Competitive Information Provision Among Internet Routing Nodes

Joshua H. Seaton, Sena Hounsinou, Gedare Bloom, and Philip N. Brown

***Abstract*— In proposed *path-aware* designs for the Internet, end hosts can select which path their packets use. What criteria should the end hosts use to select paths? Recent work has proposed path-aware access control frameworks in which routing nodes publicly report their knowledge of the security postures of other nodes; end hosts can then base their routing choices on these reports. However, nothing is known regarding the nodes' incentives to report their knowledge *truthfully*. In this paper, we consider the case in which each network node is strategic, and seeks to craft its public reports to manipulate traffic patterns in its own favor. In the context of a simple selfish routing problem with two strategic nodes, we show that for a wide swath of the parameter space, each node has a dominant reporting strategy, meaning that its individually optimal strategy does not depend on the strategy of the other node. These dominant strategies are generally *not* truthful. At the resulting dominant-strategy Nash equilibrium, we show that the expected social cost is (often considerably) higher than that achieved when both nodes are completely truthful. Nonetheless, we prove that these equilibrium reporting strategies are never *perverse,* meaning that their resulting social cost is never worse than if traffic were uninformed as to network state.**

I. Introduction

In today's Internet, end users have essentially no control over the paths used by their data; packet routing is governed almost entirely by protocols executed locally at each node [1]. This lack of end host routing control can negatively impact the overall efficiency of the Internet along a variety of dimensions; for instance, leading to conflicts between latency and bandwidth. In addition, it has security and privacy implications; for example, end users cannot certify that their data has been in trusted hands or in permitted jurisdictions all along its route. To address these and other shortfalls, much recent work has proposed new *path-aware* Internet architectures which enable path selection and path authentication by end hosts [2]–[5].

However, in a future path-aware Internet it remains an open question as to precisely *how* end hosts should obtain information about which paths satisfy their security and quality-of-service requirements, and how they should select paths accordingly. As an initial step to answer this question, in [6] we propose a framework which allows routing nodes to provide reports on the security postures or trustworthiness of other nodes, and allows end hosts to condition access to data resources on the basis of these aggregated reports.

This work was supported in part by the National Security Agency under Grant Number H98230-21-1-0155, by Colorado SB18-086, and by the National Science Foundation under Grant ECCS-2013779.

The authors are with the Department of Computer Science, University of Colorado Colorado Springs, {jseaton,shoueto,gbloom,pbrown2}@uccs.edu

In the real world, complex economic relationships exist between routing nodes, so that the individual objectives of each node may not be aligned with that of the system planner. Thus, if our framework were to be implemented, it would be desirable to ensure that nodes are properly incentivized to report their full knowledge truthfully. Simultaneously, each node would need to solve an *information design* problem [7] in order to determine what information to report or withhold to optimize its own objectives.

In the present paper, we initiate a study on the above information design problem by posing a simplified model and determining how each node should report on others in order to obtain a competitive advantage. To do this, we apply analysis techniques from *game theory* in the context of a *selfish routing* model with competitive *Bayesian persuasion*. To situate our paper in the broader research context, we briefly introduce each of these concepts before summarizing our contributions.

Selfish routing models how network traffic might be expected to select routes in the absence of centralized coordination [8]. Traditionally, the main application for selfish routing models has been transportation networks [9]–[13]. However, a path-aware Internet would explicitly allow end hosts to select routes, which would immediately make the Internet a selfish routing problem. Some recent work has begun to consider the implications of this [14]–[16], showing that it has nontrivial consequences which likely require novel incentive mechanisms.

Bayesian persuasion describes a scenario in which one agent is informed about some state of the world, and strategically reveals (or does not reveal) information to a second agent in order to influence its behavior [17], [18]. One of the main conclusions of the Bayesian persuasion literature is that the informed agent can often gain an advantage by employing an optimal revelation policy, and this policy is often nontrivial. Among many other applications, this concept has been used to study the problem of influencing driver behavior in transportation networks by coupling it with selfish routing models [19]–[22].

In the present paper, we pose a simple selfish routing model with competitive Bayesian persuasion. In our model, traffic must select between two routes on the basis of latency and the perceived security of each route. Each of the two routes is associated with a strategic agent we call a *signaler*; each signaler desires to maximize the expected traffic choosing its route. Each day, a security incident may independently afflict one or both routes rendering the route either secure or insecure; any traffic using an insecure route experiences this insecurity as an additional cost. The traffic

knows only the probability of a security incident on each path. Both signalers are fully informed as to the realized security state, and each signaler gives a public report about its competitor's security state.

