

Copyright

by

Edmund Liangfei Wong

2013

The Dissertation Committee for Edmund Liangfei Wong  
certifies that this is the approved version of the following dissertation:

## **Raising the BAR in Dependable Cooperative Services**

Committee:

---

Lorenzo Alvisi, Supervisor

---

Michael Dahlin

---

Thomas Moscibroda

---

Vitaly Shmatikov

---

Thomas Wiseman

**Raising the BAR in  
Dependable Cooperative Services**

by

**Edmund Liangfei Wong, B.S. C.S.; B.S. E.C.E.**

**DISSERTATION**

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

**DOCTOR OF PHILOSOPHY**

**THE UNIVERSITY OF TEXAS AT AUSTIN**

August 2013

To my advisor, Lorenzo Alvisi, and all the wonderful people—of which there are too many to list—that I worked with, had fun with, and supported me during some of the most arduous (and fun!) years of my life.

Thank you.

# Raising the BAR in Dependable Cooperative Services

Edmund Liangfei Wong, Ph.D.

The University of Texas at Austin, 2013

Supervisor: Lorenzo Alvisi

Cooperative services—a term which includes any system that relies on the resources and participation of its clients to function—have proven to be a popular, naturally scalable means to disseminate content, distribute computational workloads, or provide network connectivity. However, because these services critically depend on participants that are not controlled by a single administrative domain, these services must be designed to function in environments where *no* participant—because of failure or selfishness—will necessarily follow the specified protocol.

This thesis addresses the challenge of establishing and maintaining cooperation in cooperative services by (1) advancing our understanding of the limits to what our services can guarantee in the presence of failure, (2) demonstrating the critical role that correct participants can play in the incentives provided by the service, and (3) proposing a new notion of equilibrium that, unlike traditional notions, provides both rigorous yet practical guarantees in the presence of collusion. Furthermore, we demonstrate that our ideas can be applied to practice by designing and implementing Seer, a system that provides a scalable, reliable, and robust method for disseminating content even if participants may fail arbitrarily or deviate selfishly as a coalition.

# Contents

<b>Abstract</b>	<b>v</b>
<b>List of Tables</b>	<b>ix</b>
<b>List of Figures</b>	<b>x</b>
<b>Chapter 1 Introduction</b>	<b>1</b>
1.1 Byzantine . . . . .	4
1.2 Acquiescent . . . . .	6
1.3 Rationally Colluding . . . . .	8
1.4 System . . . . .	11
1.5 Organization of thesis . . . . .	12
<b>Chapter 2 Background</b>	<b>14</b>
<b>Chapter 3 Byzantine</b>	<b>19</b>
3.1 Model . . . . .	22
3.2 Regret freedom . . . . .	22
3.2.1 $(k, t)$ -robustness in communication games . . . . .	24
3.2.2 What if we know who is Byzantine? . . . . .	26
3.2.3 What if we know how Byzantine nodes behave? . . . . .	29
3.3 Regret bravery . . . . .	32
3.3.1 $t$ -maximin equilibrium . . . . .	35
3.3.2 Bayes equilibrium . . . . .	37

3.4	Summary . . . . .	41
<b>Chapter 4 Acquiescent</b>		<b>42</b>
4.1	Setup . . . . .	46
4.2	The need for acquiescence . . . . .	51
4.3	Acquiescence to the rescue . . . . .	53
4.3.1	When does a rational player pester? . . . . .	57
4.3.2	When does a rational player contribute? . . . . .	59
4.3.3	Complete formal results . . . . .	60
4.4	Characterizing the equilibrium . . . . .	81
4.5	Summary . . . . .	85
4.6	Appendix: Infinite-horizon games . . . . .	86
4.6.1	The unbounded model . . . . .	86
4.6.2	Equilibria of interest . . . . .	86
4.6.3	What if a Byzantine player acts destitute? . . . . .	87
4.6.4	What if a Byzantine player never contributes? . . . . .	90
<b>Chapter 5 Rationally Colluding</b>		<b>95</b>
5.1	Disincentivizing coalitions . . . . .	98
5.1.1	Can trusted third parties limit equilibria? . . . . .	100
5.1.2	What if nodes may fail? . . . . .	104
5.1.3	Do nodes want to punish one another? . . . . .	107
5.1.4	What other issues are there? . . . . .	113
5.2	Accepting coalitions . . . . .	114
5.2.1	Coalition-indistinguishable equilibria . . . . .	116
5.2.2	From indistinguishability to stability . . . . .	117
5.2.3	Examples of equilibria . . . . .	119
5.3	Summary . . . . .	123
<b>Chapter 6 System</b>		<b>124</b>
6.1	Principles . . . . .	127
6.2	System model . . . . .	128

6.3	Overview of Seer . . . . .	129
6.3.1	How does Seer encourage use of P2P? . . . . .	131
6.3.2	How does Seer incentivize adherence? . . . . .	133
6.3.3	What role does the server play? . . . . .	134
6.3.4	Seer: a close-up . . . . .	136
6.4	Incentives in Seer . . . . .	139
6.4.1	Abstracting away uncertainty . . . . .	140
6.4.2	What do clients do if they use P2P? . . . . .	142
6.4.3	When do clients choose to use P2P? . . . . .	153
6.4.4	The effects of Byzantine failure . . . . .	155
6.4.5	The effects of collusion . . . . .	156
6.4.6	Discussion . . . . .	158
6.5	Implementation and evaluation . . . . .	161
6.5.1	Selecting block size . . . . .	163
6.5.2	Scalability . . . . .	165
6.6	Summary . . . . .	168
<b>Chapter 7 Related Work</b>		<b>170</b>
7.1	Incentives in the presence of failure . . . . .	170
7.2	Leveraging acquiescent participants . . . . .	171
7.3	Incentives in the presence of collusion . . . . .	173
7.4	Other P2P and hybrid P2P systems . . . . .	175
<b>Chapter 8 Conclusion</b>		<b>178</b>
<b>Bibliography</b>		<b>180</b>

# List of Tables

2.1	List of symbols used in this thesis and their general meaning. .	18
4.1	Summary of symbols that define the payoffs of various actions.	48
4.2	The maximum number of times player 1 contributes for each of the thresholds shown in Figure 4.3. Simulation run with $\gamma_{\uparrow c}/\gamma_{\downarrow p} = 2$ , $R = 20$ , and $\mu_1(\mathbf{B}) = 0.1$ . . . . .	84
6.1	Parameters used for various pieces of content (§6.4.3). . . . .	153
6.2	Average throughput per client vs. block size. . . . .	163

# List of Figures

4.1	Sufficient initial beliefs for a rational player 2 in player 1 being acquiescent to incentivize player 2 to pester ( $y$ -axis, shaded area) for varying amounts of acquiescent generosity ( $x$ -axis); network loss or $\rho$ (top/bottom plots); and $b - \gamma_{\downarrow c}$ (left/right plots). Simulation run with $\gamma_{\uparrow p} = 1$ , $R = 20$ , and $\mu_2(\mathbf{B}) = 0.5$ .	82
4.2	Sufficient initial beliefs for a rational player 2 in player 1 being acquiescent to incentivize player 2 to pester ( $y$ -axis, shaded area) for varying amounts of acquiescent generosity ( $x$ -axis); network loss or $\rho$ (top/bottom plots); and $\mu_2(\mathbf{B})$ (left/right plots). Simulation run with $\gamma_{\uparrow p} = 1$ , $R = 20$ , and $b - \gamma_{\downarrow c} = 10^5$ .	83
4.3	Player 1's belief thresholds. Solid lines represent $\rho = 0.05$ ; dotted lines represent $\rho = 0.25$ . Simulation run with $\gamma_{\uparrow c}/\gamma_{\downarrow p} = 2$ and $R = 20$ .	84
5.1	In the simple secret-sharing game (Definition 5.7), the minimum benefit needed for a $k$ -resilient equilibrium where nodes attempt to reconstruct the secret.	107
6.1	The messages involved in initiating a block exchange.	130
6.2	The messages involved in the main block exchange loop.	132
6.3	The messages involved in finishing a block exchange. The italicized fields are optional.	133

6.4	An illustration of how various symbols relate to a particular exchange. . . . .	141
6.5	The minimum average effective channel bandwidth, expressed as a fraction of the channel bandwidth, that a client must experience to prefer using P2P over the free server tier. For values of the discount factor ( $x$ -axis) that do not have corresponding points, the client never prefers using P2P. . . . .	153
6.6	The minimum average effective channel bandwidth, expressed as a fraction of the channel bandwidth, that a client must experience to prefer using P2P over the paid server tier. For values of the discount factor ( $x$ -axis) that do not have corresponding points, the client never prefers using P2P. . . . .	154
6.7	Proportion of data served by the server to a flash crowd with varying block sizes. . . . .	164
6.8	Comparison of throughput, normalized to the theoretical maximum of 10 Mbps, between Seer, BitTorrent, and client/server with a varying number of clients in a flash crowd. . . . .	166
6.9	Proportion of received data served by the server with a varying number of clients in a flash crowd. . . . .	167
6.10	Comparison of throughput, normalized to the theoretical maximum of 10 Mbps, between Seer, BitTorrent, and client/server with a varying number of clients that download at random times.	168
6.11	Proportion of received data served by the server with a varying number of clients that download at random times. . . . .	169

# Chapter 1

## Introduction

*How does one reason about and build systems with provable properties in an environment where no participant is guaranteed to follow the specified protocol?*

While such an environment may seem implausible in reality, this is precisely the environment in which services that span multiple administrative domains must function. In such services—which include applications such as content dissemination (e.g., [4]), file backup (e.g., [16]), volunteer computing (e.g., [15]), multihop wireless networking (e.g., [12]), and Internet routing—resources are not under the control of a single administrative domain, so the necessary cooperation cannot simply be achieved by fiat. Instead, it is imperative that the service be structured so that nodes—which are administered by different, potentially selfish entities—have an incentive to help sustain it. Such issues are not simply the concern of researchers secluded in an ivory tower: much evidence suggests that a large number of peers will free-ride or deviate from the assigned protocol if it is in their interest to do so (e.g., [10, 22, 65, 96]). Yet, ensuring cooperation is critical to the dependability, if not the survival, of these aptly-named *cooperative services*, which rely on their participants' resources to provide their offerings. It is this challenge that this thesis addresses: establishing and maintaining cooperation in cooperative ser-

vices between participants that may arbitrarily fail or selfishly deviate when doing so is in their best interest.

There has been much previous work that has attempted to address this challenge. Traditional game theory has been heavily leveraged to provide a rigorous basis for designing a cooperative service (e.g., [4, 18, 43, 51, 70]) where nodes do not want to deviate (an *equilibrium*). This approach—while rigorous when all participants are rational—has shortcomings that stem from game theory’s general assumption that all participants act rationally, whereas failures—which may cause participants to act “irrationally”<sup>1</sup>—are a reality in any distributed system. This is particularly the case in cooperative services, where the nodes themselves are often unreliable personal machines riddled with malware and other exploits [1, 26]. Meanwhile, traditional fault tolerance techniques can handle rational deviations by modeling them as failures, but the limitations of this approach are obvious: since basic distributed computing primitives such as consensus and reliable broadcast cannot be implemented if more than one third of the nodes are Byzantine [75], fault tolerance techniques alone are unlikely to be able to handle the potentially large number of selfish participants who may find it advantageous to deviate.

Previous work has combined principles from both game theory and fault-tolerant distributed computing to provide a model that overcomes the limitations of the individual approaches. One approach is to draw inspiration from traditional Byzantine fault tolerance techniques and require that rational nodes prefer the specified equilibrium regardless of how up to  $t$  failures occur [19, 21, 49]. In this thesis, we refer to this general approach as  $(k, t)$ -robustness (from [19]). Another general model, formalized by Aiyer et al. [24], is known as BAR for the Byzantine (faulty), acquiescent (correct),<sup>2</sup> and rational (or selfish) participants that are explicitly modeled.

Both approaches leave much to be desired.  $(k, t)$ -robustness, while theo-

---

<sup>1</sup>One can model arbitrarily faulty peers as rational peers who follow an unknown utility function. Unfortunately, doing so does not simplify the problem.

<sup>2</sup>These nodes were originally known as “altruistic”, but we have renamed them to better capture the idea that these peers are obedient rather than irrationally generous [20].

retically elegant, has overly strict requirements in what guarantees an equilibrium provides with respect to Byzantine failures and coalitions. As we show in §3.2 and §5.2, these requirements make  $(k, t)$ -robustness unachievable in many real-world scenarios. On the other hand, the BAR model has been applied to real systems [24, 78, 79], but these systems, as well as other work that has used similar models [86], have:

- Generally required that rational participants believe the worst regarding Byzantine failure,
- Assumed acquiescent nodes do not exist or models them as rational, and
- Do not explicitly handle collusion, a phenomenon that has been observed in practice [80].

So what should be the basis for a rigorous treatment of cooperative services? How do we model the participants, and what types of guarantees should we aim for? Most importantly, *can we apply our model and guarantees in practice in a real system?* As the aforementioned shortcomings with  $(k, t)$ -robustness and BAR illustrate, there is a tension between theory and practice: theoretically robust and elegant notions may prove impossible to realize in real systems, but sprinkling a system with incentives whose rationale is rooted in intuition and common sense provides only a modicum of protection that is often defeated when exposed to more than casual strategic behavior [10, 68, 76, 81, 93, 99].

This thesis raises the bar in both the theory and practice of cooperative systems by first describing three theoretical contributions that address the aforementioned shortcomings in the BAR model while advancing the modeling and understanding of all three types of nodes: Byzantine, acquiescent, and rational(ly colluding). Importantly, we show how these contributions, while interesting theoretically, are applicable to a real system.

In summary, this thesis shows that:

- Notions of equilibrium inspired by traditional Byzantine fault-tolerant techniques, such as  $(k, t)$ -robustness, do not admit any useful equilibria even under a very general model of cooperative services. As a result, any practical equilibria must take into account a rational node’s beliefs regarding how Byzantine failure will occur (§3).
- Acquiescent nodes—the nodes that enable selfish nodes to free-ride to begin with—can actually be a boon for the system. As we will show, acquiescent nodes are not only sufficient, they are often necessary to provide incentives for rational nodes to cooperate in cooperative services when nodes know when the service will come to an end (§4).
- Traditional methods for analyzing protocols in the presence of collusion are likely to bear little fruit in real systems. We propose new notions of equilibria that provide rigorous guarantees of stability even when nodes may collude (§5).
- The aforementioned theoretical contributions are useful in practice. We design and implement a new hybrid content distribution system, Seer, that applies our theoretical insights towards building a scalable, robust, and dependable method for distributing content, thus demonstrating that we can have our (theoretical) cake and eat it (in practice) too (§6).

In the remainder of this chapter, we delve into the contributions we make to each type of node as well as the details of the system.

## 1.1 Byzantine

In traditional Byzantine fault tolerance, as long as the number of Byzantine nodes does not exceed some threshold  $t$ , the system is guaranteed to provide its safety properties *independent of who the  $t$  Byzantine nodes are and how they behave*. It is appealing to aim for a notion of equilibrium based on the same principles, in which rational nodes—either unilaterally or as a part of

a coalition—cannot improve their utility by deviating *independent of who the  $t$  Byzantine nodes are and of how they behave*. This approach, formalized as a part of  $(k, t)$ -robustness [19, 21] and other notions such as fault-tolerant Nash equilibrium [49], is in principle very attractive: at equilibrium, peers will always be free of regret, as their chosen strategy is guaranteed to prove a best response regardless of who and how  $t$  nodes fail.

We show that, despite its appeal, any notion of equilibrium that guarantees *regret freedom* is fundamentally unable to yield non-trivial equilibria in any game that captures three key characteristics of achieving some desired functionality in many fault-tolerant distributed systems:

- Some nodes need to communicate.
- Bandwidth is not free.
- The desired functionality can be achieved despite  $t$  Byzantine failures.

Furthermore, we argue that requiring regret freedom in any manner is simply too much to ask. In particular, we find that even if we weaken  $(k, t)$ -robustness by requiring regret freedom in only one dimension—namely, that rational nodes do not regret their decision even if they knew (1) who would fail but not how they fail, or (2) how nodes fail but not who would fail—still only achieves regret-free equilibria under very limited circumstances.

To overcome this impasse, we argue instead that nodes must brave through the possibility of regret. This approach, which we call *regret braving*, is motivated by the observation that rational agents that operate under uncertainty about the strategy of other players (as is the case when players are Byzantine) are often willing to cooperate without requiring absolute regret freedom, as long as cooperation is *expected* to yield the highest payoff. For instance, when stock traders buy or sell shares, they are well aware of the possibility of regretting their actions. Nonetheless, they follow a particular strategy as long as they cannot improve their utility with respect to their expectation about their environment—the worth of the traded asset, their comfort with

risk, and what they believe will be the trends in the market—by deviating. Similarly, we consider notions of equilibrium in which rational nodes aim to best respond to their expectations regarding Byzantine failures: the chosen strategy guarantees no regret only to the extent that such expectations prove correct.

We will describe two notions of equilibria: in the first, rational nodes play a maximin strategy that guarantees the best worst-case outcome despite any possible Byzantine failure; in the other, rational nodes assign probabilities to various possible faulty behaviors and aim for a Bayesian equilibrium. We then show that regret-braving equilibria admit simple and intuitive equilibria for communication games where even the weakened versions of  $(k, t)$ -robustness could not.

## 1.2 Acquiescent

There exists a sizable fraction of acquiescent peers in many environments [22]. Yet, despite their ubiquity, it is not obvious what impact these participants can have on cooperative services and their incentives. On one hand, services often rely on the existence of acquiescent peers—and the unselfishness codified in the protocol they obediently follow—to continue providing service. In particular, these acquiescent peers enable selfish participants to leech off the service without contributing their fair share. On the other, heavily relying on such peers to maintain the service by making up for free-riders is a risky proposition: even well-meaning peers, if blatantly taken advantage of, may give in to the temptation of joining the ranks of the selfish, leading in turn to more defections and to the service’s collapse. Because of the difficulty in finding a way to rely on acquiescent peers without having rational peers simply freeload off of them, acquiescent nodes have been largely ignored in the design of cooperative service protocols.<sup>3</sup>

---

<sup>3</sup>Acquiescent nodes can be modeled as Byzantine [19] or as rational nodes by assigning them the rational strategy. Neither classification is particularly satisfying.

In this thesis, we show that not only are acquiescent nodes not antithetical to rational cooperation, but that, in a fundamental way, rational cooperation can only be achieved in the presence of acquiescent nodes. We demonstrate how acquiescent nodes can induce cooperation by distilling the issue of whether cooperation can be maintained in a cooperative service to one critical exchange: the *last exchange*, the only exchange that, by definition, cannot rely on the promise of future benefit, and the exchange which provides incentives that form the cornerstone of cooperation in prior interactions.

We show that, without requiring oft-used assumptions on never-ending services or ignorance on a node’s part regarding when the service ends—assumptions that are often simply not true in practice and, worse, may still be insufficient to incentivize cooperation given a lossy network and Byzantine failure—the presence of acquiescent peers is sufficient and, in many ways, necessary to motivate rational peers to contribute. In particular, in the context of a model of the last exchange where one node wants to acquire content from its peer and has the capability to pester its peer (which induces cost on both the node and its peer), we prove that:

- There exists no equilibrium strategy where rational peers contribute if all peers are either rational or Byzantine. This result essentially holds even if we allow for an infinite number of pestering rounds.<sup>4</sup>
- The presence of acquiescent peers is sufficient to transform pestering into a credible threat. Intuitively, if rational peers have sufficiently high beliefs that they may be interacting with an acquiescent peer, they are motivated to pester, making it in turn preferable for rational peers to contribute.

The fraction of acquiescent nodes sufficient to sustain rational contribution depends on several system parameters, including the probability of

---

<sup>4</sup>In §4.6, we precisely define the conditions under which our result holds given an infinite number of rounds; intuitively, we assume players do not want to be punished and randomize at most a finite number of times.

network loss, the fraction of Byzantine peers in the system, and the behavior that rational peers expect from acquiescent and Byzantine peers. Exploring this space through a simulator, we find that:

- Acquiescent peers make rational cooperation easy to achieve under realistic conditions. In particular, we find that even if less than 10% of the population is acquiescent, rational peers are incentivized to cooperate in a system where the network drops 5% of all packets and Byzantine peers make up over 50% of the remainder of the population.
- Confirming our prior intuition, overly acquiescent peers—those that contribute every time they are pestered—make it more difficult to achieve rational cooperation: we cannot always achieve rational cooperation; when we do, it requires an implausibly high fraction of acquiescent peers. This is good news: the less foolishly generous is the acquiescent behavior sufficient to incentivize rational contribution, the more feasible it is to design systems with a sustainable population of acquiescent peers.
- The uncertainty introduced by network loss is both a bane and a boon. On the one hand, it significantly complicates the analysis of a peer’s optimal strategy because each peer does not know what the other has observed. On the other, it lowers the threshold for rational cooperation by leaving open some possibility that the other peer may be acquiescent, even when the observed behavior suggests otherwise.

### 1.3 Rationally Colluding

The social nature of cooperative services suggests that nodes will develop, or may have already established, a rich web of relationships (e.g., based on friendship or on belonging to the same organization), which may cause *coalitions* of nodes to collude and deviate together [80]. The literature offers two approaches to guarantee that deviations resulting from collusion do not affect

the incentives provided to rational nodes. The first approach is to model collusion as a fault and colluding nodes as Byzantine. While possible—after all, Byzantine failure includes any arbitrary failure—this approach suffers the same limitations that modeling rational deviations as Byzantine does: it introduces an artificially low cap on the number of colluders. The second approach—taken by *strong Nash* [31], *k-resilient* equilibria [19, 21], and *coalition-proof* Nash equilibria [33], to name a few—is to deny any benefit to colluders: if the equilibrium is a best response not just to every individual, but also to every possible coalition, then collusion poses no harm to the equilibrium’s stability, since nodes gain no benefit by colluding.

Our work is motivated by what we believe to be a critical flaw of the second approach: its requirement that every node prefer the same course of action, despite whom it may be colluding with. After all, nodes that collude are likely to trust each other more and, more generally, be able to hold stronger assumptions about one another. Since stronger assumptions typically lead to more efficient protocols, identifying a single strategy that is a best response both inside and outside of every possible coalition is very difficult and is unlikely to yield a useful, practical foundation for building dependable services.

To overcome this challenge, this dissertation introduces a fundamentally different approach to dealing with coalitions. The key observation is that the fundamental property an equilibrium provides is *stability* (in that nodes do not want to deviate), and that while finding a single best response between all nodes is sufficient, it is not *necessary* to achieve stability. Leveraging this observation, we introduce two new notions of equilibrium that achieve stability through a simple observation: coalitions (including the trivial singleton coalition of one non-colluding node) will not deviate from an equilibrium as long as the equilibrium specifies a best-response strategy for every *coalition*. Thus, the strategy a node follows may depend on whom the node is colluding with, thereby enabling the equilibrium to explicitly model the advantages that coalition members have while providing assurances that nodes do not deviate from the specified equilibrium.

The first notion of equilibrium, *k-indistinguishability*, provides an attractively simple guarantee: the strategy a node plays may depend on whom it is colluding with, but its participation in any coalition has no effect on its actions towards peers the node is not colluding with. Thus, in a *k*-indistinguishable equilibrium, nodes cannot tell whether another node, with whom they are not colluding, is itself part of some other coalition (of at most *k* nodes).

*k*-indistinguishable equilibrium, while simple, provides stronger than necessary guarantees. The second notion, *k-stability*, instead adheres to the conditions necessary for stability: like *k*-indistinguishability, *k*-stable equilibria specify a strategy per coalition that is a best response to the strategies played by all other possible coalitions; unlike *k*-indistinguishability, the actions that a node takes as a part of a *k*-stable equilibrium may be informative about whether it is colluding and with whom. Finally, we introduce *strategy functions*, a new construct that enables us to express the strategies a node may play as a function of whom a node is colluding with.

In summary, we make the following contributions to how coalitions are dealt with in cooperative services:

- We illustrate the limits of generalizing Nash equilibria to coalitions, which requires that a single strategy be a best response for every node regardless of whether it is colluding, and show that is too strong a requirement to admit equilibria in several common cooperative service scenarios.
- We distill the fundamental stability property that defines an equilibrium and introduce (a) two new notions of equilibria—*k*-indistinguishability and *k*-stability—that continue to guarantee stability while allowing nodes to benefit from their coalitions, and (b) a new construct, strategy functions, for expressing a node’s strategy as a function of its coalitions.
- We demonstrate the applicability of our equilibria by showing their use in scenarios for which previous equilibria did not exist.

## 1.4 System

Putting it all together, we demonstrate the utility of our ideas and insights by applying them to the design and implementation of Seer, a hybrid P2P content distribution system.

Currently, content is predominantly delivered via servers in large data centers, which provide a trusted, reliable source of content. The increasing popularity of these services and richness of the content has resulted in the rapid deployment of new data centers to offset the load (e.g., [38, 39]). The cost of connecting and running these data centers can quickly add up: for instance, Netflix, who streamed roughly 4 billion hours worth of content in Q1 2013 [91], pays over \$400 million a year for bandwidth alone (assuming it costs Netflix roughly \$0.05 per 2 hours of content [36]).

Cooperative services, such as those that use peer-to-peer (P2P) technology, provide a cheaper alternative: because participants are responsible for distributing content, the number of content distributors inherently scales with the number of clients. However, such scalability does not come for free: in addition to the aforementioned challenges in designing incentives and dealing with failures, P2P systems provide no guarantee that clients will acquire the content they desire.

In this thesis, we describe Seer, a hybrid P2P-client/server that provides a reliable, scalable, and robust content distribution service with strong incentives for cooperation. Like previous hybrid P2P systems—deployed by companies such as Spotify [17], Blizzard [5], and PPTV [14, 64] and studied and used in much prior work (e.g., [89, 92, 100])—Seer leverages the strengths of both approaches on which it is based: the scalability and low cost of P2P services combined with the reliability of a server or CDN. Unlike other hybrid P2P systems—which have largely assumed that clients will not free-ride off the server (e.g., by changing the settings in the client’s software preferences, using a firewall to block P2P traffic, or modifying the client itself) or used informal arguments about their incentives (which could lead to potential exploits

[68, 10, 76, 81, 93, 99])—Seer leverages both our theoretical contributions and the trusted server to provide robust guarantees that peers will be incentivized to disseminate content despite the presence of Byzantine, acquiescent, and rational participants. Thus, Seer ensures that it will not simply devolve to a client/server architecture where the server provides all the content to the clients.

Seer is proof that one can design a provably dependable cooperative service under assumptions significantly more realistic than previously achieved: more than ever before, Seer is not only rigorous, but *rigorous in practice*. In particular, Seer does not assume that clients are (1) exceedingly risk-averse with respect to Byzantine failure, (2) blissfully unaware of when the last exchange is and that they may receive no future benefit from the service, and (3) unable to collude; furthermore, rational nodes in Seer account for network loss and latency. Despite these weaker assumptions, we are able to rigorously prove that clients want to help disseminate content in Seer. We accomplish this in part by basing design decisions on what we are able to prove about the system, thereby ensuring that Seer’s policies and mechanisms are backed by strong theoretical guarantees. Yet, as we show, rigor does not have to come at the price of performance: Seer neither sacrifices the scalability of P2P services nor the reliability of client/server services. Evaluating our implementation of Seer, we find that Seer can support significantly more clients than a traditional client/server service and can outperform BitTorrent, a popular traditional P2P service, by over 20%.

## 1.5 Organization of thesis

We introduce some background and common terminology and concepts that we will use in the rest of the thesis in §2. We delve into our theoretical contributions in §3 (Byzantine), §4 (acquiescent), and §5 (rational). We describe Seer in §6. Finally, we describe related work in §7, and conclude in §8.

The theoretical contributions in this thesis have appeared in previous

publications [104, 106, 107]. Seer has been published as a technical report [105].

# Chapter 2

## Background

Before describing our contributions in detail, we provide some general background, define some useful notions and terminology, and give a high-level introduction on how we model and reason about cooperative services. We define relevant notions and terminology as needed in subsequent chapters.

We model cooperative services as a game played by a set of  $n$  players  $N = \{1, \dots, n\}$ . These players represent participants in the cooperative service (which we also refer to as “nodes” or “peers”). In our game,

- Nodes communicate by sending and receiving messages.
- A node does not benefit directly from the act of sending a message. Thus, a node does not receive benefit from sending a message if the outcome of the node (e.g., the messages or credits it receives) or system (e.g., the outcome of consensus) does not change.
- Doing more (e.g., sending or receiving more or larger messages) incurs more cost.
- A node’s payoff solely depends on what benefits it earns and what costs it incurs.

- The actions a node can choose to take do not decrease as a result of receiving a message. In particular, a node can always act as if it had not received the message.<sup>1</sup>

We will define concrete instantiations of this model in subsequent chapters.

A *strategy*  $\sigma_x$  is a complete description of the *actions* some node  $x$  takes at any point in the game; in the context of this thesis, a strategy is effectively the protocol a node uses. We will mostly focus on *pure strategies*, which specify exactly one action at every point where  $x$  must make a decision, but we will also discuss *mixed strategies*, which may randomly select an action among some set of actions at some point in the game.

We refer to the service-assigned protocol as the *assigned strategy*. A *strategy profile*  $\sigma = (\sigma_x)_{x \in N}$  assigns a strategy  $\sigma_x$  to each node  $x$ . We refer to the actions that a node  $x$  has performed or observed in the past as the node's *history* and denote  $x$ 's history at some time  $r$  as  $h_x^r$ .

A *utility function*  $U$  defines the preferences of all nodes. If there is no randomness in the system or environment (e.g., packets are not randomly dropped because of network loss), then given a strategy profile  $\sigma$ , every node  $x$  earns *payoff* or *utility*  $U_x(\sigma)$ . Rational nodes prefer and select strategies that increase their payoffs as specified by the utility function. We refer to the payoff a node earns from the middle of the game (e.g., from some history  $h_x^r$ ) on as a node's *continuation payoff*. We denote the amount of payoff some node  $x$  earns from some strategy profile starting from some history  $h_x^r$  as  $U_x(\sigma|h_x^r)$ .

We denote “everyone but  $x$ ” as  $-x$ ; indicate the combination of multiple strategies into a strategy profile using parentheses, e.g.,  $\sigma = (\sigma_x, \sigma_{-x})$ ; and drop parentheses when the meaning is clear. For example,  $U_x(\sigma'_x, \sigma_{-x})$  denotes the payoff that  $x$  earns from playing  $\sigma'_x$  while everyone else plays  $\sigma_{-x}$ . We use the same notation for sets of nodes as well, e.g., for some set of nodes  $K$ ,  $-K$  represents “everyone but nodes in  $K$ .”

---

<sup>1</sup>This assumption may seem obvious in the context of computer science, but we state this explicitly to make clear that we are not considering any “unusual” games where this property may not hold.

Generally, we strive to prove that a particular protocol is an *equilibrium*. When there is no randomness in the system and environment, an *equilibrium* is a set of strategies (typically a strategy profile) in which every node prefers to play its assigned strategy. For example, the celebrated Nash equilibrium achieves this stability by ensuring that the strategy  $\sigma_x^*$  of any given node  $x$  is a *best response* (i.e., it maximizes  $x$ 's payoff) to everyone else following  $\sigma_{-x}^*$ . Thus, no node has any incentive to unilaterally deviate, making the strategy profile *stable*.

DEFINITION 2.1. A strategy profile  $\sigma^*$  is a Nash equilibrium if for all  $x \in N$ , there does not exist some strategy  $\sigma'_x$  such that

$$U_x(\sigma'_x, \sigma_{-x}^*) > U_x(\sigma^*)$$

Definition 2.1 is an example of a notion of equilibrium, or a *solution concept*, in that it defines a set of conditions that describe when a set of strategies is considered an equilibrium. We will define other solution concepts in subsequent chapters.

We focus on environments in which neither trusted hardware nor trusted third-parties are used to monitor *all* communication between peers.<sup>2</sup> Although such a monitor is useful, it is often impractical or even infeasible to provide one, and in practice few cooperative systems leverage trusted hardware to prove communication. We express this reality in the following assumption:

ASSUMPTION 2.2. A node that sent a message  $m$  cannot unilaterally prove that it sent  $m$ .

**BAR.** In this thesis, we consider three different types of nodes:

---

<sup>2</sup>Note that this does not preclude the existence of a trusted third-party such as the server in Seer (§6), which does not observe every message between every pair of nodes.

- *Byzantine*: these nodes play an arbitrary strategy.
- *Acquiescent*: these nodes follow the assigned strategy.
- *Rational*: these nodes follow a strategy if and only if deviating does not increase their payoff.

Each node  $x$  has a *type*  $\theta_x$  that distinguishes it as being Byzantine, acquiescent, or rational. For simplicity, we will generally assume that all rational nodes are of the same type (denoted  $\mathbf{R}$ ); we assume the same of acquiescent nodes (denoted  $\mathbf{A}$ ), who would anyway follow any strategy assigned to them. On the other hand, a Byzantine node may potentially play one of many different strategies. When we need to distinguish the different strategies a Byzantine node may play, we will denote a Byzantine node  $x$ 's type by its strategy  $\tau_x$ . When we do not (and we simply need to know the expected Byzantine strategy), we will denote the Byzantine type as  $\mathbf{B}$ .

We assume that a player's type is only known a priori to the player itself and any peers it may be colluding with (if collusion is possible). Thus, players must choose their strategies given incomplete information about the environment. The remainder of the thesis describes how we deal with and, at times, leverage this uncertainty.

**Summary of symbols used.** This thesis makes heavy use of a variety of symbols. For the reader's convenience and reference, Table 2.1 lists many of the symbols that we repeatedly use and how we generally use them in this thesis. Note that some of the symbols listed in Table 2.1 have not yet been introduced.

Symbol	General meaning
$b$	Benefit
$h_x^r$	Player $x$ 's history at some time $r$
$i, j$	Generic index variable
$k$	Size of coalition
$n$	Number of players
$r$	Variable related to time or the round
$t$	Maximum number of failures handled
$u, v$	Function or variable related to payoff
$x, y, z$	Index variable for players
$K$	A coalition of size $k$
$N$	Set of players $\{1, \dots, n\}$
$R$	Number of rounds (for finitely-repeated games)
$T$	Set of failed nodes
$U_x$	Utility function for player $x$
$V$	Function or variable related to payoff
$\gamma$	Cost of communication
$\delta$	Discount factor (for infinitely-repeated games)
$\theta_x$	Type of player $x$
$\mu / \mu_x$	Set of beliefs / player $x$ 's belief
$\sigma / \sigma_x$	Strategy (protocol) profile / player $x$ 's strategy
$\tau_x$	Strategy for (and type of) Byzantine player $x$
<b>B, A, R</b>	Byzantine, acquiescent, and rational types

**Table 2.1:** List of symbols used in this thesis and their general meaning.

# Chapter 3

## Byzantine

We start off by describing the type that is most familiar to those in distributed computing and systems. Because real distributed systems often fail in unexpected ways, there has been considerable work in developing techniques to tolerate Byzantine, i.e., arbitrary, failure [75].

In cooperative services, simply being able to tolerate failure is insufficient: non-faulty participants that are selfish may deviate for their own self-interest. Moreover, failure complicates any incentives that a service provides for cooperation, as these incentives must be robust to the possibility that some participants fail and act “irrationally.” It is useful, then, to provide a rigorous basis for analyzing such services and incentives by devising a solution concept that admits, as equilibrium, strategies (protocols) that rational nodes continue to follow despite the possibility of failure. But how does one devise such a solution concept?

A natural approach is to draw inspiration from traditional Byzantine fault-tolerant computing. In traditional Byzantine fault-tolerant systems, as long as the number of Byzantine nodes does not exceed a threshold  $t$ , the system is guaranteed to provide its safety properties *independent of who the  $t$  Byzantine nodes are and how they behave*. Similarly, it is appealing to aim for a notion of equilibrium in which rational nodes—either unilaterally or as a part of a coalition—cannot improve their utility by deviating *independent of who*

the  $t$  Byzantine nodes are and of how they behave. This approach, elegantly formalized in the notion of  $(k, t)$ -robustness [19, 21] (as well as in fault-tolerant Nash equilibrium [49]), is in principle very attractive: at equilibrium, peers will never have reason to regret their chosen strategy, which is guaranteed to prove a best response to any Byzantine strategy, independent of the identities of Byzantine nodes.

In this chapter, we show that, despite its appeal, a solution concept that guarantees *regret freedom* is fundamentally unable to yield non-trivial equilibria in games (which we name *communication games*) that capture three key characteristics of many practical fault-tolerant distributed systems:

- To achieve some desired functionality, some nodes need to communicate.
- Bandwidth is not free.
- The desired functionality can be achieved despite  $t$  Byzantine failures.

Moreover, we find that weakening  $(k, t)$ -robustness, even considerably, does not help. For example, suppose that, magically, all rational nodes in a communication game knew precisely the identity of all Byzantine nodes (but not their strategy); or, alternatively, that they knew their strategy (but not their identities). We find that in both cases a regret-free equilibrium can be achieved only under very limited circumstances.

