$F = \frac{S_k + (n - k)}{n - k + 1}.$$

Since we know that there is a unique solution, we can run the following linear algorithm to compute the flow in the planner equilibrium. We iterate for each $k = 0, 1, \dots, n - 1$ until finding k such that $\alpha_{k+1} \geq 1 - \frac{S_k + (n - k)}{n - k + 1}$. After such k is found, we return $\frac{S_k + (n - k)}{n - k + 1}$.

The value of each x_i can be computed directly from F using Equation 33 in constant time. Note that computing S_k from scratch in every iteration would result in a quadratic implementation. However, for each k we can simply update its value in constant time.

H Robustness Against Random Defections

Up to this point, most of our results rely on the ability of planners to discipline deviations by switching to a fully selfish policy that drives the system toward a planner equilibrium whenever a planner or a user defects. While analytically convenient, such punishments are unlikely to be efficient in practice. In large-scale systems, some deviations are essentially unavoidable: even when following the prescribed policy is optimal in the long run, a small fraction of users will occasionally ignore recommendations due to unawareness, urgency, mistakes, or other factors. To capture this phenomenon in a discrete model, we can study a *perturbed* version of the game in which, in each round, every user independently deviates from the proposed route with probability $\epsilon > 0$, in the spirit of trembling-hand perturbations and trembling-hand perfect equilibrium.

Recall that the strategy profiles constructed in Sections C and D have a simple structure: planners propose a socially optimal flow, and upon a deviation they switch to a planner-equilibrium “punishment phase” for a fixed number of rounds. In the distorted model, this logic breaks down.

Even if the profile satisfies desiderata (1), (2), and (4), random trembles will trigger punishments repeatedly, making it likely that play spends (asymptotically) almost all rounds in the planner-equilibrium phase. In particular, the construction no longer guarantees convergence to an optimal flow and thus fails desideratum (3). In this section, we show how to modify the constructions of Sections C and D to retain (1), (2), and (4) while making the dynamics more robust to persistent, low-rate random defections. Our refinement relies on an additional transparency assumption: there is no private information, so planners observe which users deviate in each round and observe one another's recommendations.⁵ For clarity we present the refinement in a discrete model; extending it to continuous flows is conceptually straightforward, but requires additional technical machinery.

The construction is parameterized by constants c , k , and N , and proceeds as follows. In each stage, every planner recommends a karma-based implementation of a socially optimal flow. If a *planner* deviates, then all planners enter a punishment phase and play a planner equilibrium for the next N rounds. If instead a *user* i under planner p deviates, then p permanently deducts c units of karma from i and issues i a *ticket*. This ticket remains active until i 's karma, relative to the average karma of the remaining users, returns to its pre-ticket level. Finally, while any user accumulates k or more tickets, then the planners play the planner equilibrium.

The parameters c , k , and N serve distinct roles. The constant c is chosen so that receiving a ticket forces the user to be routed on the more costly edges for several consecutive future rounds, making a ticket strictly undesirable from the user's perspective. The horizon N is chosen as in Sections C and D to deter planners from defecting. The threshold k is then set so that, under the distorted model where each user trembles with probability ϵ , it is exceedingly unlikely that any user reaches k tickets due to random mistakes alone.

Arguments analogous to those in Sections C and D show that the resulting strategy profile satisfies desiderata (2) and (4).⁶ To verify desideratum (1), observe that a user has no incentive to deviate under this scheme. By construction, receiving a ticket is strictly undesirable: if user i returns to following recommendations, the karma deduction ensures that i is routed on the longer path for several consecutive rounds; if instead i continues to defect, then after accumulating k tickets the system triggers a punishment phase in which all planners play a planner equilibrium until i recovers part of its karma.