This setup gives rise to a hierarchical game: the signalers competitively design their reporting policies to maximize their expected share of the traffic; given the signalers' policies, the traffic plays a nonatomic Bayesian game to determine the allocation of traffic to the two routes. While competitive Bayesian persuasion is not new, to the best of our knowledge, this is the first work in which the competitive signalers are associated with network links; that is, the signalers have conflicting goals as to the *network flows* themselves rather than to subscriber counts [23].

Our analysis presents several results regarding both the equilibrium behavior of the signalers and traffic as well as the quality of the resulting equilibria. First, we show that when the *a priori* expected cost of a security incident is comparable with the network's worst-case latency costs, the game played between the signalers is trivial: both signalers have a dominant-strategy revelation policy. That is, each signaler can compute its individually-optimal strategy *without knowing* the strategy of the other signaler.

Second, we show that the expected *social cost* (i.e., the sum of latency and security costs experienced by traffic) improves with the truthfulness of the signalers' policies; this is in contrast to other related work in which more information can actually harm traffic [22]. Thus, the social cost at the dominant-strategy signalling equilibrium is no worse than the uninformed social cost, meaning that the *perversity index* of competitive signaling is unity [24]. However, we show numerically that the equilibrium social cost can be substantially worse than the optimal (i.e., fully-informed) social cost. Thus, the *price of anarchy* of competitive signalling may be large [25].

Finally, we present the results of some numerical experiments which demonstrate our results and we include a link to Python code for these experiments.

II. Model and Metrics

The network model studied in this paper is depicted in Figure 1. The model consists of four nodes: two nodes represent the source and the target of a non-atomic unit of traffic, and the remaining two nodes represent possibly insecure nodes through which network traffic is routed. The path $(s, 0, t)$ will be referred to as path 0 or node 0, dependent upon the context. Path $(s, 1, t)$ will likewise be referred to as node 1 or path 1. Each of the these two nodes may be in one of two security states from state space $\Theta = \{\mathrm{S}, \mathrm{I}\}$. Path i is in the secure state S with probability q_i and is otherwise in the insecure state I.

For each of the two nodes $i \in \{0, 1\}$, there is a strategic agent—which we call the *signaler*—that knows the security state $\theta_j, i \neq j$ of the opposing signaler's node, and they seek to maximize the flow x_i of path i. A report from signaler i about the security state of path j is $r_j \in \{\mathrm{S}, \mathrm{I}\}$. A report r is an ordered pair (r_0, r_1) such that $r \in R = \{\mathrm{SS}, \mathrm{SI}, \mathrm{IS}, \mathrm{II}\}$.

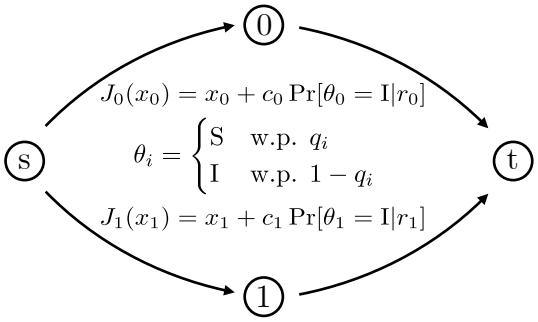

Fig. 1: Diagram of the network studied in this paper. A single unit of (non-atomic) selfish traffic routes from node s to node t, selecting its route based on the latency on each road (modeled by x_i) and the reports given by nodes 0 and 1. Note that Node 0 delivers a report on its competitor 1, and node 1 delivers a report on its competitor 0.

Each signaler i chooses a reporting policy that seeks to maximize traffic on path i. These reports are issued pursuant to the signaler's policy such that:

$$r_j = \begin{cases} \mathrm{S} & \text{w.p. } p_j \text{ if } \theta_j = \mathrm{S} \\ \mathrm{I} & \text{w.p. } 1 - p_j \text{ if } \theta_j = \mathrm{I} \text{ and w.p. } 1 \text{ if } \theta_j = \mathrm{I}. \end{cases} \tag{1}$$

We assume that signaler i always reports truthfully when the security state θ_j is I in order to simplify the network model.

A. Routing Game

We model traffic as a continuum of infinitely many infinitesimally small *packets*. These packets know the probability q_i with which a node i will be insecure. They also know the security cost c_i. Thus, the unconditioned cost of path i is

$$J_i(x_i) = \begin{cases} x_i & \text{w.p. } q_i \\ x_i + c_i & \text{w.p. } 1 - q_i. \end{cases} \tag{2}$$