These results are not interesting because of their proofs, which are straightforward, but because they show that in fault-tolerant distributed systems, conditioning rational cooperation on the expectation of regret freedom may be fundamentally too much to ask. Furthermore, the limitations of this approach appear hard to fix, since they are rooted in the universal quantifiers (e.g., “for all strategies” or “for all sets of  $t$  Byzantine nodes”) that are at the very essence of regret freedom.

The second part of this chapter describes how we overcome this impasse, using an approach we call *regret braving*. Regret braving is motivated by the observation that rational agents that operate under uncertainty about

the strategy of other players (as is the case when players are Byzantine) are often willing to cooperate without requiring absolute regret freedom, as long as cooperation is expected to yield the highest payoff. For instance, when stock traders buy or sell shares, they are well aware of the possibility of regretting their actions. Nonetheless, they follow a particular strategy as long as they cannot improve their utility with respect to their expectation about their environment—the worth of the traded asset, their comfort with risk, and what they believe will be the trends in the market—by deviating. Similarly, we consider solution concepts in which rational nodes aim to best respond to their expectations regarding Byzantine failures: the chosen strategy guarantees no regret only to the extent that such expectations prove correct.

We find that regret-braving solution concepts admit simple and intuitive equilibria for communication games where even the weakened versions of  $(k, t)$ -robustness could not. We will describe two solution concepts: in the first, rational nodes play a maximin strategy that guarantees the best worst-case outcome despite any possible Byzantine failure; in the other, rational nodes assign probabilities to various possible faulty behaviors and aim for a Bayesian equilibrium. We do *not* suggest that these solution concepts are the “right” ones or that they can be directly applied to every BAR-tolerant system. What these preliminary results do show, however, is that regret-braving solution concepts are not subject to the fundamental limitations inherent to regret freedom.

**Organization of chapter.** §3.1 formalizes how we model players and introduces the communication game that we use to compare solution concepts. §3.2 explores the land of the (regret) free, showing why equilibria that base rational cooperation on regret freedom are fundamentally hard to achieve. §3.3 describes instead the home of the (regret) brave: we discuss two models of rational beliefs that admit useful equilibria in an instantiation of the communication game.

## 3.1 Model

In this chapter, we model a fault-tolerant system in which communication is not free and at least some nodes must communicate to achieve the desired functionality using the following game. This game mirrors many of characteristics described in §2.

DEFINITION 3.1. A communication game consists of some set of nodes  $N = \{1, \dots, n\}$  in which

- Communication incurs some cost and does not generate direct benefit to the sender.
- Communication incurs some cost to the receiver.
- The actions a node can choose to take does not decrease as a result of receiving a message.
- Benefit is obtained from functionality that (a) can be achieved in the presence of up to  $t < n$  Byzantine failures and (b) requires communication between some pair of nodes.

For simplicity, in this chapter, we use the same communication cost  $\gamma > 0$  for both sending and receiving, we assume that messages are never lost, and we assume that there are no acquiescent nodes. We focus on *non-trivial* strategy profiles, in which some positive utility is expected for at least one node; this implies that some communication must occur.

## 3.2 Regret freedom

In Byzantine fault-tolerant systems, safety properties hold regardless of how Byzantine failures occur. Ideally, one would like rational cooperation to be achieved under similarly strong guarantees.  $(k, t)$ -robustness [19, 21] is an el-

egant solution concept that captures this attractive intuition. A  $(k, t)$ -robust equilibrium is completely impervious to the actions of Byzantine nodes: rational nodes will never have to second-guess their decision even if the identities and strategies of the Byzantine nodes become known.

$(k, t)$ -robustness offers three key properties. We define all three properties here for completeness, but our focus will be on the best-response condition, which is the condition that provides regret freedom.

The first property, *t-immunity* [19], captures the intuition that nodes following a strategy profile should not be adversely affected by Byzantine failures, as long as there are at most  $t$ .

DEFINITION 3.2. A strategy profile  $\sigma$  is *t-immune* if, for all  $T \subseteq N$  such that  $|T| \leq t$ , all strategy profiles  $\tau$ , and  $x \notin T$ ,

$$U_x(\sigma_{-T}, \tau_T) \geq U_x(\sigma)$$

Note that *t-immunity* is *not* equivalent to Byzantine fault tolerance, as *t-immunity* does not specify that a strategy profile  $\sigma$  must provide any sort of desirable safety or liveness properties despite  $t$  faults. In fact, any  $\sigma$ , fault-tolerant or not, is *t-immune* if it specifies actions so terrible that Byzantine nodes, playing anything other than  $\sigma$ , cannot hurt a player’s utility.

The second, *k-resilience* [19], addresses the possibility of collusion and is effectively a generalization of the Nash equilibrium for coalitions of up to size  $k$ : a *k-resilient* equilibrium guarantees that a coalition of size at most  $k$  cannot deviate in a way that benefits every member.<sup>1</sup>

---

<sup>1</sup>Abraham et al. also define a strong version of collusion resilience in which there must not exist a deviation in which even *one* coalition member can do better [19, 21]. While we explore the issue of collusion in §5, the impossibility results we derive in this chapter are orthogonal to how we handle collusion. Consequently, we will use the weak version as Abraham et al. do in [21]; since any strongly *k-resilient* equilibrium is (weakly) *k-resilient*, our impossibility results hold in both versions.

DEFINITION 3.3. A strategy profile  $\sigma^*$  is a  $k$ -resilient equilibrium if, for all  $K \subseteq N$  such that  $|K| \leq k$ , there does not exist an alternate strategy  $\sigma'_K$  such that for all  $x \in K$ ,

$$U_x(\sigma'_K, \sigma_{-K}^*) > U_x(\sigma^*)$$

The  $(k, t)$ -robustness solution concept is the combination of  $t$ -immunity,  $k$ -resilience, and *regret freedom with respect to Byzantine failure*: regardless of how Byzantine failures occurs,  $(k, t)$ -robustness guarantees that no coalition of at most  $k$  nodes can ever do better than following the equilibrium strategy. It is this last condition—the best-response condition—that provides regret freedom and which we focus on in this chapter.

DEFINITION 3.4. A strategy profile  $\sigma^*$  is a  $(k, t)$ -robust equilibrium if  $\sigma^*$  is  $t$ -immune and, for all (a)  $K, T \subseteq N$  such that  $K \cap T = \emptyset$ ,  $|K| \leq k$ , and  $|T| \leq t$ , and (b) strategy profiles  $\tau$ , there does not exist an alternate strategy  $\sigma'_K$  such that for all  $x \in K$ ,

$$U_x(\sigma'_K, \tau_T, \sigma_{-\{K \cup T\}}^*) > U_x(\sigma_{-T}^*, \tau_T)$$

Note that there is no notion of expectation here, because rational nodes, whoever they may be, must be willing to follow the protocol no matter who the Byzantine nodes are.

### 3.2.1 $(k, t)$ -robustness in communication games

We show that the very property that makes  $(k, t)$ -robustness so appealing—regret freedom regardless of how Byzantine failures occur—makes it infeasible in many real-world systems. The reason, fundamentally, is that *communication always incurs cost but could potentially yield no benefit if one is communicating with a Byzantine node*. In other words, a rational node may realize in hindsight that it could have reduced its costs without affecting its benefits by avoiding all

communication with Byzantine nodes, thus improving its utility. As any node can be Byzantine, this implies that the only possible  $(k, t)$ -robust equilibrium is one in which no node communicates.

**THEOREM 3.5.** There exists no non-trivial  $(k, t)$ -robust equilibrium in any communication game.

*Proof.* Consider some non-trivial  $(k, t)$ -robust strategy  $\sigma^*$ . There must exist some node  $x$  which, with positive probability  $\alpha$  under  $\sigma^*$ , sends a message to some other node  $z$  before receiving any other messages. Suppose that  $z$  is Byzantine. Since  $\sigma^*$  is  $(k, t)$ -robust,  $x$  must not be able to do better with some alternate strategy, regardless of who has failed and what a failed node will do. In particular, for all alternate strategies  $\sigma'_x$  for  $x$  and Byzantine strategies  $\tau_z$  for  $z$ , it must be that

$$U_x(\sigma_{-z}^*, \tau_z) \geq U_x(\sigma'_x, \tau_z, \sigma_{-\{x,z\}}^*) \quad (3.1)$$

Suppose  $\tau_z$  is the strategy in which  $z$  crashes immediately, i.e.,  $z$  never sends any messages. Let  $\sigma'_x$  be the strategy in which  $x$  plays the same actions with the same probability as in  $\sigma_x^*$ , except  $x$  sends nothing to  $z$ . By Assumption 2.2,  $x$  cannot prove that it communicated with  $z$ ; it thus follows that  $(\sigma'_x, \tau_z, \sigma_{-\{x,z\}}^*)$  has the same functionality as  $(\sigma_{-z}^*, \tau_z)$  and is indistinguishable to any node in  $N \setminus \{x, z\}$ . Clearly, if  $z$  follows  $\tau_z$ ,  $x$  can do better by never communicating with  $z$ :  $x$ 's outcome will not change (since  $z$  never communicates with anyone), and  $x$ 's communication costs are lower. Formally,

$$U_x(\sigma'_x, \tau_z, \sigma_{-\{x,z\}}^*) = U_x(\sigma_{-z}^*, \tau_z) + \alpha\gamma > U_x(\sigma_{-z}^*, \tau_z)$$

which contradicts inequality (3.1). □

More broadly, Theorem 3.5 suggests that it may be hard to build non-trivial  $(k, t)$ -robust equilibria for any game where a player's actions incur cost.

Indeed, in all the games for which Abraham et al. derive  $(k, t)$ -robust equilibria [19, 21], a node’s utility depends only on the game’s outcome (e.g., in a secret-sharing game based on Shamir’s scheme, utility depends on whether a node can learn the secret) and is independent of how much communication is required to reach that outcome.

It follows from Theorem 3.5 that  $(k, t)$ -robustness—which requires regret freedom along two axes: *who* the Byzantine nodes are and *how* they behave—may be too strong to require in practice. However, one may wonder whether regret freedom can still be applied along one axis. In particular, can we achieve regret freedom in communication games if we know exactly who the Byzantine nodes are, but not how they behave? What if we do not know who is Byzantine, but we know how they behave?

### 3.2.2 What if we know who is Byzantine?

Let us assume that we know *exactly* who all the Byzantine players are before the game begins. This may already appear a strong assumption, but it is necessary, since if the identity of even one Byzantine node were unknown, Theorem 3.5 would still apply. We show that, even with this strong assumption, a solution concept that is regret-free with respect to the strategies of Byzantine nodes is possible only to the extent that it defines away the problem: the only possible equilibria are those in which rational nodes communicate only among themselves, completely excluding Byzantine nodes from the system. Furthermore, we show that many interesting communication games do not yield a regret-free equilibrium even if one takes the drastic step of excluding Byzantine nodes. In communication games where Byzantine nodes may take actions that can *affect* a rational node’s utility by more than the cost of sending a *single* message, there exists no regret-free equilibrium, even if the identity of all Byzantine nodes are known a priori.

We first define the equivalent of  $t$ -immunity (Definition 3.2) and  $(k, t)$ -robustness (Definition 3.4) for a fixed set  $T$  of Byzantine nodes.

DEFINITION 3.6. A strategy profile  $\sigma$  is  $T$ -strategy-immune if for all strategy profiles  $\tau$  and  $x \notin T$ ,

$$U_x(\sigma_{-T}, \tau_T) \geq U_x(\sigma)$$

DEFINITION 3.7. A strategy profile  $\sigma^*$  is  $(k, T)$ -strategy-robust with respect to  $T \subseteq N$  iff  $\sigma^*$  is  $T$ -strategy-immune and for all  $K \subseteq N \setminus T$  such that  $|K| \leq k$  and all strategy profiles  $\tau$ , there does not exist some  $\sigma'$  such that for all  $x \in K$ ,

$$U_x(\sigma'_K, \tau_T, \sigma^*_{-(K \cup T)}) > U_x(\sigma^*_{-T}, \tau_T)$$

A  $(k, T)$ -strategy-robust equilibrium need only be a best response to the specified set  $T$  of Byzantine nodes. The following theorem shows that no  $(k, T)$ -strategy-robust equilibrium is possible unless rational nodes “blacklist” all nodes in  $T$ .

THEOREM 3.8. In a communication game, there does not exist any  $(k, T)$ -strategy-robust equilibrium  $\sigma^*$  where any  $x \notin T$  communicates with any  $z \in T$ .

*Proof.* This is similar to the proof of Theorem 3.5. Assume there exists such an equilibrium  $\sigma^*$  and, in  $\sigma^*$ , some node  $x$ , with positive probability  $\alpha$ , sends a message to some other node  $z \in T$ . Suppose the strategy  $\tau_z$  that some Byzantine node  $z$  employs is the crash strategy: it never communicates.

Consider an alternate strategy  $\sigma'_x$  in which some rational node  $x$  plays the same actions with the same probability as in  $\sigma^*$ , except  $x$  does not communicate with  $z$ . Since  $\sigma^*$  is a  $(k, T)$ -strategy-robust equilibrium,  $U_x(\sigma^*_{-z}, \tau_z) \geq U_x(\sigma'_x, \tau_z, \sigma^*_{-\{x, z\}})$ . Yet, by Assumption 2.2,  $x$  cannot prove it communicated with  $z$ . It follows that  $(\sigma^*_{-\{x, z\}}, \sigma'_x, \tau_z)$  has the same functional-

ity as  $(\sigma_{-z}^*, \tau_z)$  and is indistinguishable to any node in  $N \setminus \{x, z\}$ . Thus,

$$U_x(\sigma_{-\{x,z\}}^*, \sigma'_x, \tau_z) \geq U_x(\sigma_{-z}^*, \tau_z) + \alpha\gamma > U_x(\sigma_{-z}^*, \tau_z)$$

Contradiction. □

Although Theorem 3.8 does not rule out all  $(k, T)$ -strategy-robust equilibria, Theorem 3.9 proves that these equilibria, which must be regret-free for *any* Byzantine strategy, only exist in limited circumstances. Intuitively, if different Byzantine failures can affect a rational node's payoff by even just the cost of communication and these failures can be triggered by whether a rational node communicates or not, a rational node may find that ignoring Byzantine nodes may not be optimal in hindsight (as required by regret freedom).

**THEOREM 3.9.** No communication game can yield a  $(k, T)$ -strategy-robust equilibrium for any set  $T \subseteq N$  of Byzantine nodes if for some  $x \notin T$ ,

- $x$  has at least one opportunity to send a message to some  $z \in T$ ,
- Members of  $T$  can freely coordinate their strategies,<sup>a</sup> and
- For any strategy profile  $\sigma$  in which  $x$  does not send any message to any member of  $T$ , there exist two Byzantine strategies  $\tau_T^h$  and  $\tau_T^\ell$  such that  $\tau_T^h$  and  $\tau_T^\ell$  are the same until one of  $x$ 's opportunities to send a message to some  $z \in T$  and

$$U_x(\sigma_{-T}, \tau_T^h) - U_x(\sigma_{-T}, \tau_T^\ell) > \gamma$$

<sup>a</sup>Note that this trivially holds if  $|T| = 1$ .

*Proof.* By contradiction. Fix  $\sigma^*$  to be some  $(k, T)$ -strategy-robust equilibrium. We know by Theorem 3.8 that if  $\sigma^*$  is  $(k, T)$ -strategy-robust, then any rational node  $x$  following  $\sigma^*$  never chooses to send to any member of  $T$ .

By assumption, we know that, given  $\sigma^*$ , there exists two Byzantine strategies  $\tau_T^h$  and  $\tau_T^\ell$  such that they affect  $x$ 's utility by more than  $\gamma$  during some opportunity that  $x$  has to communicate with some  $z \in T$ . We compose a strategy  $\tau_T$  where:

- If  $x$  sends a message at this time to  $z$ ,  $T$  plays actions from  $\tau_T^h$  as if  $x$  had not sent a message.
- $T$  otherwise plays actions from  $\tau_T^\ell$ .

It is obvious then that  $x$  prefers to send a message to  $z$ . More formally, consider some alternate strategy profile  $\sigma'$  that is the same as  $\sigma^*$  except  $x$  chooses to communicate in its first interaction with any member of  $T$ ;  $x$  then plays the same actions as if it had not communicated in  $\sigma_x^*$ . It follows that

$$\begin{aligned} U_x(\sigma'_x, \tau_T, \sigma_{-\{x\} \cup T}^*) &= U_x(\sigma_{-T}^*, \tau_T^h) - \gamma \\ &> U_x(\sigma_{-T}^*, \tau_T^\ell) \\ &= U_x(\sigma_{-T}^*, \tau_T) \end{aligned}$$

This contradicts the assumption that  $\sigma^*$  is a  $(k, T)$ -strategy-robust equilibrium.  $\square$

Theorem 3.9—unlike Theorem 3.5—provides conditions under which no  $(k, t)$ -strategy-robust equilibria exist, whether trivial or not. Since  $(k, t)$ -strategy-robust equilibria are a superset of  $(k, t)$ -robust equilibria, it naturally follows from Theorem 3.9 that no  $(k, t)$ -robust equilibria exist under the same conditions.

### 3.2.3 What if we know how Byzantine nodes behave?

Let us now consider a solution concept that assumes that the strategy played by every Byzantine node is known a priori and yields equilibria that are regret-free with respect to who the Byzantine nodes are.

DEFINITION 3.10. The strategy profile  $\sigma^*$  is a  $(k, t, \tau)$ -type-robust equilibrium iff  $\sigma^*$  is  $t$ -immune<sup>a</sup> and for all  $K, T \subseteq N$  such that  $K \cap T = \emptyset$ ,  $|K| \leq k$ , and  $|T| \leq t$ , there does not exist some  $\sigma'$  such that for all  $x \in K$ ,

$$U_x(\sigma'_K, \tau_T, \sigma^*_{-(K \cup T)}) > U_x(\sigma^*_{-T}, \tau_T)$$

---

<sup>a</sup>While we could have defined a weaker notion of immunity with respect to  $\tau$ , since our focus is on regret freedom and not  $t$ -immunity, we use  $t$ -immunity for simplicity in order to avoid introducing another notion of immunity.

Despite the strong assumption on which they rely,  $(k, t, \tau)$ -type-robust equilibria are impossible to achieve for many Byzantine behaviors. In particular, it follows immediately from Theorem 3.5 that no such equilibrium is possible if the known Byzantine strategy calls for any Byzantine node to crash at the very beginning of the game.

THEOREM 3.11. There exist no non-trivial  $(k, t, \tau)$ -type-robust equilibria in the communication game in which a Byzantine node  $z$ , following  $\tau_z$ , crashes at the beginning of the game.

*Proof.* Same as proof of Theorem 3.5. □

In general, we can show that, if time is discrete (and thus can be divided into periods<sup>2</sup>) and there exists some period after which a Byzantine node becomes “unresponsive,” i.e., the node’s behavior becomes independent of how the game has been played so far (e.g., the node crashes or starts flooding all other nodes with messages), there do not exist any non-trivial  $(k, t, \tau)$ -type-robust equilibria that guarantees that nodes choose the best response at every given point in the game.<sup>3</sup>

---

<sup>2</sup>As the game is not necessarily repeated, we use the word “period” rather than “round” to try to avoid any confusion.

<sup>3</sup>In game theory parlance, there do not exist any  $(k, t, \tau)$ -type-robust equilibrium that rely on credible threats, are subgame-perfect equilibrium, or are perfect Bayes equilibrium.

**THEOREM 3.12.** There exists no non-trivial  $(k, t, \tau)$ -type-robust equilibria in the communication game in which

- Any node  $z$ , following  $\tau_z$ , (a) plays as if it were playing  $\sigma_z^*$  before some period  $r$ , and (b) is unresponsive after period  $r$ ; and
- Every node is always best-responding to every other node.

*Proof.* By contradiction. Assume that  $\sigma^*$  is a non-trivial  $(k, t, \tau)$ -type-robust equilibrium as described in the theorem statement. Consider some period  $r$ , and suppose, for now, there exists some communication between some nodes after period  $r$ . Consider the first period in which, under  $\sigma^*$ , a node  $x$  sends a message to another node  $z$  with positive probability (if there are multiple such nodes, arbitrarily choose one).

Suppose now that  $z$  is Byzantine and is playing  $\tau_z$ . As the equilibrium strategy cannot depend on who the Byzantine nodes are,  $x$ 's choice to send a message must not depend on whether  $z$  is Byzantine. However, if  $z$  turns out to be unresponsive,  $x$  is clearly better off not sending to  $z$  (using the same argument as in the proof of Theorem 3.5).

Suppose then that  $\sigma^*$  does not have any nodes communicating after period  $r$ . We can then prove that in any equilibrium, communication never occurs with positive probability at any period in the game. We prove this using backwards induction on the period; we use  $i$  as the induction variable.

Base case:  $i = r + 1$  (or any period after  $r$ ). As proven above, no node communicates in the equilibrium with positive probability after time  $r$ .

Inductive step. Assume true for all  $i > r_0$ ; we now prove it to be true for period  $r_0$ . If no communication occurs in period  $r_0$ , then we are already done; thus, assume that some rational node  $y$ , following  $\sigma_y^*$ , sends a message at period  $r_0$  with probability  $\alpha > 0$  (again, if there are multiple such nodes, choose one arbitrarily). Let  $h_y^{r_0}$  be some history that leads to this point in the game and let  $\sigma'_y$  be some alternate strategy in which  $y$  does not send anything at

or after period  $r_0$ . By the induction hypothesis, we know that there is no communication that occurs after period  $r_0$ .

Thus, it must be the case that

$$U_y((\sigma'_y, \sigma^*_{-y})|h_y^{r_0}) = U_y(\sigma^*|h_y^{r_0}) + \alpha\gamma > U_y(\sigma^*|h_y^{r_0})$$

contradicting that communicating at period  $r_0$  was in  $y$ 's best interest.  $\square$

Note that the second-half of the proof above can be applied to any equilibrium that requires a best-response at any point in the game, not just  $(k, t, \tau)$ - or  $(k, t)$ -robust equilibrium.

### 3.3 Regret bravery

Finding a single strategy that is a best response against all possible Byzantine strategies or all possible  $t$ -sized subsets of Byzantine nodes (or both) appears fundamentally hard: regret-free solution concepts, for which rational cooperation depends on finding such a strategy, seem unlikely to provide a viable theoretical framework for many BAR-tolerant systems.

*Regret bravery*, the alternative we explore in this section, explicitly forgoes seeking a “universal” best response. Instead, it makes rational cooperation dependent on identifying a strategy that is a best response to the Byzantine behavior that rational nodes *expect* to be exposed to. Before we proceed to look at examples of regret-braving equilibria, we answer some natural questions.

*Is aiming for a best response towards only a subset of all possible Byzantine behaviors in effect abdicating the general claims (and benefits) of Byzantine fault tolerance?* No. Any BAR-tolerant protocol, independent of the underlying solution concept, must be a strategy that guarantees Byzantine fault tolerance. The choice of a solution concept is not about fault tolerance; rather, it specifies under which conditions rational nodes will be willing to follow a given strategy, fault-tolerant or not. Regret-braving solution concepts are motivated by the observation that rational nodes may be willing to cooperate even without the

guarantee that the considered strategy will, in *all* circumstances, prove to be a best response.

*Do regret-braving solution concepts limit how Byzantine node can behave?* No more than a threshold  $t$  on the number of Byzantine faults limits a system to experience, in reality, more than  $t$  faults. Regret braving asks rational nodes to build a model of expected Byzantine behavior, but of course Byzantine failures are in no way bound to follow that model. If Byzantine behavior does not match the expectation of rational nodes, then a regret-braving equilibrium strategy may not, in hindsight, prove to be a best response.

*What is the right set of expectations when it comes to Byzantine behavior?* It all depends on the application being considered. We discuss below two concrete examples inspired by approaches (maximin and Bayes equilibria) that have been extensively studied in the economics literature, but we do not claim that these solution concepts model “realistic” expectations for all distributed systems. For example, the maximin approach produces a best response to the expectation that the system always includes exactly  $t$  Byzantine nodes, when it may instead often be reasonable to expect that the actual number of Byzantine faults will be lower.<sup>4</sup> Indeed, we believe that the challenge of finding equilibrium strategies under more flexible solution concepts is an extremely exciting research opportunity.

**The threshold communication game.** To show the viability of regret-brave solution concepts in a communication game, we consider a concrete communication game: a *threshold game*, which models protocols, such as secret-sharing [98], replicated state machines [74], and terminating reliable broadcast [58] in which functionality is achieved if and only if some sufficiently large subset of peers cooperate with a node.

---

<sup>4</sup>A worst-case attitude is actually not uncommon when designing fault-tolerant systems, even for benign failures. For instance, non-early stopping protocols for synchronous terminating reliable broadcast always run for  $t + 1$  rounds, even in executions that experience no failures.

DEFINITION 3.13. A (synchronous) threshold game is an infinitely-repeated communication game where

- There are at least 3 nodes ( $n \geq 3$ ).
- The game repeats indefinitely. In every round, for every pair of nodes  $x, y \in \mathbf{N}$  such that  $x \neq y$ ,  $x$  decides whether to send a message (“contribute”) or not (“snub”) to  $y$ .
- At the end of the round, every  $x \in \mathbf{N}$  simultaneously (1) observes who contributed to it and (2) receives its payoff.<sup>a</sup>  $x$  incurs a cost of  $\gamma$  for each node  $x$  contributes to and for each node that contributes to  $x$ ;  $x$  incurs no cost for snubbing or being snubbed.  $x$  realizes a positive benefit of  $b > 2n\gamma$  in any round where  $q > 0$  *other* nodes contribute to  $x$ , where  $q$  is the threshold.
- The total payoff is the  $\delta$ -discounted sum of each individual round’s payoff, where  $0 < \delta < 1$ .<sup>b</sup>

---

<sup>a</sup>In game theory parlance, the game is a simultaneous game; in distributed systems, synchronous.

<sup>b</sup> $\delta$ -discounting is a commonly-accepted way of handling utility in infinite-horizon games (see any game theory text, e.g., [54, 87]). This models the reality that earning benefit (incurring cost) now is better (worse) than doing so later. For example, it is often preferable to have a dollar now rather than later, since money can be invested and can earn interest in the meantime.

When up to  $t$  failures can occur, we will generally focus our attention only on threshold games where  $q < n - t$ , i.e., where even  $t$  failures cannot prevent a node from achieving benefit.

In the subsequent sections, we consider two concrete regret-braving solution concepts for the threshold game. We do not claim that these solution concepts are new; instead, our goal is to illustrate that taking a node’s expectations into account (for which we provide two examples) allow us to bypass the difficulties with dealing with Byzantine failures that regret freedom faced.

In a paper we previously published [107], we show how these solution concepts could be extended to explicitly consider collusion via  $k$ -resilience. Because the choice of how Byzantine failures are modeled is orthogonal to how coalitions are modeled and because of recent advances in how coalitions should be dealt with (§5), we omit these details here.

### 3.3.1 $t$ -maximin equilibrium

In the first, rational nodes best-respond to fearing the worst, i.e., they follow a maximin strategy with respect to Byzantine failures. This is the notion of equilibrium used in previous BAR systems (e.g., [24, 79]).

DEFINITION 3.14. The strategy profile  $\sigma^*$  is a  $t$ -maximin equilibrium iff for all  $x \in \mathbf{N}$ , there does not exist an alternate strategy profile  $\sigma'$  such that

$$\min_{\substack{T \subseteq \mathbf{N} \setminus \{x\}: \\ |T| \leq t}} \min_{\tau} U_x(\sigma'_x, \tau_T, \sigma^*_{\{-x\} \cup T}) > \min_{\substack{T \subseteq \mathbf{N} \setminus \{x\}: \\ |T| \leq t}} \min_{\tau} U_x(\sigma^*_{-T}, \tau_T)$$

To demonstrate that  $t$ -maximin yields equilibria in communication games, we show the following simple  $t$ -maximin equilibrium in the threshold game.

THEOREM 3.15. Let the strategy profile  $\sigma^*$  be defined as follows: any  $x \in N$  following  $\sigma^*_x$  contributes to some  $y \neq x$  iff  $x$  and  $y$  have always contributed to each other in the past and  $x$  has been snubbed by at most  $t$  different nodes. Then  $\sigma^*$  is a  $t$ -maximin equilibrium in the threshold game if  $q < n - t$  and

$$\frac{b}{\gamma} \geq \frac{2(n-1) - (1-\delta^2)(q+t)}{\delta^2} \quad (3.2)$$

*Proof.* Consider some rational node  $x$ . If all rational nodes follow  $\sigma^*$ , then each node will receive a threshold number of contributions regardless of Byzantine

behavior and the worst a rational node can do is incur the cost of contributing and being contributed to by every node in the system. Thus, a rational node, following  $\sigma^*$ , will earn at least

$$V^* = \frac{1}{1 - \delta}(b - 2(n - 1)\gamma) \quad (3.3)$$

Suppose instead that in round  $r$ ,  $x$  snubs some node  $y$  that  $x$  was supposed to contribute to after history  $h_x^r$ . If  $h_x^r$  is not “expected” to occur—i.e., given  $\sigma^*$  and the  $T$  and  $\tau$  that minimize  $\sigma^*$ ,  $h_x^r$  occurs with zero probability<sup>5</sup>—then the expected change in utility as a result of snubbing is 0, and the proof is trivially complete. Otherwise, suppose  $h_x^r$  is expected to occur. If  $y$  is rational and  $t$  Byzantine nodes, in addition to  $x$ , snub  $y$  by round  $r$ ,  $y$  will snub all nodes from round  $r + 1$  at latest. If  $y$  and all  $t$  Byzantine nodes snub every node by round  $r + 1$ , then all other nodes snub every node from, at latest, round  $r + 2$  on. Therefore,  $x$  earns at most  $b - (q + t)\gamma$  in rounds  $r$  and  $r + 1$  and 0 in subsequent rounds for a total payoff of

$$V' = (1 + \delta)(b - (q + t)\gamma)$$

Given inequality (3.2),  $V^* \geq V'$  and thus deviating is not worthwhile.

Now consider if  $x$  contributes to some node  $y$  that  $x$ , under  $\sigma^*$ , was supposed to snub after some history  $h_x^r$ . As we just showed, rational nodes never snub unless they were snubbed first. Because  $q < n - t$ , a rational node will be guaranteed to get a threshold number of contributions regardless of how Byzantine nodes play. Thus, the worst damage a Byzantine node can inflict on a rational node  $x$  is to have  $x$  contribute in every round to it and to contribute in every round back. As a result,  $h_x^r$  must occur with zero probability, implying that deciding to contribute after such a history changes the expected utility by 0 and is thus not a profitable deviation.  $\square$

---

<sup>5</sup>This is similar to being “off the equilibrium path” in traditional game theory.

### 3.3.2 Bayes equilibrium

One advantage of using the  $t$ -maximin solution concept is its simplicity: because we need only consider the worst possible case,  $t$ -maximin equilibria are simple to analyze. Unfortunately, although a rational node playing a  $t$ -maximin equilibrium may receive a safe, steady amount of utility, Byzantine failures are unlikely to always occur in the worst possible way, and a rational node willing to take a risk and deviate from the prescribed strategy may be able to do better in expectation.

In this section, we show how a standard solution concept, Bayes equilibrium, can be used to take a node's expectation regarding failure into account by having rational nodes maintain beliefs that represent the probability of various failure scenarios occurring. More specifically, each node  $x$  has some beliefs  $\mu_x((\mathbf{R}_{-T}, \tau_T) | \mathbf{R}_x)$  that represents its belief that all nodes  $z \in T$  are Byzantine and of type (i.e., playing strategy)  $\tau_z$  and all nodes  $y \notin T$  are rational (i.e., of type  $\mathbf{R}$ ), given that  $x$  is rational. For notational simplicity, we denote it as  $\mu_x(\tau_T)$ .

DEFINITION 3.16. The strategy profile/belief tuple  $(\sigma^*, \mu^*)$  is a Bayes equilibrium iff for all  $x \in \mathbf{N}$ , there does not exist an alternate strategy profile  $\sigma'$  such that

$$\mathbf{E}^{(\sigma'_x, \sigma_{-x}^*), \mu^*} [U_x] > \mathbf{E}^{\sigma^*, \mu^*} [U_x]$$

where

$$\mathbf{E}^{\sigma, \mu} [U_x] = \sum_{T \subseteq \mathbf{N} \setminus \{x\}} \sum_{\tau} \mu_x(\tau_T) U_x(\sigma_{-T}, \tau_T)$$

In the remainder of this section, we demonstrate that the Bayesian approach provides flexibility in how Byzantine nodes are modeled by rational nodes by showing equilibria given two different sets of beliefs. Our goal in these examples is to simply illustrate the existence of Bayesian equilibria, not

to derive tight bounds for when these equilibria exist. Thus, for simplicity of exposition, we will be extremely optimistic about the utility earned by deviating and pessimistic about the utility earned by cooperating.

We first show a Bayes equilibrium in a simple scenario that roughly models the one used in the proof of Theorem 3.9: Byzantine nodes are expected to either crash or threaten to inflict communication costs unless rational nodes contribute.

**THEOREM 3.17.** For any  $x \in N$ , let  $T_x^i$  be the set of nodes who have snubbed  $x$  in round  $i$  and let  $t$  be the expected number of Byzantine failures. Define strategy profile  $\sigma^*$  as follows.  $x$ , playing  $\sigma_x^*$ , (a) in round 0, contributes to all nodes, and (b) in round  $r > 0$ , if  $|T_x^0| > t$  or there exists some round  $i < r$  with  $T_x^i \not\subseteq T_x^0$ , then snub all nodes; otherwise, contribute to all nodes in  $N \setminus T_x^0$ .

Let  $\mu^*$  be some set of beliefs which place positive probability only on the following Byzantine strategies: (a) snub everyone (the crash strategy); and (b) snub everyone in the first round, and, in any subsequent round  $r$ , snub a node  $y$  iff  $y$  previously contributed to it.

Let  $\psi$  be the joint probability (based on  $\mu^*$ ) that the environment has exactly  $t$  Byzantine nodes and that a node, picking a peer at random, selects a rational one. Then  $(\sigma^*, \mu^*)$  is a Bayes equilibrium in the threshold game if  $q < n - t$ ,  $\psi > 0$ , and

$$\frac{b}{\gamma} \geq \frac{1}{\delta^2 \psi} (2(n-1) - q(1 - \delta^2 \psi)) \quad (3.4)$$

*Proof.* Consider some rational node  $x$ . As shown in the proof of Theorem 3.15, if  $x$  follows  $\sigma^*$ , the worst  $x$  can do is receive a threshold number of contributions in every round and incur communication costs with everyone, resulting in a payoff of  $V^*$ , where  $V^*$  is defined as in equation (3.3).

Consider the ways that  $x$  may deviate.

Case 1:  $x$  snubs some set of nodes  $L \subset N$  in round 0. We optimistically assume that (a)  $x$  is only hurt if there exists some rational node  $y \in L$  that  $x$  snubbed and there are exactly  $t$  Byzantine nodes, which occurs with probability at least  $\psi$ , and (b)  $x$  earns the maximum round payoff that it can (i.e.,  $b - q\gamma$  if it receives a threshold number of contributions and 0 if not) when deviating. Thus, if  $L$  contains only Byzantine nodes, then  $x$  earns  $b - q\gamma$  in all rounds. Otherwise,  $x$  earns at most  $b - q\gamma$  in rounds 0 and 1 and 0 in all subsequent rounds.  $x$ 's total expected payoff from deviating is then

$$V' = \psi(1 + \delta)(b - q\gamma) + (1 - \psi)\frac{1}{1 - \delta}(b - q\gamma)$$

Inequality (3.4) ensures that  $V^* \geq V'$ , and thus  $x$  does not deviate in this fashion.

Case 2:  $x$  snubs a node that  $x$  is supposed to contribute to in round  $r > 0$ . In any subsequent round  $r > 0$ ,  $x$  knows exactly who the Byzantine nodes are. Thus, the second case, in which a rational node that deviates by snubbing some (rational) node  $y$  that has never snubbed it before, results in  $y$  snubbing everyone in round  $r + 1$ , causing all nodes to snub  $x$  by round  $r + 2$ . In this case,  $x$  earns at most  $(1 + \delta)(b - q\gamma)$  in rounds  $r$  and  $r + 1$  and 0 in all subsequent rounds. The same argument used in the previous case implies that  $x$  does not deviate in this way either.

Case 3:  $x$  contributes to a node that  $x$  is supposed to snub in round  $r > 0$ . Suppose  $x$  contributes to a node  $z$  that  $x$  is supposed to snub in round  $r$ . Regardless of the reason, contributing to  $z$  incurs cost on  $x$  yet does not affect  $z$ 's strategy in any way. It is obvious then that this deviation is never in  $x$ 's best interest.  $\square$

We now demonstrate a Bayes equilibrium in a crash-failure scenario that is similar to that used in many of the theorems in this chapter.