Finally, this modification is substantially more robust to random defections. In the original construction, a single spurious deviation could repeatedly trigger punishment and keep the system near the planner equilibrium. Here, isolated trembles are absorbed locally through tickets: global punishment is activated only if a user deviates frequently enough to accumulate many tickets in a short period, an event that is exceedingly unlikely under the ϵ -distortion.

⁵One can interpret this as planners pooling information to improve outcomes for all parties.

⁶In discrete settings with a large number of users, these properties hold up to a negligible approximation error.

For a more formal comparison, we can bound the probability that a *single* user triggers a punishment phase under the two constructions. Under the strategy profiles of Sections C and D, any deviation immediately induces a punishment phase, so the triggering probability in a given round is simply ϵ . In the ticket-based construction, punishment occurs only if the user accumulates k tickets. Let T_c be an upper bound on the number of rounds required for a ticket to expire after a deduction of c karma (i.e., the time needed for the user's karma to recover relative to the rest). Then the probability that the user reaches k tickets at a given round is at most

$$\epsilon(1 - (1 - \epsilon)^{T_c})(1 - (1 - \epsilon)^{2T_c}) \dots (1 - (1 - \epsilon)^{(k-1)T_c}).$$

To justify this bound, note that if the user has k active tickets at some round t , then it must defect at t (probability ϵ), and it must also have defected at least once in each of the preceding windows of lengths $T_c, 2T_c, \dots, (k-1)T_c$: the first ticket can persist for at most T_c rounds, but each additional ticket received before the previous ones expire can extend the time to clear all tickets by at most another T_c rounds.

Using Bernoulli's inequality $(1 - \epsilon)^n \geq 1 - \epsilon n$, we obtain

$$1 - (1 - \epsilon)^{mT_c} \leq \epsilon mT_c,$$

and therefore

$$\epsilon \prod_{m=1}^{k-1} (1 - (1 - \epsilon)^{mT_c}) \leq \epsilon \prod_{m=1}^{k-1} (\epsilon mT_c) = \epsilon^k T_c^{k-1} (k-1)!.$$

The expression above can be rewritten as

$$\epsilon^k T_c^{k-1} (k-1)! = \epsilon \cdot \left(\epsilon^{\sqrt{k}} T_c^{\sqrt{k}} (\sqrt{k})! \right) \cdot \prod_{m=\sqrt{k}+1}^{k-1} (\epsilon T_c m).$$

If we choose k so that $\epsilon k T_c < 1$, then for all $m \leq k-1$ we have $\epsilon T_c m \leq \epsilon T_c k < 1$, and hence

$$\prod_{m=\sqrt{k}+1}^{k-1} (\epsilon T_c m) \leq 1.$$

Therefore, we get that

$$\epsilon^k T_c^{k-1} (k-1)! \leq \epsilon \cdot (\epsilon T_c \sqrt{k})^{\sqrt{k}}.$$

Since $\sqrt{\epsilon T_c k} < 1$, we can further upper bound the above expression by

$$\epsilon \cdot (\epsilon T_c \sqrt{k})^{\sqrt{k}} < \epsilon \cdot (\sqrt{\epsilon T_c})^{\sqrt{k}},$$

and thus the probability that a given user triggers a planner-equilibrium phase is many orders of magnitude smaller than ϵ whenever $\epsilon T_c \ll 1$.

This condition is typically easy to satisfy. For example, in the canonical Pigou network one can choose c so that a deviating user is forced to traverse the top edge for three consecutive rounds, i.e., $T_c = 3$. If $\epsilon = 0.05$, taking $k = 6$ yields the bound

$$\epsilon \cdot (\sqrt{\epsilon T_c})^{\sqrt{k}} \approx 0.0049,$$

an order-of-magnitude reduction relative to the original triggering probability ϵ . If instead $\epsilon = 0.01$, choosing $k = 30$ gives

$$\epsilon \cdot (\sqrt{\epsilon T_c})^{\sqrt{k}} \approx 6.75 \times 10^{-7},$$

which is more than 10^4 times smaller.