However, the report acts to augment the packets' information about the state of the nodes. Since they know both q, the probabilities that the nodes are insecure, and p, the reporting frequencies of the signalers, packets seek to optimize cost by choosing the path that minimizes the cost function:

$$J_i^r(x_i) = x_i + c_i \Pr\left[\theta_i = \mathrm{I}|r\right]. \tag{3}$$

We call an interim traffic state $x^r(p)$ an *interim Wardrop equilibrium* for joint report r if it satisfies:

$$x_i^r > 0 \implies J_i^r(x_i) \leq J_j^r(x_j), i \neq j. \tag{4}$$

The expected flow across the set of all reports is:

$$x_i(p) = \sum_{r \in R} \Pr[r] x_i^r(p) \tag{5}$$

B. Game Between the Signalers

Signaler i seeks to select its reporting frequency p_j to solve the optimization problem:

$$\max_{p_j \in [0,1]} x_i(p). \tag{6}$$

A pair of reporting frequencies (p_0^*, p_1^*) is a Nash Equilibrium (NE) when:

$$x_0(p_0^*, p_1^*) \geq x_0(p_0^*, p_1) \tag{7}$$

and

$$x_0(p_0^*, p_1^*) \geq x_0(p_0^*, p_1) \tag{8}$$

for all p_0 and p_1.

We wish to clarify that each signaler i emits a report $r_j, i \neq j$, but they do not directly select which report to emit. The report from signaler i is based upon the probability $p_j, j \neq i$, the policy chosen by signaler i. Therefore, traffic must make an inference about the security state on path j. Therefore, each signaler $i \in \{0, 1\}$ desires to set a policy which maximizes the flow of traffic across path i. Thus, signaler i selects a p_j given an opposing policy p_i.

C. System-Level Cost Metrics

The system-level costs of the systems are measured in a variety of ways. These costs are measured both for a given report as well as an expected value across the set of reports.

For a given report, the flow of traffic will induce an interim Wardrop equilibrium. The interim total latency of the system for a given report r is:

$$L(x^r) = \sum_{i \in \{0,1\}} (x_i^r)^2 \tag{9}$$

The expected total latency of the network is:

$$\mathcal{L}(p) = \sum_{r \in R} L\left(x^r(p)\right) P\left[r\right] \tag{10}$$

For a given report r the interim total security cost for the system is:

$$S(x^r) = \sum_{i \in \{0,1\}} x_i^r c_i \Pr\left[\theta_i = \mathrm{I} | r\right] \tag{11}$$

The expected total security costs of a game is:

$$\mathcal{S}(p) = \sum_{r \in R} S\left(x^r(p)\right) P\left[\theta_i = \mathrm{I} | r\right] \tag{12}$$

The social cost for a given report r is:

$$C(x^r) = L(x^r) + S(x^r) \tag{13}$$

The expected social cost is:

$$\mathcal{C}(p) = \mathcal{L}(p) + \mathcal{S}(p) \tag{14}$$

III. Analytical Contributions

This paper's analytical results focus mainly on the case in which $c_i(1 - q_i) \leq 1$ for both links $i \in \{0, 1\}$. That is, the *a priori* expected cost of a security incident on a network link is no greater than the largest latency possible on the network. In other words, this case models a scenario with security incidents which are either not too frequent or not too costly. The high-security-cost parameter regime does not generally admit pure Nash equilibria between the signalers, and we postpone its study for future work.

A. The Dominant-Strategy Equilibrium

To concisely express these results, we frequently write $e_i := c_i(1 - q_i)$ to denote the *a priori* expected security cost of edge i. When $c_i(1 - q_i) \leq 1$ for both $i \in \{0, 1\}$, a dominant-strategy Nash equilibrium emerges between the signalers, as we show in Theorem 3.1. For many parameter values, this equilibrium fixes the interim expected security cost $c_i \Pr\left[\theta_i = \mathrm{I} \mid r_i = \mathrm{I}\right]$ at exactly 1 for each link.

Theorem 3.1: Let $c_i(1 - q_i) \leq 1$ for both $i, j \in \{0, 1\}$. Then agent i has a dominant strategy to select

$$p_j^* = \begin{cases} 1 & \text{if } c_j < 1, \\ \frac{1 - c_j(1 - q_j)}{q_j} & \text{if } c_j \in \left[1, \frac{1 - q_j/2}{1 - q_j}\right], \\ 0.5 & \text{otherwise,} \end{cases} \tag{15}$$

where $j \neq i$. At this equilibrium, it holds that the conditional (interim) expected security cost of link $i \in \{0, 1\}$ is given by

$$c_i \Pr\left[\theta_i = \mathrm{I} \mid r_i = \mathrm{I}\right] = \begin{cases} c_i & \text{if } p_i^* = 1, \\ 1 & \text{if } p_i^* = \frac{1 - c_i(1 - q_i)}{q_i}, \\ c_i(1 - q_i) & \text{if } p_i^* = 0.5. \end{cases} \tag{16}$$

Before presenting the proof, we note that this result exhibits several interesting features. First, (15) indicates that the dominant strategy of each agent depends only on parameters associated with its opponent's link, and that each agent's truthfulness is nondecreasing in its opponent's link cost c_i. Intuitively, it is easier to lie about very risky links than about very safe links.