**THEOREM 3.18.** Let  $t$  be the expected number of Byzantine failures. Define the strategy profile  $\sigma^*$  such that any  $x \in N$ , following  $\sigma_x^*$ , contributes to any  $y \neq x$  iff  $x$  and  $y$  have always contributed to each other in the past and  $x$  has been snubbed by at most  $t$  peers.

Let  $\tau$  be defined as the random  $t$ -crash strategy: in any given round, a node  $z$  playing  $\tau_z$  has some positive probability  $\rho$  of crashing. Define the set of beliefs  $\mu^*$  such that for all  $x \in \mathbf{N}$ ,

- $\mu_x^*(\tau_T) = 0$  for any  $T$  such that  $|T| \neq t$ , and
- $\mu_x^*(\tau_{T_1}) = \mu_x^*(\tau_{T_2}) > 0$  for any  $T_1, T_2 \subseteq \mathbf{N} \setminus \{x\}$  such that  $|T_1| = |T_2| = t$ .

Then  $(\sigma^*, \mu^*)$  is a Bayes equilibrium if  $q < n - t$  and

$$\frac{b}{\gamma} \geq \frac{2n - q - 2}{\delta^2 \rho^t} \frac{n - 1}{n - t - 1} + q \quad (3.5)$$

*Proof.* Consider the perspective of some rational node  $x$ . We optimistically assume that if  $x$  deviates in round  $r$ , it only loses utility if  $t$  nodes crash on or before round  $r$ , which occurs with probability at least  $\rho^t$ .

As shown in the proof of Theorem 3.15, if  $x$  follows  $\sigma^*$ , the worst  $x$  can do is receive a threshold number of contributions in every round and incur communication costs with everyone, resulting in a payoff of  $V^*$ , where  $V^*$  is defined as in equation (3.3).

Suppose that  $x$  snubs some node  $y$ . Since the probability that a node is rational is uniform across all nodes,  $y$  is rational with probability at least  $1 - t/(n - 1)$ , and with probability at least  $\rho^t$ ,  $y$  will observe  $t$  other nodes snub it by round  $r$ .  $y$  then snubs everyone starting in round  $r + 1$ , all nodes snub everyone starting in round  $r + 2$ , and  $x$  earns at most 0 in every round starting from round  $r + 2$ . Otherwise, we assume  $x$  earns the maximum round

payoff  $b - q\gamma$ . Thus, deviating results in a payoff of

$$V' = \psi(1 + \delta)(b - q\gamma) + (1 - \psi) \frac{1}{1 - \delta}(b - q\gamma)$$

where  $\psi = \rho^t(1 - t/(n - 1))$ . Given inequality (3.5),  $V^* \geq V'$ ; thus, deviating is not in  $x$ 's best interest.  $\square$

### 3.4 Summary

To formally reason about cooperative services, we need a solution concept that provides rigorous guarantees for rational cooperation without sacrificing real-world applicability. This chapter argues that solution concepts based on regret freedom, despite their intuitive correspondence to the traditional guarantees of fault-tolerant distributed computing, are unlikely to provide the basis for a viable theoretical framework for real-world systems. In particular, we believe that any practical solution concept should be able to admit equilibria in games where a rational node's payoff is not based simply on the outcome but also on the cost of the actions required to achieve said outcome. While our discussion here has focused on communication costs, other costs should be included, such as computational costs [61]. We believe that regret-brave solution concepts provide a rigorous and realistic framework for games that account for these costs.

# Chapter 4

## Acquiescent

In §3, we discussed how incentives in cooperative services must be resilient to faulty peers. However, incentives must also be robust against a more subtle threat: an overabundance of good will from the unselfish peers who simply follow the protocol run by the service. It is, after all, the unselfishness of correct peers—as codified in the protocol they obediently follow—that allows selfish peers to continue receiving service without contributing their fair share. Yet, the efforts of well-meaning peers alone may be insufficient to sustain the service. Further, asking these peers to increase their contribution to make up for free-riders may backfire: even well-meaning peers, if blatantly taken advantage of, may give in to the temptation of joining the ranks of the selfish, leading in turn to more defections and to the service’s collapse.

Although real cooperative services include a sizable fraction of correct and unselfish peers [22], their impact on the incentive structure of cooperative services is not well understood. Existing BAR-tolerant systems have sidestepped the challenge of dealing with acquiescent nodes by designing protocols that neither depend on nor leverage the presence of acquiescent peers.<sup>1</sup>

In this chapter, we ask the following question: *can we leverage the good*

---

<sup>1</sup>Gossip-based BAR-tolerant streaming protocols [78, 79] do rely on a trusted source to seed the stream; however, the nodes we are focusing on are those nodes that participate in the gossip protocol, which are modeled as either rational or Byzantine.

*will of acquiescent nodes and still motivate rational participants to cooperate?*

We find that not only is acquiescence not antithetical to rational cooperation, but that, in a fundamental way, rational cooperation can only be achieved in the presence of acquiescent nodes. To do so, we distill the issue to a rational peer’s *last* opportunity to cooperate.

**The last exchange.** Rational peers are induced to cooperate with another peer (or, more generally, with a service) by the expectation that, if they cooperate, they will receive future benefit. However, in most cases, interaction with a particular peer or with the service itself eventually comes to an end. In this last exchange, rational peers do not have incentive to contribute, as doing so incurs cost without any future benefit. Unfortunately, rational cooperation throughout the protocol often hinges on this critical last exchange: the lack of incentive to cooperate at the end may, in a sort of reverse domino effect, demotivate rational peers from cooperating in *any* prior exchange.

Most current systems address this problem in one of three ways (or some combination of them). Some systems [24, 79] assume that rational peers interact with the service forever, and thus future incentives always exist; others [76, 77, 78] assume rational peers deviate only if their increase in utility is above a certain threshold; others, finally, try to threaten rational peers with the possibility of losing utility if they deviate. For instance, in BAR Gossip [79], peers that do not receive the data they expect *pester* the guilty peer by repeatedly requesting the missing contribution.

Unfortunately, each of these approaches relies on somewhat unrealistic assumptions. Few relationships in life are infinite in length; worse, as we will show later in this paper, with a lossy network and the possibility of Byzantine peers it may be impossible to incentivize cooperation even in an infinite-length protocol. The real possibility of penny-pinching peers can undermine any system that assumes no deviation unless their expected gain is “large enough.” Finally, threats such as pestering are effective only when they are credible: to feel threatened, a peer must believe that it will be rational for the other

peer to pester. Since pestering incurs cost for the initiator as well as for the receiver, it is surprisingly hard to motivate rational peers to pester in the first place. For example, pestering in BAR Gossip is credible only under the rather implausible assumption that a peer, even when faced with enduring silence, will never give up on an unresponsive peer and forever continue to attribute a peer’s lack of contribution to the unreliability of the network [79].

We model the last-exchange problem as a finite-round game between two players 1 and 2; neither player expects to interact with the other beyond this exchange. We assume player 1 holds a contribution (e.g., some information) that is of value to player 2; however, contributing yields no expectation of further benefit for player 1. We are interested in studying whether player 2 can nonetheless induce a selfish player 1 to contribute by threatening to pester it if player 1 fails to do so. Pestering is an attractive threat because it is simple and does not require the involvement of a third party. We want to determine whether it can be made a credible threat under realistic system assumptions, unlike in BAR Gossip [79]. In each round, player 1 is given a choice whether to contribute or not; in response, player 2 may pester player 1. Players communicate through a lossy channel and therefore do not necessarily share the same view of the ongoing game. For instance, player 1 may have contributed, but player 2 may not have received the contribution.

**Our contributions.** We show that, without requiring implausible network assumptions or the specter of never-ending pestering, the presence of acquiescent peers is both necessary and sufficient to make pestering a credible threat and motivate rational peers to contribute. In particular:

- We prove that there exists no equilibrium strategy where rational peers contribute if all peers are either rational or Byzantine—even if we allow for an infinite number of pestering rounds.
- We show that the presence of acquiescent peers is sufficient to transform pestering into a credible threat. Intuitively, if rational peers have

sufficiently high beliefs that they may be interacting with an acquiescent peer, they are motivated to pester, making it in turn preferable for rational peers to contribute.

The fraction of acquiescent peers sufficient to sustain rational contribution depends on several system parameters, including the probability of network loss, the fraction of Byzantine peers in the system, and the behavior that rational peers expect from acquiescent and Byzantine peers. Exploring this space through a simulator we find that:

- Acquiescent peers make rational cooperation easy to achieve under realistic conditions. In particular, we find that even if less than 10% of the population is acquiescent, rational peers are incentivized to cooperate in a system where the network drops 5% of all packets and Byzantine peers make up over 50% of the remainder of the population.
- Prodigious acquiescent peers do harm rational cooperation: if acquiescent peers contribute every time they are pestered, then we cannot always achieve rational cooperation; when we do, it requires an implausibly high fraction of acquiescent peers. This is good news: the less foolishly generous is the acquiescent behavior sufficient to incentivize rational contribution, the more feasible it is to design systems with a sustainable population of acquiescent peers.
- The uncertainty introduced by network loss is both a bane and a boon. On the one hand, it significantly complicates the analysis of a peer's optimal strategy because each peer does not know what the other has observed. On the other, it lowers the threshold for rational cooperation by leaving open some possibility that the other peer may be acquiescent, even when the observed behavior suggests otherwise.

**Organization of chapter.** After presenting in §4.1 the game theoretic framework used to analyze the last exchange problem, we show in §4.2 that rational

cooperation is impossible in the absence of acquiescent peers. We proceed to derive, in §4.3, conditions under which acquiescence is sufficient to elicit rational cooperation in the last exchange and, in §4.4, use simulations to study the implications of these conditions on the design of cooperative services.

## 4.1 Setup

In this chapter, we consider cooperative services that can be modeled as a collection of peer-to-peer pairwise exchanges, in which two players communicate over unreliable channels. In particular, we focus solely on the last exchange between these two players and are interested in studying the conditions under which one player can induce another to contribute with the threat of pestering and without any exogenous incentives or entities such as a server.

We model this last exchange as a two-player,  $(R + 1)$ -round stochastic sequential game, which is similar to a repeated game except that it allows players' payoffs to change as the game progresses. This flexibility is critical to model the intuition that player 2 benefits from player 1's contribution only the first time player 2 receives it. In each round, player 1 moves first by choosing between two actions: contribute (denoted by  $c$ ) or do nothing ( $o$ ). Player 2 follows by choosing between two actions: pester ( $p$ ) or do nothing ( $o$ ). Since our analysis of the game often relies on the number of rounds remaining rather than on the round number, we label the first round as round  $R$  and the last round as round 0.

Doing nothing has neither cost nor benefit. Player 1 incurs a cost  $\gamma_{\uparrow c}$  in every round in which it contributes and a cost  $\gamma_{\downarrow p}$  in every round in which it is pestered; player 2 incurs a cost  $\gamma_{\downarrow c}$  in every round it receives a contribution and a cost  $\gamma_{\uparrow p}$  in every round it pesters player 1.<sup>2</sup> A non-Byzantine player 2 starts off being *destitute*, i.e., player 2 does not have the contribution that player 1 can provide. A destitute player 2 receives a one-time benefit  $b \gg \gamma_{\downarrow c} + \gamma_{\uparrow p}$  the

---

<sup>2</sup>One easy way to remember which symbol corresponds to what is to remember that  $\uparrow$  corresponds to uploading (or sending) and  $\downarrow$  corresponds to downloading (or receiving).

first time it receives player 1's contribution; now no longer destitute, player 2 gains no further benefit from receiving further copies of the contribution.

**Network loss, signals, and utilities.** To model the unreliable channel through which players 1 and 2 communicate, we adopt from game theory the concept of *private signals*: for every action  $a$  played by some player, both players privately observe some (possibly different) resulting *signal*. Specifically, let  $\rho$ ,  $0 < \rho < 1$ , be the rate of network loss, which we assume to be common knowledge. When player  $x$  plays  $a$ , player  $x$  observes  $a$ , and player  $-x$  observes  $a$  with probability  $1 - \rho$  and  $o$  otherwise. Thus, players do not always observe their peer's actions accurately and cannot rely on their peer accurately observing their own actions.

As mentioned in §2, a node's history describes what a player has performed or observed in the past. In this chapter, we denote the history that results from the sequence of signals observed by player  $x$  until player  $y$ 's turn in round  $r$  as  $h_x^{r,y}$ . At the beginning of the game,  $h_x^{R,1}$  is the empty sequence. During player 1's turn in round  $r < R$ , player  $x$ 's history  $h_x^{r,1}$  is obtained by appending player  $x$ 's previous history  $h_x^{r+1,2}$  with some signal  $\omega_x^{r+1,2}$  observed by player  $x$  corresponding to player 2's action in round  $r+1$ :  $h_x^{r,1} = (h_x^{r+1,2}, \omega_x^{r+1,2})$ . Similarly, during player 2's turn in round  $r \leq R$ , player  $x$ 's history  $h_x^{r,2}$  is obtained by appending  $h_x^{r,1}$  with some signal  $\omega_x^{r,1}$  observed by player  $x$  corresponding to player 1's action in round  $r$ :  $h_x^{r,2} = (h_x^{r,1}, \omega_x^{r,1})$ . For simplicity, we drop the second superscript (e.g., the  $y$  in  $h_x^{r,y}$  and  $\omega_x^{r,y}$ ) when it is obvious whose turn it is. We use the subscript to denote different histories (e.g.,  $h_1^r$ ,  $h_{1,c}^r$ ) and the superscript to denote the prefix of a history (e.g., for  $i \geq r$ ,  $h_1^{i,x}$  is the first  $i$  rounds of  $h_1^r$ , including player 1's signal in round  $i$  if  $x = 2$ ).

We define  $u$  to be a mapping from the signals that a player observes to

Symbol	Meaning
$\gamma_{\uparrow p}$	Cost of sending a pester
$\gamma_{\downarrow p}$	Cost of receiving a pester
$\gamma_{\uparrow c}$	Cost of contributing
$\gamma_{\downarrow c}$	Cost of receiving a contribution
$b$	Benefit of receiving a contribution

**Table 4.1:** Summary of symbols that define the payoffs of various actions.

the payoff they receive:

$$\begin{aligned}
 u_1(C, \hat{P}) &= - \left( |\hat{P}| \gamma_{\downarrow p} + |C| \gamma_{\uparrow c} \right) \\
 u_2(P, \hat{C}) &= H[|\hat{C}| - 1] b - \left( |P| \gamma_{\uparrow p} + |\hat{C}| \gamma_{\downarrow c} \right)
 \end{aligned}$$

where  $C$  and  $\hat{P}$  are the sets of rounds in which player 1 respectively contributed and observed player 2 pester;  $P$  and  $\hat{C}$  are the sets of rounds in which player 2 respectively pestered and observed player 1 contribute; and  $H[n]$  is the unit step function.<sup>3</sup> We can define the utility function  $U$  as a function of the players' strategies, but we stick to  $u$  in this chapter for simplicity.

**Strategies, types, beliefs, and equilibrium.** As noted earlier, our game is stochastic: the payoffs of a non-Byzantine player 2 change depending on whether it is destitute (and wants player 1 to contribute) or not (and wants player 1 to do nothing). In this chapter, for convenience, in addition to **A** and **R** denoting acquiescent and rational types (§4.1), we abuse notation and introduce two additional types, **D** and  $\neg\mathbf{D}$ , to characterize the state of a non-Byzantine player 2, depending on whether or not it is destitute. If player 1 contributes, then with probability  $(1 - \rho)$ , a player 2 of type **D** observes  $c$  and hence change to type  $\neg\mathbf{D}$ .

In this chapter, we model rational players as entities that maximize their payoffs ex-ante. As described in §2, every player  $x$  starts with some

---

<sup>3</sup> $H[i] = 0$  if  $i < 0$ ; else  $H[i] = 1$ .

initial beliefs  $\mu_x(\theta)$  representing the probabilities that player  $x$  assigns to the statement that player  $-x$  is of type  $\theta$ .<sup>4</sup> We assume that, for all  $x \in \{1, 2\}$ ,  $\mu_x(\theta)$  equals an initial value  $\mu(\theta)$ , which is common knowledge, and that the beliefs of a rational player  $x$ 's evolve based on the history  $h_x^r$  it has observed; we use  $\mu_x(\theta|h_x^r)$  to denote player  $x$ 's conditional beliefs.

For a given set of beliefs, a rational player's strategy depends on the specific strategy that it expects its peer to adopt—which, in turn, depends on the peer's type. A rational player expects an acquiescent player  $x$ 's strategy  $\sigma_{\mathbf{A},x}$  to be identical to the initially assigned protocol and a rational player to follow  $\sigma_{\mathbf{R},x}$  (assuming that it is a best response). A player's strategy may depend on its observed history  $h_x^r$ ; we use  $\sigma_{\theta,x}(a|h_x^r)$  to denote the conditional probability that  $a$  is played by player  $x$  of type  $\theta$  given  $h_x^r$ .

If player  $x$  is Byzantine, its strategy can in principle be arbitrary, significantly complicating the task of identifying a rational player's best response ex-ante. In this chapter, we address this difficulty by limiting our attention to a particular set of beliefs that a rational player can hold vis-à-vis Byzantine behaviors: we assume that a rational player does not expect to be able to influence a Byzantine peer's strategy through its actions, and the rational beliefs we consider only put positive probability on Byzantine strategies of this form. More formally, a rational player  $-x$  expects to observe a Byzantine peer player  $x$  do nothing in round  $r$  with some probability  $\tau_x(o|h_{-x}^r) \geq \rho$  that at most depends on player  $x$ 's signals in player  $-x$ 's current history  $h_{-x}^r$ . While this restriction sacrifices the generality of our results—our results do not hold if a rational player believes it can influence Byzantine failures<sup>5</sup>—we believe it captures a large and realistic set of beliefs that models the reasonable distrust that a rational player is likely to harbor towards a Byzantine peer's threats and promises.

As a player only cares about the expected Byzantine strategy, we denote

---

<sup>4</sup>Technically, a player's belief is also a function of its own type, i.e.,  $\mu_x(\theta_{-x}|\theta_x)$ , but since we are only considering the beliefs of a single-type rational player, we simplify our notation.

<sup>5</sup>Note that this restriction is on a rational player's beliefs and expectations. A Byzantine failure may still occur in ways that defy a rational player's expectation (as in §3.3).

the Byzantine type as  $\mathbf{B}$  and, for notational consistency, denote  $\tau_x$  as  $\sigma_{\mathbf{B},x}$ . This allows us to use  $\sigma_x = \{\sigma_{\mathbf{B},x}, \sigma_{\mathbf{A},x}, \sigma_{\mathbf{R},x}\}$  to denote the strategies that a rational player expects player  $x$  to adopt, depending on player  $x$ 's type;  $\sigma = (\sigma_1, \sigma_2)$  to denote the strategy profile that describes the (expected) strategies for players 1 and 2; and  $\mu = (\mu_1, \mu_2)$  to denote the belief profile that describes the beliefs  $\mu_1$  and  $\mu_2$  held by rational players 1 and 2.

In this chapter, we are interested in perfect Bayes equilibrium: a strategy profile and set of beliefs  $(\sigma^*, \mu^*)$  such that for all  $x \in \{1, 2\}$ ,  $\mu_x^*(\theta|h_x^r)$  is computed using Bayes rule whenever  $h_x^r$  is reached via a signal that may be observed with positive probability; and for all histories  $h_x^r$  and strategies  $\sigma'_{\mathbf{R},x}$ :

$$\mathbb{E}^{(\sigma_{\mathbf{R},x}^*, \sigma_{-x}^*), \mu^*} [u_x | h_x^r] \geq \mathbb{E}^{(\sigma'_{\mathbf{R},x}, \sigma_{-x}^*), \mu^*} [u_x | h_x^r]$$

where  $\mathbb{E}^{(\sigma_{\mathbf{R},x}, \sigma_{-x}^*), \mu^*} [u_x | h_x^r]$  is a rational player  $x$ 's expected utility from playing  $\sigma_{\mathbf{R},x}$  with beliefs  $\mu_x^*$ , with both strategy and beliefs conditional on  $h_x^r$ , while its peer player  $-x$  plays  $\sigma_{-x}^*$ . To lighten the already substantial notation, we drop  $\mu^*$  when it is obvious that we are referring to that particular set of beliefs and almost always refer to  $\sigma_{\mathbf{R},x}$  as  $\sigma_x$  unless it is not obvious we are referring to the rational strategy.<sup>6</sup>

Finally, we often refer to a player  $x$ 's expected utility given that its peer  $-x$  is of a specific type  $\theta$ . We denote this expected utility as  $\mathbb{E}^\sigma [u_x | \theta]$  and the expected continuation utility given some history  $h_x^r$  as  $\mathbb{E}^\sigma [u_x | \theta, h_x^r]$ .

We assume that all players are limited to actions in the strategy space. This can be accomplished in practice if actions outside of the strategy space generate a proof of misbehavior [24, 59] and if the associated punishments (e.g., financial penalties) are sufficient to deter rational players. Finally, we assume that a rational player 1 does not try to avoid pestering by severing its network connection: if losing a fraction of bandwidth from pestering is undesirable, disconnecting and losing all of it is even less desirable.

---

<sup>6</sup>This allows us to rewrite the best-response condition as  $\mathbb{E}^{\sigma^*} [u_x | h_x^r] \geq \mathbb{E}^{(\sigma'_x, \sigma_{-x}^*)} [u_x | h_x^r]$ .

## 4.2 The need for acquiescence

Acquiescence is not only sufficient to incentivize cooperation, it is necessary. In this section, we assume that there are no acquiescent nodes, and we show that, as a result, rational players never pester or contribute.

LEMMA 4.1. There exists no equilibrium where a rational player 2 pesters with any positive probability if a rational player 1 will not subsequently contribute.

*Intuition.* Player 2 incurs cost by pestering—with no chance of future contribution from player 1.

*Proof.* Suppose such an equilibrium  $(\sigma^*, \mu^*)$  exists where after some history  $h_2^r$ , player 2 pesters with probability  $\psi > 0$  during some round  $r$ , but player 1 will never subsequently contribute. Consider an alternate strategy  $\sigma'_2$  in which player 2 plays exactly as in  $\sigma_2^*$  until round  $r$ , after which player 2 never pesters again. Since player 1 will never contribute again, then

$$\begin{aligned} \mathbb{E}^{\sigma^*}[u_2|h_2^r] &\leq \psi(-\gamma_{\uparrow p} + \mathbb{E}^{\sigma^*}[u_2|(h_2^r, p)]) + (1 - \psi)(\mathbb{E}^{\sigma^*}[u_2|(h_2^r, o)]) \\ &< 0 = \mathbb{E}^{(\sigma'_2, \sigma_1^*)}[u_2|(h_2^r, o)] \end{aligned}$$

Following  $\sigma'_2$  instead of  $\sigma_2^*$  improves player 2's utility. Contradiction.  $\square$

LEMMA 4.2. There exists no equilibrium where a rational player 1 contributes with any positive probability if a rational player 2 will not subsequently pester.

*Intuition.* Player 1 incurs cost by contributing, yet there is no threat of pestering from player 2.

*Proof.* Suppose such an equilibrium  $(\sigma^*, \mu^*)$  exists where, after some history  $h_1^r$ , player 1 contributes with probability  $\psi > 0$  during some round  $r$ , but

player 2 will never subsequently pester. Consider an alternate strategy  $\sigma'_1$  in which player 1 plays exactly as in  $\sigma_1^*$  until round  $r$ , after which player 1 never contributes again. Since player 1 is not pestered again from round  $r$  on,

$$\begin{aligned} \mathbb{E}^{\sigma^*}[u_1|h_1^r] &\leq \psi(-\gamma_{\uparrow c} + \mathbb{E}^{\sigma^*}[u_1|(h_1^r, c)]) + (1 - \psi)\mathbb{E}^{\sigma^*}[u_1|(h_1^r, o)] \\ &< \mathbb{E}^{(\sigma'_1, \sigma_2^*)}[u_1|(h_1^r, o)] \end{aligned}$$

Following  $\sigma'_1$  instead of  $\sigma_1^*$  improves player 1's utility. Contradiction.  $\square$

**THEOREM 4.3.** There exists no equilibrium in which rational players 1 and 2 contribute and pester, respectively.

*Proof.* Suppose such an equilibrium  $(\sigma^*, \mu^*)$  exists. Then there exists some rounds  $r_c$  and  $r_p$  such that players 1 and 2 contribute and pester, respectively, with some positive probability for the last time. By Lemma 4.2, a rational player 1 never contributes after round  $r_p$  and so  $r_c \geq r_p$ .<sup>7</sup> However, by Lemma 4.1, a rational player 2 only pesters until round  $r_c + 1$ ; thus,  $r_p > r_c$ . Finally, it is obvious that contributing or pestering to a Byzantine peer is never in either player's best interest, since a Byzantine peer will play a strategy independent of what the player does. Contradiction.  $\square$

Theorem 4.3 only holds when the game lasts for a finite number of rounds. When there exists no bound on the number of rounds, a weaker, yet in practice still crippling, result holds. We summarize a simplified version of the main result here; the model of the infinitely-repeated game and details of the results are in §4.6.

**THEOREM 4.4.** In the infinitely-repeated game, suppose a non-destitute player 2 always prefers to do nothing and rational players expect that there exists some positive fraction of Byzantine peers that either (a) when

---

<sup>7</sup>Recall that, in this chapter, we count rounds in reverse.

playing as player 1, never contributes; or (b) when playing as player 2, plays the same strategy played by a destitute rational player 2. Then there exists no finitely mixed equilibrium<sup>a</sup> in which rational players 1 and 2 contribute and pester, respectively.

<sup>a</sup>In other words, an equilibrium in which only a finite number of histories can be reached with positive probability

*Proof.* (Sketch) If some Byzantine player 2 pesters as if playing the destitute rational strategy, despite player 1's contributions, then player 1's belief that player 2 is Byzantine eventually grows arbitrarily close to 1. Similarly, if a Byzantine player 1 never contributes despite player 2's incessant pestering, then player 2 becomes increasingly certain player 1 is Byzantine. It can be shown that a player's belief in its peer being Byzantine eventually grows sufficiently high such that the expected utility of contributing (in the first case) or pestering (in the second) is lower than that of doing nothing. By showing a bound of the number of rounds in which a rational player 1 contributes or player 2 pesters, it follows, using an argument similar to the finitely-repeated game, that this bound must be 0.  $\square$

### 4.3 Acquiescence to the rescue

We now show that acquiescence is sufficient to incentivize rational peers to, respectively, pester and contribute by constructing a cooperative strategy profile and proving that it is an equilibrium. We start by specifying the *acquiescent strategy*.

DEFINITION 4.5. The *acquiescent strategy*  $\sigma_{\mathbf{A}}^*$  is the following:

- $\sigma_{\mathbf{A},1}^*$ : Player 1 contributes during round  $R$ . During round  $r < R$ , player 1 contributes, only if pestered in the previous round  $r + 1$ ,

with probability  $(1 - \alpha)/(1 - \rho)^2$ , where  $\alpha$  is a known parameter such that  $0 < (1 - \alpha)/(1 - \rho)^2 \leq 1$ .<sup>a</sup>

- $\sigma_{\mathbf{A},2}^*$ : For any round  $r > 0$ , player 2 pesters if and only if it is destitute.

<sup>a</sup>Hence, if player 2 pesters an acquiescent player 1 during round  $r$ , player 2 expects to observe a contribution in round  $r - 1$  with probability  $1 - \alpha$ .

In practice, all players are initially given the acquiescent strategy. Although we cannot guarantee that a rational player 1 will follow  $\sigma_{\mathbf{A},1}^*$ , we prove that, under the expectation that its peer player  $-x$  of type  $\theta$  plays  $\sigma_{\theta,-x}^*$ , a rational player  $x$  will play the following *rational strategy*.

DEFINITION 4.6. The *rational strategy*  $\sigma_{\mathbf{R}}^*$  is the following:

- $\sigma_{\mathbf{R},1}^*$ : During round  $r$ , player 1 contributes if and only if being pestered is sufficiently expensive to overcome the cost of contributing, i.e.,  $r > \bar{r}$ , where

$$\bar{r} = \left\lceil \frac{1}{(1 - \rho)^2} \frac{\gamma_{\uparrow c}}{\gamma_{\downarrow p}} \right\rceil - 1 \quad (4.1)$$

player 1 knows that player 2, if non-Byzantine, is destitute; and player 1's belief that player 2 is destitute exceeds some threshold  $\bar{\mu}_1^r$ .

- $\sigma_{\mathbf{R},2}^*$ : Same as  $\sigma_{\mathbf{A},2}^*$ .

We prove that rational players follow  $\sigma_{\mathbf{R}}^*$  under the following set of assumptions:

ASSUMPTION 4.7. A destitute rational player 2 always prefers pestering a known acquiescent player 1:

$$\alpha \leq 1 - \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}} \quad (4.2)$$

ASSUMPTION 4.8. For all histories  $h_2^r$  that a rational player 2 may observe which do not contain a contribution, its belief  $\mu_2^*(\mathbf{A}|h_2^r)$  that player 1 is acquiescent satisfies the following condition:

$$\mu_2^*(\mathbf{A}|h_2^r) > \frac{\gamma_{\uparrow p}}{\alpha^{r-1}(1-\alpha)(b-\gamma_{\downarrow c}) + (1-\alpha^{r-1})\gamma_{\uparrow p}} \quad (4.3)$$

ASSUMPTION 4.9.

$$R < \frac{1-\rho+\rho^2\gamma_{\uparrow c}}{\rho^2(1-\rho)^2\gamma_{\downarrow p}} \quad (4.4)$$

For consistency, we denote the Byzantine strategy as  $\sigma_{\mathbf{B}}^*$ .

In a perfect Bayes equilibrium, whether a rational player deviates or not depends on its beliefs for all histories, both those on and off the equilibrium path. In our desired equilibrium, almost every history has some positive probability of being observed: a rational player 2 expects that an acquiescent player 1 contributes with positive probability (if player 2 pestered); destitute player 2 always pester; and, as a result of network loss, doing nothing is always observable with positive probability from either player. Thus, for most histories, Bayes' rule can be applied to calculate a rational player's beliefs.

However, there are still signals that are never expected from a rational or acquiescent player; whether they are observable from a Byzantine peer depends on the expected Byzantine strategy. In this section, all sets of beliefs that we consider use the following set of rules to specify beliefs on and off the equilibrium path:

DEFINITION 4.10. Our rules for updating beliefs are as follows:

- For any action that is expected as a result of  $\sigma^*$ , update using Bayes rule.
- Player 2 observes contribution in round  $r < R$  from player 1 unex-

pectedly. We assume at this point that player 2 believes player 1 to be Byzantine with probability 1.

- Player 1 observes pestering in round 0. At this point, the game is over.

Note that the second case has no effect on player 2's strategies, since a non-destitute player 2, following  $\sigma_{\mathbf{A},2}^*$  or  $\sigma_{\mathbf{R},2}^*$ , does nothing for the remainder of the game anyway.

It is possible for player  $x$  to reach off-equilibrium information sets as a result of its own play. While player  $x$ 's signals themselves do not affect player  $x$ 's belief about player  $-x$ , player  $-x$ 's response to said deviation does. We simply update a player's beliefs using the rules above.

Generally, there may exist multiple strategies that result in a cooperative equilibrium. All our lemmas and theorems here apply only to our particular cooperative equilibrium (and thus, all our results should be prefaced by, "In our cooperative equilibrium. . ."). We believe that our rational strategy represents a sensible design point: incentivizing a rational player 1 to contribute in every round would require player 1 to start with an unrealistically low belief in player 2 being Byzantine. Fortunately, this is unnecessary: we show in §4.4 that the rational strategy results in player 1 often contributing multiple times.

We will proceed as follows. For clarity of exposition, we defer the detailed formal results and proofs to §4.3.3. We instead provide a summary of why it is in a rational player 2's best interest to pester in §4.3.1 and why it is in a rational player 1's best interest to contribute in §4.3.2. This will subsequently allow us to prove the following main result:

**THEOREM 4.11.** The strategy profile and set of beliefs  $(\sigma^*, \mu^*)$  make up a perfect Bayes equilibrium.

*Proof.* See page 81.

We also prove the following, which gives the rational player 1’s strategy a flavor of the acquiescent player 1’s strategy. We find that player 1 effectively only contributes in some round  $r$  when pestered in previous round  $r + 1$ . The only exception is if player 1 “accidentally” did nothing in round  $r + 1$  when contributing was the best response.

**THEOREM 4.12.** Let  $r < R$  be the current round, and suppose that player 1 played its best-response action in round  $r + 1$ . Then, if player 1 observes player 2 do nothing in round  $r + 1$ , then player 1 does nothing in round  $r$ .

*Intuition.* The belief that player 2 is destitute is strictly non-decreasing when player 1 observes player 2 do nothing, and the number of expected pesters also drops as the number of remaining rounds decreases. Thus, if it was in player 1’s best interest to do nothing in the prior round, then player 1 is better served by not contributing in the current round as well.

*Proof.* See page 81.

### 4.3.1 When does a rational player pester?

In this section, we consider the incentives of a rational player 2 and its choice of actions.<sup>8</sup> Intuitively, a rational player 2 pesters only if it is destitute *and* believes that player 1 is sufficiently acquiescent (and thus willing to contribute, even in the final rounds). Naturally, player 2’s strategy also depends on a rational player 1’s strategy; in this section, we assume a rational player 1 plays  $\sigma_{\mathbf{R},1}^*$ .

We start by making two simple observations that are easy to prove.

**LEMMA 4.13.** If (a rational) player 2 is non-destitute, player 2 does nothing.

---

<sup>8</sup>Thus, all our results in this section apply only to a rational player 2.

*Intuition.* If player 2 already has the contribution, player 2 receives no further benefit from receiving another contribution. In fact, pestering and receiving another contribution only incurs cost.

*Proof.* See page 61.

We now show that player 2 is no less likely to get a contribution from a rational player 1 if player 2 pesters more frequently.

LEMMA 4.19. Player 2 is as likely to receive a contribution from a rational player 1 if player 2 pesters instead of doing nothing.

*Intuition.* It is obvious that an acquiescent player 1 is more likely to contribute if pestered, and a Byzantine player 1 will contribute independent of what player 1 does. Finally, since a rational player 1 expects a destitute player 2 to always pester whereas Byzantine player 2 may not, a rational player 1 is more convinced that player 2 is destitute if pestering is observed.

*Proof.* See page 70.

From these lemmas (along with others; see §4.3.3), we can prove that player 2 pesters in every round  $r > 0$ .

THEOREM 4.20. A destitute player 2 pesters in all rounds  $r > 0$ .

*Intuition.* Since a Byzantine player 1 is expected to play independently, an acquiescent player 1 will contribute only when pestered, and a rational player 1 is never discouraged from contributing by pestering (by Lemma 4.19), a destitute player 2 is willing to incur the minor cost of pestering to get a contribution.

*Proof.* See page 71.

### 4.3.2 When does a rational player contribute?

We now consider the incentives of a rational player 1. In every round, player 1 must make a choice:

- Pay the cost of contributing now ( $\gamma_{\uparrow c}$ ), hoping to stop a non-Byzantine player 2 from pestering in the future. The savings are a function of the remaining rounds and the beliefs about player 2.
- Delay contributing, at the risk of being pestered (with cost at most  $(1 - \rho)\gamma_{\downarrow p}$ ), hoping to glean more about player 2's type.

Procrastination has its lure. Since we are considering strategies where a non-Byzantine player 2 always pesters (minus the last round) whereas a Byzantine player 2 may not, every additional signal can drastically affect player 1's expected utility and possibly save player 1 the cost of contributing. Moreover, doing nothing now does not preclude player 1 from contributing in the future.

In this section, we find that if player 1 has sufficiently strong belief that player 2 is destitute, procrastination is something best put off until tomorrow: for every round sufficiently removed from the end of the game, there exists a belief threshold above which contributing yields a higher expected utility for player 1. We prove this under the assumption that a rational player 2 will play  $\sigma_{\mathbf{R},2}^*$ .

We start off our proof that the rational strategy is in a rational player 1's best interest by proving an obvious result: player 1 never contributes when the threat of pestering does not offset the cost of contributing.

LEMMA 4.21. Player 1 does nothing for rounds  $r \leq \bar{r}$ , where  $\bar{r}$  is defined as in condition (4.1).

*Proof.* See page 73.

We can further show that player 1 is never better off trying to contribute unless player 1 knows for certain that player 2, if non-Byzantine, is destitute.

LEMMA 4.14. Let  $r < R$  be the current round,  $h_1^r = h_1^{r,1}$  be the current history during player 1's turn in round  $r$ , and  $\mu_1$  be player 1's beliefs. If  $\mu_1(\neg \mathbf{D} | h_1^r) > 0$ , i.e., player 1 has contributed in the past and has not been pestered since, then player 1 does nothing in round  $r$ .

*Intuition.* Given a sufficiently low  $\rho$ , player 1 is better off not being overly anxious: player 1's contribution has a sufficiently high probability of reaching player 2, so player 1 is thus better off waiting for a definitive signal that player 2 is still destitute (if non-Byzantine).

*Proof.* See page 61.

Given these lemmas (along with other lemmas in §4.3.3), we can show that playing  $\sigma_{\mathbf{R},1}$  is in player 1's best interest.

THEOREM 4.23. Let  $h_1^r$  be the history of player 1 in some round  $r$  such that  $\bar{r} < r \leq R$ , where  $\bar{r}$  is defined by condition (4.1), and let  $\mu_1$  be player 1's beliefs such that  $\mu_1(\neg \mathbf{D} | h_1^r) = 0$ . Then there exists some threshold  $\bar{\mu}_1^r$  such that if  $\mu_1(\mathbf{D} | h_1^r) \geq \bar{\mu}_1^r$ , player 1 contributes; otherwise, player 1 does nothing.

*Intuition.* Player 1, who has sufficiently strong belief that its peer is destitute to contribute, will contribute if its belief is even stronger.

*Proof.* See page 75.

### 4.3.3 Complete formal results

In this subsection, we provide full proofs of the results we have in §4.3.1 and §4.3.2. As a reminder to the reader, as in any typical proof of equilibrium, we will assume that, when considering a player  $x$ 's behavior, player  $x$  is rational and believes player  $-x$  is following the specified strategy.

LEMMA 4.13. If (a rational) player 2 is non-destitute, player 2 does nothing.

*Proof.* Consider some alternate strategy  $\sigma'_2$  in which player 2 pesters with some probability  $\psi > 0$  after observing some history  $h_2^r$  which contains at least one  $c$ . Since  $\sigma_2^*$  specifies player 2 to do nothing starting from history  $h_2^r$ , player 2's expected difference in utility between  $\sigma_2^*$  versus  $\sigma'_2$  is at least

$$\mathbb{E}^{\sigma_2^*}[u_2|h_2^r] - \mathbb{E}^{(\sigma'_2, \sigma_1^*)}[u_2|h_2^r] \geq \psi\gamma_{\uparrow p} > 0$$

and thus player 2 does strictly better following  $\sigma_2^*$ . □

LEMMA 4.14. Let  $r < R$  be the current round,  $h_1^r = h_1^{r,1}$  be the current history during player 1's turn in round  $r$ , and  $\mu_1$  be player 1's beliefs. If  $\mu_1(\neg \mathbf{D}|h_1^r) > 0$ , i.e., player 1 has contributed in the past and has not been pestered since, then player 1 does nothing in round  $r$ .

*Proof.* By induction on  $r$ , i.e., backwards induction on time.

Base case:  $r = 0$ . If player 1 does nothing, it earns 0 since player 2 never pesters in the last round. If it contributes, it earns at most  $-\gamma_{\uparrow c} < 0$ .

Inductive step. Assume true for all rounds  $r \leq r_0$ ; we prove the inductive step ( $r = r_0 + 1$ ) by contradiction. Let  $m_c$  and  $m_p$  be the number of rounds elapsed since player 1 has last contributed or observed pestering, respectively, i.e., player 1 has done nothing in the past  $m_c$  rounds and observed player 2 do nothing in the past  $m_p$  rounds. By assumption,  $m_p > m_c \geq 0$ . Thus,  $h_1^r$  is of the form:

$$h_1^r = (h_1^{r+m_c+1,1}, c, o, o, o, \dots, o, o)$$

where  $h_1^{r+m_c+1,1}$  is some history from round  $r + m_c + 1$ .

Let  $\sigma_{1,c}$  be some strategy such that player 1 contributes after some history  $h_1^r$ , i.e., a history in which player 1 was not pestered since its last contribution. Given a type- $\theta$  player 2, let  $V(\theta)$  be the continuation payoff starting from player 2's turn in round  $r$ , i.e., after player 1 contributes in round  $r$ . Construct an alternate strategy  $\sigma_{1,o}$  where player 1 does nothing in round  $r$  but for the remaining rounds is identical to  $\sigma_{1,c}$  (i.e., following  $\sigma_{1,c}$  as if player 1 had contributed in round  $r$ ). Thus, the continuation payoff from playing  $\sigma_{1,o}$  after doing nothing in round  $r$  is also  $V(\theta)$ . Note that  $V(\neg\mathbf{D}) = 0$  since a non-destitute player 2 never pesters, and, by the inductive hypothesis, player 1 never contributes in later rounds unless player 1 has been pestered since its last contribution.

Following  $\sigma_{1,c}$  and contributing after  $h_1^r$  results in an expected payoff of

$$\begin{aligned} & -\gamma_{\uparrow c} + \mu_1(\mathbf{D}|(h_1^r, c))V(\mathbf{D}) + \mu_1(\mathbf{B}|(h_1^r, c))V(\mathbf{B}) \\ = & -\gamma_{\uparrow c} + \rho\mu_1(\mathbf{D}|h_1^r)V(\mathbf{D}) + \mu_1(\mathbf{B}|(h_1^r, c))V(\mathbf{B}) \end{aligned}$$

whereas following  $\sigma_{1,o}$  and doing nothing earns

$$\begin{aligned} & \mu_1(\mathbf{D}|(h_1^r, o))V(\mathbf{D}) + \mu_1(\mathbf{B}|(h_1^r, o))V(\mathbf{B}) \\ = & \mu_1(\mathbf{D}|h_1^r)V(\mathbf{D}) + \mu_1(\mathbf{B}|(h_1^r, o))V(\mathbf{B}) \end{aligned}$$

Observe that  $\mu_1(\mathbf{B}|(h_1^r, o)) = \mu_1(\mathbf{B}|(h_1^r, c))$ . Thus, player 1 contributes only if

$$\gamma_{\uparrow c} \leq -(1 - \rho)\mu_1(\mathbf{D}|h_1^r)V(\mathbf{D}) \quad (4.5)$$

Note that in the above condition, both sides are positive ( $V(\mathbf{D}) \leq 0$ ).

By Bayes rule, we have

$$\begin{aligned} & \mu_1(\mathbf{D}|h_1^r) \\ = & \frac{\mu_1(\mathbf{D}|h_1^{r+m_c+1})\rho^{m_c+2}}{\mu_1(\mathbf{D}|h_1^{r+m_c+1})(\rho^{m_c+2} + (1 - \rho)) + \mu_1(\neg\mathbf{D}|h_1^{r+m_c+1}) + \mu_1(\mathbf{B}|h_1^{r+m_c+1})\beta} \end{aligned}$$

where

$$\beta = \prod_{i=0}^{m_c} (1 - (1 - \rho)\tau_2(p|h_1^{r-i+m_c+1,2}))$$

Intuitively, the numerator represents the probability that non-Byzantine player 2, destitute after  $h_1^{r+m_c+1}$ , remains destitute after one contribution (with probability  $\rho$ ) and that player 1 does not observe any pesters from round  $r + m_c + 1$  to round  $r + 1$ , inclusive of both endpoints (with probability  $\rho^{m_c+1}$ ).

We can further evaluate the above fraction, giving us

$$\begin{aligned} \mu_1(\mathbf{D}|h_1^r) &\leq \frac{\mu_1(\mathbf{D}|h_1^{r+m_c+1})\rho^{m_c+2}}{\mu_1(\mathbf{D}|h_1^{r+m_c+1})(\rho^{m_c+2} + (1 - \rho))} \\ &= \frac{\rho^{m_c+2}}{\rho^{m_c+2} + 1 - \rho} \\ &\leq \frac{\rho^2}{1 - \rho + \rho^2} \end{aligned}$$

The above fraction gives us an upper bound to  $\mu_1(\mathbf{D}|h_1^r)$ . We can also derive a lower bound on  $V(\mathbf{D}) \leq 0$  (and thus an upper bound on  $-V(\mathbf{D}) \geq 0$ ) by observing that, in any optimal strategy against a destitute player 2, player 1 can do no worse than being pestered for the entire game. It follows then that  $V(\mathbf{D}) \geq -R(1 - \rho)\gamma_{\downarrow p}$ . Plugging these bounds into condition (4.5), it follows that player 1 contributes only if

$$\gamma_{\uparrow c} \leq \frac{\rho^2(1 - \rho)^2}{1 - \rho + \rho^2} R\gamma_{\downarrow p}$$

but this contradicts condition (4.4).  $\square$

LEMMA 4.15. Let  $h_1^r$  and  $h_1^{r'}$  be two histories that player 1 could observe and  $\mu_1$  be player 1's beliefs. Then  $\mu_1(\mathbf{D}|h_1^r) \geq \mu_1(\mathbf{D}|h_1^{r'})$  if

1. For every round, player 2's signals in  $h_1^r$  and  $h_1^{r'}$  are the same;
2. The number of contributions in  $h_1^r$  is greater than or equal to the

number in  $h_{1'}^r$ ; and

3. The last contribution in both  $h_1^r$  and  $h_{1'}^r$  are followed by a pester at some point.

*Proof.* Since pestering has been observed after the last contribution and a non-destitute player 2 never pesters,  $\mu_1(\neg \mathbf{D} | h_1^r) = \mu_1(\neg \mathbf{D} | h_{1'}^r) = 0$ . Letting  $m_c$  and  $m'_c$  be the number of rounds in which player 1 contributed in  $h_1^r$  and  $h_{1'}^r$ , and  $\hat{P}$  and  $\hat{O}$  be the rounds in which player 2 is observed to pester and do nothing,

$$\begin{aligned} \mu_1(\mathbf{D} | h_1^r) &= \frac{\mu_1(\mathbf{D})(1-\rho)^{|\hat{P}|} \rho^{|\hat{O}|+m_c}}{\mu_1(\mathbf{D})(1-\rho)^{|\hat{P}|} \rho^{|\hat{O}|+m_c} + \mu_1(\mathbf{B})\beta} \\ &\geq \frac{\mu_1(\mathbf{D})(1-\rho)^{|\hat{P}|} \rho^{|\hat{O}|+m'_c}}{\mu_1(\mathbf{D})(1-\rho)^{|\hat{P}|} \rho^{|\hat{O}|+m'_c} + \mu_1(\mathbf{B})\beta} \\ &= \mu_1(\mathbf{D} | h_{1'}^r) \end{aligned}$$

where

$$\beta = \prod_{i \in \hat{P}} (1-\rho)\tau_2(p|h_1^{i,2}) \prod_{i \in \hat{O}} (1-(1-\rho)\tau_2(p|h_1^{i,2}))$$

□

LEMMA 4.16. Suppose that player 1 follows a threshold strategy (i.e., contribute iff its belief that player 2 is destitute is above some threshold) starting from some round  $r-1$  and history  $h_1^{r-1}$ . Let  $\mu_1$  be player 1's beliefs and let  $h_{1,p}^i$  and  $h_{1,o}^i$  be two continuation histories from player 1's perspective such that  $i < r$  and

$$\begin{aligned} h_{1,p}^i &= (h_1^r, \omega_1^{r,1}, p, \omega_{1,p}^{r-1,1}, \omega_1^{r-1,2}, \dots, \omega_{1,p}^{i+1,1}, \omega_1^{i+1,2}) \\ h_{1,o}^i &= (h_1^r, \omega_1^{r,1}, o, \omega_{1,o}^{r-1,1}, \omega_1^{r-1,2}, \dots, \omega_{1,o}^{i+1,1}, \omega_1^{i+1,2}) \end{aligned}$$

In other words,  $h_{1,p}^i$  and  $h_{1,o}^i$  differ in player 1's perception of player 2's round  $r$  signal and potentially every signal from player 1 after round  $r$ .

Then either:

1. There has not been a pester since the last contribution in  $h_{1,o}^i$  and  $h_{1,p}^i$  contains at least as many  $c$ 's as  $h_{1,o}^i$ : for  $i < j < r$ ,  $|\{\omega_{1,p}^{j,1} | \omega_{1,p}^{j,1} = c\}| = |\{\omega_{1,o}^{j,1} | \omega_{1,o}^{j,1} = c\}|$ ;
2.  $\mu_1(\mathbf{D} | h_{1,p}^i) \geq \mu_1(\mathbf{D} | h_{1,o}^i)$  and  $h_{1,p}^i$  contains at least as many  $c$ 's as  $h_{1,o}^i$ ; or
3.  $h_{1,p}^i$  contains more  $c$ 's than  $h_{1,o}^i$ : for  $i < j < r$ ,  $|\{\omega_{1,p}^{j,1} | \omega_{1,p}^{j,1} = c\}| > |\{\omega_{1,o}^{j,1} | \omega_{1,o}^{j,1} = c\}|$ .

*Proof.* By backwards induction on  $i$ .

Base case:  $i = r - 1$ . Let  $h_1^{r,2} = (h_1^r, \omega_1^{r,1})$ ,  $h_{1,p}^i = (h_1^{r,2}, p)$  and  $h_{1,o}^i = (h_1^{r,2}, o)$ . Then since player 1 expects that a destitute player 2 always pesters, whereas a Byzantine player 2 may not, then if player 1 observes pestering, then, by Bayes rule, its belief that player 2 is destitute equals

$$\begin{aligned} \mu_1(\mathbf{D} | h_{1,p}^i) &= \frac{(1 - \rho)\mu_1(\mathbf{D} | h_1^{r,2})}{(1 - \rho)(\mu_1(\mathbf{D} | h_1^{r,2}) + \tau_2(p | h_1^{r,2})\mu_1(\mathbf{B} | h_1^{r,2}))} \\ &= \frac{\mu_1(\mathbf{D} | h_1^{r,2})}{\mu_1(\mathbf{D} | h_1^{r,2}) + \tau_2(p | h_1^{r,2})\mu_1(\mathbf{B} | h_1^{r,2})} \end{aligned}$$

The numerator in the above expression represents the probability of observing pester from a destitute player 2, whereas the bottom represents the total probability of observing pester. The  $(1 - \rho)$  term is the possibility of observing the signal given network loss (which cancels out).

On the other hand, if player 1 observes nothing, then, by Bayes rule, its belief that player 2 is destitute equals

$$\begin{aligned} \mu_1(\mathbf{D} | h_{1,o}^i) &= \frac{\rho\mu_1(\mathbf{D} | h_1^{r,2})}{\rho\mu_1(\mathbf{D} | h_1^{r,2}) + \mu_1(\neg\mathbf{D} | h_1^{r,2}) + (1 - (1 - \rho)\tau_2(p | h_1^{r,2}))\mu_1(\mathbf{B} | h_1^{r,2})} \end{aligned}$$

From these two expressions, it can be shown that

$$\mu_1(\mathbf{D}|h_{1,p}^i) \geq \mu_1(\mathbf{D}|h_{1,o}^i)$$

and both  $h_{1,p}^i$  and  $h_{1,o}^i$  have the same number of  $c$ 's, thus satisfying case 2.

Inductive step. Assume true for all  $r_0 \leq i < r$ ; we prove the lemma for  $i = r_0 - 1$  by starting from the inductive hypothesis for round  $i + 1$  and showing that one of the above cases must hold when extending the history into round  $i$ .

Case 1: There has not been a pester since the last contribution in  $h_{1,o}^{i+1}$  and  $h_{1,p}^{i+1}$  contains at least as many  $c$ 's as  $h_{1,o}^{i+1}$ . By Lemma 4.14, we know that player 1 will never contribute following  $h_{1,o}^{i+1}$ . Therefore,

- If player 1 contributes after  $h_{1,p}^{i+1}$ , then case 3 holds.
- If player 1 does nothing after  $h_{1,p}^{i+1}$  and subsequently observes player 2 doing nothing, then case 1 holds.
- If player 1 does nothing after  $h_{1,p}^{i+1}$  and subsequently observes player 2 pestering, then it follows from Lemma 4.15 that  $\mu(\mathbf{D}|(h_{1,p}^{i+1}, o, p)) \geq \mu(\mathbf{D}|(h_{1,p}^{i+1}, o, p))$ , thus satisfying case 2.

Case 2:  $\mu_1(\mathbf{D}|h_{1,p}^{i+1}) \geq \mu_1(\mathbf{D}|h_{1,o}^{i+1})$  and  $h_{1,p}^{i+1}$  contains at least as many  $c$ 's as  $h_{1,o}^{i+1}$ . If player 1 contributes after  $h_{1,o}^{i+1}$  because  $\mu_1(\mathbf{D}|h_{1,o}^{i+1})$  exceeds some threshold, then  $\mu_1(\mathbf{D}|h_{1,p}^{i+1})$  must also exceed the threshold by assumption, so player 1 contributes after  $h_{1,p}^{i+1}$ . It is obvious that the number of contributions in  $(h_{1,p}^{i+1}, c, \omega_1^{i+1,2})$  continues to be at least that in  $(h_{1,o}^{i+1}, c, \omega_1^{i+1,2})$ . Furthermore,

- If player 1 observes player 2 pester in round  $i + 1$  ( $\omega_1^{i+1,2} = p$ ), then, by Lemma 4.15,  $\mu_1(\mathbf{D}|(h_{1,p}^{i+1}, c, p)) \geq \mu_1(\mathbf{D}|(h_{1,o}^{i+1}, c, p))$ , thus satisfying case 2.
- If player 1 observes player 2 do nothing in round  $i + 1$  ( $\omega_1^{i+1,2} = o$ ), then case 1 holds.

Case 3:  $h_{1,p}^{i+1}$  contains more  $c$ 's than  $h_{1,o}^{i+1}$ . If player 1 contributes after  $h_{1,p}^{i+1}$  or does nothing after  $h_{1,o}^{i+1}$ , then case 3 holds. Otherwise, if player 1 contributes after  $h_{1,o}^{i+1}$  but not after  $h_{1,p}^{i+1}$ , then it is obvious that the number of contributions in  $(h_{1,p}^{i+1}, o, \omega_1^{i+1,2})$  continues to be at least that in  $(h_{1,o}^{i+1}, c, \omega_1^{i+1,2})$ . Furthermore,

- If player 1 observes player 2 pester in round  $i + 1$  ( $\omega_1^{i+1,2} = p$ ), then, by Lemma 4.15,  $\mu_1(\mathbf{D}|(h_{1,p}^{i+1}, o, p)) \geq \mu_1(\mathbf{D}|(h_{1,o}^{i+1}, c, p))$ , thus satisfying case 2.
- If player 1 observes player 2 do nothing in round  $i + 1$  ( $\omega_1^{i+1,2} = o$ ), then case 1 holds.

□

LEMMA 4.17. Suppose that player 1 is playing a threshold strategy starting from some round  $r - 1$  and history  $h_1^{r-1}$ . Let  $h_{1,c}^i$  and  $h_{1,o}^i$  be two continuation histories from player 1's perspective such that

$$\begin{aligned} h_{1,c}^i &= (h_1^r, c, \omega_1^{r,2}, \omega_{1,c}^{r-1,1}, \omega_1^{r-1,2}, \dots, \omega_{1,c}^{i+1,1}, \omega_{1,c}^{i+1,2}) \\ h_{1,o}^i &= (h_1^r, o, \omega_1^{r,2}, \omega_{1,o}^{r-1,1}, \omega_1^{r-1,2}, \dots, \omega_{1,o}^{i+1,1}, \omega_{1,o}^{i+1,2}) \end{aligned}$$

In other words,  $h_{1,c}^i$  and  $h_{1,o}^i$  differs in how player 1 plays in round  $r$  and may also differ in how player 1 plays in subsequent rounds. Then either:

1. There has not been a pester since the last contribution in  $h_{1,o}^i$  and  $h_{1,c}^i$  contains at least as many  $c$ 's as  $h_{1,o}^i$ : for  $i < j < r$ ,  $|\{\omega_{1,c}^{j,1} | \omega_{1,c}^{j,1} = c\}| + 1 = |\{\omega_{1,o}^{j,1} | \omega_{1,o}^{j,1} = c\}|$ ;
2.  $\mu_1(\mathbf{D}|h_{1,c}^i) \geq \mu_1(\mathbf{D}|h_{1,o}^i)$  and  $h_{1,c}^i$  contains at least as many  $c$ 's as  $h_{1,o}^i$ ;  
or
3.  $h_{1,c}^i$  contains more  $c$ 's than  $h_{1,o}^i$ : for  $i < j < r$ ,  $|\{\omega_{1,c}^{j,1} | \omega_{1,c}^{j,1} = c\}| + 1 > |\{\omega_{1,o}^{j,1} | \omega_{1,o}^{j,1} = c\}|$ .

*Proof.* By backwards induction on  $i$ .

Base case:  $i = r - 1$ . Player 1 has one more contribution by construction.

Inductive step. Assume true for all  $r_0 \leq i < r$ ; we prove the lemma for  $i = r_0 - 1$  by starting from the inductive hypothesis for round  $i + 1$  and showing that one of the above cases must hold when extending the history into round  $i$ .

Case 1: There has not been a pester since the last contribution in  $h_{1,o}^{i+1}$  and  $h_{1,c}^{i+1}$  contains at least as many  $c$ 's as  $h_{1,o}^{i+1}$ . By Lemma 4.14, player 1 will never contribute after  $h_{1,o}^{i+1}$ . Therefore,

- If player 1 contributes after  $h_{1,c}^{i+1}$ , then case 3 holds.
- If player 1 does nothing after  $h_{1,c}^{i+1}$  and subsequently observes player 2 doing nothing, then case 1 holds.
- If player 1 does nothing after  $h_{1,c}^{i+1}$  and subsequently observes player 2 pester in round  $i + 1$ , it follows from Lemma 4.15 that  $\mu(\mathbf{D}|(h_{1,c}^{i+1}, o, p)) \geq \mu(\mathbf{D}|(h_{1,o}^{i+1}, o, p))$ , thus satisfying case 2.

Case 2:  $\mu_1(\mathbf{D}|h_{1,c}^{i+1}) \geq \mu_1(\mathbf{D}|h_{1,o}^{i+1})$  and  $h_{1,c}^{i+1}$  contains at least as many  $c$ 's as  $h_{1,o}^{i+1}$ . If player 1 contributes after  $h_{1,o}^{i+1}$  because  $\mu_1(\mathbf{D}|h_{1,o}^{i+1})$  exceeds some threshold, then  $\mu_1(\mathbf{D}|h_{1,c}^{i+1})$  must also exceed the threshold, so player 1 contributes after  $h_{1,c}^{i+1}$ . It is obvious that the number of contributions in  $(h_{1,c}^{i+1}, c, \omega_1^{i+1,2})$  continues to be at least that in  $(h_{1,o}^{i+1}, c, \omega_1^{i+1,2})$ . Furthermore,

- If player 1 observes player 2 pester in round  $i + 1$  ( $\omega_1^{i+1,2} = p$ ), then, by Lemma 4.15,  $\mu_1(\mathbf{D}|(h_{1,c}^{i+1}, c, p)) \geq \mu_1(\mathbf{D}|(h_{1,o}^{i+1}, c, p))$ , thus satisfying case 2.
- If player 1 observes player 2 do nothing in round  $i + 1$  ( $\omega_1^{i+1,2} = o$ ), then case 1 holds.

Case 3:  $h_{1,c}^{i+1}$  contains more  $c$ 's than  $h_{1,o}^{i+1}$ . If player 1 contributes after  $h_{1,c}^{i+1}$  or does nothing after  $h_{1,o}^{i+1}$ , then case 3 holds. Otherwise, if player 1 contributes

after  $h_{1,o}^{i+1}$  but not after  $h_{1,c}^{i+1}$ , then it is obvious that the number of contributions in  $(h_{1,c}^{i+1}, o, \omega_1^{i+1,2})$  continues to be at least that in  $(h_{1,o}^{i+1}, c, \omega_1^{i+1,2})$ . Furthermore,

- If player 1 observes player 2 pester in round  $i + 1$  ( $\omega_1^{i+1,2} = p$ ), then, by Lemma 4.15,  $\mu_1(\mathbf{D}|(h_{1,c}^{i+1}, o, p)) \geq \mu_1(\mathbf{D}|(h_{1,o}^{i+1}, c, p))$ , thus satisfying case 2.
- If player 2 observes player 2 do nothing in round  $i + 1$  ( $\omega_1^{i+1,2} = o$ ), then case 1 holds.

□

The following lemma states that, given that a non-Byzantine player 2 is known to be destitute (and thus player 1 is willing to consider contributing), there exist (infinitely-many) prior beliefs which, after the same sequence of signals, map to arbitrarily-close posterior beliefs.

LEMMA 4.18. Let  $h_1^r$  be the current history and  $\mu_1$  be player 1's beliefs such that  $\mu_1(\mathbf{D}|h_1^r) > 0$  and  $\mu_1(\neg\mathbf{D}|h_1^r) = 0$ . For all  $\epsilon > 0$ , there exists some  $\delta > 0$  such that:

1. For all histories  $h_{1'}^r$  and associated beliefs  $\mu_1'(\theta|h_{1'}^r)$  where  $\mu_1'(\neg\mathbf{D}|h_{1'}^r) = 0$  and  $0 \leq \mu_1'(\mathbf{D}|h_{1'}^r) - \mu_1(\mathbf{D}|h_1^r) < \delta$ ; and
2. For any history  $h_1^i$  and  $h_{1'}^i$ , that contain  $h_1^r$  and  $h_{1'}^r$ , respectively, as a prefix such that:

- Both continuations contain the same signals:

$$\begin{aligned} h_1^i &= (h_1^r, \omega_1^{r,1}, \omega_1^{r,2}, \dots, \omega_1^{i+1,1}, \omega_1^{i+1,2}) \\ h_{1'}^i &= (h_{1'}^r, \omega_1^{r,1}, \omega_1^{r,2}, \dots, \omega_1^{i+1,1}, \omega_1^{i+1,2}) \end{aligned}$$

- Player 1 knows a non-Byzantine player 2 is destitute after either  $h_1^i$  and  $h_{1'}^i$ :

$$\mu_1(\neg\mathbf{D}|h_1^i) = \mu_1'(\neg\mathbf{D}|h_{1'}^i) = 0$$

Then the following holds:  $0 \leq \mu'_1(\mathbf{D}|h_{1'}^i) - \mu_1(\mathbf{D}|h_1^i) < \epsilon$ .

*Proof.* Let  $\hat{P}$  and  $\hat{O}$  be the rounds in which pestering and doing nothing, respectively, are observed in  $\omega_1^{r,1}, \omega_1^{r,2}, \dots, \omega_1^{i+1,1}, \omega_1^{i+1,2}$ ; similarly, let  $m_c$  be the number of rounds in which contribution occurs in this sequence. Letting  $\zeta = (1 - \rho)^{|\hat{P}|} \rho^{|\hat{O}| + m_c}$  and  $\beta = \prod_{i \in \hat{P}} (1 - \rho) \tau_2(p|h_1^{i,2}) \prod_{i \in \hat{O}} (1 - (1 - \rho) \tau_2(p|h_1^{i,2}))$ , we have

$$\begin{aligned}
& \mu'_1(\mathbf{D}|h_{1'}^i) - \mu_1(\mathbf{D}|h_1^i) \\
&= \frac{\mu'_1(\mathbf{D}|h_{1'}^r) \zeta}{\mu'_1(\mathbf{D}|h_{1'}^r) \zeta + \mu'_1(\mathbf{B}|h_{1'}^r) \beta} - \frac{\mu_1(\mathbf{D}|h_1^r) \zeta}{\mu_1(\mathbf{D}|h_1^r) \zeta + \mu_1(\mathbf{B}|h_1^r) \beta} \\
&= \frac{(\mu'_1(\mathbf{D}|h_{1'}^r) \mu_1(\mathbf{B}|h_1^r) - \mu_1(\mathbf{D}|h_1^r) \mu'_1(\mathbf{B}|h_{1'}^r)) \beta \zeta}{(\mu'_1(\mathbf{D}|h_{1'}^r) \zeta + \mu'_1(\mathbf{B}|h_{1'}^r) \beta)(\mu_1(\mathbf{D}|h_1^r) \zeta + \mu_1(\mathbf{B}|h_1^r) \beta)} \\
&= \frac{(\mu'_1(\mathbf{D}|h_{1'}^r) - \mu_1(\mathbf{D}|h_1^r)) \zeta \beta}{(\mu'_1(\mathbf{D}|h_{1'}^r) \zeta + \mu'_1(\mathbf{B}|h_{1'}^r) \beta)(\mu_1(\mathbf{D}|h_1^r) \zeta + \mu_1(\mathbf{B}|h_1^r) \beta)}
\end{aligned}$$

It is obvious that the above expression is non-negative if  $\mu'_1(\mathbf{D}|h_{1'}^r) - \mu_1(\mathbf{D}|h_1^r) \geq 0$ . Moreover, since the denominator in the above expression is positive and less than 1 and  $0 \leq \beta \zeta < 1$ , it can be verified that the above expression is strictly less than  $\epsilon$  if  $\mu'_1(\mathbf{D}|h_{1'}^r) - \mu_1(\mathbf{D}|h_1^r) < \epsilon = \delta$ .  $\square$

## Full results for player 2

LEMMA 4.19. Player 2 is as likely to receive a contribution from a rational player 1 if player 2 pesters instead of doing nothing.

*Proof.* Let  $r$  be the current round and let  $h_1^r$  and  $h_2^r$  be the histories that player 1 and 2 have observed. If player 2 pesters and player 1 does not observe it, then regardless of whether player 2 pesters or not, player 1 starts from history  $(h_1^r, o)$  and plays exactly the same. As a result, the likelihood of player 1 contributing and player 2 receiving said contribution is exactly the same.

Thus, suppose that player 1 observes the history  $(h_1^r, p)$  if player 2 pesters and  $(h_1^r, o)$  otherwise. Consider any two complete histories  $h_{1,p}$  and  $h_{1,o}$  from player 1's perspective such that

$$\begin{aligned} h_{1,p} &= (h_1^r, p, \omega_{1,p}^{r-1,1}, \omega_1^{r-1,2}, \dots, \omega_{1,p}^{0,1}, \omega_1^{0,2}) \\ h_{1,o} &= (h_1^r, o, \omega_{1,o}^{r-1,1}, \omega_1^{r-1,2}, \dots, \omega_{1,o}^{0,1}, \omega_1^{0,2}) \end{aligned}$$

Let  $m_p$  and  $m_o$  be the number of rounds from round  $r - 1$  until the end of the game in which player 1 observes player 2 pestering and doing nothing. By Lemma 4.16,  $h_{1,p}$  must contain at least as many  $c$ 's as  $h_{1,o}$ ; let  $m_{c,p}$  and  $m_{c,o}$  be the number of rounds in which player 1 contributes after  $h_{1,p}^r$  and  $h_{1,o}^r$ , respectively, where  $m_{c,p} \geq m_{c,o}$ .

Player 1 expects that a destitute player 2 will pester every round except the last. Moreover, note that (a rational) player 1 never mixes: based on the history it has observed, player 1 either contributes or does nothing with probability 1. Thus, given that player 2 never receives the contribution, the probability that  $h_{1,p}$  and  $h_{1,o}$  result from  $(h_1^r, p)$  and  $(h_1^r, o)$ , respectively, is the same:  $(1 - \rho)^{m_p} \rho^{m_o - 1}$  (note the  $m_o - 1$  is due to the fact that player 2 always does nothing in the last round). The probability that player 2 never receives a contribution given  $h_{1,p}$  is  $\rho^{m_{c,p}}$ . Since the probability that player 2 never receives a contribution given  $h_{1,o}$  is  $\rho^{m_{c,o}} \geq \rho^{m_{c,p}}$ , player 1 is as likely to receive a contribution in history  $h_{1,p}$  as in  $h_{1,o}$ .

The probability of player 2 not getting the contribution starting from  $(h_1^r, p)$  and  $(h_1^r, o)$  is simply the sum of the probabilities of all possible complete histories  $h_{1,p}$  and  $h_{1,o}$  that start with  $(h_1^r, p)$  and  $(h_1^r, o)$ , respectively. It then follows that player 2 is no less likely to get the contribution starting from  $(h_1^r, p)$  instead of  $(h_1^r, o)$ .  $\square$

**THEOREM 4.20.** A destitute player 2 pesters in all rounds  $r > 0$ .

*Proof.* By contradiction. Assume player 2 prefers to do nothing despite condition (4.3), i.e., there exists some strategy  $\sigma_{2,o}$  in which player 2 does nothing in round  $r$  despite having beliefs which satisfy condition (4.3). Construct an alternate strategy  $\sigma_{2,p}$  in which:

1. Player 2 pesters after  $h_2^r$ .
2. If player 2 receives a contribution in round  $r - 1$ , player 2 does nothing for the remainder of the game.
3. Otherwise, player 2 plays  $\sigma_{2,o}$  as if it did nothing in round  $r$ , i.e.,  $\sigma_{2,o}$  and  $\sigma_{2,p}$  are identical from round  $r - 1$  on if player 2 does not receive a contribution in round  $r$ .

Consider player 2's difference in expected utility between playing  $\sigma_{2,o}$  and  $\sigma_{2,p}$ . If player 1 is Byzantine, the expected difference in utility between  $\sigma_{2,o}$  and  $\sigma_{2,p}$  is  $\gamma_{\uparrow p}$ . If player 1 is rational, by Lemma 4.19, player 2 is as likely to receive a contribution if player 2 pesters in round  $r$  (versus doing nothing); hence, the expected difference in utility between  $\sigma_{2,p}$  and  $\sigma_{2,o}$  is at most  $\gamma_{\uparrow p}$ . Finally, if player 1 is acquiescent, then the expected utility, starting from player 2's turn in round  $r - 1$ , of playing  $\sigma_{2,o}$  or  $\sigma_{2,p}$  as a destitute player is the same; let  $V(\mathbf{A}, r - 1)$  represent this utility. Thus, the expected difference in utility between  $\sigma_{2,o}$  and  $\sigma_{2,p}$  when facing an acquiescent player 1 is then

$$\mathbb{E}^{(\sigma_{2,o}, \sigma_1^*)}[u_2|h_2^r, \mathbf{A}] - \mathbb{E}^{(\sigma_{2,p}, \sigma_1^*)}[u_2|h_2^r, \mathbf{A}] = \gamma_{\uparrow p} - (1 - \alpha)(b - \gamma_{\downarrow c} - V(\mathbf{A}, r - 1))$$

Note that we subtract away  $V(\mathbf{A}, r - 1)$  with probability  $1 - \alpha$  since, with probability  $1 - \alpha$ , player 2 observes the contribution and, for the remainder of the game, will do nothing along with an (acquiescent) player 1.

By Assumption 4.7, pestering an acquiescent player 1 until player 2 gets the contribution or  $r = 0$  is in player 2's best interest, and thus for  $i < r$ ,  $V(\mathbf{A}, i) \leq -\gamma_{\uparrow p} + (1 - \alpha)(b - \gamma_{\downarrow c}) + \alpha V(\mathbf{A}, i - 1)$ , where  $V(\mathbf{A}, 0) = 0$ . Solving

the recursion, we have

$$V(\mathbf{A}, r-1) \leq \frac{1-\alpha^{r-1}}{1-\alpha} (-\gamma_{\uparrow p} + (1-\alpha)(b-\gamma_{\downarrow c}))$$

Using condition (4.3) we get

$$\begin{aligned} & \mathbb{E}^{(\sigma_{2,o}, \sigma_1^*)}[u_2|h_2^r] - \mathbb{E}^{(\sigma_{2,p}, \sigma_1^*)}[u_2|h_2^r] \\ & \leq \gamma_{\uparrow p} - \mu_2^*(\mathbf{A}|h_2^r)(\alpha^{r-1}(1-\alpha)(b-\gamma_{\downarrow c}) + (1-\alpha^{r-1})\gamma_{\uparrow p}) \\ & < 0 \end{aligned}$$

and thus player 2 prefers to pester. Contradiction.  $\square$

### Full proofs for player 1

LEMMA 4.21. Player 1 does nothing for rounds  $r \leq \bar{r}$ , where  $\bar{r}$  is defined as in condition (4.1).

*Proof.* By induction.

Base case:  $r = 0$ . If player 1 does nothing, it earns 0 since player 2 never pesters in the last round. If it contributes, it earns at most  $-\gamma_{\uparrow c} < 0$ .

Inductive step. Assume true for all rounds  $r$  such that  $0 \leq r < r_0 \leq \bar{r}$ . We now prove it to be true for  $r = r_0$ . Let  $h_1^r$  be player 1's current history. If player 1 does nothing for the remainder of the game, its expected utility is at least

$$-\mu_1^*(\mathbf{D}|h_1^r)r(1-\rho)\gamma_{\downarrow p} + \mu_1^*(\mathbf{B}|h_1^r) \mathbb{E}[u_1|h_1^r, \mathbf{B}]$$

since, by the induction hypothesis, player 1 does nothing for the remaining rounds.

If player 1 instead contributes, its expected utility is at most

$$-\gamma_{\uparrow c} - \mu_1^*(\mathbf{D}|h_1^r)r\rho(1-\rho)\gamma_{\downarrow p} + \mu_1^*(\mathbf{B}|h_1^r) \mathbb{E}[u_1|h_1^r, \mathbf{B}]$$

for similar reasons.