Second, (16) shows that strategic signaling can severely curtail the quality of information that reaches traffic. Specifically, each signaler attempts to signal in such a way as to fix the interim expected security cost (conditioned on a signal of "insecure") of the opponent's link at exactly 1.

The proof proceeds in the following steps:

1) In Lemma 3.2, we characterize the interim Wardrop equilibria.
2) In Lemma 3.3, we show that when $e_i \leq 1$ for $i \in \{0, 1\}$, both signalers have a dominant strategy.
3) Finally, we combine Lemmas 3.2 and 3.3 to complete the proof of Theorem 3.1.

Throughout the proof, we perform all analysis from the perspective of link 0, whose objective is to select p_1 to maximize $x_0(p_0, p_1)$. For notational convenience, we write x^r to mean $x(p; r)$, the interim equilibrium flow given reports r; note that the dependence on p is always implied but often suppressed for brevity. For example, x^{SS} means $x(p; \mathrm{SS})$. In addition, we write $c_i^r := c_i \Pr(\theta_i = \mathrm{I} \mid r)$ to denote these interim expected costs.

Lemma 3.2 begins with a description of the interim equilibrium network flows.

Lemma 3.2: For any set of parameters, the interim equi-

librium flows on link 0 satisfy the following:

$$x_0^{\mathrm{SS}} = 1/2. \tag{17}$$

$$x_0^{\mathrm{SI}} = \begin{cases} 1 & \text{if } c_1^{\mathrm{SI}} \geq 1 \\ \frac{1}{2}\left(1 + c_1^{\mathrm{SI}}\right) & \text{otherwise.} \end{cases} \tag{18}$$

$$x_0^{\mathrm{IS}} = \begin{cases} 0 & \text{if } c_0^{\mathrm{IS}} \geq 1 \\ \frac{1}{2}\left(1 - c_0^{\mathrm{IS}}\right) & \text{otherwise.} \end{cases} \tag{19}$$

$$x_0^{\mathrm{II}} = \begin{cases} 0 & \text{if } c_0^{\mathrm{II}} - c_1^{\mathrm{II}} \geq 1 \\ 1 & \text{if } c_1^{\mathrm{II}} - c_0^{\mathrm{II}} \geq 1 \\ \frac{1}{2}\left(1 + c_1^{\mathrm{II}} - c_0^{\mathrm{II}}\right) & \text{otherwise.} \end{cases} \tag{20}$$

To aid exposition, Figure 2 depicts the inequalities on c_i^r found in (18)-(20), and illustrates the general form of the interim equilibrium in all cases.

Proof: First, consider the case that the joint report is $r = \mathrm{SS}$. Due to the signaling structure, the traffic knows with certainty that both links have $\theta_i = \mathrm{S}$, so $c_i^{\mathrm{SS}} = 0$. Hence, in this case the interim equilibrium is simply equal to (17) to satisfy (4).

Next, if the joint report is $r = \mathrm{SI}$, link 0 is known to be secure, but the interim expected cost of link 1 is nonzero: $c_1^{\mathrm{SI}} = c_1(1 - q_1)/(1 - q_1 p_1)$. Hence, the $r = \mathrm{SI}$ interim equilibrium satisfying (4) is equal to (18).

Similarly, if the joint report is $r = \mathrm{IS}$, link 1 is known to be secure, but the interim expected cost of link 0 is nonzero: $c_0^{\mathrm{IS}} = c_0(1 - q_0)/(1 - q_0 p_0)$. Hence, the $r = \mathrm{IS}$ interim equilibrium satisfying (4) is equal to (19).

Finally, if the joint report is $r = \mathrm{II}$, both links have a nonzero interim expected security cost. Hence, the $r = \mathrm{II}$ interim equilibrium satisfying (4) now has the 3 cases given in (20). ∎

Our next lemma shows that the strategic problem faced by the signalers is very simple and is driven entirely by the inequalities in (18) and (19).

Lemma 3.3: Let $c_i(1 - q_i) \leq 1$ for both $i \in \{0, 1\}$. Then $x_0(p)$ is nonincreasing in p_1 if $c_1^{\mathrm{SI}} \geq 1$ and nondecreasing in p_1 if $c_1^{\mathrm{SI}} < 1$. Similarly, $x_1(p)$ is nonincreasing in p_0 if $c_0^{\mathrm{IS}} \geq 1$ and nondecreasing in p_0 if $c_0^{\mathrm{IS}} < 1$.