Thus, doing nothing is strictly better if

$$-\mu_1^*(\mathbf{D}|h_1^r)r(1-\rho)^2\gamma_{\downarrow p} > -\gamma_{\uparrow c}$$

This is satisfied given condition (4.1).  $\square$

LEMMA 4.22. Let  $r < R$  be the current round and  $h_1^r$  be the current history, and suppose that player 1 follows a threshold strategy  $\sigma_1$  starting from round  $r$  and history  $h_1^r$ . Then

$$-\gamma_{\uparrow c} + \mathbb{E}^{(\sigma_1, \sigma_2^*)}[u_1|(h_1^r, c), \mathbf{B}] \leq \mathbb{E}^{(\sigma_1, \sigma_2^*)}[u_1|(h_1^r, o), \mathbf{B}]$$

*Proof.* Starting from history  $(h_1^r, o)$  and  $(h_1^r, c)$ , consider any two complete histories of the following form:

$$\begin{aligned} h_{1,c} &= (h_1^r, c, \omega_1^{r,2}, \omega_{1,c}^{r-1,1}, \omega_1^{r-1,2}, \dots, \omega_{1,c}^{0,1}, \omega_1^{0,2}) \\ h_{1,o} &= (h_1^r, o, \omega_1^{r,2}, \omega_{1,o}^{r-1,1}, \omega_1^{r-1,2}, \dots, \omega_{1,o}^{0,1}, \omega_1^{0,2}) \end{aligned}$$

such that  $\omega_{1,c}^{j,1}$  and  $\omega_{1,o}^{j,1}$  is the action specified by  $\sigma_1$  given history  $h_{1,c}^j$  and  $h_{1,o}^j$  (for  $0 \leq j < r$ ).

Given the fact that  $\sigma_1$  is a threshold strategy and that player 2 is Byzantine and is thus expected to play actions independent of what it observes from player 1, (1)  $h_{1,c}$  and  $h_{1,o}$  occur with equal probability, (2) by Lemma 4.17,  $h_{1,c}$  contains at least as many  $c$ 's as  $h_{1,o}$ , and (3) the cost incurred by pestering in both histories is exactly the same.

It follows then that, starting from history  $(h_1^r, c)$  and history  $(h_1^r, o)$ , player 1 will contribute at most once more in continuation starting from  $(h_1^r, o)$  with no difference in the cost incurred by pestering. Thus,

$$\mathbb{E}^{\sigma^*}[u_1|(h_1^r, c), \mathbf{B}] - \mathbb{E}^{\sigma^*}[u_1|(h_1^r, o), \mathbf{B}] \leq \gamma_{\uparrow c}$$

as needed. □

**THEOREM 4.23.** Let  $h_1^r$  be the history of player 1 in some round  $r$  such that  $\bar{r} < r \leq R$ , where  $\bar{r}$  is defined by condition (4.1), and let  $\mu_1$  be player 1's beliefs such that  $\mu_1(\neg \mathbf{D}|h_1^r) = 0$ . Then there exists some threshold  $\bar{\mu}_1^r$  such that if  $\mu_1(\mathbf{D}|h_1^r) \geq \bar{\mu}_1^r$ , player 1 contributes; otherwise, player 1 does nothing.

*Proof.* By induction on  $r$ , i.e., backwards induction on time.

Base case:  $r = \bar{r} + 1$ . player 1 contributes iff

$$-\gamma_{\uparrow c} + \mathbb{E}^{(\sigma_{\mathbf{R},1}, \sigma_2^*)}[u_1|(h_1^r, c)] \geq \mathbb{E}^{(\sigma_{\mathbf{R},1}, \sigma_2^*)}[u_1|(h_1^r, o)]$$

Since, by Lemma 4.21, player 1 does nothing in subsequent rounds regardless of player 2's actions, we have

$$-\gamma_{\uparrow c} \geq \mathbb{E}^{(\sigma_{\mathbf{R},1}, \sigma_2^*)}[u_1|(h_1^r, o)] - \mathbb{E}^{(\sigma_{\mathbf{R},1}, \sigma_2^*)}[u_1|(h_1^r, c)] = -(1 - \rho)^2 \mu_1(\mathbf{D}|h_1^r) \bar{r} \gamma_{\downarrow p}$$

Solving for  $\mu_1(\mathbf{D}|h_1^r)$ , we have

$$\mu_1(\mathbf{D}|h_1^r) \geq \frac{1}{(1 - \rho)^2 \bar{r}} \frac{\gamma_{\uparrow c}}{\gamma_{\downarrow p}} = \bar{\mu}_1^{\bar{r}+1}$$

Inductive step. Assume true for all  $r$ ,  $\bar{r} < r \leq r_0$ ; we prove  $r = r_0 + 1$  by contradiction, i.e., there exists some alternate strategy  $\sigma_1'$  such that, in round  $r$ , there are beliefs (in player 2 being destitute) in which player 1 is (weakly) better off contributing and higher beliefs in which player 1 is better off doing nothing; in subsequent rounds, by the induction hypothesis, player 1 is best off playing the threshold strategy. Formally, there must exist some  $\eta_1^r < 1$ ,  $\delta > 0$ , and some belief  $\mu_1(\theta|h_1^r)$  such that  $\mu_1(\neg \mathbf{D}|h_1^r) = 0$  and either:

1. For  $\mu_1(\mathbf{D}|h_1^r) = \eta_1^r$ , player 1 contributes, and for  $\mu_1(\mathbf{D}|h_1^r)$  such that  $\eta_1^r < \mu_1(\mathbf{D}|h_1^r) < \eta_1^r + \delta$ , player 1 does nothing; or

2. For  $\mu_1(\mathbf{D}|h_1^r)$  such that  $\eta_1^r - \delta < \mu_1(\mathbf{D}|h_1^r) < \eta_1^r$ , player 1 contributes, and for  $\mu_1(\mathbf{D}|h_1^r) = \eta_1^r$ , player 1 does nothing.

We only consider the first case, as the proof of the other case is very similar. Suppose that  $\mu_1(\neg\mathbf{D}|h_1^r) = 0$  and  $\mu_1(\mathbf{D}|h_1^r) = \eta_1^r$  for some history  $h_1^r$ . By the inductive hypothesis, there exists a threshold  $\bar{\mu}_1^i$  for all rounds  $i < r$ . For any  $i \leq r$ , let  $H_1^i$  represent the set of all possible histories that start with  $h_1^r$  in which player 1 knows a non-Byzantine player 2 is destitute:

$$H_1^i = \{h_1^i \mid h_1^i = (h_1^r, \omega_1^{r,1}, \omega_1^{r,2}, \dots, \omega_1^{i+1,1}, \omega_1^{i+1,2}), \mu_1(\neg\mathbf{D}|h_1^i) = 0\}$$

Then we let:

$$\epsilon = \min_{0 \leq i < r, h_1^i \in H_1^i} \{\bar{\mu}_1^i - \mu_1(\mathbf{D}|h_1^i) \mid \bar{\mu}_1^i > \mu_1(\mathbf{D}|h_1^i)\}$$

If  $\mu_1(\mathbf{D}|h_1^i) \geq \bar{\mu}_1^i$  for all  $h_1^i$ , then we set  $\epsilon = 1$ .  $\epsilon$  represents the minimum difference, if one exists, between player 1's belief and the threshold at any future round when player 1's belief in player 2 being destitute is less than that round's threshold.

By Lemma 4.18, we know that we can find some history  $h_1^r$ , and associated belief  $\mu_1'(\theta|h_1^r)$  such that  $\mu_1'(\neg\mathbf{D}|h_1^r) = 0$  and for  $0 \leq \mu_1'(\mathbf{D}|h_1^r) - \mu_1(\mathbf{D}|h_1^r) < \delta$  and  $0 \leq i < r$ ,

$$0 \leq \mu_1'(\mathbf{D}|h_1^i) - \mu_1(\mathbf{D}|h_1^i) < \epsilon$$

for all  $h_1^i \in H_1^i$  (where  $H_1^i$  is defined similarly to  $H_1^i$ ) and  $h_1^i \in H_1^i$ .

Thus, for any round  $i < r$  where  $\mu_1(\neg\mathbf{D}|h_1^r) = \mu_1'(\neg\mathbf{D}|h_1^r) = 0$ ,  $\mu_1(\mathbf{D}|h_1^i) < \bar{\mu}_1^i$  iff  $\mu_1'(\mathbf{D}|h_1^i) < \bar{\mu}_1^i$ . By Lemma 4.14, player 1 is better off doing nothing if its belief that player 2 is non-destitute is positive. It follows that player 1 playing  $\sigma_1'$  with belief  $\mu_1(\theta|h_1^r)$ , upon observing some non-empty sequence of signals after  $h_1^r$ , plays the same action as if player 1 held the belief  $\mu_1'(\theta|h_1^r)$  and observed the same non-empty sequence of signals after  $h_1^r$ . Given that player 2

is of type  $\theta$ , player 1's expected utility of playing action  $a_1^r$  followed by the threshold strategy with either belief  $\mu_1(\theta|h_1^r)$  or  $\mu'_1(\theta|h_{1'}^r)$  must be equal; let  $V(a_1^r, \theta)$  be this expected continuation utility.

By assumption,  $\mu_1(\neg\mathbf{D}|h_1^r) = \mu'_1(\neg\mathbf{D}|h_{1'}^r) = 0$  and so  $\mu_1(\mathbf{B}|h_1^r) = 1 - \mu_1(\mathbf{D}|h_1^r)$  and  $\mu'_1(\mathbf{B}|h_{1'}^r) = 1 - \mu_1(\mathbf{D}|h_{1'}^r)$ . However, since  $\mu_1(\mathbf{D}|h_1^r) = \eta_1^r < \mu'_1(\mathbf{D}|h_{1'}^r) < \eta_1^r + \delta$ , then given belief  $\mu'_1(\theta|h_{1'}^r)$ , player 1 prefers to do nothing during round  $r + 1$ :

$$\begin{aligned} -\gamma_{\uparrow c} + \mathbb{E}^{(\sigma'_1, \sigma_2^*), \mu}[u_1|(h_1^r, c)] &\geq \mathbb{E}^{(\sigma'_1, \sigma_2^*), \mu}[u_1|(h_1^r, o)] \\ -\gamma_{\uparrow c} + \mathbb{E}^{(\sigma'_1, \sigma_2^*), \mu'}[u_1|(h_{1'}^r, c)] &< \mathbb{E}^{(\sigma'_1, \sigma_2^*), \mu'}[u_1|(h_{1'}^r, o)] \end{aligned}$$

We know that

$$\begin{aligned} &\mathbb{E}^{(\sigma'_1, \sigma_2^*), \mu}[u_1|(h_1^r, \omega_1^{r,1})] \\ &= \mu_1(\mathbf{B}|h_1^r)V(\omega_1^{r,1}, \mathbf{B}) + \mu_1(\mathbf{D}|h_1^r)V(\omega_1^{r,1}, \mathbf{D}) + \mu_1(\neg\mathbf{D}|h_1^r)V(\omega_1^{r,1}, \neg\mathbf{D}) \\ &\mathbb{E}^{(\sigma'_1, \sigma_2^*), \mu'}[u_1|(h_{1'}^r, \omega_1^{r,1})] \\ &= \mu'_1(\mathbf{B}|h_{1'}^r)V(\omega_1^{r,1}, \mathbf{B}) + \mu'_1(\mathbf{D}|h_{1'}^r)V(\omega_1^{r,1}, \mathbf{D}) + \mu'_1(\neg\mathbf{D}|h_{1'}^r)V(\omega_1^{r,1}, \neg\mathbf{D}) \end{aligned}$$

Since a non-Byzantine player 2 stops pestering upon receiving a contribution, then by Lemma 4.14, player 1 never contributes unless pestered:  $V(a_1^r, \neg\mathbf{D}) = 0$ . Combining the last two groups of expressions and moving terms around, we get

$$\begin{aligned} \mu_1(\mathbf{B}|h_1^r)(V(o, \mathbf{B}) - V(c, \mathbf{B}) + \gamma_{\uparrow c}) &\leq \mu_1(\mathbf{D}|h_1^r)(-\gamma_{\uparrow c} + \rho V(c, \mathbf{D}) - V(o, \mathbf{D})) \\ \mu'_1(\mathbf{B}|h_{1'}^r)(V(o, \mathbf{B}) - V(c, \mathbf{B}) + \gamma_{\uparrow c}) &> \mu'_1(\mathbf{D}|h_{1'}^r)(-\gamma_{\uparrow c} + \rho V(c, \mathbf{D}) - V(o, \mathbf{D})) \end{aligned}$$

In rounds  $r + 1$  on, we know by the inductive hypothesis that player 1 is better off playing a threshold strategy. By Lemma 4.22, we know that  $V(o, \mathbf{B}) - V(c, \mathbf{B}) + \gamma_{\uparrow c} \geq 0$ . If  $V(o, \mathbf{B}) - V(c, \mathbf{B}) + \gamma_{\uparrow c} = 0$ , an immediate

contradiction arises:

$$\mu'_1(\mathbf{D}|h_1^r)(-\gamma_{\uparrow c} + \rho V(c, \mathbf{D}) - V(o, \mathbf{D})) < 0 \leq \mu_1(\mathbf{D}|h_1^r)(-\gamma_{\uparrow c} + \rho V(c, \mathbf{D}) - V(o, \mathbf{D}))$$

Thus, assuming that  $V(o, \mathbf{B}) - V(c, \mathbf{B}) + \gamma_{\uparrow c} > 0$ , we have

$$\frac{\mu_1(\mathbf{B}|h_1^r)}{\mu_1(\mathbf{D}|h_1^r)} \leq \frac{-\gamma_{\uparrow c} + \rho V(c, \mathbf{D}) - V(o, \mathbf{D})}{V(o, \mathbf{B}) - V(c, \mathbf{B}) + \gamma_{\uparrow c}} < \frac{\mu'_1(\mathbf{B}|h_1^r)}{\mu'_1(\mathbf{D}|h_1^r)} \quad (4.6)$$

However, by assumption,  $\mu'_1(\mathbf{D}|h_1^r) > \mu_1(\mathbf{D}|h_1^r)$ , so

$$\mu_1(\mathbf{B}|h_1^r) = 1 - \mu_1(\mathbf{D}|h_1^r) > 1 - \mu'_1(\mathbf{D}|h_1^r) = \mu'_1(\mathbf{B}|h_1^r)$$

thus contradicting condition (4.6).  $\square$

LEMMA 4.24. Suppose there exists a threshold strategy  $\sigma_1$  which specifies a best response for player 1 at every point in the game. Let  $r$  be the current round such that  $0 < r < R$  and  $R$  satisfies condition (4.4), and let  $h_1^r$  be the current history. If player 1, following  $\sigma_1$ , does nothing in round  $r$  after history  $h_1^r$  and observes player 2 doing nothing, then player 1, following  $\sigma_1$ , also does nothing in round  $r - 1$  following history  $(h_1^r, o, o)$ .

*Proof.* Since player 1 never contributes starting from round  $\bar{r}$  (as defined by condition (4.1)), we prove the lemma for  $r = r_0 + 1 > \bar{r}$ .

For the sake of contradiction, suppose that there exists a threshold strategy  $\sigma_{1,o}$  that specifies a best response for player 1 at every point in the game and in which player 1, following  $\sigma_{1,o}$ , does nothing after history  $h_1^r$  but contributes after history  $h_1^{r-1} = (h_1^r, o, o)$ . By Lemma 4.16,

$$\mu_1(\mathbf{D}|h_1^{r-1}) = \mu_1(\mathbf{D}|(h_1^r, o, o)) \leq \mu_1(\mathbf{D}|(h_1^r, o, p))$$

It follows that, since player 1 is playing a threshold strategy, player 1 will also

contribute following  $(h_1^r, o, p)$ .

Consider, instead, an alternate strategy  $\sigma_{1,c}$ , which is exactly the same as  $\sigma_{1,o}$  except player 1 contributes after history  $h_1^r$  but does nothing after history  $h_1^{r-1} = (h_1^r, c, o)$ . The subsequent actions that player 1 plays, following  $\sigma_{1,c}$ , are the same as if player 1 had faithfully followed  $\sigma_{1,o}$  the entire time (i.e., as if player 1 contributed in round  $r - 1$  and not the prior round  $r$ ). Intuitively,  $\sigma_{1,c}$  differs from  $\sigma_{1,o}$  in that player 1 contributes one round early in  $\sigma_{1,c}$ .

To complete the contradiction, we show the following:

$$E^{(\sigma_{1,c}, \sigma_2^*)}[u_1 | h_1^r] > E^{(\sigma_{1,o}, \sigma_2^*)}[u_1 | h_1^r]$$

Consider what player 1's expected continuation utility is given these two different strategies. First, consider if player 2 is non-destitute. By Lemma 4.14, player 1 only contributes if it has been subsequently pestered. Since a non-destitute player 2 never pesters, the continuation utility against a non-destitute player 2 is the same  $(-\gamma_{\uparrow c})$  regardless of whether player 1 is playing  $\sigma_{1,c}$  or  $\sigma_{1,o}$ .

Next, consider if player 2 is destitute. Then, starting from  $h_1^r$  and given that player 2 is destitute, player 1 expects to earn

$$E^{(\sigma_{1,o}, \sigma_2^*)}[u_1 | h_1^r, \mathbf{D}] = -\gamma_{\uparrow c} + (1 - \rho)(-\gamma_{\downarrow p} + \rho E^{(\sigma_{1,o}, \sigma_2^*)}[u_1 | (h_1^r, o, p, c), \mathbf{D}]) + \rho(\rho E^{(\sigma_{1,o}, \sigma_2^*)}[u_1 | (h_1^r, o, o, c), \mathbf{D}])$$

following  $\sigma_{1,o}$ . The  $(1 - \rho)(\cdot)$  term (the  $\rho(\cdot)$  term) represents the expected payoff if player 1 observes pestering (nothing) from player 2 in round  $r$ . The  $\rho$  multiplying the  $E[\cdot]$  expression in both of the aforementioned terms represents the likelihood that the contribution in round  $r$  is missed by player 2, who remains destitute as a result. Again, since a non-destitute player 2 never pesters, by Lemma 4.14, player 1's expected continuation utility from round  $r_0$  on for any destitute player that receives a contribution is 0.

On the other hand, player 1 expects to earn

$$\begin{aligned} \mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|h_1^r, \mathbf{D}] &= -\gamma_{\uparrow c} + \rho((1 - \rho)(-\gamma_{\downarrow p} + \mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|(h_1^r, c, p, o), \mathbf{D}]) + \\ &\quad \rho \mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|(h_1^r, c, o, o), \mathbf{D}]) \end{aligned}$$

following  $\sigma_{1,c}$ . The  $\rho(\cdot)$  term represents the expected payoff if the contribution in round  $r$  is missed by player 2; the expression inside of the parenthesis represents the expected payoff given that player 1 observes pestering with probability  $1 - \rho$  and nothing with probability  $\rho$ ; and, again, player 1's expected continuation utility from round  $r_0$  on is 0 from any destitute player that receives a contribution and thus becomes non-destitute.

We note that

$$\mathbb{E}^{(\sigma_{1,o}, \sigma_2^*)}[u_1|(h_1^r, o, p, c), \mathbf{D}] = \mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|(h_1^r, c, p, o), \mathbf{D}]$$

as  $\sigma_{1,o}$  and  $\sigma_{1,c}$  do not differ in rounds after  $r$ , and a destitute player 2 has seen the same history from its perspective (i.e., no contributions) and will thus play the same actions in continuation from either history.

By the same argument,

$$\mathbb{E}^{(\sigma_{1,o}, \sigma_2^*)}[u_1|(h_1^r, o, o, c), \mathbf{D}] = \mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|(h_1^r, c, o, o), \mathbf{D}]$$

Using these to compare  $\mathbb{E}^{(\sigma_{1,o}, \sigma_2^*)}[u_1|h_1^r, \mathbf{D}]$  and  $\mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|h_1^r, \mathbf{D}]$  gives us

$$\mathbb{E}^{(\sigma_{1,o}, \sigma_2^*)}[u_1|h_1^r, \mathbf{D}] - \mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|h_1^r, \mathbf{D}] \leq -(1 - \rho)^2 \gamma_{\downarrow p} < 0$$

Finally, it can be easily verified through similar arguments that against a Byzantine player, the expected utility of playing either  $\sigma_{1,c}$  and  $\sigma_{1,o}$  is exactly the same:

$$\mathbb{E}^{(\sigma_{1,o}, \sigma_2^*)}[u_1|h_1^r, \mathbf{B}] = \mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|h_1^r, \mathbf{B}]$$

We thus have  $\mathbb{E}^{(\sigma_{1,c}, \sigma_2^*)}[u_1|h_1^r] > \mathbb{E}^{(\sigma_{1,o}, \sigma_2^*)}[u_1|h_1^r]$  which implies that player 1 is better off contributing in round  $r$ , thus completing the contradiction.  $\square$

**THEOREM 4.11.** The strategy profile and set of beliefs  $(\sigma^*, \mu^*)$  make up a perfect Bayes equilibrium.

*Proof.* Lemmas 4.14, 4.13, and 4.21 and Theorems 4.20 and 4.23 show that  $\sigma^*$  specifies a best response given the set of beliefs  $\mu^*$ . Furthermore,  $\mu^*$ , by construction, is the result of using Bayes rule in continuation whenever possible.  $\square$

**THEOREM 4.12.** Let  $r < R$  be the current round, and suppose that player 1 played its best-response action in round  $r + 1$ . Then, if player 1 observes player 2 do nothing in round  $r + 1$ , then player 1 does nothing in round  $r$ .

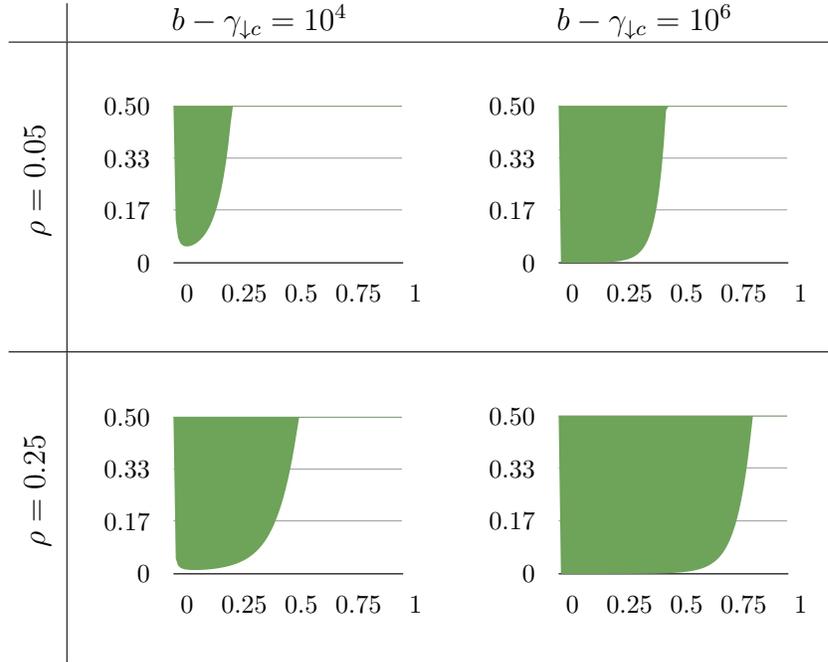
*Proof.* By Theorem 4.23, we know that it is in player 1's best interest to play a threshold strategy. By assumption, player 1 observes player 2 do nothing in round  $r + 1$ . If player 1's best-response action was to contribute in round  $r + 1$  then by Lemma 4.14, player 1 does nothing in round  $r$ . If player 1's best-response action was to do nothing in round  $r + 1$ , then by Lemma 4.24, player 1 does nothing in round  $r$ .  $\square$

## Auxiliary results

### 4.4 Characterizing the equilibrium

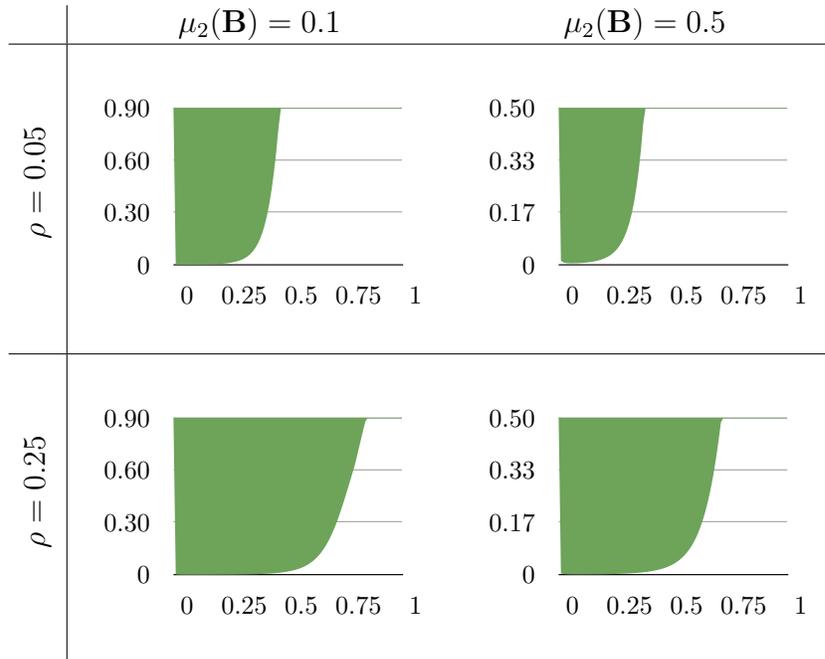
To understand the implications of §4.3 on the design of cooperative services, we explore, through simulation, the parameter space for which our cooperative equilibrium holds. We ask the following questions:

**1. What fraction of acquiescent peers suffices to motivate a player to pester?** The shaded areas in Figures 4.1 and 4.2 show (for different rates of network loss, different initial beliefs about the likelihood of player 1 being



**Figure 4.1:** Sufficient initial beliefs for a rational player 2 in player 1 being acquiescent to incentivize player 2 to pester ( $y$ -axis, shaded area) for varying amounts of acquiescent generosity ( $x$ -axis); network loss or  $\rho$  (top/bottom plots); and  $b - \gamma_{\downarrow c}$  (left/right plots). Simulation run with  $\gamma_{\uparrow p} = 1$ ,  $R = 20$ , and  $\mu_2(\mathbf{B}) = 0.5$ .

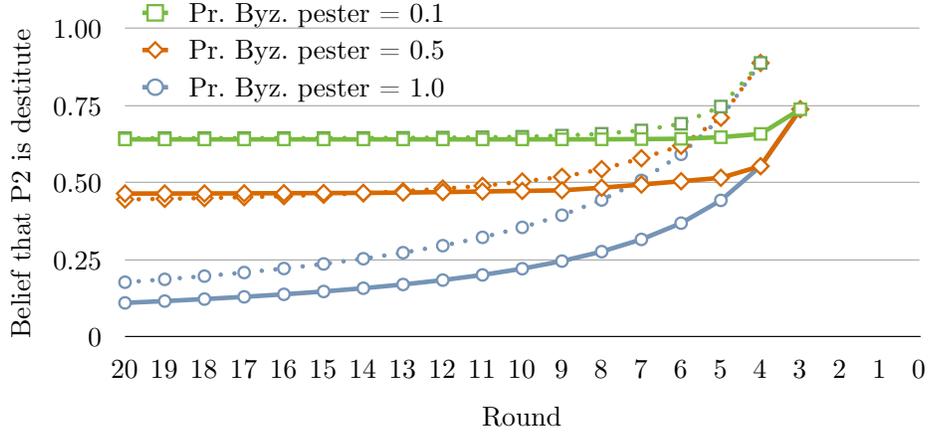
Byzantine, and different worth of receiving a contribution) the fraction of acquiescent peers that suffices to trigger player 2’s pestering, as a function of the probability  $((1 - \alpha)/(1 - \rho)^2)$  that an acquiescent player 1 will contribute if pestered, which we refer to here as *acquiescent generosity*. We assume that player 2 believes an acquiescent player 1 follows the acquiescent strategy; a Byzantine player 1 never contributes; and a rational player 1 only contributes in round  $R$ . This is a conservative estimate on the fraction of acquiescent peers sufficient to motivate player 2; in practice, the actual amount is likely to be lower than we report. As expected, for a given level of generosity, it is easier to incentivize player 2 if the value of the contribution increases and the likelihood of player 1 being Byzantine decreases. Given a highly lossy network, player 2



**Figure 4.2:** Sufficient initial beliefs for a rational player 2 in player 1 being acquiescent to incentivize player 2 to pester ( $y$ -axis, shaded area) for varying amounts of acquiescent generosity ( $x$ -axis); network loss or  $\rho$  (top/bottom plots); and  $\mu_2(\mathbf{B})$  (left/right plots). Simulation run with  $\gamma_{\uparrow p} = 1$ ,  $R = 20$ , and  $b - \gamma_{\downarrow c} = 10^5$ .

is also more willing to continue pestering, as it is more willing to attribute to network loss its failure to receive a contribution.

**2. How do player 1’s beliefs and the rate of network loss affect player 1’s willingness to contribute?** Player 1 may contribute only if its belief that player 2 is destitute is above a certain threshold. We show in Figure 4.3 how that threshold changes over the course of a game in which  $R = 20$ . For six configurations, with two different rates of network loss and three different probabilities that a Byzantine player 2 will pester, we plot the belief threshold and report the number of times that player 1 contributes in Table 4.2. For a given round, we assign player 1 some initial belief that player 2 is destitute and construct the game tree to determine whether that initial belief is sufficient to



**Figure 4.3:** Player 1’s belief thresholds. Solid lines represent  $\rho = 0.05$ ; dotted lines represent  $\rho = 0.25$ . Simulation run with  $\gamma_{\uparrow c}/\gamma_{\downarrow p} = 2$  and  $R = 20$ .

Pr. Byz. pester	$\rho = 0.05$	$\rho = 0.25$
0.1	14	17
0.5	5	9
1.0	2	3

**Table 4.2:** The maximum number of times player 1 contributes for each of the thresholds shown in Figure 4.3. Simulation run with  $\gamma_{\uparrow c}/\gamma_{\downarrow p} = 2$ ,  $R = 20$ , and  $\mu_1(\mathbf{B}) = 0.1$ .

motivate player 1 to contribute in that round; we use binary search to approximate the threshold value. As expected, when the game has only few rounds left and the cost from being pestered is not enough to overcome the cost of contributing, there is no threshold above which player 1 contributes. Note also that increasing  $\rho$  increases the belief threshold required to convince player 1 to contribute (as it reduces the expected threat from pestering) but also makes player 1 more likely to contribute when pestered, since past contributions are more likely to have been dropped. Also, the belief threshold increases as the likelihood of a Byzantine player 2 pestering decreases, since it becomes more in player 1’s interest to delay contribution, waiting to see whether player 2 will pester. However, when player 1 observes pestering, its belief that player 2

is destitute increases, and player 1 becomes more more willing to contribute. Finally, decreasing the relative cost of contributing ( $\gamma_{\uparrow c}/\gamma_{\downarrow p}$ ) has an obvious effect on player 1's likelihood to contribute (not shown).

**3. Too much generosity?** An intriguing conclusion from Figure 4.1 and 4.2 is that acquiescent generosity can make it much harder to motivate player 2 to pester. The reason is that the more generous acquiescent peers are, the easier it is for a rational player 2 to determine, from observed signals, whether player 1 is acquiescent or not, which in turn affects whether player 2 continues to pester. Figures 4.1 and 4.2 show the effects that an acquiescent peer's generosity has on cooperation. For higher levels of acquiescent generosity, we can only guarantee cooperation if such generosity is offset by a high  $\rho$  or  $b - \gamma_{\downarrow c}$ . Acquiescent generosity becomes a more obvious discriminant if a Byzantine player 1 never contributes, but it becomes less conspicuous with higher rates of network loss, which affects the observed generosity from player 2's perspective. As expected, player 2 is more willing to pester given a more valuable contribution.

## 4.5 Summary

Despite the presence of acquiescent peers in real-world cooperative systems, little attention has been given to their role in establishing rational cooperation. In this paper, we take the first step in understanding their function by showing that altruism is necessary and sufficient to motivate rational cooperation in the crucial last exchange between peers. Our results suggest that, while a small fraction of acquiescent peers is sufficient to spur rational peers into action even in systems with a large fraction of Byzantine peers, overly generous acquiescent peers can irreparably harm rational cooperation.

## 4.6 Appendix: Infinite-horizon games

### 4.6.1 The unbounded model

We focus on the differences between this model and the one presented in §4.1. As in §4.2, we assume that there are only Byzantine and rational participants. As we are dealing with an infinite-length game, we count the rounds in order, i.e., the game starts at round 0 and increases from there. Another consequence of an infinite-length game is that we must introduce discount factors on payoffs in future rounds; otherwise, payoffs may diverge as it becomes possible to achieve infinitely positive (or negative) payoffs. Thus, the utility functions for players 1 and 2 are:

$$u_1(C, \hat{P}) = - \left( \sum_{r \in \hat{P}} \delta^r 1_{\hat{P}}(r) \gamma_{\downarrow p} + \sum_{r \in C} \delta^r 1_C(r) \gamma_{\uparrow c} \right)$$

and

$$u_2(P, \hat{C}) = \delta^{\min(\hat{C})} H[|\hat{C}| - 1] b - \left( \sum_{r \in P} \delta^r 1_P(r) \gamma_{\uparrow p} + \sum_{r \in \hat{C}} \delta^r 1_{\hat{C}}(r) \gamma_{\downarrow c} \right)$$

where  $\delta$  is the discount factor, assumed to be common knowledge, and  $1_X(r) = 1$  if  $r \in X$  and is 0 otherwise. Note that these formulas are simply generalizations of the utility functions as defined in §4.2.

### 4.6.2 Equilibria of interest

The equilibria we focus on in this section have the following properties:

1. They are *finitely mixed*: they place positive probability only on a finite number of histories. We leave (arbitrarily) mixed strategies as future work.
2. A destitute player 2 believes it can earn positive payoff from the last

exchange. More formally, a destitute player 2 expects that following any equilibrium strategy  $\sigma^*$  after some history  $h_2^r$  with a rational player 1 results in non-negative utility:

$$E^{\sigma^*}[u_2|h_2^r, \mathbf{R}] \geq 0 \quad (4.7)$$

3. A non-destitute player 2 always prefers not to pester. In other words, we do not consider “unusual” equilibria in which player 2 might pester to dissuade player 1 from contributing.

### 4.6.3 What if a Byzantine player acts destitute?

In this section, we assume that if player 2 is Byzantine, it plays the following strategy.

DEFINITION 4.25. The *destitute strategy* is a strategy in which player 2 plays the rational strategy as though no contribution were ever observed.

LEMMA 4.26. Suppose there exists some positive probability that player 2 is a Byzantine player. Furthermore, suppose that Byzantine players follow the destitute strategy. Then for any round  $r$  and any history  $h_1^r$ , if  $\mu_1(\mathbf{B}|h_1^r) \geq \bar{\mu}_1$ , where

$$\bar{\mu}_1 = 1 - \frac{\gamma_{\uparrow c}}{\gamma_{\uparrow c} + \gamma_{\downarrow p}} \frac{1 - \delta}{1 - \rho} < 1 \quad (4.8)$$

then player 1 is better off doing nothing.

*Proof.* Let  $\sigma_1^*$  denote a strategy in which player 1 contributes after  $h_1^r$  given beliefs  $\mu_1(\mathbf{B}|h_1^r) \geq \bar{\mu}_1$  and maximizes player 1’s continuation payoff. Let  $\sigma_1'$  be the same strategy as  $\sigma_1^*$  except player 1 does nothing after  $h_1^r$  and then plays

as if it had contributed from round  $r + 1$  on.

Denoting  $V$  as the  $\delta$ -discounted continuation utility of playing  $\sigma_1^*$  with a (still) destitute player 2 after contributing, i.e.,

$$V = E^{\sigma_1^*}[u_1|(h_1^r, c), \mathbf{D}] \leq 0$$

then we have

$$E^{(\sigma_1^*, \sigma_2^*)}[u_1|h_1^r] = -\gamma_{\uparrow c} + \mu_1(\mathbf{B}|(h_1^r, c))V + \rho(1 - \mu_1(\mathbf{B}|(h_1^r, c)))V$$

whereas  $E^{(\sigma_1', \sigma_2^*)}[u_1|h_1^r] = V$ . It can be easily verified that  $E^{(\sigma_1', \sigma_2^*)}[u_1|h_1^r] > E^{(\sigma_1^*, \sigma_2^*)}[u_1|h_1^r]$  and thus player 1 is better off doing nothing if  $V = 0$  or

$$\mu_1(\mathbf{B}|h_1^r) > \bar{\mu}_1 \geq \frac{V(1 - \rho) + \gamma_{\uparrow c}}{V(1 - \rho)}$$

□

LEMMA 4.27. Suppose there exists some positive probability that player 2 is a Byzantine player. Furthermore, suppose that a Byzantine player 2 follows the destitute strategy. Then there exists no equilibrium (with the properties specified in §4.6.2) where a rational player 1 contributes with positive probability an unbounded number of times.

*Proof.* Suppose there exists some equilibrium  $(\sigma^*, \mu^*)$  in which player 1 contributes with positive probability an unbounded number of times. This implies that there must exist some sequence of signals

$$(\omega_1^{0,1}, \omega_1^{0,2}, \omega_1^{1,1}, \omega_1^{1,2}, \dots)$$

such that for any number of contributions  $m_c \geq 0$ , there exists a finite prefix of this sequence  $h_1^r$  that occurs with positive probability (given  $\sigma^*$ ) and that

(a) contains at least  $m_c$  contributions and (b) has a pester following the last contribution.<sup>9</sup>

Note that, while destitute and non-destitute rational player 2 are distinguishable, destitute and Byzantine player 2 are not. Since pestering has been observed in  $h_1^r$ ,  $\mu_1(\mathbf{D}|h_1^r) = 0$ . It follows that

$$\mu_1(\mathbf{B}|h_1^r) = \frac{\mu_1(\mathbf{B})}{\mu_1(\mathbf{B}) + \rho^{m_c}(1 - \mu_1(\mathbf{B}))}$$

It can be verified that for any finite continuation  $h_1^i$  of  $h_1^r$ ,  $\mu_1(\mathbf{B}|h_1^i) \geq \mu_1(\mathbf{B}|h_1^r)$ , as any signals that are observed by player 1 at most affect the second term in the denominator (the  $1 - \mu_1(\mathbf{B})$  term). Furthermore,

$$\lim_{m_c \rightarrow \infty} \frac{\mu_1(\mathbf{B})}{\mu_1(\mathbf{B}) + \rho^{m_c}(1 - \mu_1(\mathbf{B}))} = 1$$

so by selecting a sufficiently large  $m_c$  and corresponding finite prefix  $h_1^{\bar{r}_c}$  that has  $m_c$  contributions and a pester following them, we have

$$\mu_1(\mathbf{B}|h_1^{\bar{r}_c}) > \bar{\mu}_1$$

after which player 1 never contributes again. Contradiction.  $\square$

**THEOREM 4.28.** Suppose there exists some positive probability that player 2 is a Byzantine player. Furthermore, suppose that a Byzantine player 2 follows the destitute strategy. Then there exists no equilibrium (with the properties specified in §4.6.2) where a rational player 1 contributes or rational player 2 pesters.

*Proof.* Suppose there exists an equilibrium  $(\sigma^*, \mu^*)$  in which a rational player

---

<sup>9</sup>Case (a) holds by assumption; if case (b) did not hold, then this would imply that player 2, destitute or not, never pesters after  $m_c$  contributions. It follows then that player 1 would have been better off contributing only  $m_c - 1$  times.

1 contributes and player 2 pesters. By Lemma 4.27, there exists no equilibrium where a rational player 1 contributes with positive probability an unbounded number of times. Because  $\sigma^*$  is finitely mixed, we can consider all possible histories that follow from  $\sigma^*$  with positive probability and find the last round in which player 1 contributes and after which player 1 never contributes again. Denote this last round as  $\bar{r}_c$ . As in the proof of Theorem 4.3, it follows that a rational player 2 never pesters with positive probability starting from round  $\bar{r}_c$  in  $\sigma^*$  and therefore a rational player 1 is better off not contributing in round  $\bar{r}_c$ . Contradiction.  $\square$

#### 4.6.4 What if a Byzantine player never contributes?

In this section, we assume if player 1 is Byzantine, there is some probability it has crashed.

DEFINITION 4.29. The *crash strategy* is a strategy in which player 1 never contributes.

We denote the subset of Byzantine nodes that have crashed using a new type:  $\mathbf{S}$ .

LEMMA 4.30. Let  $(\sigma^*, \mu^*)$  be some equilibrium and let  $\phi(h_2^r)$  be the probability that player 2, following  $\sigma^*$ , will receive at least one contribution from a rational player 1 starting from some history  $(h_2^r, p)$ . If

$$\phi(h_2^r) < \frac{1}{\mu_2^*(\mathbf{R}|h_2^r)} \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}} \quad (4.9)$$

then player 2 does nothing in round  $r$ .

*Proof.* Suppose that there exists some equilibrium  $\sigma^*$  such that for some history  $h_2^r$ , condition (4.9) holds but player 2 pesters with positive probability. Letting  $\sigma_2'$  be the strategy in which player 2 does nothing in round  $r$  but follows

$\sigma_2^*$  as if it had pestered for the remainder of the game, we know that

$$-\gamma_{\uparrow p} + \delta \mathbf{E}^{\sigma^*}[u_2|(h_2^r, p)] \geq \delta \mathbf{E}^{(\sigma_{-2}^*, \sigma_2')}[u_2|(h_2^r, o)] \quad (4.10)$$

Since Byzantine players are expected to play independently of player 2's actions, we know that

$$\mathbf{E}^{\sigma^*}[u_2|(h_2^r, p), \mathbf{B}] = \mathbf{E}^{(\sigma_{-2}^*, \sigma_2')}[u_2|(h_2^r, o), \mathbf{B}]$$

Let  $V(p)$  represent player 2's expected utility given that a rational player 1 observed player 2 pester in round  $r$ ; define  $V(o)$  similarly. Then we have

$$\begin{aligned} \mathbf{E}^{\sigma^*}[u_2|(h_2^r, p), \mathbf{R}] &= (1 - \rho)V(p) + \rho V(o) \\ \mathbf{E}^{(\sigma_2', \sigma_{-2}^*)}[u_2|(h_2^r, o), \mathbf{R}] &= V(o) \end{aligned}$$

Plugging these into condition (4.10), we have

$$-\gamma_{\uparrow p} + \delta \mu_2^*(\mathbf{R}|h_2^r) ((1 - \rho)V(p) + \rho V(o)) \geq \delta \mu_2^*(\mathbf{R}|h_2^r) V(o)$$

By condition (4.7),  $V(o) \geq 0$ . In the best case, player 2 receives the contribution in the following round and incurs no additional costs, i.e.,  $\phi(h_2^r)(b - \gamma_{\downarrow c}) \geq V(p)$ :

$$-\gamma_{\uparrow p} + \delta \mu_2^*(\mathbf{R}|h_2^r) (1 - \rho) \phi(h_2^r) (b - \gamma_{\downarrow c}) \geq \delta (1 - \rho) \mu_2^*(\mathbf{R}|h_2^r) V(o) \geq 0$$

Moving around terms results in

$$\phi(h_2^r) \geq \frac{1}{\mu_2^*(\mathbf{R}|h_2^r)} \frac{1}{\delta(1 - \rho)} \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}} \geq \frac{1}{\mu_2^*(\mathbf{R}|h_2^r)} \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}}$$

a contradiction. □

LEMMA 4.31. Suppose there exists some positive probability that player 1 is a Byzantine player that has crashed and never contributes. Then there exists no equilibrium (with the properties specified in §4.6.2) where a destitute player 2 pesters with positive probability an unbounded number of times.

*Proof.* Suppose there exists some equilibrium  $(\sigma^*, \mu^*)$  in which a destitute player 2 pesters with positive probability an unbounded number of times. Formally, there exists some sequence of signals

$$(\omega_2^{0,1}, \omega_2^{0,2}, \omega_2^{1,1}, \omega_2^{1,2}, \dots)$$

such that any finite prefix of this sequence is reached with positive probability (given  $\sigma^*$ );  $\omega_1^{i,1} = o$  for all  $i$ ; and there does not exist some  $\bar{r}_p$  such that  $\omega_1^{i,2} = o$  for all  $i \geq \bar{r}_p$ . We denote any finite prefix of this sequence as  $h_2^r = (\omega_2^{0,1}, \omega_2^{0,2}, \dots, \omega_2^{r+1,1})$ .

Let  $\bar{\sigma}_{\mathbf{R},1}(\omega_2^{r,1} | h_2^r)$  represent player 2's expectation on the probability of observing  $\omega_2^{r,1}$  from a rational player 1 given that player 2 has observed  $h_2^r$ . We first show that

$$\lim_{R \rightarrow \infty} \prod_{r \in [0, R)} \bar{\sigma}_{\mathbf{R},1}(o | h_2^r) = 0 \tag{4.11}$$

Suppose not; suppose instead that the limit converges to some  $L > 0$ . Since  $\prod_{r \in [0, R)} \bar{\sigma}_{\mathbf{R},1}(o | h_2^r)$  is monotonically non-increasing with respect to  $R$ , this implies that there exists some round  $r_0$  such that for all  $r_1 \geq r_0$ ,

$$\prod_{r \in [0, r_1)} \bar{\sigma}_{\mathbf{R},1}(o | h_2^r) - L < \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c} - \gamma_{\uparrow p}} L$$

Moving terms around, we have

$$\lim_{R \rightarrow \infty} \prod_{r \in [0, R]} \bar{\sigma}_{\mathbf{R},1}(o|h_2^r) = L > \left(1 - \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}}\right) \prod_{r \in [0, r_1]} \bar{\sigma}_{\mathbf{R},1}(o|h_2^r)$$

This implies that

$$\lim_{R \rightarrow \infty} \prod_{r \in [r_1, R]} \bar{\sigma}_{\mathbf{R},1}(o|h_2^r) > 1 - \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}} \quad (4.12)$$

Since player 2 pesters an unbounded number of times, there must exist some round  $r_2 \geq r_1$  where player 2 pesters ( $\omega_2^{r_2,2} = p$ ). As in Lemma 4.30, let  $\phi(h_2^{r_2})$  be the probability that player 2 will receive at least one contribution from a rational player 1 starting from some history  $(h_2^{r_2}, p)$ . By definition,

$$\phi(h_2^{r_2}) = 1 - \lim_{R \rightarrow \infty} \prod_{r \in (r_2, R)} \bar{\sigma}_{\mathbf{R},1}(o|h_2^r)$$

It follows, from this expression and condition (4.12), that

$$\phi(h_2^{r_2}) = 1 - \lim_{R \rightarrow \infty} \prod_{r \in (r_2, R)} \bar{\sigma}_{\mathbf{R},1}(o|h_2^r) \leq 1 - \lim_{R \rightarrow \infty} \prod_{r \in (r_1, R)} \bar{\sigma}_{\mathbf{R},1}(o|h_2^r) < \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}}$$

However, given this condition, Lemma 4.30 states that player 2 should not have pestered after  $h_2^{r_2}$ . This contradiction completes the proof of condition (4.11).

Given condition (4.11), it follows that over time, if player 2 continues to hear nothing from player 1, then the belief that player 2 has about player 1 being a crashed Byzantine node approaches 1:

$$\lim_{R \rightarrow \infty} \mu_2^*(\mathbf{B}|h_2^r) = \lim_{R \rightarrow \infty} \frac{\mu_2^*(\mathbf{S}) + \mu_2^*(\mathbf{B} \setminus \mathbf{S})\beta}{\mu_2^*(\mathbf{S}) + \mu_2^*(\mathbf{B} \setminus \mathbf{S})\beta + \mu_2^*(\mathbf{R}) \prod_{r \in [0, R]} \bar{\sigma}_{\mathbf{R},1}(o|h_2^r)} = 1$$

where

$$\beta = \prod_{r \in [0, R]} (1 - (1 - \rho)\tau_1(c|h_2^r))$$

By the definition of the limit, there exists some round  $\bar{r}_p$  such that for all rounds  $r \geq \bar{r}_p$ ,

$$\mu_2^*(\mathbf{R}|h_2^r) = 1 - \mu_2^*(\mathbf{B}|h_2^r) < \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}}$$

and thus

$$\frac{1}{\mu_2^*(\mathbf{R}|h_2^r)} \frac{\gamma_{\uparrow p}}{b - \gamma_{\downarrow c}} > 1 \geq \phi(h_2^r)$$

By Lemma 4.30, it follows that in every round after  $\bar{r}_p$ , player 2 does nothing. Contradiction. □

**THEOREM 4.32.** Suppose there exists some positive probability that player 1 is a Byzantine player that has crashed and never contributes. There exists no equilibrium (with the properties specified in §4.6.2) where a rational player 1 contributes or player 2 pesters.

*Proof.* Suppose there exists an equilibrium  $(\sigma^*, \mu^*)$  in which a rational player 1 contributes and player 2 pesters. By Lemma 4.31, there exists no equilibrium where a rational player 2 pesters with positive probability an unbounded number of times. Because  $\sigma^*$  is finitely mixed, we can consider all possible histories that follow from  $\sigma^*$  with positive probability and find the last round in which player 2 pesters and after which player 2 never pesters again. Denote this last round as  $\bar{r}_p$ . As in the proof of Theorem 4.3, it follows that a rational player 1 never contributes with positive probability starting from round  $\bar{r}_p + 1$  in  $\sigma^*$  and therefore a rational player 2 is better off not pestering in round  $\bar{r}_p$ . Contradiction. □

# Chapter 5

## Rationally Colluding

In the previous chapters, we described how to guarantee that rational individuals will not deviate in the presence of Byzantine, acquiescent, and rational nodes. Preventing *individual* deviations, however, is unlikely to be sufficient to build robust cooperative services. The social nature of these services suggests that nodes will develop, or may have already established, a rich web of relationships (based, for instance, on friendship or on belonging to the same organization), which may cause *coalitions* of nodes to collude and deviate together [80]. We submit that cooperative services that ignore the possibility of collusion do so at their own peril. That most cooperative services still choose to do so is a testament to how hard it is to address the threat posed by collusion to the stability of an equilibrium.

The literature offers two approaches to address this threat. The first is to model collusion as a fault and colluding nodes as Byzantine [24, 49, 86]. The limitations of this approach are obvious: since basic distributed computing primitives such as consensus and reliable broadcast cannot be implemented if more than one third of the nodes are Byzantine [75], modeling colluding nodes as Byzantine imposes a cap on the number and size of coalitions that is both artificial (since it lacks a game theoretic basis) and dangerously low.

The second approach is to deny any benefit to colluders. If the equilibrium is a best response not just to every individual, but also to every possible

coalition, then collusion poses no harm to the equilibrium’s stability, since nodes gain no benefit by colluding. This is the aim of solution concepts such as strong Nash [31] and  $k$ -resilient equilibria [19, 21], which offer this guarantee, respectively, for all conceivable coalitions and for arbitrary coalitions of size at most  $k$ . Coalition-proof Nash equilibria [33] similarly ensure that nodes cannot gain any benefit from colluding and deviating in a self-enforcing way (such that there cannot be further profitable deviations from sub-coalitions).

Our work is motivated by what we believe to be a critical flaw of the second approach: its inability to account for the role played by social factors that are impossible to completely capture a priori (such as friendships or shared participation in social groups) in determining whether a node will consider a strategy to be a best response. Intuitively, nodes in coalitions formed on the basis of social “side channels” are likely to know more about each other, trust each other more, and in general be able to hold stronger assumptions about one another than about non-coalition members. Since stronger assumptions typically lead to more efficient protocols, techniques that aim to deny benefits to coalitions face a fundamentally uphill battle: as we show in §5.1, identifying a single strategy that is a best response both inside and outside every possible coalition is very hard.

To overcome this impasse, this paper introduces and begins to explore a fundamentally different approach to dealing with coalitions. The key observation is that the fundamental property provided by an equilibrium is *stability*—in that nodes do not want to deviate—and that while finding a single best response between all nodes is sufficient to achieve stability, it is not *necessary*: insisting on this requirement as the means to providing stability puts the cart (i.e., best responding) before the horse (i.e., stability). As a first concrete step in this new direction, we introduce two new solution concepts that do not require fighting the strong headwinds of social relationships to guarantee stable cooperative services; instead, they explicitly model the advantages that coalition members have while ensuring that nodes do not want to deviate from the specified equilibrium. Both solution concepts achieve stability

through a simple observation: coalitions (including the trivial singleton coalition of one non-colluding node) will not deviate from an equilibrium as long as the equilibrium specifies a best-response strategy for every *coalition*. Thus, the strategy a node follows depends on whom the node is colluding with, allowing the equilibrium to specify how nodes can benefit from their coalitions.

The first solution concept, *k-indistinguishability*, achieves stability through a guarantee that, while stronger than necessary, is attractively simple. In a *k*-indistinguishable equilibrium, the actions performed by a node within its coalition may depend on who belongs to the coalition, but the actions towards those with whom that participant is not colluding are unaffected. Thus, in a *k*-indistinguishable equilibrium, nodes cannot tell whether another peer, with whom they are not colluding, is itself part of some other coalition (of at most *k* nodes). The second solution concept, *k-stability*, instead adheres to the conditions necessary for stability: like *k*-indistinguishability, *k*-stable equilibria specify a strategy per coalition that is a best response to the strategies played by all other possible coalitions; unlike *k*-indistinguishability, the actions that a node takes as a part of a *k*-stable equilibrium may be informative about whether it is colluding and with whom. Finally, because *k*-stability and *k*-indistinguishability allow nodes to change their strategies depending on whom they are colluding with, strategy profiles—traditionally used by equilibria to specify a single best-response strategy per node—cannot capture the range of strategies that a node may play. Instead, we use *strategy functions*, a new construct that lets us express a node’s strategy as a function of the coalition the node belongs to.

**Our contributions.** In summary, our new contributions to the treatment of collusion in cooperative services are as follows:

- We illustrate the limits of generalizing Nash equilibria that prevent colluding nodes from receiving any benefit. Specifically, we show that requiring that a single strategy be a best response for every node, regardless of whether it is colluding, does not admit an equilibrium in several scenarios

that commonly arise in cooperative services.

- We decouple the fundamental property that defines an equilibrium—stability—from the requirement that a single strategy be a best-response. This requirement, while sufficient, is not necessary when nodes may collude. We take a first step at leveraging this separation by introducing (a) a new construct, strategy functions, that allows us to describe, for each node and each possible coalition it may be part of, the strategies the node will play, and (b) two new solution concepts,  $k$ -indistinguishability and  $k$ -stability, that admit a strategy function as an equilibrium if no coalition wants to deviate from its specified strategy.
- We demonstrate the applicability and utility of specifying a strategy per coalition by showing how our solution concepts admit useful equilibria in the same scenarios where traditional solution concepts could not.

**Organization of chapter.** §5.1 demonstrates the limits of generalizing traditional equilibria in the context of several common scenarios encountered in many cooperative services. §5.2 defines our two new solution concepts— $k$ -indistinguishability and  $k$ -stability—and demonstrates how these solution concepts overcome challenges faced by traditional approaches.

## 5.1 Disincentivizing coalitions

Solution concepts such as strong Nash equilibria and  $k$ -resilience specify, for each node, a single best response in which a node’s actions towards a peer do not depend on whether the two are colluding. However, if coalition members trust each other more than other nodes, the practical applicability of these solution concepts are fundamentally limited. To illustrate this point, we describe techniques and scenarios likely to occur in cooperative services where the stronger assumptions that insiders can rely on when dealing with one another hamper the ability to achieve  $k$ -resilience. These examples are by no means

comprehensive; rather, our goal is to provide a taste of the larger challenges faced by solution concepts that aim to discourage coalition formation.

As a reminder, before we proceed, we repeat the definition of  $k$ -resilience here. Recall that  $k$ -resilience generalizes the Nash equilibrium (Definition 2.1) by requiring that the strategy profile be a best response (i.e., admit no profitable deviations) not only for every individual node (as required by a Nash equilibrium) but also for any coalition of up to size  $k$ .

DEFINITION 3.3. A strategy profile  $\sigma^*$  is a  $k$ -resilient equilibrium if, for all  $K \subseteq N$  such that  $|K| \leq k$ , there does not exist an alternate strategy  $\sigma'_K$  such that for all  $x \in K$ ,

$$U_x(\sigma'_K, \sigma^*_{-K}) > U_x(\sigma^*)$$

As a Nash equilibrium is simply a 1-resilient equilibrium, we generally focus on  $k$ -resilient equilibria where  $k \geq 2$ . Note that a strong Nash equilibrium is a  $n$ -resilient equilibrium.

As noted earlier in §3.2, we use the weak version of  $k$ -resilience to prove our negative results; our results therefore apply to stronger notions of  $k$ -resilience that guarantee stability even if coalitions are willing to deviate for less [19, 21]. Our negative results also do not rely on coalition members being able to “cheap talk”, i.e., communicate at no cost, during the game [44, 50].

We can analogously define a Bayesian notion of  $k$ -resilient and strong Nash equilibrium similar to a Bayes (Nash) equilibrium.

DEFINITION 5.1. A strategy profile and set of beliefs  $(\sigma^*, \mu^*)$  is a  $k$ -resilient Bayes equilibrium if for all  $K \subseteq N$  such that  $|K| \leq k$ , there does not exist some strategy  $\sigma'_K$  such that for all  $x \in K$ ,

$$E^{(\sigma'_K, \sigma^*_{-K}), \mu^*} [U_x] > E^{\sigma^*, \mu^*} [U_x]$$

where  $E^{\sigma, \mu}[U_x]$  represents  $x$ 's expected payoff from the strategy profile  $\sigma$  with belief  $\mu_x$ , given that  $x \in K$ .

It is important to note that all the solution concepts and equilibria we discuss in this paper are notions from *non-cooperative* game theory. There has also been extensive work in *cooperative* game theory (see any game theory text, e.g., [87], for a survey of related work) that explicitly studies the formation of coalitions in games where players are trying to work together. Cooperative and non-cooperative game theory significantly differ in focus: cooperative game theory focuses on interactions *within* a coalition—how and which coalitions form (players join a coalition based on the benefit the coalition offers) and how payoffs are allocated among coalition members (based on each member's value to the coalition)—whereas non-cooperative game theory focuses instead on the interactions *between* competing players (which, in our case, consist of exogenously-determined coalitions and non-colluding nodes).

### 5.1.1 Can trusted third parties limit equilibria?

Cooperative services often rely on a trusted third party to incentivize cooperation among nodes. This type of trust, which in some cases is indispensable (e.g., to implement fair exchange [72, 90]), is unnecessary among coalition members; indeed, perhaps surprisingly, it can actually render  $k$ -resilient equilibria impossible to achieve.

We illustrate this point through the following game, which models the fundamental choice that each node makes in P2P cooperative services: should I contribute my fair share?

For simplicity, we assume, in this example, that all nodes are rational.

**DEFINITION 5.2.** The mediated pairwise-exchange game is a  $R$ -repeated game where, in each round  $r \in \{1, \dots, R\}$ , each node  $x \in N$ :

1. Decides (simultaneously) on some set of peers  $M_x^r \subseteq N \setminus \{x\}$  to use a mediator with.
2. Observes which peers are using a mediator with  $x$ .
3. Decides on some set of peers  $\Gamma_x^r \subseteq N \setminus \{x\}$  to contribute to; any other peer is snubbed.
4. Receives a contribution from a peer  $y$  if  $y$  contributed to  $x$  and either (a)  $y$  did not use a mediator with  $x$ , or (b)  $x$  contributed to  $y$ . Denote the set of all such  $y$  as  $C_x^r$ , i.e.,  $y \in C_x^r$  iff  $x \in \Gamma_y^r \wedge (x \notin M_y^r \vee y \in \Gamma_x^r)$ .

$x$  pays  $\gamma$  per peer that  $x$  contributes to and  $\epsilon$  per peer that  $x$  uses a mediator with.  $x$  earns  $b > 2\gamma + \epsilon$  per received contribution, for a round payoff of  $v_x^r = |C_x^r|b - |\Gamma_x^r|\gamma - |M_x^r|\epsilon$ . A node's total payoff is the sum of all round payoffs:  $\sum_{r=1}^R v_x^r$ .

While this game resembles a finitely-repeated prisoner's dilemma, the mediator, who can serve as a trusted third party and ensure a fair pairwise exchange, enables the existence of Nash equilibria in which contribution occurs (without the mediator, no such equilibrium exists).

**THEOREM 5.3.** Let  $\sigma^*$  be a strategy profile in the mediated pairwise-exchange game in which a node  $x$ , following  $\sigma_x^*$ :

- Contributes to a peer  $y$ , using a mediator only in round  $R$ , iff (a)  $x$  and  $y$  have never snubbed each other in the past and (b)  $x$  and  $y$  have not used a mediator in any round other than  $R$ .
- Snubs a peer  $y$  without a mediator otherwise.

Then  $\sigma^*$  is a Nash equilibrium.

*Proof.* Same as the backwards-induction half of the proof of Theorem 5.15.  $\square$

The Nash equilibrium in Theorem 5.3 uses the mediator to ensure cooperation in the last round, which encourages cooperation in prior rounds without the mediator. We now prove that this same mediator precludes the existence of  $k$ -resilient equilibria. The reason, essentially, is that using the mediator, which incurs cost, is undesirable between colluding nodes (Lemma 5.5) but necessary to ensure cooperation between two non-colluding nodes (Lemma 5.4). This tension makes it impossible for a single strategy to be a node's best response regardless of how it colludes (Theorem 5.6).

LEMMA 5.4. In any  $k$ -resilient equilibrium of the mediated pairwise-exchange game where some node contributes, the last time in the game that any node contributes with positive probability to a peer must always involve a mediator.

*Proof.* By contradiction. Fix some  $k$ -resilient equilibrium  $\sigma^*$ , where the last time that any node contributes with positive probability does not involve a mediator with positive probability (if there exist multiple such node/peer pairings, choose one arbitrarily). During this “last contribution,” let  $x$  be the node that contributes,  $y$  be the receiving peer, and  $\alpha$  be the probability that  $x$  contributes to  $y$  after deciding not to use a mediator with  $y$ . By assumption,  $\alpha > 0$ .

Since  $\sigma^*$  must be a best response regardless of who is colluding, suppose  $x$  and  $y$  are not colluding. Then it must be the case that, in  $\sigma^*$ ,  $y$  snubs  $x$  during the last contribution if  $x$  does not use a mediator:  $y$  expects to earn, from  $x$ 's contribution,  $\alpha b$  without incurring the cost of contributing; moreover, since this is the last time a contribution occurs with positive probability,  $y$ 's choice of whether to snub  $x$  does not negatively impact  $y$ 's continuation payoff. It follows that  $x$  could profitably deviate from  $\sigma^*$  by always snubbing  $y$  during the last contribution if  $x$  does not use a mediator: doing so would save  $x$  an expected cost of  $\alpha\gamma$  with no negative effect on  $x$ 's continuation payoff. Contradiction. □

LEMMA 5.5. In any  $k$ -resilient equilibrium of the mediated pairwise exchange game where some node contributes, the last time in the game that any node contributes with positive probability to a peer must never involve a mediator.

*Proof.* By contradiction. Fix some  $k$ -resilient equilibrium  $\sigma^*$  where the last time that any node contributes with positive probability also involves a mediator with positive probability (if there exist multiple such node/peer pairings, choose one arbitrarily). During this “last contribution,” let  $x$  be the node that contributes;  $y$  be the peer;  $\alpha > 0$  be the probability that  $x$  decides to use a mediator with  $y$ ; and  $\psi_x$  ( $\psi_y$ ) be the probability that  $y$  ( $x$ ) observes a contribution from  $x$  ( $y$ ) in expectation over all possible combinations of  $x$  and  $y$ ’s choices regarding using a mediator and contributing with one another.

Since  $\sigma^*$  must be a best response regardless of who is colluding, suppose  $x$  and  $y$  are colluding. Consider an alternate strategy profile  $\sigma'$  in which all nodes play the same actions with the same probabilities as in  $\sigma^*$ , except, during the last contribution,  $x$  and  $y$  do not use a mediator with one another,  $x$  ( $y$ ) contributes to  $y$  ( $x$ ) with probability  $\psi_x$  ( $\psi_y$ ), and  $x$  and  $y$  subsequently play actions as if  $x$  and  $y$  had instead followed  $\sigma^*$ . It follows that the payoffs for  $x$  and  $y$  are exactly the same, with the exception of the payoffs that  $x$  and  $y$  receive from one another during the last contribution, where (1)  $x$  and  $y$ ’s expected benefit remains the same, (2)  $x$ ’s expected cost is strictly lower since  $x$  contributes with the same probability in expectation without the cost of a mediator ( $\alpha\epsilon > 0$ ), and (3)  $y$ ’s expected cost is no higher (and is lower if  $y$  was using a mediator in  $\sigma^*$ ). Thus,  $x$  is better off and  $y$  is no worse off. Contradiction. □

THEOREM 5.6. There exists no  $k$ -resilient equilibrium in the mediated pairwise-exchange game.

*Proof.* Lemmas 5.4 and 5.5 imply that there exists no  $k$ -resilient equilibrium where nodes contribute. Further, a strategy profile  $\sigma$  in which all nodes snub and earn 0, while a Nash equilibrium, is not a  $k$ -resilient equilibrium. To see why, consider an alternate strategy profile  $\sigma'$  and some coalition  $K$  (such that  $|K| \geq 2$ ) where no one uses mediators and only members of  $K$  contribute to one another.  $\sigma'$  earns  $K$ 's members payoffs of  $(|K| - 1)R(b - \gamma) > 0$  each, making it a profitable deviation from  $\sigma$ .  $\square$

### 5.1.2 What if nodes may fail?

When nodes may fail, a node's best response will generally depend on the probability with which it expects other nodes may fail. Greater trust and access to more information (e.g., concerning the frequency with which fellow coalition nodes are patched) may allow nodes within a coalition to reasonably believe that fellow coalition members have a lower probability of failing than outsiders. Unfortunately, even a slightly lower failure probability can make  $k$ -resilience practically unachievable.

We illustrate this point using a simple single-shot simultaneous game that models a simplified version of secret-sharing [19, 98]. In this game, each node wants to reconstruct a secret that requires the node to request shares from its peers. These peers deliver the requested shares unless they fail (e.g., by crashing). Each node must then decide how many shares to request: requesting more shares incurs more cost, but requesting fewer shares may result in the node being unable to reconstruct the secret because of peer failures. In this example, we assume that nodes are either rational or may fail by failing to send the requested share.

DEFINITION 5.7. The simple secret-sharing game is a single-shot, simultaneous game in which every node  $x \in N$ :

1. Selects a set  $\Gamma_x \subseteq N \setminus \{x\}$  of nodes to request shares from.

2. Pays  $|\Gamma_x|\gamma$  for this request.
3. Receives shares from some set  $C_x \subseteq \Gamma_x$ .
4. Earns benefit  $b > |N|\gamma$  iff  $|C_x| \geq m$ , where  $m$  is the number of shares that  $x$  must gather from its peers before being able to reconstruct the secret.

The simple secret-sharing game is a decision theory problem: a node's choice does not affect its peers' outcomes.<sup>1</sup> This is intentional: our goal is to show that, despite the game's simplicity, it is often impossible to find  $k$ -resilient equilibria. To account for a node's beliefs regarding how likely its peers are to fail, we use  $k$ -resilient Bayes equilibria (Definition 5.1). In this game, a strategy profile  $\Gamma$  represents the peers that each node requests from. A set of beliefs  $\mu$  represents the view of each node, given the set of peers it is colluding with, regarding the likelihood that any peer will successfully deliver its share if requested. In other words,  $\mu$  represents each node  $x$ 's view of the likelihood that a peer in  $\Gamma_x$  will also be in  $C_x$ . An equilibrium in the simple secret-sharing game is some  $(\Gamma^*, \mu^*)$  where no node  $x$ , colluding with any  $(k-1)$  peers, could do any better in expectation requesting shares from some set  $\Gamma'_x \neq \Gamma_x^*$ . More formally, for all  $K \subseteq N$  such that  $|K| \leq k$ , there is no  $\Gamma'_x$  such that for all  $x \in K$ ,

$$H[|C'_x| - m]b - |\Gamma'_x|\gamma > H[|C_x^*| - m]b - |\Gamma_x^*|\gamma$$

where  $H[i]$  is the discrete unit step function.<sup>2</sup>

**THEOREM 5.8.** Let  $(\Gamma^*, \mu^*)$  be a  $k$ -resilient Bayes equilibrium of the simple secret-sharing game in which some node  $x \in N$  believes that a peer  $y$  will fail with probability  $\mu_x^*$  if  $x$  and  $y$  are not colluding and  $\mu_x^* - \epsilon$  if  $x$  and  $y$

<sup>1</sup>If the game were sequential, the choice of some node  $x$  to request a share from some peer  $y$  could inform  $y$  of whether  $x$  has failed. However, finding  $k$ -resilient equilibria is no less challenging, since (1) there is at least one node (the first node to move) that will never have such a signal and (2) even if  $x$  successfully requests a share from  $y$ ,  $x$  could subsequently fail before  $y$ 's turn.

<sup>2</sup>As a reminder,  $H[i] = 1$  if  $i \geq 0$ ; otherwise  $H[i] = 0$ .

are, where  $\epsilon > 0$ . Then either  $x$  requests secrets from no one or everyone, i.e.,  $\Gamma_x^* \in \{\emptyset, N \setminus \{x\}\}$ .

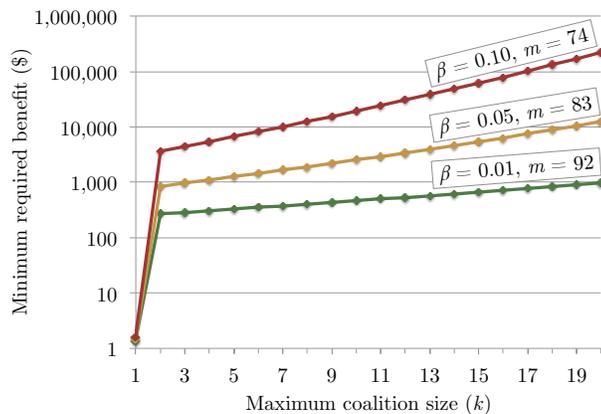
*Proof.* Suppose  $k = 2$  and  $K = \{x, y\}$ , i.e.,  $x$  and  $y$  are colluding. If  $x$  incurs more cost requesting shares than it earns in expectation from reconstructing the secret (e.g., because of high rates of failure), then  $\Gamma_x^* = \emptyset$ . Otherwise, suppose  $x$  requests shares from peers in  $\Gamma_x^* \neq \emptyset$ . It is obvious that since  $x$  believes that  $y$  will fail with probability  $\mu_x^* - \epsilon$ , which is lower than the probability of any other peer  $z \neq y$  failing ( $\mu_x^*$ ),  $x$  should always request shares from  $y$ , so any 2-resilient  $\Gamma_x^*$  must contain  $y$ . However, as  $y$  can be any peer, the only  $\Gamma_x^*$  that is guaranteed to contain all possible  $y$  is  $\Gamma_x^* = \{y \mid y \in N \wedge y \neq x\} = N \setminus \{x\}$ . Finally, as  $k$ -resilience implies 2-resilience, this result applies to  $k$ -resilience for  $k \geq 2$ .  $\square$

A node that wants to reconstruct the secret rarely wants to request shares from all of its peers, since the cost of these additional requests is not worth the slight insurance that redundant shares provide. However, in such cases, it follows from Theorem 5.8 that no  $k$ -resilient Bayes equilibrium exists. Therefore, the only scenarios in which a node wants to reconstruct the secret as a part of a  $k$ -resilient Bayes equilibrium are those in which the secret's value is sufficiently high to justify requesting shares from all peers to maximize the likelihood of success.

Figure 5.1 quantifies what this value must be, using example numbers based on a movie-streaming context:  $n = 100$  nodes; each node expects that coalition members never fail<sup>3</sup> and that non-coalition members fail with independent probability  $\beta$ ;  $m$  is set such that, given an independent failure probability of  $\beta$ , there is at least a 0.99999 chance that at least  $m$  peers, out of  $n - 1$  possible peers, will not fail; and  $\gamma$  is set to  $(1500 \text{ Kbps}) \times (2 \text{ hours}) \times$

---

<sup>3</sup>While this may seem extreme, note that this is exactly what failure-aware  $k$ -resilient solution concepts, such as  $(k, t)$ -robustness [19, 21], require: nodes do not deviate assuming that the coalition and set of faulty nodes do not overlap.



**Figure 5.1:** In the simple secret-sharing game (Definition 5.7), the minimum benefit needed for a  $k$ -resilient equilibrium where nodes attempt to reconstruct the secret.

$(\$1/\text{GB})/(m+1)$ . As  $k$  increases, the expected probability of a coalition member reconstructing the secret increases, thus making it more difficult to convince such a node to request shares from every other peer. Note that while Figure 5.1 implies that the minimum required benefit goes up as probability of failure goes up, this is an artifact of how we define  $m$ ; in reality, the minimum required benefit goes up as the probability of failure goes down, as expected.

As Figure 5.1 shows, even with coalitions of at most two nodes and beliefs that non-coalition nodes fail with probability 0.01, a 2-resilient equilibrium exists only if a node values a two-hour movie, which incurs  $\gamma(n-1) > \$1.37$  in communication costs, at over \$268.95!

### 5.1.3 Do nodes want to punish one another?

Cooperative services often incentivize nodes not to deviate by relying on the threat of punishment. In this section, we show that punishments that hurt both the enforcing and receiving nodes are never used within a coalition, and other forms of punishment will be difficult, if not impossible, to achieve in real-world scenarios. We illustrate this through a simplified version of the mediated

pairwise-exchange game.<sup>4</sup>

**DEFINITION 5.9.** The simple pairwise-exchange game is an infinitely-repeated game where, in each round  $r \geq 0$ , every node  $x \in N$ :

1. Simultaneously decides on some set of nodes  $\Gamma_x^r \subseteq N \setminus \{x\}$  that it will contribute to; any node not in  $\Gamma_x^r$  is snubbed.
2. Observes which peer  $y \neq x$  contributed to it; let  $C_x^r$  denote the set of all such  $y$ , i.e.,  $y \in C_x^r$  iff  $x \in \Gamma_y^r$ .