To simplify the proof of Lemma 3.3, we introduce the following claim, whose tedious proof is summarized in the appendix:

Claim 3.4: The following are true whenever the corresponding flow x_r satisfies $J_0(x_0^r) = J_1(x_1^r)$:

$$x_0^{\mathrm{SS}} \Pr[r = \mathrm{SS}] = \frac{1}{2} \Pr[r_0 = \mathrm{S}] \Pr[r_1 = \mathrm{S}] \tag{21}$$

$$x_0^{\mathrm{SI}} \Pr[r = \mathrm{SI}] = \frac{1}{2} \Pr[r_0 = \mathrm{S}] \left(\Pr[r_1 = \mathrm{I}] + e_1\right) \tag{22}$$

$$x_0^{\mathrm{IS}} \Pr[r = \mathrm{IS}] = \frac{1}{2} \Pr[r_1 = \mathrm{S}] \left(\Pr[r_0 = \mathrm{I}] - e_0\right) \tag{23}$$

$$x_0^{\mathrm{II}} \Pr[r = \mathrm{II}] = \frac{1}{2} \left(\Pr[r_0 = \mathrm{I}] \Pr[r_1 = \mathrm{I}] - e_0 \Pr[r_1 = \mathrm{I}] + e_1 \Pr[r_0 = \mathrm{I}]\right) \tag{24}$$

The proof of Claim 3.4 follows in a straightforward manner from the definitions of the considered quantities and is omitted for reasons of space.

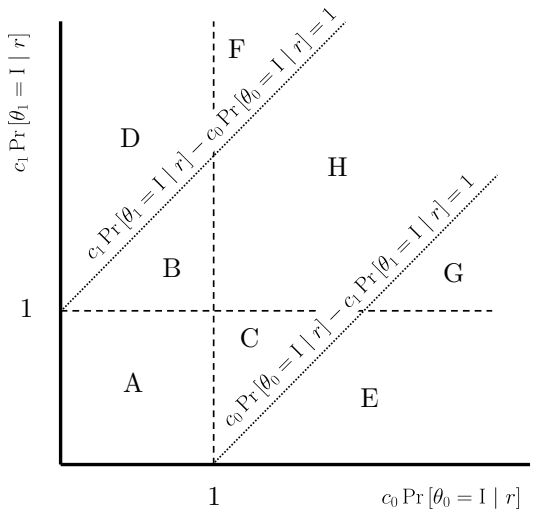

Fig. 2: Depiction of the constraints on $c_i \Pr[\theta_i = \mathrm{I} \mid r]$ which lead to various types of interim equilibria.

Proof of Lemma 3.3: First, consider region A from Figure 2. By Lemma 3.2, in this region all four joint reports yield indifferent interim flows, meaning that the overall expected flow $x_0(p)$ is simply given by the sum of (21)-(24), which can be written

$$\begin{aligned} x_0(p_0, p_1) = \frac{1}{2}\Big[& q_0 p_0 q_1 p_1 \\ & + q_0 p_0 (1 - q_1 p_1 + e_1) + q_1 p_1 (1 - q_0 p_0 - e_0) \\ & + (1 - q_1 p_1)(1 - q_0 p_0) \\ & + (1 - q_0 p_0) e_1 - (1 - q_1 p_1) e_0 \Big]. \end{aligned} \tag{25}$$

By combining terms, it can easily be shown that (25) is constant in p_1, proving Lemma 3.3 for region A.

Now, consider region B in Figure 2. Here, all flows are indifferent except for x^{SI}; in that case, Lemma 3.2 gives that $x_0^{\mathrm{SI}} = 1$, so that

$$\begin{aligned} x_0^{\mathrm{SI}} \Pr[r = \mathrm{SI}] &= \Pr[r = \mathrm{SI}] \\ &= (1 - q_1 p_1) q_0 p_0. \end{aligned} \tag{26}$$

Accordingly, substituting derived expressions from (21), from (23), and (24), we have in region B that

$$\begin{aligned} x_0(p_0, p_1) = & \frac{1}{2} q_0 p_0 q_1 p_1 \\ & + q_0 p_0 (1 - q_1 p_1) + \frac{1}{2}\Big[q_1 p_1 (1 - q_0 p_0 - e_0) \\ & + (1 - q_1 p_1)(1 - q_0 p_0) \\ & + (1 - q_0 p_0) e_1 - (1 - q_1 p_1) e_0 \Big], \end{aligned} \tag{27}$$

which can easily be shown to be nonincreasing in p_1, and strictly decreasing in p_1 whenever $q_0 p_0 > 0$.