$x$ 's round payoff is  $v_x^r = |C_x^r|b - |\Gamma_x^r|\gamma$ , where  $b > \gamma$ .  $x$ 's total payoff is the  $\delta$ -weighted sum of the round payoffs:  $\sum_{r=0}^{\infty} \delta^r v_x^r$ .

**THEOREM 5.10.** Let  $\sigma^*$  be a  $k$ -resilient equilibrium in the simple pairwise-exchange game in which some contribution occurs. In other words,  $\sigma^*$  specifies that, at some point in the game, a node  $y$  contributes to some node  $x$ , who “rewards”  $y$  if  $y$  contributes and “punishes”  $y$  if  $y$  snubs  $x$ . Then either (1)  $x$  must prefer punishing to rewarding  $y$ , and/or (2)  $x$  punishing  $y$  is not a  $k$ -resilient best response (i.e.,  $x$  may threaten to punish  $y$ , but, given the opportunity,  $x$  and  $y$  can profitably deviate by not following through).

*Proof.* Fix some  $k$ -resilient equilibrium  $\sigma^*$  in which contribution occurs and, unlike condition (1) above,  $x$  is no worse off rewarding  $y$ . We prove that condition (2) follows:  $x$  punishing  $y$  is not a  $k$ -resilient best response. Let  $r$  be the round in which this contribution occurs,  $U_x(\sigma^* | (\Gamma_x^r, C_x^r \cup \{y\}))$  denote  $x$ 's continuation payoff from rewarding  $y$ , and  $U_x(\sigma^* | (\Gamma_x^r, C_x^r \setminus \{y\}))$  denote  $x$ 's

<sup>4</sup>While we could use the mediated pairwise-exchange game to illustrate this point, we instead use a game with an infinite horizon (which enables the existence of Nash equilibria where contribution occurs) and no mediator as the mediator already makes  $k$ -resilient equilibria impossible to achieve.

continuation payoff for punishing  $y$ . We have:

$$U_x(\sigma^* | (\Gamma_x^r, C_x^r \cup \{y\})) \geq U_x(\sigma^* | (\Gamma_x^r, C_x^r \setminus \{y\})) \quad (5.1)$$

Denote  $y$ 's continuation payoff from contributing to and snubbing  $x$  as  $U_y(\sigma^* | (\Gamma_y^r \cup \{x\}, C_y^r))$  and  $U_y(\sigma^* | (\Gamma_y^r \setminus \{x\}, C_y^r))$ , respectively. As  $y$  contributes to  $x$  as a part of a  $k$ -resilient equilibrium,  $y$  must be no worse off doing so:

$$|C_y^r|b - |\Gamma_y^r \cup \{x\}| \gamma + U_y(\sigma^* | (\Gamma_y^r \cup \{x\}, C_y^r)) \geq |C_y^r|b - |\Gamma_y^r \setminus \{x\}| \gamma + U_y(\sigma^* | (\Gamma_y^r \setminus \{x\}, C_y^r))$$

Unsurprisingly, it follows that  $y$ , in continuation, is worse off being punished than being rewarded:

$$U_y(\sigma^* | (\Gamma_y^r \cup \{x\}, C_y^r)) \geq \gamma + U_y(\sigma^* | (\Gamma_y^r \setminus \{x\}, C_y^r)) > U_y(\sigma^* | (\Gamma_y^r \setminus \{x\}, C_y^r)) \quad (5.2)$$

Suppose  $K = \{x, y\}$ , and let  $\sigma'_K$  specify the same actions as in  $\sigma^*$ , except  $x$  and  $y$  play  $\sigma^*$  as if  $y$  contributed even if  $y$  snubbed  $x$ . We can see that by inequality (5.1),

$$U_x((\sigma'_K, \sigma_{-K}^*) | (\Gamma_x^r, C_x^r \setminus \{y\})) = U_x(\sigma^* | (\Gamma_x^r, C_x^r \cup \{y\})) \geq U_x(\sigma^* | (\Gamma_x^r, C_x^r \setminus \{y\}))$$

and, by inequality (5.2),

$$U_y((\sigma'_K, \sigma_{-K}^*) | (\Gamma_y^r \setminus \{x\}, C_y^r)) = U_y(\sigma^* | (\Gamma_y^r \cup \{x\}, C_y^r)) > U_y(\sigma^* | (\Gamma_y^r \setminus \{x\}, C_y^r))$$

Thus,  $x$  punishing  $y$  is not a  $k$ -resilient best response.  $\square$

To get a sense of Theorem 5.10's impact, consider, for simplicity, a  $k$ -resilient equilibrium  $\sigma^*$  that uses punishment to encourage nodes to continually contribute to one another. If all nodes continually contribute to one another,

without fail, each node earns a payoff of

$$\frac{1}{1-\delta}(n-1)(b-\gamma) \quad (5.3)$$

We now consider how Theorem 5.10 applies to many forms of punishment, including:

*Grim trigger (global).* Node  $x$  contributes to another node  $y$  iff every node has always contributed in the past. A global grim trigger can incentivize cooperation only if:

$$\frac{1}{1-\delta}(n-1)(b-\gamma) \geq (n-1)b \quad (5.4)$$

where the right-hand side represents the payoff  $x$  receives from immediately snubbing everyone.

Consider the scenario in which  $x$  has observed  $y$  snubbing it in some round  $r$ . If  $x$  administers a global grim trigger, it earns at most  $(n-2)b$  (if  $y$  only snubbed  $x$ ) in round  $r+1$  and nothing after. If  $x$  and  $y$ , as a coalition, simply pretended nothing happened and continuing contributing, then  $x$  would earn the payoff specified in expression (5.3) if  $x$  believes that  $y$  only snubbed  $x$ .<sup>5</sup> By inequality (5.4), this payoff is larger than the  $(n-2)b$  that  $x$  would have earned administering the global grim trigger.

*Grim trigger (local).* Node  $x$  contributes to another node  $y$  iff neither node has snubbed in the past. This is not  $k$ -resilient since  $x$  will earn

$$\frac{1}{1-\delta}(n-2)(b-\gamma) < \frac{1}{1-\delta}(n-1)(b-\gamma)$$

---

<sup>5</sup>A node  $x$  in the simple pairwise-exchange game can only observe how a peer  $y$  behaves towards  $x$ , so  $x$  does not know who else  $y$  may have snubbed. However, since observing snub from  $y$  is a zero-probability event, we can assign any beliefs we want (see, e.g., [87]). Moreover, if  $x$  and  $y$  are colluding, it is possible they could communicate (e.g., via cheap talk).

*Forgiving trigger.* Node  $x$  contributes to another node  $y$  iff neither node has not deviated (i.e., contributed when it was supposed to snub and vice versa) in the past  $s > 0$  rounds. This is not  $k$ -resilient since  $x$  will earn

$$\frac{1}{1-\delta}(n-2)(b-\gamma) + \frac{\delta^s}{1-\delta}(b-\gamma) < \frac{1}{1-\delta}(n-1)(b-\gamma)$$

*Tit-for-tat.* Node  $x$  contributes to another node  $y$  iff this node contributed in the previous round. For  $\delta > \gamma/b$ , this is not  $k$ -resilient since  $x$  will earn

$$\frac{1}{1-\delta^2}((n-1)b - (n-2)\gamma) + \frac{\delta}{1-\delta^2}((n-2)b - (n-1)\gamma) < \frac{1}{1-\delta}(n-1)(b-\gamma)$$

where the first half of the expression on the left is the payoff  $x$  gets in rounds where  $y$  contributes to  $x$  and  $x$  snubs  $y$ , and the second half is the payoff gets in rounds where  $x$  contributes to  $y$  and  $y$  snubs  $x$ . As typically  $b \gg \gamma$ , this means that tit-for-tat is often not  $k$ -resilient.

A  $k$ -resilient equilibrium can still use these punishments as a non-credible threat and hope that such bluffs are not called in practice.<sup>6</sup> Alternatively, any punishment in which nodes strictly prefer to punish than reward peers can be part of a  $k$ -resilient equilibrium. Here is one such punishment scheme.

*Contrite tit-for-tat* [37]. Every node maintains a simple one-bit reputation with respect to every peer and vice versa. Every node starts off in “good standing” with all of its peers. If a node  $x$  snubs a peer  $y$  that is in good standing,  $x$  falls into “bad standing” with respect to  $y$ . A node’s reputation does not change if it snubs a peer who was in bad standing. If a node  $x$  contributes to its peer  $y$ ,  $x$  returns into good standing with  $y$ .

Given this reputation, a node  $x$  snubs its peer  $y$  iff  $x$  is in good standing and  $y$  is not.

---

<sup>6</sup>Such punishments can never be part of a subgame-perfect or perfect Bayes equilibrium.

The reason this scheme is  $k$ -resilient is because a node prefers to punish its peers rather than reward them. To see why, suppose that after the first round,  $y$  is in good standing and  $x$  is not. Using notation from the proof of Theorem 5.10,  $y$ , playing contrite tit-for-tat, earns

$$\begin{aligned} U_x(\sigma^* | (\Gamma_x^r, C_x^r \setminus \{y\})) &= \frac{1}{1-\delta}(n-1)(b-\gamma) + \gamma \\ &> \frac{1}{1-\delta}(n-1)(b-\gamma) \\ &= U_x((\sigma'_K, \sigma^*_{-K}) | (\Gamma_x^r, C_x^r \setminus \{y\})) \end{aligned}$$

However, network loss in real-world environments may cause complications to using non-credible threats and punishment strategies where the punishing node prefers to punish. Network loss may result in a node falsely believing that it was snubbed by its peer, even if contributing is the only action that should have been played. Such false deviations are not differentiable from true deviations and result in histories where punishment is supposed to be played. If a node (rationally) reneges on its non-credible threat, then, in the absence of punishment, other nodes are unlikely to be incentivized to contribute, causing the collapse of any  $k$ -resilient equilibrium that relies on these punishments to encourage contribution.

Moreover, network loss may enable nodes to frivolously punish other nodes under the false pretense of being snubbed. This is because, given network loss, a node that contributed does not know whether its peer observed the contribution, and a node that observes snubbing does not know whether its peer contributed. Because it is impossible to distinguish true deviations from false ones and legitimate and illegitimate claims of deviation, a node that prefers to punish may profit from falsely claiming to be snubbed.

While we could try to circumvent both of these problems by ensuring that a node is indifferent between punishing and rewarding its peer, this is still not a  $k$ -resilient best response, since the guilty node's strict preference to being rewarded (inequality (5.2)) implies that both nodes can still profitably

deviate by avoiding punishment (e.g., via  $\sigma'_K$ ).

#### 5.1.4 What other issues are there?

Finally, we briefly describe two commonly-used techniques that are often not  $k$ -resilient.

**Digital signatures.** Digital signatures, which guarantee non-repudiation (the signer of a message cannot later deny sending it), are useful in adversarial environments, but their cost (in bandwidth and computation) is hard to justify within a coalition where members trust each other. More generally, digitally signing messages is part of a  $k$ -resilient protocol only if not doing so may affect the outcome of the protocol.

One straightforward example in which this may occur is if this message is passed around to more than  $k$  nodes that check the signature. However, this is not enough. Observe that the overhead of expending extra bandwidth to forward a signed message is unnecessary until the message is actually forwarded outside the coalition. Since coalition members trust each other, the coalition can simply exchange the keys needed to sign any message that leaves the coalition. By doing this, we can construct a strategy profile that is both indistinguishable to nodes outside of the coalition and saves the coalition in bandwidth costs.

Although exchanging the signing keys may have some positive cost which offsets the savings of not exchanging signed messages, note that the size of a digital signature is often comparable to the size of the signing key.<sup>7</sup> In addition, if coalition members sign for one another, one node may end up incurring more cost because it may have to generate signatures for larger messages. However, if the additional cost is sufficiently low, the savings from not sending the signature may be enough to cover this additional cost. Moreover, nodes often exchange more than one message; these future exchanges provide

---

<sup>7</sup>In the case of DSA and its variants, the savings of not signing one message subsumes the cost of exchanging keys; in the case of RSA [67], two is enough.

opportunities for these nodes to level out any computational disparities that may exist.

**Junk.** Junk, i.e., semantically meaningless data, has been used (e.g., [24, 79, 109]) as a form of payment to ensure that nodes contribute their fair share to the cooperative service. For instance, if a node is required to send data but has nothing useful to send, it may instead send protocol-specified “junk.” By making junk more expensive to transfer than useful content, junk transfers discourage free-riding by incentivizing nodes to send real content whenever possible. However, junk transfers incur bandwidth costs on the sender and receiver while providing no benefit to the receiver; nodes that trust each other have no incentive to perform them. It follows that no protocol that relies on junk transfers is  $k$ -resilient.

## 5.2 Accepting coalitions

The scenarios in §5.1 suggest that it is difficult for a single strategy profile to specify strategies that a node, colluding with up to  $(k - 1)$  peers, prefers over all possible deviations, as required by  $k$ -resilience. Yet, we believe these scenarios are symptomatic of a more general problem: the ability for colluding nodes to hold stronger beliefs and assumptions about fellow coalition members (and potentially about the system as a whole) often results in more efficient protocols. As a result, we believe there are likely very few scenarios in which  $k$ -resilience will bear fruit.

In this section, we show that the insight to overcome this impasse is to recognize that denying benefits to nodes that belong to a coalition, while sufficient for stability, is not necessary. We propose a fundamentally different notion of equilibrium: instead of specifying a single best-response strategy to each node, our equilibria map each node to possibly multiple strategies, depending on whom it colludes with. By effectively mapping each possible coalition to a strategy, our equilibria can specify, as a part of the strategy, the

efficiencies that a coalition can leverage among its members. Despite this flexibility, our equilibria guarantee that the strategies specified for every coalition is a best response to what other nodes play, despite how they collude.

**Specifying coalitional strategies.** Because our equilibria specify a strategy per coalition, the strategies that the nodes, within each coalition, follow may depend on whom they are colluding with. Our equilibria cannot use strategy profiles used by traditional equilibria because they specify only a single strategy per node. Our equilibria instead use a novel construct, a *strategy function*, to specify a node’s strategy based on whom the node is colluding with. We formally represent how nodes collude by a partition  $P$  of  $N$ , in which two nodes  $x$  and  $y$  are colluding if there exists some element (a coalition)  $K \in P$  such that  $x, y \in K$ . Intuitively, each partition represents one way that nodes can collude. We use  $\mathbb{P}^k = \{P : \forall K \in P, |K| \leq k\}$  to denote the space of all partitions that contains no coalition larger than size  $k$ .

DEFINITION 5.11. A strategy function  $\mathcal{S}$  is a mapping from a partition (representing a particular way that nodes have chosen to collude) to a strategy profile (which specifies the strategies that these nodes will play as a result) such that if there exists some coalition  $K$  that is in  $P$  and  $P'$ ,  $\mathcal{S}$  maps the same strategy to  $K$  in  $P$  and  $P'$ , i.e., if  $K \in P$  and  $K \in P'$ ,  $\mathcal{S}_K(P) = \mathcal{S}_K(P')$ , where  $\mathcal{S}_K(P)$  and  $\mathcal{S}_K(P')$  denote the strategies deployed by  $K$  given partitions  $P$  and  $P'$ .

Note that a node’s strategy does not depend on how nodes outside of its coalition collude, which a node may not know. We define  $\mathcal{M}$  as the membership function:  $\mathcal{M}(x, P) = K$  if, in partition  $P$ ,  $K$  is the coalition that  $x$  is a part of, i.e.,  $K \in P$  and  $x \in K$ . With respect to a node  $x$  in coalition  $K$ , all nodes in  $K$  are *insiders*, and all others are *outsiders*.

### 5.2.1 Coalition-indistinguishable equilibria

Where  $k$ -resilience makes coalitions futile,  $k$ -indistinguishability makes them invisible; where  $k$ -resilience fundamentally aims to deny coalitions any claim of exceptionalism and sees a system as a collection of individual nodes,  $k$ -indistinguishability sees a system as a collection of coalitions, some of which may contain a single node; where  $k$ -resilience ensures that every node best responds to every other node,  $k$ -indistinguishability ensures that every coalition best responds to every other coalition: in both equilibria, nodes that belong to different coalitions interact with each other as if no coalition existed.

DEFINITION 5.12. Two strategy profiles  $\sigma$  and  $\sigma'$  are indistinguishable with respect to some node  $x$ , denoted as  $\sigma \stackrel{x}{=} \sigma'$ , if all histories resulting from  $\sigma$  and  $\sigma'$ , as observed by  $x$ , occur with equal probability and  $U_x(\sigma) = U_x(\sigma')$ .

DEFINITION 5.13.  $\mathcal{S}^*$  is a  $k$ -indistinguishable equilibrium if:

- For any  $P, P' \in \mathbb{P}^k$ , any coalition  $K$  such that  $K \in P$  and  $K \in P'$ , and any  $x \in K$ ,  $\mathcal{S}^*(P) \stackrel{x}{=} \mathcal{S}^*(P')$ .
- For all  $P \in \mathbb{P}^k$  and all  $K \in P$ , there does not exist a strategy  $\sigma'_K$  such that for all  $x \in K$ ,

$$U_x(\sigma'_K, \mathcal{S}_{-K}^*(P)) \geq U_x(\mathcal{S}^*(P))$$

and, for some  $y \in K$ , the inequality is strict.

Intuitively, the first condition (indistinguishability) requires that a node cannot distinguish whether an outsider is itself colluding with others; the second condition (best response) requires that in any partition, there exists some node in every coalition that prefers the equilibrium-specified strategy to any

coalitional deviation. Note that while we defined best response to be consistent with the definition of  $k$ -resilience, weaker or stronger notions could have been used instead. Also, observe that the best-response condition of  $k$ -indistinguishable equilibria must hold for all possible partitions. Therefore, like  $k$ -resilient equilibria, a  $k$ -indistinguishable equilibrium consists of strategies that make up a best response for *all* possible coalitions of up to size  $k$ , not just one particular coalition or set of coalitions.

Every  $k$ -resilient and Nash equilibrium  $\sigma^*$  has an equivalent  $k$ -indistinguishable equilibrium  $\mathcal{S}^*$  in which  $\mathcal{S}^*(P) = \sigma^*$  for all  $P$ . However, by allowing nodes to base their strategies on whom they collude with,  $k$ -indistinguishable equilibria circumvent the challenges described in §5.1 while ensuring that no coalition will deviate from its specified strategy (§5.2.3). Moreover, similar to  $k$ -resilience, any service that uses a protocol which is the non-colluding strategy in a  $k$ -indistinguishable equilibrium is guaranteed to be supported and maintained by nodes, even if they may collude. Although  $k$ -indistinguishability cannot guarantee that the exact protocol will be followed to the letter by a node when interacting with a fellow insider,  $k$ -indistinguishability does guarantee that any actions that a node takes when interacting with an outsider is the same as those specified by the service’s protocol. Thus, from the service’s perspective, every node is effectively running the service’s protocol and supporting the service.

### 5.2.2 From indistinguishability to stability

Although indistinguishability is an attractive guarantee, it may in practice prove too stringent for some applications. For example, a content-distribution service in which colluding nodes freely exchange content with one another may not be  $k$ -indistinguishable because non-colluding nodes may be able to detect the presence of a coalition simply by observing that colluding nodes statistically have more content at any given time than everyone else.  $k$ -stable equilibria do away with indistinguishability, focusing only on the conditions

necessary for stability.

DEFINITION 5.14.  $\mathcal{S}^*$  is a  $k$ -stable equilibrium if for all  $P \in \mathbb{P}^k$  and all  $K \in P$ , there does not exist a strategy  $\sigma'_K$  such that for all  $x \in K$ ,

$$U_x(\sigma'_K, \mathcal{S}_{-K}^*(P)) \geq U_x(\mathcal{S}^*(P))$$

and, for some  $y \in K$ , the inequality is strict.

As in  $k$ -indistinguishable equilibria, a  $k$ -stable equilibrium requires a best response for all possible coalitions of up to size  $k$ , and every  $k$ -resilient and Nash equilibrium has a  $k$ -stable equivalent. Moreover, every  $k$ -indistinguishable equilibrium is also  $k$ -stable. However,  $k$ -stable equilibria do not guarantee that a colluding node's strategy is indistinguishable from that of a non-colluding node. In other words, it is possible that the strategy of a colluding node  $x$  provides outsiders with information about whether  $x$  is colluding, with whom  $x$  is colluding, etc. In addition, if  $x$  chooses to collude,  $x$ 's coalition may affect the payoffs of peers both inside and outside of  $x$ 's coalition. Nevertheless, a  $k$ -stable equilibrium still guarantees that, for any coalition, the specified strategy is a best response to the strategies played by all outsiders, regardless of how these other nodes may collude.

**Other  $k$ -stable solution concepts.**  $k$ -stability is a very general notion that, we believe, provides a useful basis for developing new solution concepts that guarantee stability in the presence of collusion.  $k$ -indistinguishability is one such solution concept, the result of adding indistinguishability to  $k$ -stability. Another requirement that one may desire is some notion of self-enforcement (no profitable deviation by sub-coalitions), e.g., a solution concept could require that, in equilibrium, nodes prefer to be with their respective coalitions over working alone ( $k$ -stability and  $k$ -indistinguishability do not have any such requirement). Alternatively, one could devise a Bayesian version of  $k$ -stability that guarantees an expected best response for each coalition based

on the likelihood that certain coalitions will form. Yet another interesting direction would be to devise a version of  $k$ -stability that bounds the “price of collusion,” i.e., how much a node’s payoff is affected when outsiders choose to collude (similar to the notion of a safety-net guarantee used in [109]). We leave exploring these and other notions of equilibrium to future work.

### 5.2.3 Examples of equilibria

In this section, we show the applicability of  $k$ -stability and  $k$ -indistinguishability by showing that such equilibria exist in the scenarios described in §5.1, where  $k$ -resilient equilibria did not exist before.

**$k$ -stability and  $k$ -indistinguishability in the mediated pairwise-exchange game.** It is simple to prove that there exists a  $k$ -indistinguishable equilibrium in the mediated pairwise-exchange game (Definition 5.2). Because  $k$ -indistinguishable and  $k$ -stable equilibria allow nodes to base their play on whom they are colluding with, a node, as a part of a  $k$ -indistinguishable equilibrium, can use the mediator with outsiders (as in Theorem 5.3) and leverage the trust provided by the coalition with insiders.

**THEOREM 5.15.** Let  $\mathcal{S}^*$  be a strategy function such that, for any partition  $P \in \mathbb{P}^k$  and any  $x \in N$ ,  $\mathcal{S}_x^*(P)$  specifies that

- For  $y \in \mathcal{M}(x, P)$  such that  $y \neq x$ ,  $x$  never uses a mediator and always contributes.
- For  $y \notin \mathcal{M}(x, P)$ ,  $x$  contributes to  $y$ , using a mediator only in round  $R$ , iff (1)  $x$  and  $y$  have never snubbed each other in the past and (2)  $x$  and  $y$  have not used a mediator in any round other than  $R$ . Otherwise,  $x$  snubs  $y$  without a mediator.

Then  $\mathcal{S}^*$  is a  $k$ -indistinguishable equilibrium.

*Proof.* Without loss of generality, fix some partition  $P$ ,<sup>8</sup> and consider the interactions of some node  $x$  with some peer  $y$ .<sup>9</sup> Suppose that  $y$  is an insider, i.e.,  $y \in \mathcal{M}(x, P) = K$ . Let  $R_s$  be the set of rounds in which  $x$  snubs  $y$  and  $R_m$  be the set of rounds in which  $x$  uses a mediator with  $y$ . In each round in  $R_s$ ,  $x$  gains  $\gamma$ , but  $y$  loses  $b$ . In each round in  $R_m$ ,  $x$  loses  $\epsilon$ ;  $y$ 's payoff is unaffected. Any deviation in which  $R_s \neq \emptyset$  or  $R_m \neq \emptyset$  is then not in  $K$ 's best interest.

Suppose instead that  $y$  is an outsider, i.e.,  $y \notin \mathcal{M}(x, P)$ . We can show that by following  $\mathcal{S}_x^*(P)$  with respect to  $y$  is  $x$ 's best response by backwards induction.

Base case: round  $R$  (the last round). We first show that  $\mathcal{S}_x^*(P)$  is a best response for  $x$  with respect to  $y$  by considering the following two cases:

- $x$  and  $y$  have always contributed to one another. If  $x$  deviates by snubbing and/or not using a mediator,  $x$  saves at most  $\gamma + \epsilon$ . However, since  $y$  is using a mediator,  $x$  loses benefit  $b$  it would have received from  $y$  otherwise. Since  $b > 2\gamma + \epsilon > \gamma + \epsilon$  by assumption (Definition 5.2),  $x$  is clearly worse off.
- $x$  and/or  $y$  have snubbed one another in the past. If  $x$  deviates by contributing to  $y$  or using a mediator,  $x$  is obviously worse off:  $x$  must pay at least  $\min(\gamma, \epsilon) > 0$  but receives no additional benefit.

Inductive step. Assume that for all rounds following some round  $r_0 > 1$ ,  $\mathcal{S}_x^*(P)$  is a best response for  $x$  with respect to  $y$ . We now prove the inductive step— $\mathcal{S}_x^*(P)$  is a best response for  $x$  with respect to  $y$  in round  $r_0$ —in a similar fashion by considering the following two cases:

- $y$  has always contributed to  $x$ . If  $x$  deviates by using a mediator,  $x$  is at least  $\epsilon$  worse off in round  $r_0$ . If  $x$  deviates by snubbing  $y$ ,  $x$  saves  $\gamma$  in

---

<sup>8</sup>As our proof makes no assumptions about  $P$ , it follows that our proof holds for all possible partitions  $P \in \mathbb{P}^k$ .

<sup>9</sup>We can safely do this because each interaction between any two pairs of nodes in  $\mathcal{S}^*$  is independent.

round  $r_0$ . Regardless,  $y$  will snub  $x$  in every subsequent round, resulting in  $x$  losing at least  $b - (\gamma + \epsilon)$  per round.  $x$  is then worse off since the net change in  $x$ 's payoff is at least  $\gamma - (b - (\gamma + \epsilon)) = -b + 2\gamma + \epsilon < 0$ .

- $y$  has snubbed  $x$ . If  $x$  deviates by contributing or using a mediator,  $x$  is worse off, as argued in the base case.

Thus,  $\mathcal{S}_x^*(P)$  is a best response for  $x$ . □

The mediated pairwise-exchange game, as defined in Definition 5.2, involves every node  $x$  *privately* observing which peers use a mediator with or contribute to  $x$ ;  $x$  does not know what other peers have chosen with respect to one another. If such choices were publicly observable (e.g., if the mediator published a list describing which pairs of nodes it would mediate for),  $\mathcal{S}^*$  would no longer be a  $k$ -indistinguishable equilibrium, since non-colluding nodes would be able to observe that coalition members never use a mediator with one another. However, because nodes, regardless of whom they collude with, are still better off following the strategies specified in  $\mathcal{S}^*$ ,  $\mathcal{S}^*$  would remain a  $k$ -stable equilibrium.

**$k$ -stability in the simple secret-sharing game.** Likewise, it is straightforward to show that the simple secret-sharing game (Definition 5.7) has a  $k$ -stable equilibrium. In particular, a node, depending on whom it is colluding with, can choose the exact set of peers to request secrets from that the node expects will maximize its payoff.

**$k$ -stability and  $k$ -indistinguishability in the simple pairwise-exchange game.** We can incorporate the punishments in §5.1.3 into a protocol that is  $k$ -stable or  $k$ -indistinguishable in the simple pairwise-exchange game (Definition 5.9). As an example, we demonstrate how a local grim-trigger punishment can be used here.

THEOREM 5.16. Let  $\mathcal{S}^*$  be the following strategy function: for any partition  $P \in \mathbb{P}^k$  and for any  $x \in N$ ,  $\mathcal{S}_x^*(P)$  specifies the following action for  $x$ :

- For  $y \in \mathcal{M}(x, P)$  such that  $y \neq x$ , contribute to  $y$ .
- For  $y \notin \mathcal{M}(x, P)$ , contribute to  $y$  iff  $r = 0$  or  $x$  and  $y$  have always contributed to one another.

Then  $\mathcal{S}^*$  is a subgame-perfect  $k$ -indistinguishable equilibrium (i.e., at every point in the game, nodes play a  $k$ -indistinguishable best response) if

$$\frac{b}{\gamma} \geq \frac{1}{\delta} \tag{5.5}$$

*Proof.* Without loss of generality, fix  $P$ . For any  $K \in P$  in which  $|K| > 1$ ,  $\mathcal{S}_K^*(P)$  is a best response when interacting with fellow insiders. To see why, observe that following  $\mathcal{S}_K^*(P)$  in each round earns a round payoff of  $(n - 1)(b - \gamma)$ . Deviating by snubbing an insider improves one node's payoff by  $\gamma$  but causes a loss of  $b > \gamma$  to another's; the coalition as a whole earns  $(n - 2)(b - \gamma) < (n - 1)(b - \gamma)$  as a result in that round, so someone in the coalition must be worse off.

Now consider any two nodes  $x, y$  that are not colluding, i.e.,  $y \notin \mathcal{M}(x, P)$ . If  $x$  and  $y$  have always contributed to each other and  $x$  snubs  $y$ ,  $x$  gains  $\gamma$  in the current round but loses at least  $(b - \gamma)$  in every subsequent round. This is profitable only if

$$\frac{\delta(b - \gamma)}{1 - \delta} < \gamma$$

which is never the case given inequality (5.5). Finally, if  $y$  has snubbed  $x$  and  $x$  deviates by contributing to (rather than snubbing)  $y$ ,  $x$  incurs an additional cost of  $\gamma$ ; this is clearly not in  $x$ 's best interest.  $\square$

Similar to the previous example,  $\mathcal{S}^*$  as defined in Theorem 5.16 would remain a  $k$ -stable equilibrium (but would not be indistinguishable at every point in the game) if a node's choices of whom to contribute to were publicly observable.

**$k$ -stability and  $k$ -indistinguishability with digital signatures and junk.**

Mechanisms such as digital signatures or junk transfers fit naturally within a  $k$ -stable or  $k$ -indistinguishable equilibrium. The equilibrium may specify that these mechanisms are used between outsiders and bypassed between insiders when unneeded.

### 5.3 Summary

Trying to identify strategies that eliminate all incentives to collude, as traditional approaches attempt to do, is difficult, possibly futile, and fundamentally unnecessary. This paper introduces a new approach to handle the challenge posed by collusion: accept that coalitions will form, allow coalitions to benefit among themselves, and aim for stability by ensuring that the strategies or protocols specified for every *coalition*, not just every *node*, are best responses. While we are only beginning to explore the space of solution concepts and equilibria allowed by this new approach, we believe our initial results are encouraging: our proposed framework offers rigorous guarantees to both colluding and non-colluding nodes in cooperative services where traditional approaches are often provably unable to yield an equilibrium.

# Chapter 6

## System

So far, this thesis has described what theoretical guarantees practical cooperative services should strive for in the presence of Byzantine and rational nodes and described one way acquiescent nodes can be leveraged to encourage rational cooperation. After all, while simply sprinkling a system with incentives whose rationale is rooted in intuition and common sense may provide a modicum of protection, these schemes are typically easily defeated once exposed to more than casual strategic behavior [68, 10, 76, 81, 93, 99]. At the same time, as we have seen in the previous chapters, overly strong equilibrium notions may not even be achievable in fault-tolerant distributed systems.

While the previous three chapters have worked on spanning the classic theory/practice divide by making practical theoretical advances, this chapter focuses on spanning the divide in the other direction: we demonstrate how to design a system that provides robust guarantees, using many of the ideas we described in the previous chapters, with far more realistic assumptions than previously achieved. In particular, this chapter will describe the design and implementation of Seer, a cooperative content distribution service that provides provably robust incentives to individual clients for faithfully disseminating content, even if their peers may fail or try to game the system by colluding. content distribution

Seer is a hybrid peer-to-peer (P2P) protocol, in the style of popular

commercial services including Spotify [17], Blizzard [5], PPTV [14, 64], and Akamai [2] and various prior work (e.g., [89, 92, 100]). Like these hybrid P2P services, Seer leverages trusted servers to offer its users the best of both worlds: the scalability and low cost of P2P dissemination combined with the reliability that a trusted set of servers or CDN can provide. Seer share these systems' practical concerns for scalability and performance, but fundamentally departs from other hybrid P2P services in how it ensures its clients' cooperation. Where commercial hybrid P2P services rely on largely ad-hoc solutions to dissuade clients from free-riding off the server (such as changing the settings in the client's software preferences, using a firewall to block P2P traffic, or modifying the client itself), Seer takes a principled approach: it applies game-theoretic techniques to incentivize clients to *want* to help disseminate content.

Seer is not the first research system to aim at a rigorous treatment of incentives [24, 78, 79], but it manages to achieve its dual goals of performance and provable guarantees under fundamentally more realistic assumptions than any of these prior systems: more than ever before, Seer is not only rigorous, but *rigorous in practice*. In particular, Seer does *not* rely any of the following frequently-made assumptions:

- Clients do not collude, or, if they do, they can be modeled as Byzantine. Seer recognizes that clients that develop trust in one another (because of real-life connections outside the systems or from interacting in the system) may well choose to collude to improve their own standing.
- Clients do not know when their interaction with other clients will end: hence, the promise of future benefits will always be an effective incentive to cooperation. Seer recognizes that in practice clients are likely to infer when a peer is about to leave and leverages the trusted server to incentivize cooperation even in end-of-game situations.
- Clients are uniformly and exceedingly risk-averse when it comes to incentives. Seer does not assume that clients will simply be frightened

into compliance, but recognizes that when clients interact their behavior, rather than by ancestral fear, is going to be determined by more nuanced expectations.

**Our contributions.** In summary, this chapter makes the following contributions:

- We present Seer, a robust hybrid P2P service that advances the state-of-the-art in dependable cooperative services. Seer, unlike prior systems, makes few assumptions about the environment and participants and incentivizes cooperation despite the possibility of arbitrary failure or collusion.
- We rigorously prove, under a realistic set of assumptions, that clients disseminate content quickly in Seer. We accomplish this in part by basing design decisions on what we are able to prove about the system, thereby ensuring that Seer’s policies and mechanisms are backed by strong theoretical guarantees.
- We evaluate Seer and show that rigor can go hand-in-hand with performance: Seer neither sacrifices the scalability of P2P services nor the reliability of client/server services. Evaluating our implementation of Seer, we find that Seer can support significantly more clients than a traditional client/server service and can outperform BitTorrent, a popular traditional P2P service, by over 20%.

Naturally, Seer’s guarantees still rely on certain key assumptions to hold, and Seer is still far from the final word in how cooperative systems should be built. Nonetheless, we believe that the design and principles behind Seer represent a significant contribution towards building a more robust bridge over the theory/practice divide.

**Organization of chapter.** We discuss some of the principles underlying Seer in §6.1. We describe the system model in §6.2, provide a description of Seer in §6.3, and describe the incentives that underlie Seer in §6.4. We present an implementation and evaluation of Seer in §6.5 and summarize in §6.6.

## 6.1 Principles

There are several important principles that underlie the design and implementation of Seer. While we do not claim to be the first to recognize these principles, we do believe that the strategic application of the *combination* of these principles enables us to build a dependable cooperative service with much stronger guarantees and fewer assumptions than previous systems.

*Leverage impatience.* Clients not only want content, they want it now. Leverage this impatience to induce clients to contribute their resources in order to accelerate their own acquisition of content.

*Align incentives of participants.* Downloaders want their content fast; uploaders want to do as little as possible as slowly as possible. Design the incentives so that they are working towards a common goal.

*Balance the role of the server.* In many P2P applications, a trusted server exists, but a P2P mechanism is deployed to reduce the cost of serving content. Leverage the server to help the reliability of the service, but limit its involvement in the common case to preserve scalability.

*Keep assumptions about rationality to a minimum.* Individual clients' take on what constitutes maximizing utility is subjective and depends on how they value their time, bandwidth, and the content, as well as how they view their peers. To maximize the applicability of the system's guarantees, strive to minimize what the system assumes about a node's motivations.

## 6.2 System model

**Principals.** In Seer, there are two types of principals: servers, which are run by some *content provider* in order to serve some content, and nodes, which are interested in acquiring and consuming the content. While there may be multiple servers, we model the set of servers as a single trusted entity for simplicity. The server has a known IP address and public key.

Nodes, on the other hand, may be arbitrarily faulty (*Byzantine*), correct (*acquiescent*), or selfish and deviate if doing so is in its best interest (*rational*). We do not restrict the behavior of Byzantine nodes except that we assume that no node, Byzantine or otherwise, can subvert cryptographic primitives. We do assume that rational nodes are somewhat pessimistic when dealing with a known Byzantine peer. We assume no Sybil attacks and ignore malicious denial-of-service attacks, which any system is susceptible to. We discuss these assumptions more rigorously in §6.4. Note that we do allow rational nodes to collude and deviate as a coalition if doing so is weakly better for the coalition.