Now, consider region C in Figure 2. Here, all flows are indifferent except for x^{IS}; in that case, Lemma 3.2 gives that $x_0^{\mathrm{SI}} = 0$. Thus, substituting derived expressions

from (21), (22), and (24), we have in region C that

$$x_0(p_0,p_1) = \frac{1}{2}[q_0p_0q_1p_1 + q_0p_0(1-q_1p_1+e_1) + (1-q_1p_1)(1-q_0p_0) + (1-q_0p_0)e_1 - (1-q_1p_1)e_0]. \tag{28}$$

It can be verified that (28) is nondecreasing in p_1 whenever

$$e_0/(1-q_0p_0) \geq 1. \tag{29}$$

Here, (29) is equivalent to writing $c_1^{\mathrm{SI}} \geq 1$, which is always satisfied by definition in the C region. Hence, (28) is nondecreasing in p_1 in the C region.

Regions D through G proceed in a similar manner; $x_0(p_0,p_1)$ is nonincreasing in p_1 in D and F, and nondecreasing in p_1 in E and G. We omit the detailed derivation for reasons of space.

Finally, consider region H in Figure 2. Here, Lemma 3.2 gives that flows x^{SS} and x^{II} are indifferent, but $x_0^{\mathrm{SI}} = 1$ and $x_0^{\mathrm{IS}} = 0$. Thus, substituting derived expressions from (21) and (24), we have in region H that

$$x_0(p_0,p_1) = \frac{1}{2}q_0p_0q_1p_1 + q_0p_0(1-q_1p_1) + \frac{1}{2}\Big[(1-q_1p_1)(1-q_0p_0) + (1-q_0p_0)e_1 - (1-q_1p_1)e_0\Big]. \tag{30}$$

It can be shown that (30) is nonincreasing in p_1 if and only if $c_0(1-q_0) \leq 1$.

Taking the above all together, we have shown that whenever $c_0(1-q_0) \leq 1$, $x_0(p)$ is nonincreasing in p_1 if $c_1^{\mathrm{SI}} \geq 1$ (i.e., in regions B, D, F, G, and H), and $x_0(p)$ is nondecreasing in p_1 if $c_1^{\mathrm{SI}} < 1$ (i.e., in regions A, C, and E). Due to the symmetry of the problem, a complementary statement holds for $x_1(p)$ with respect to p_0, completing the proof of Lemma 3.3. ∎

Proof of Theorem 3.1: Consider that by Bayes' rule,

$$c_1 \Pr[\theta_1 = \mathrm{I} \mid \mathrm{SI}] = \frac{c_1(1-q_1)}{1-q_1p_1}. \tag{31}$$

It follows from the above that

$$c_1 \Pr[\theta_1 = \mathrm{I} \mid \mathrm{SI}] \geq 1 \iff p_1 \geq \frac{1-c_1(1-q_1)}{q_1}. \tag{32}$$

Let $p_1^* = (1-c_1(1-q_1))/q_1$. In light of Lemma 3.3, this means that link 0 always maximizes $x_0(p_0,p_1)$ by selecting p_1 *in the direction of* p_1^*, yielding (15) for the case of p_1; the case of p_0 is given by the symmetry of the problem.

Finally, the first case of (16) holds since if $p_i = 1$, the security state of link i is always known with certainty. The middle case of (16) holds due to (32), and the third case of (16) holds because if $p_i = 0.5$, no additional information is revealed at interim so $\Pr[\theta_i = \mathrm{I} \mid r_i = \mathrm{I}] = \Pr[\theta_i = \mathrm{I}]$. ∎

Next, we examine the equilibrium which emerges when both signalers employ their dominant strategies.

Proposition 3.5: When $p_i^* = \frac{1-c_i(1-q_i)}{q_i}$ for both $i \in \{0,1\}$, the dominant-strategy equilibrium satisfies

$$x^{\mathrm{SS}}(p^*) = (1/2, 1/2) \tag{33}$$
$$x^{\mathrm{SI}}(p^*) = (1, 0) \tag{34}$$
$$x^{\mathrm{IS}}(p^*) = (0, 1) \tag{35}$$
$$x^{\mathrm{II}}(p^*) = (1/2, 1/2). \tag{36}$$

Here, (33)-(36) show that in the case when both signalers can "control the message," the equilibrium reduces to a very simple form.

Proof: The proof is immediate from Lemma 3.2 and (16) in Theorem 3.1. ∎

B. Characterizing Social Cost

Unlike some similar models of information provision in selfish routing (such as [13], [22]) our model predicts that more information always benefits the traffic from the standpoint of social cost.

Theorem 3.6: For all possible parameter values for c and q, the expected social cost $\mathcal{C}(p)$ is nonincreasing in both p_0 and p_1. Thus, it holds that strategic behavior by the signalers does not harm social cost relative to the uninformed case:

$$\mathcal{C}(p^*) \leq \mathcal{C}((1/2, 1/2)). \tag{37}$$

Theorem 3.6 shows at least that truthfulness is well-aligned with social cost; thus, future work can focus on shaping the equilibrium strategies among the nodes to incentivize truthfulness. However, as can be observed from the expressions for $\mathcal{C}(p)$ given in the proof of Theorem 3.6 and in the figures presented in Section IV, it is frequently true that (37) holds *with equality*, meaning that despite the lack of perversity, the traffic experiences very little benefit at equilibrium relative to being uninformed.

The proof of this theorem proceeds in a straightforward manner by examining the social cost in each of the regions depicted in Figure 2. For the sake of brevity, however, the proof will be omitted from this manuscript.

IV. Numerical Experiments

In this section, we illustrate our results for total latency and numerically using our Python code.[1] Figure 3a depicts the equilibrium total latency as a function of p_0 and p_1 for the given parameters. Note that the total latency is lower in the uninformed state than in the equilibrium state (marked with a large $+$). Furthermore, for these parameters, the fully-informed total latency is *also* lower than at equilibrium.

Similarly, Figure 3b depicts the equilibrium social cost as a function of p_0 and p_1, illustrating that the informed social cost is minimal, and the uninformed is maximal as in Theorem 3.6. However, it is also clear from the figure that the equilibrium social cost is no better than that of the uninformed case.

[1] Our Python code is available at `https://github.com/descon-uccs/strategic-nodes`

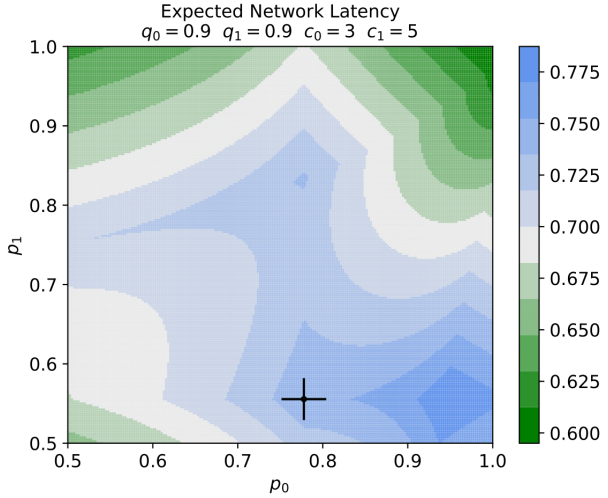

(a) Expected total latency with respect to p_0 and p_1 for given parameter values. The Nash equilibrium p^* is marked with a large black $+$ symbol. Note that the Nash equilibrium total latency is considerably higher than both the uninformed (lower-left) and fully-informed (upper-right) total latency.

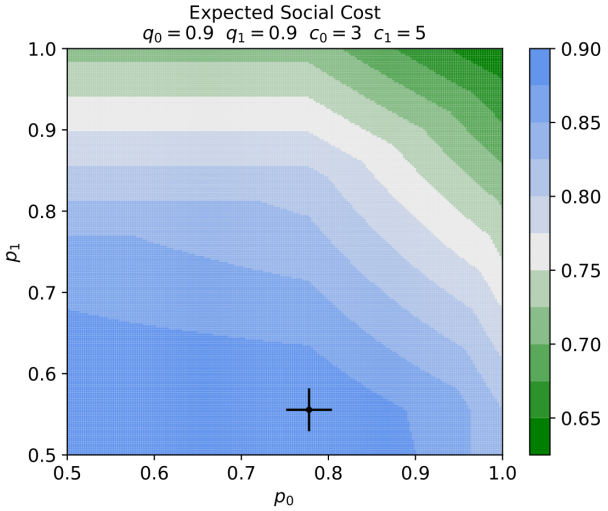

(b) Expected social cost with respect to p_0 and p_1 for given parameter values. The Nash Equilibrium p^* is marked with a large black $+$ symbol. Note that the Nash equilibrium social cost is considerably higher than the fully-informed (upper-right) social cost, but no worse than the uninformed (lower-left) social cost.

Fig. 3: Cost metrics with respect to p_0 and p_1 for given parameter values.

References

[1] L. Gao and J. Rexford, "Stable Internet routing without global coordination," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 681–692, 2001.

[2] T. Anderson, K. Birman, R. Broberg, M. Caesar, D. Comer, C. Cotton, M. J. Freedman, A. Haeberlen, Z. G. Ives, A. Krishnamurthy, W. Lehr, B. T. Loo, D. Mazières, A. Nicolosi, J. M. Smith, I. Stoica, R. van Renesse, M. Walfish, H. Weatherspoon, and C. S. Yoo, "The NEBULA Future Internet Architecture," in *The Future Internet*, pp. 16–26, 2013.

[3] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," *ACM SIGCOMM Computer Communication Review*, vol. 39, pp. 111–122, aug 2009.

[4] B. Raghavan, P. Verkaik, and A. Snoeren, "Secure and Policy-Compliant Source Routing," *IEEE/ACM Transactions on Networking*, vol. 17, pp. 764–777, jun 2009.

[5] M. Legner, T. Klenze, M. Wyss, C. Sprenger, and A. Perrig, "EPIC: Every Packet Is Checked in the Data Plane of a Path-Aware Internet," in *29th USENIX Security Symposium (USENIX Security 20)*, pp. 541—-558, 2020.

[6] J. H. Seaton, S. Hounsinou, T. Wood, S. Xu, P. N. Brown, and G. Bloom, "Poster: Toward Zero-Trust Path-Aware Access Control," in *Proceedings of the 27th ACM on Symposium on Access Control Models and Technologies*, (New York, NY, USA), pp. 267–269, ACM, jun 2022.

[7] D. Bergemann and S. Morris, "Information Design: A Unified Perspective," *Journal of Economic Literature*, vol. 57, pp. 44–95, mar 2019.

[8] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[9] F. Farokhi and K. H. Johansson, "A Study of Truck Platooning Incentives Using a Congestion Game," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 2, pp. 581–595, 2015.

[10] F. Benita, V. Bilò, B. Monnot, G. Piliouras, and C. Vinci, "Data-Driven Models of Selfish Routing: Why Price of Anarchy Does Depend on Network Topology," in *WINE*, pp. 252–265, Springer International Publishing, 2020.

[11] P. N. Brown and J. R. Marden, "Can Taxes Improve Congestion on All Networks?," *IEEE Transactions on Control of Network Systems*, vol. 7, pp. 1643–1653, dec 2020.

[12] B. T. Gould and P. N. Brown, "On Partial Adoption of Vehicle-to-Vehicle Communication: When Should Cars Warn Each Other of Hazards?," *2022 American Control Conference*, pp. 627–632, jul 2022.

[13] B. T. Gould and P. N. Brown, "Information design for vehicle-to-vehicle communication," *Transportation Research Part C: Emerging Technologies*, vol. 150, p. 104084, 2023.

[14] S. Scherrer, M. Legner, A. Perrig, and S. Schmid, "Incentivizing stable path selection in future Internet architectures," *Performance Evaluation*, vol. 144, p. 102137, dec 2020.

[15] S. Scherrer, A. Perrig, and S. Schmid, "The Value of Information in Selfish Routing," in *SIROCCO 2020: Structural Information and Communication Complexity*, Information Security and Cryptography, (Cham), pp. 366–384, Springer International Publishing, 2020.

[16] S. Scherrer, M. Legner, A. Perrig, and S. Schmid, "An axiomatic perspective on the performance effects of end-host path selection," *Performance Evaluation*, vol. 151, p. 102233, nov 2021.

[17] E. Kamenica and M. Gentzkow, "Bayesian Persuasion," *American Economic Review*, vol. 101, pp. 2590–2615, oct 2011.

[18] R. Alonso and O. Câmara, "Bayesian persuasion with heterogeneous priors," *Journal of Economic Theory*, vol. 165, pp. 672–706, sep 2016.

[19] M. Wu, S. Amin, and A. E. Ozdaglar, "Value of Information in Bayesian Routing Games," *Operations Research*, vol. 69, pp. 148–163, jan 2021.

[20] M. Castiglioni, A. Celli, A. Marchesi, and N. Gatti, "Signaling in Bayesian Network Congestion Games: the Subtle Power of Symmetry," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 5252–5259, may 2021.

[21] Y. Zhu and K. Savla, "Information Design in Nonatomic Routing Games With Partial Participation: Computation and Properties," *IEEE Transactions on Control of Network Systems*, vol. 9, pp. 613–624, jun 2022.

[22] B. L. Ferguson, P. N. Brown, and J. R. Marden, "Avoiding Unintended Consequences: How Incentives Aid Information Provisioning in Bayesian Congestion Games," in *61st IEEE Conference on Decision and Control*, pp. 3781–3786, apr 2022.

[23] H. Tavafoghi, A. Shetty, K. Poolla, and P. Varaiya, "Strategic Information Platforms in Transportation Networks," in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 816–823, IEEE, sep 2019.

[24] P. N. Brown, "When Altruism is Worse than Anarchy in Nonatomic Congestion Games," in *2021 American Control Conference (ACC)*, pp. 4503–4508, IEEE, may 2021.

[25] C. Papadimitriou, "Algorithms, games, and the internet," in *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pp. 749–753, 2001.