We call a node that is downloading content a *consumer*, and a node that is offering or uploading content to other peers a *distributor*. A node may, of course, be a consumer for some content and a distributor for other content.

**Incentives.** Rational nodes are trying to maximize their own utility; in order to do this, they try to maximize their benefits while minimizing their costs. In Seer, rational nodes have some notion of impatience; a node's benefit is derived from being able to acquire some content of interest more quickly. A node's costs come from participating in the service, whether from downloading some content or from helping distribute. §6.4 formalizes these notions in more detail.

**Network.** We assume that the network may be lossy and asynchronous. Nodes have certain beliefs regarding how lossy and asynchronous the network might be. While not strictly necessary, for simplicity of analysis, we assume

that a node believes that any connection to the server is reliable and synchronous. We assume nodes have full-duplex connections in that serving and downloading content are independent with respect to network bandwidth. As we describe later in §6.3.2, a node’s upstream and downstream bandwidth are split into a number of upstream and downstream capped-bandwidth channels.

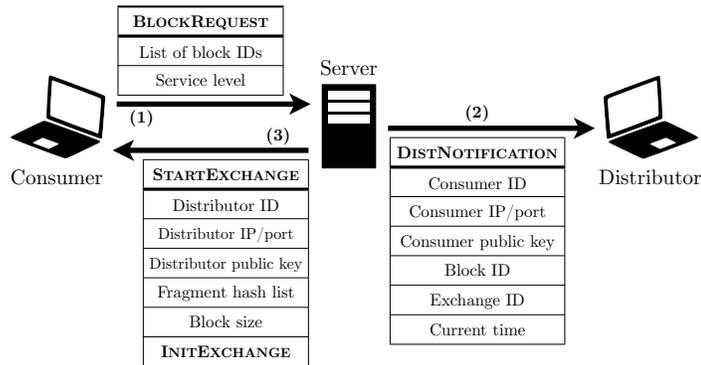
**Timing assumptions.** We assume nodes in the service have synchronized clocks. Because our incentives depend on the impatience of nodes, which inherently is based on time, an unsynchronized clock may affect whether a node chooses to cooperate or not. Ultimately, the server serves as the authoritative time source, and when initiating a P2P exchange between two nodes, the server sends some messages that include a timestamp, which nodes use to roughly synchronize their clocks.

**Content.** Similar to other P2P services, content in Seer is broken up into blocks. For a particular piece of content, every block (with the possible exception of the last block in the content) is the same size and is further broken up into fixed-sized fragments. We pad content to be a multiple of the fragment size. The size of the block and fragment, while constant for a particular piece of content, may differ depending on the content. The block and fragment sizes present various tradeoffs that we discuss further in §6.5, but, for now, we generally assume that blocks will be large and fragments small.

## 6.3 Overview of Seer

Seer is a hybrid P2P service designed for bulk transfer of content. The core of the protocol and of the incentive structure focus on encouraging clients to help exchange blocks.

At a high-level, consumers acquire content in Seer by (1) purchasing/requesting the content from the server, which provides the list of block identifiers that make up the content; (2) for each block, consumers can choose



**Figure 6.1:** The messages involved in initiating a block exchange.

to download the block form either (a) the server (which will in turn assign the task to a server-backed distributor) or (b) a peer distributor; and (3) downloading the block via the chosen method, falling back on the server if needed.

Whether they choose to download from the server or a peer, consumers download the content on a block-by-block basis. A consumer performs the following steps to download a block (illustrated by Figures 6.1, 6.2, and 6.3):

1. The consumer sends to the server a **BLOCKREQUEST** message to get distributors for the blocks it wants to download, indicating whether the distributor should be a peer or a server.
2. The server selects distributors for each block and notifies them via a **DISTNOTIFICATION** message. For blocks for which the consumer selected to download from the server, the server selects a server-backed distributor.
3. Upon confirmation from a distributor, the server sends the consumer a **STARTEXCHANGE** message.
4. The consumer starts the block exchange by sending the distributor the **INITEXCHANGE** message included in the the **STARTEXCHANGE** message received from the server.
5. The distributor encrypts the block using a block-specific key (the *block key*). The distributor sends each block fragment to the consumer inside

of a `FRAGMENT` message and waits for the consumer to acknowledge the fragment.

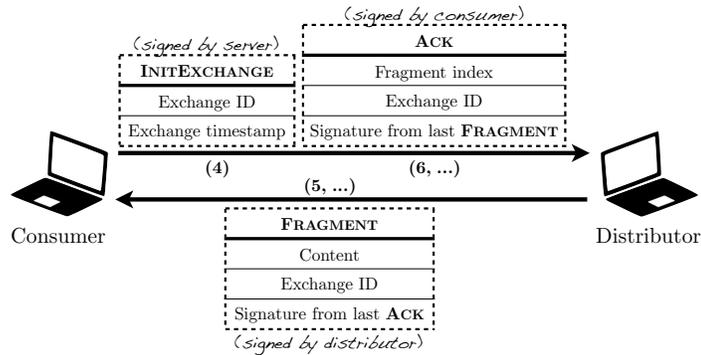
6. For every fragment, the consumer sends a cumulative acknowledgment message (`ACK`) back to the distributor.
7. The consumer and distributor independently report back to the server using `CONREPORT` and `DISTREPORT` messages, respectively, when the exchange completes or the consumer or distributor give up on the transfer. If the consumer chose to download the block from the server, the server already knows that the exchange has completed, and this step can be skipped.
8. The server sends to the consumer a `CONSFINISH` message that includes the block key needed to decrypt the block and, if necessary, the means to download from a server-backed distributor any fragment missing from the block—a process that simply involves performing steps 4, 5, and 6).

Before serving a particular block, distributors must register for the blocks they want to serve. Seer ensures that distributors only register for blocks that they possess. §6.3.4 describes this process in more detail.

Note that in Seer, consumers and distributors ultimately have complete freedom over what exchanges they want to participate in and how long they want to participate. In the following subsections, we describe the key challenges in Seer and how we leverage our principles to solve them.

### **6.3.1 How does Seer encourage use of P2P?**

Because in hybrid P2P services like Seer a trusted server provides clients with a reliable source for content, a peer does not *need* other peers to receive the content it is interested in. Hence, the prospect of receiving content from another peer is insufficient to motivate rational peers to cooperate in the process of content distribution. Seer instead leverages the clients' impatience: by



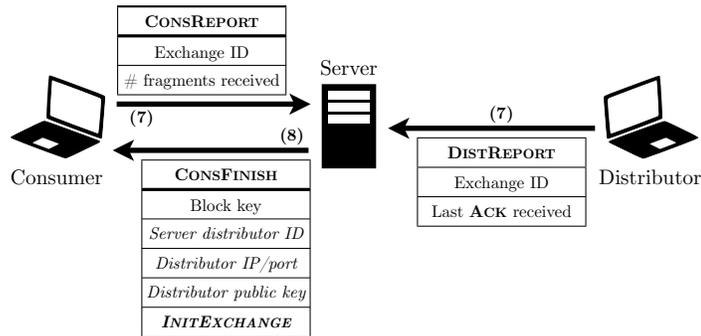
**Figure 6.2:** The messages involved in the main block exchange loop.

contributing their resources to the system, clients can accelerate their own acquisition of content.

We put in place this in a straightforward manner by using the trusted server as a bank that issues *credits* to clients. Clients earn credits by distributing content and spend credits to get content more quickly. Credits are not tied to a particular piece of content, so a client can earn credits for serving one piece of content and use them to download another. Thus, a client’s incentives are not tied to a particular piece of content.

When requesting the server for blocks, a consumer has a choice of three different levels of service. At the lowest level, the *free tier*, clients download content slowly from the server for free. This enables the client to acquire content even if it does not have any credit. There are certain limitations to using this level of service. For example, a consumer can only use this level of service on a small number of downstream channels and cannot simultaneously use its other downstream channels for downloading content.

Otherwise, a consumer can use all of its downstream channels in the *paid tier*, either in P2P or server mode. In P2P mode, a channel is used to download content from a distributor; in server mode, a channel is used to download content from the server. Both modes incur a positive cost; the *server rate* (the cost of using the server) costs more credit than the *P2P rate* (the cost of using P2P). A client may use any number of channels in P2P or server



**Figure 6.3:** The messages involved in finishing a block exchange. The italicized fields are optional.

mode, and a consumer may begin an exchange on a channel using P2P and complete it from the server.

In Seer, clients are not actually issued any form of actual token or coin; instead, the server tracks all the credits that clients have. Since all P2P requests go through the server first, a server will only pair up a distributor with a consumer if the latter has enough credit to pay for the download. By using the server, we avoid issues with forging currency and double-spending.

### 6.3.2 How does Seer incentivize adherence?

While its lower cost encourages clients to use the P2P tier, clients are ultimately free to choose the level of service they want to use. If clients do not faithfully participate in the P2P protocol, it is likely that the performance of using P2P will suffer, resulting in clients preferring other levels of service.

Thus, we must ensure that P2P exchanges occur in a timely fashion. The challenge is that, in P2P exchanges, the two parties involved in the exchange have incentives that are not naturally aligned: distributors want to do as little as possible as slowly as possible while still earning credit for the exchange, while the consumer wants to acquire the content as quickly as possible.

Seer aligns the incentives of consumer and distributor by providing incentives for both parties to complete their exchange in a timely fashion. For the

consumer, the server initially charges it the server rate regardless of whether the consumer chose to use P2P or (paid) server service. The server later credits the consumer back after it receives a report with the number of fragments the consumer successfully received using P2P. Thus, for the consumer, there is no ambiguity: a consumer wants to finish its exchange as quickly as possible, both to get the content and its credit back. For the distributor, the server credits the distributor only after it successfully provides the content; as a result, a distributor also prefers serving content more quickly.

However, even if both clients want to finish an exchange as quickly as possible, a client's view of what it means to be timely may differ from that of its peer. Pairing a consumer and distributor with different expectations may result in both parties being disappointed and ultimately choosing to opt out of P2P services. To ensure common expectations in the service, all block exchanges in Seer (P2P or client/server) are performed over a *capped-bandwidth (data) channel* that handles transmitting content fragments at a rate that does not exceed some system-defined bandwidth (the *channel bandwidth*).

Clients are assigned a particular number of downstream and upstream channels based on periodically-repeated bandwidth tests that the server conducts. As each channel is assigned to a single block exchange at any given time, channels are also a limited resource that clients prefer using as efficiently as possible. As we discuss in the following subsection, how and when channels can be allocated is dictated by the server and depends on the manner in which the previous exchange completed.

### **6.3.3 What role does the server play?**

The server plays a key role in dictating the incentives in Seer, both as a mediator of P2P exchanges and as a backup distributor. For every P2P exchange, the server is a matchmaker: every consumer—with sufficient credit and a free downstream channel—is matched up with a distributor—with a free upstream channel and able to offer the block in question. An exchange begins when

the server stamps the current time as the exchange timestamp on the INITEXCHANGE message (the server-provided message used to initiate a block exchange).

In addition to its role in pairing clients, the server also acts as the authoritative clock in the service and has the final say as to when, at the end of an exchange, a consumer and distributor get to reuse their channels and when they get charged/credited for the exchange. Given the capped bandwidth of the channel and the timestamp on the INITEXCHANGE, the *expected time* to transfer a given number of fragments is known to all parties involved (the consumer, distributor, and server). Using this notion of expected time, the server enforces a set of policies regarding how much credit a client receives and when it receives credit or can reuse its channel.

When a consumer ends an exchange and reports to the server, the server waits until the expected time to transfer the reported number of fragments has elapsed (if it has not already). It is only after this period elapses that the server provides the block key for decrypting the block and enables the consumer to download the fragments (if any) that it is missing from the server-backed distributor (through a capped-bandwidth channel). Besides helping to enforce the channel bandwidth, waiting provides some incentives to consumers to stick to the assigned distributor.

When a consumer successfully downloads all fragments, the server then allows the consumer to reuse the associated channel and credits the consumer the difference between the server rate and P2P rate for each fragment the consumer downloaded via P2P. Note that as a direct consequence, a consumer has no access to the content block or the channel and credits tied up in this exchange until after the expected time for transferring all the fragments.

When a distributor reports to the server, as with the consumer, the server waits until the expected time for the reported number of fragments has elapsed (if it has not already) before allowing the distributor to reuse the associated upstream channel. However, when it comes to crediting the distributor for the fragments it did send. Server wait until the time elapsed is

the larger between the current time and the expected time to transfer the fragments reported as sent by the distributor *plus* the expected time needed to transfer any of the fragments the distributor did not successfully send.

Finally, the server also acts as a distributor as well. The server provides a server-backed distributor either when there are no available distributors or if a consumer requests for service from the server, either as a part of the free tier, the paid tier, or to help it finish downloading fragments it is missing from a previous block exchange with a client distributor. Because clients can always fallback onto the server, this enables clients to demand a certain level of service from client distributors, lest they simply switch to the server distributor, depriving client distributors of credit.

#### 6.3.4 Seer: a close-up

This section provides additional details on the inner workings of Seer.

**Pairing clients (Figure 6.1).** For the server to add a peer as distributor, the peer must register with the server the blocks it wants to serve. Seer requires that clients prove to have the blocks for which they want to be listed as distributors. This condition could be enforced using cryptographic puzzles, but our prototype opts for a straightforward solution: we allow clients to only to become distributors for blocks that they have previously downloaded through Seer.

A consumer that wants to download a particular piece of content sends to the server a `BLOCKREQUEST` message that contains the list of blocks it wants to download and the level of service it wants to use (free, paid P2P/server). If the consumer chooses to use P2P, then for every requested block, the server matches up the consumer with a distributor offering the block. The probability that a particular distributor is selected is equal to the ratio between that distributor's free upstream channel and the total number of free upstream channels available at all distributors offering this block. The server pairs the consumer up with as many distributors as the consumer has

downstream channels. If there are no distributors that have free channels for a particular block, the server provides a server-backed distributor to serve the block.

After selecting a distributor, the server assigns a unique identifier to this particular P2P block exchange and sends the distributor a `DISTNOTIFICATION` message that notifies the distributor that it has been selected for an impending P2P exchange with a particular consumer and block. The distributor may choose to refuse to serve this request; if so, the server repeats the random distributor selection without any distributor that previously refused this request.

After finding a distributor willing to serve the consumer, the server sends the consumer a `STARTEXCHANGE` message that notifies the consumer of the selected distributor. This message also contains a list of fragment hashes to enable the consumer to verify the fragments it is receiving (to limit the amount of damage Byzantine distributors can cause) along with other relevant details needed to communicate with the distributor. In addition, the server gives the consumer a `INITEXCHANGE` message that the consumer uses to initiate the P2P exchange with the distributor.

**Exchanging content (Figure 6.2).** Upon receiving the server's response, the consumer decides whether to download from the assigned distributor. If the consumer chooses to proceed, it sends `INITEXCHANGE` to the distributor. Upon verifying `INITEXCHANGE`, the distributor retrieves the encrypted version of the block (either by retrieving the encrypted version or re-encrypting the block with the block key, both of which it received at the earlier time when it was a consumer of this block) and breaks the block up into fragments. For each fragment, the distributor decides whether it wants to continue the exchange; if so, it sends a `FRAGMENT` message to the consumer with the encrypted content and awaits an acknowledgment before repeating this process.

Upon receiving a `FRAGMENT` message, the consumer verifies that it matches the list of hashes provided by the server. If it does, the consumer

responds with a `ACK` message that acts as a cumulative acknowledgment.

**Finishing up an exchange (Figure 6.3).** When either the exchange is completed or a consumer or distributor unilaterally aborts the exchange, both the consumer and distributor report on how the exchange went using the `CONREPORT` and `DISTREPORT` messages. The `CONREPORT` message contains how many fragments the consumer received; similarly, the `DISTREPORT` message reports the last `ACK` received in the exchange, which indicates how many fragments the distributor served. Upon receiving the reports, the server proceeds to credit the participants for the exchange and goes on to pair them with other peers as described in §6.3.3. If the consumer aborted the exchange early, the server includes information similar to the `STARTEXCHANGE` message that provides, among other things, a `INITEXCHANGE` message to obtain the content from a server-provided distributor itself.

Clients are only allowed to send one exchange report for a given exchange. When a consumer and distributor send differing reports, the server reconciles them as follows. If the consumer reports one more fragment than the distributor, the server assumes that either the consumer neglected to send an `ACK` message or the `ACK` message was dropped; thus, the server charges the consumer as if the consumer had reported the number of fragments listed in the distributor report. Note that we do this instead of crediting the distributor to encourage consumers to send `ACK` messages (thereby overcoming any end-game problems within a particular block exchange). Otherwise, one client has underreported the number of fragments exchanged. While the server could potentially detect which client lied if the messages contained slightly more information (e.g., the `CONREPORT` message contained information about the last `FRAGMENT` message received), note that underreporting is never in a client's best interest. As a result, we chose to have the server ignore this and simply credits the consumer and distributor based on their own individual reports.

Content mostly consists of fixed-sized blocks, but when content is not

evenly divisible by the block size, the last block may be smaller than the rest. For this block, the amounts of all charges and credits are directly scaled to the number of fragments this block contains, as is the time before the channel is eligible for reuse. However, the time at which credits are processed, both for the consumer and distributor, is as if the block were a full-sized block.

**Security.** To ensure that clients cannot spoof traffic, we use secure channels for all control traffic. While we could have used signed messages with sequence numbers to ensure authenticity and prevent replay, control messages are generally small, so securing the entire channel is unlikely to cause significant overhead. The server is used to act as a key distributor or certificate authority to distribute the necessary keys or certificates.

On the other hand, because the content being sent over the channel is potentially large, Seer uses a combination of signatures to ensure authenticity and integrity of the message. To prevent arbitrary clients from connecting to a client’s channel, the consumer sends a nonce along with the `INITEXCHANGE` message on the secure control channel that the distributor presents when connecting to a consumer’s channel, and the consumer checks to be sure the nonce came from the correct IP address.

Any deviation that the server can verify (either as a first- or third-party) serves as a proof of misbehavior and results in a client being banned from the service and losing any credits it may have accumulated.

## 6.4 Incentives in Seer

In this section, we describe in more detail how Seer is able to achieve the needed incentives to guarantee that clients help disseminate content quickly.

Rational clients aim to maximize their utility, which may depend on payoffs that the client receives now and in the future. Since clients are impatient, they always prefer receiving payoff earlier. We formally model a client’s patience using a discount factor ( $\delta$ ). We use the standard formulation of dis-

counting utility exponentially, i.e.,  $\delta^r v$  represents utility of  $v$  received at time  $r$  by some client with a discount factor of  $\delta$ .

Because of the key role that timeliness plays in a client's experience, a client that is maximizing its utility is likely to maintain some notion of its peer's timeliness. We refer to a client's view of how timely its peer is expected to be as a peer's *expected timeliness*. Every client may have a different view of a peer's expected timeliness. We say a client's expected timeliness is better if the client is expected to be no less likely to be timely in the future.

We say an exchange is (or is expected to be) *timely* if, at the current time (or some future time), a client has received (or is expected to receive) credit for at least the expected number of fragments, given bandwidth  $w$ . We say that the exchange at some future time is expected to be at least as timely if, from now until that future time, a client expects to receive credit for at least the expected number of fragments. Finally, we refer to a client as being timely if it is at least as timely at all future times and refer to a node being "more" timely, "better" off, etc. in a non-strict sense, i.e., a node is no less timely, no worse off, etc.

We discuss bandwidth in terms of fragments per second, as this is the minimal unit of data that either a consumer or distributor can earn credit for.

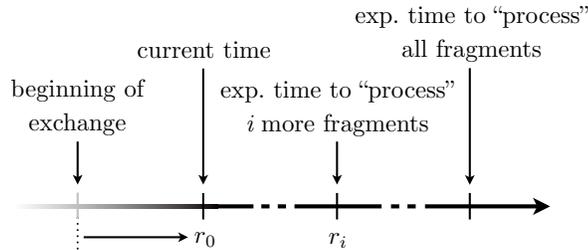
### 6.4.1 Abstracting away uncertainty

A client's beliefs represent its view on the unknown. A client must come up with some expectation over these unknowns before it can decide its (expected) best course of action. However, this becomes extremely hairy in any real-world distributed system, where clients cannot observe all the interactions that occur in the service and must maintain beliefs about what its peers may have observed, what its peers may believe about what the client has observed, etc.<sup>1</sup>

In this paper, we abstract away these details by simply representing

---

<sup>1</sup>In other words, the interactions between peers is private information.



**Figure 6.4:** An illustration of how various symbols relate to a particular exchange.

three key parameters that we believe captures the relevant details that a client considers when deciding whether to continue exchanging fragments with a peer. Note that clients may have a different value of these parameters for different peers and that a client may have a different value depending on whether it is serving as a consumer or distributor.

*What is the client’s expected marginal net payoff to get additional fragments from the peer?* This factor represents how much a client values getting credit for additional fragments (for a distributor, through a cumulative acknowledgment; for a consumer, through the fragments themselves), how much a client values its network bandwidth (the cost of getting these additional fragments), and how the client views the reliability of the network or its peer (e.g., network loss could increase the cost per fragment). We use  $\pi$  to represent the marginal net payoff for a node. For simplicity, we assume that, at any moment in time within a particular block exchange,  $\pi$  is constant, so the expected marginal net payoff is linear with respect to the number of additional fragments.

We describe how introducing Byzantine clients affects  $\pi$  in §6.4.4. Finally, we assume clients cannot collude in the first few subsections; we describe how collusion affects our results in §6.4.5.

*How much longer does the client expect additional fragments will take with this peer?* This factor incorporates a client’s expectation of how slow this peer or the network is or how likely a failure has occurred, potentially based

on previous history with this peer. We use  $r_i$  to denote a client's expectation regarding how long it will take to receive credit for  $i$  fragments, relative to the beginning of the exchange.  $r_0$  then represents the current time, i.e., the amount of time that has currently elapsed in the current exchange, and we let  $i_0$  be the number of fragments a client has already received credit for. Figure 6.4 illustrates these notions of time. Because a client cannot receive credit for  $(i + i_0)$  fragments until the client receives those additional fragments (expected to occur at time  $r_i$ ) or until the expected time for transferring those fragments elapses  $((i + i_0)/w)$ , whichever is later, a client often considers this time instead; we denote this time as  $\bar{r}_i = \max(r_i, (i + i_0)/w)$ .

*How much does the client expect to earn using this channel after this exchange?* This utility, known as a *continuation utility*, incorporates a client's expectation a client's beliefs about other peers and the environment as a whole, including the two previous factors. We denote the expected continuation utility of using the system as  $V$  and assume that it is non-negative. A negative continuation utility implies that a client believes it will be worse off participating in Seer than not doing anything at all, which would make it difficult to sustain cooperation in Seer, or any system for that matter.

### 6.4.2 What do clients do if they use P2P?

The main result that we prove is that clients want to be timely when using P2P in Seer.

**THEOREM 6.1.** If a client chooses to participate in a block exchange with a peer it is not colluding with, a client is better off being timely.

#### One single block exchange

We first discuss conditions under which, ignoring the effects on a node's expected timeliness, a client chooses to terminate an exchange before it is com-

pleted. The conditions are effectively thresholds that bound how long a client is willing to wait to receive credit for  $i$  more fragments. In this section, we ignore how a client's actions affect its perceived timeliness; we consider these effects in §6.4.2.

**Consumers.** Consumers may always choose to download from the server, but the server is expected to cost more than using P2P. Thus, consumers only choose to use the server if the cost of acquiring the content through the peer is expected to be slower and/or more costly than going through the server.

To get an idea of why Seer's design provides the incentives it does, we first consider an intermediate lemma, which describes conditions under which a consumer prefers to continue a block exchange rather than terminating it. As a reminder, recall that  $\pi$  is the expected marginal net payoff per fragment,  $V$  is the expected continuation payoff for the consumer,  $i_0$  is the number of fragments a consumer currently has,  $r_0$  is the current time relative to the exchange timestamp,  $w$  is the channel bandwidth in fragments per unit time,  $r_i$  is the expected time in which a client (in this case, a consumer) expects to get  $i$  additional fragments, and  $\bar{r}_i = \max(r_i, (i + i_0)/w)$ .

LEMMA 6.2. Ignoring the effects on expected timeliness, a consumer continues to download from a distributor in some block exchange iff there exists some  $i > 0$  such that at the time  $r_i$  the consumer expects to get  $i$  additional fragments,

1. The exchange is timely,
2. The exchange is at least as timely, and/or
3. The following condition holds:

$$r_i \leq \max\left(r_0, \frac{i_0}{w}\right) + \frac{i}{w} + \mathcal{C}_i \quad (6.1)$$

where

$$\mathcal{C}_i = -\log_\delta \left( 1 + \frac{\pi i}{\pi i_0 + V} \right)$$

*Proof.* A consumer continues downloading from a distributor iff there exists some  $i > 0$  such that

$$v_{\mathcal{C}}(i_0 + i, r_i) \geq v_{\mathcal{C}}(i_0, r_0)$$

where  $v_{\mathcal{C}}(x, r)$  is the payoff that a consumer expects to earn in an exchange if it chooses to download a block by downloading  $j$  chunks from a distributor and downloading, starting at time  $r$ ,  $(m - j)$  fragments of a block from the server; recall that time is measured relative to the beginning of a particular block's transfer.

How do we define  $v_{\mathcal{C}}$ ? A consumer that downloads  $j$  fragments using P2P expects to earn a payoff of  $(\pi j + V)$ . A consumer expects that it can redeem this benefit at the later of  $(r + (m - j)/w)$  (the amount of time it will take to download the remainder of the  $(m - j)$  fragments at time  $r$  from the server at bandwidth  $w$ ) and  $\tau = m/w$  (the beginning of the redemption period). Given the current time is  $r_0$ , we have,

$$v_{\mathcal{C}}(j, r) = \delta^{(m-j)/w + \max(r, j/w) - r_0} (\pi j + V) \quad (6.2)$$

By plugging this into the condition at the beginning of this proof and moving terms around, the condition becomes

$$\bar{r}_i \leq \bar{r}_0 + \frac{i}{w} + \mathcal{C}_i \quad (6.3)$$

where  $\mathcal{C}_i$  is defined as above.

We prove that condition (6.3) holds iff at least one of the three cases in Lemma 6.2 holds. We first prove the “if” direction.

*Case (1):* the consumer expects the exchange to be timely when it receives  $i$  more fragments, i.e.,  $r_i \leq (i_0 + i)/w$ . Then  $\bar{r}_i = (i_0 + i)/w$ ; since

$\bar{r}_0 \geq i_0/w$ , we have

$$\bar{r}_i \leq \bar{r}_0 + i/w$$

To ensure condition (6.3) holds, it is sufficient to ensure that  $\mathcal{C}_i \geq 0$ , i.e.,

$$1 + \frac{\pi i}{\pi i_0 + V} \geq 1$$

It is straightforward to show that this condition is true given that  $V \geq 0$  (by assumption).

*Case (2):* the consumer expects the exchange to be at least as timely, i.e.,  $r_i \leq r_0 + i/w$ . Assume also that  $r_i > (i_0 + i)/w$ ; otherwise, condition (6.3) is satisfied using the argument made in case (1). This implies that  $\bar{r}_0 = r_0$  (the exchange is currently untimely); since  $\bar{r}_i = r_i \leq r_0 + i/w$ ,  $\bar{r}_i \leq \bar{r}_0 + i/w$ . The proof of case (2) then proceeds using the same argument as the one used in case (1).

*Case (3):* condition (6.1) holds. The only difference between condition (6.1) and (6.3) is a  $r_i$  in the former and a  $\bar{r}_i$  in the latter. If  $r_i \leq \max((i_0 + i)/w, r_0 + i/w)$ , then either case (1) or (2) proves that condition (6.3) holds. However, if, on the other hand,  $r_i > \max((i_0 + i)/w, r_0 + i/w)$ ,  $\bar{r}_i = r_i$  as needed.

We prove the “only if” direction by using the contrapositive. If none of the cases are true, it must be the case that  $r_i > \max((i_0 + i)/w, r_0 + i/w)$ ,  $r_i = \bar{r}_i$ , and condition (6.1) is false. This directly contradicts condition (6.3).  $\square$

*Observations from Lemma 6.2:* There are several interesting properties that can be gleaned from Lemma 6.2. First, a consumer always prefers to continue an exchange if it is timely or is expected to be more timely in the future.<sup>2</sup> Recall that a consumer is immediately charged server rate even if it requests a peer. Thus, a consumer that gives up on an exchange early and submits a CONREPORT not only receives the missing fragments at the same

---

<sup>2</sup>While true, note that a node never expects this to be the case given the capped-bandwidth channel.

channel bandwidth that its peer distributor would have served at, but also receives less credit because it received fewer fragments via P2P.

Second, assuming the exchange is currently untimely (i.e.,  $\max(r_0, i_0/w) = r_0$ ),  $\mathcal{C}_i$  represents how much longer a client is willing to wait beyond the expected  $i/w$  duration it should take to transfer the  $i$  fragments. If  $\pi > 0$ , then  $\mathcal{C}_i > 0$ ; the larger  $V$  then is, the less a consumer is willing to tolerate untimeliness. If instead  $\pi < 0$ , then  $\mathcal{C}_i < 0$ , and the consumer is unwilling to continue the exchange.

Third, the  $\max(\cdot)$  term is a side-effect of Seer's policy of not releasing a channel until the expected time for a given number of fragments (in this case,  $i_0$ ) elapses. If Seer had instead allowed immediate channel reuse, a distributor that sends at channel bandwidth on average but sends faster earlier on and slower later on could actually make the consumer more likely to abort an exchange early to cash in on being ahead.

**Distributors.** Like the consumer, a distributor is only willing to serve a particular consumer if it believes it will always be behind schedule and the amount the distributor can earn from switching is worth it.

LEMMA 6.3. Ignoring effects on expected timeliness, a distributor continues serving a consumer in some block exchange iff

$$\bar{r}_i \leq \bar{r}_0 + \frac{i}{w} + \mathcal{D}_i \quad (6.4)$$

where

$$\mathcal{D}_i = -\log_\delta \left( \frac{\delta^{(m-i_0)/w} \pi (i + i_0) + \delta^{i/w} V}{\delta^{(m-i_0)/w} \pi i_0 + V} \right)$$

*Proof.* A distributor serves a consumer iff there exists some  $i > 0$  such that

$$v(i_0 + i, r_i) \geq v(i_0, r_0)$$

where  $v(j, r)$  is the payoff that a distributor expects to earn in a particular channel if it has a cumulative acknowledgment for  $j$  fragments at time  $r$  and chooses to terminate the exchange early.

To define  $v$ , consider that a distributor that terminates an exchange at time  $r$  with an acknowledgment that covers  $j$  fragments will (a) receive a payment of  $\pi j$  after the expected transfer time of those  $j$  fragments has elapsed if it has not already  $(\max(r, j/w) - r_0)$  plus an additional delay equal to the expected time to transfer the missing fragments  $((m - j)/w)$  and (b) an expected utility of  $V$  when this channel is reused and allocated to other exchanges at time  $\max(r, j/w) - r_0$  from now. This gives us the following definition of  $v$ :

$$v(j, r) = \delta^{\max(r, j/w) - r_0} (\delta^{(m-j)/w} \pi j + V)$$

It can be seen that given this definition, the condition at the beginning of the proof is satisfied given condition (6.4).  $\square$

Corollary 6.4 describes sufficient conditions similar to those of Lemma 6.2 under which a distributor continues to serve a consumer.

**COROLLARY 6.4.** Ignoring effects on expected timeliness, a distributor continues serving a consumer in some block exchange if the expected payoff of serving this consumer is at least as good as serving any consumer on average, i.e.,

$$\frac{\delta^{m/w}}{1 - \delta^{m/w}} m\pi \geq V$$

and there exists some  $i > 0$  such that at the time  $r_i$  the distributor expects an acknowledgment for  $i > 0$  additional fragments,

1. The exchange is timely, or
2. the exchange is at least as timely.

*Proof.* We prove that condition (6.4) holds given one of the two aforementioned conditions.

*Case (1):* the distributor expects the exchange to be timely when it receives an acknowledgment for  $i$  more fragments, i.e.,  $r_i \leq (i_0 + i)/w$ . This implies that  $\bar{r}_i = (i_0 + i)/w$ ; since  $\bar{r}_0 \geq i_0/w$ , we have

$$\bar{r}_i \leq \bar{r}_0 + i/w$$

To ensure condition (6.4) holds, it is sufficient to ensure that  $\mathcal{D}_i \geq 0$ , i.e.,

$$\delta^{(m-i_0)/w} \pi(i_0 + i) + \delta^{i/w} V \geq \delta^{(m-i_0)/w} \pi i_0 + V$$

Moving terms around, we have

$$V \leq \frac{\delta^{(m-i_0)/w}}{1 - \delta^{i/w}} \pi i$$

Since  $V \leq \delta^{m/w} / (1 - \delta^{m/w}) \pi m$  and  $m / (1 - \delta^{m/w}) \leq i / (1 - \delta^{i/w})$  for  $i \leq m$ , we have

$$V \leq \frac{\delta^{m/w}}{1 - \delta^{m/w}} \pi m \leq \frac{\delta^{m/w}}{1 - \delta^{i/w}} \pi i \leq \frac{\delta^{(m-i_0)/w}}{1 - \delta^{i/w}} \pi i$$

as needed.

*Case (2):* the distributor expects the exchange to be at least as timely, i.e.,  $r_i \leq r_0 + i/w$ . Assume also that  $r_i > (i_0 + i)/w$ ; otherwise, condition (6.4) is satisfied using the argument made in case (1). This implies that  $\bar{r}_0 = r_0$  (the exchange is currently untimely); since  $\bar{r}_i = r_i \leq r_0 + i/w$ ,  $\bar{r}_i \leq \bar{r}_0 + i/w$ . The proof of case (2) can then be shown using the same argument as the one used in case (1).  $\square$

## From one single block exchange to all block exchanges in a channel

The previous section relies on a node's expectation of how long it would take to receive credit for  $i$  additional fragments ( $r_i$ ). This expectation depends on a peer's perceived timeliness, which likely depend on that peer's previous actions

with respect to that node, both in previous exchanges and even the current exchange. Moreover, in addition to timeliness, a node’s best response depends on other factors as well. For instance, a node may be judged on the basis of when it chooses to terminate an exchange or the expected cost of receiving credit for a fragment ( $\pi$ ) from this node. We collectively call these factors a node’s *reputation*.

A node’s reputation can greatly affect a node’s willingness to interact with a peer. In an extreme example, a peer may blacklist any node that terminates an exchange early; this may cause a node to not terminate an untimely exchange even if the conditions in §6.4.2 would say otherwise. The effect of  $\pi$  is also explicit in Lemmas 6.6 and 6.7; §6.4.4 will further discuss how Byzantine nodes affect  $\pi$ .

In this paper, we largely take the approach of keeping our results simple and general and leave it to nodes themselves how to calculate  $r_i$  or  $\pi$  or how to update a peer’s reputation based on observed behavior. Note that as the designer of the protocol, we can “correlate” how all non-Byzantine nodes choose to (initially) update their reputations. As long as this update rule is in a node’s best interest given how other nodes operate, then nodes will continue to use it. For instance, we can choose to have nodes maintain reputations which only keep track of how timely a node is and does not penalize a node for terminating an exchange early. We will make such decisions in the implementation of Seer (§6.5).

However, the one important restriction we make is that a node expects that being timely will never worsen its expected timeliness. Without this assumption, a client may be encouraged to be untimely if peers believed that timeliness now was sign of untimeliness in the future.

ASSUMPTION 6.5. A client being timely in the current exchange with any non-Byzantine peer will not negatively affect the peer’s view of the client’s expected timeliness in the future.

By using Assumption 6.5, we can prove that, even considering the effects on a node's expected timeliness, a node is no worse off being timely in the current exchange.

LEMMA 6.6. In a particular exchange, a consumer expects to be no worse off being timely for the remainder of an exchange than not.

*Proof.* First, assume that a distributor never cuts off the exchange early. Consider the consumer's best response at any instant  $r_0$  when it has  $i_0$  blocks. Given Lemma 6.2, a consumer cuts off an exchange as specified by Lemma 6.2, resulting in an expected payoff of

$$\bar{v}_C = \max_{i \geq 0} v_C(i_0 + i, r_i)$$

Observe that  $dv_C/dt \leq 0$ , i.e., for a given amount of payoff that a consumer gets after cutting off an exchange after a given number of additional fragments, a consumer's expected payoff is non-decreasing with respect to the time it takes to get those fragments (more specifically, it is decreasing if the exchange is untimely and constant if the exchange is already timely). As a result,  $d\bar{v}_C/dt \leq 0$ , which implies a consumer expects to be no worse off being timely.

Consider now if a distributor may cut off an exchange early if the exchange is expected to be sufficiently untimely. The effect this has on a consumer's payoff is that the values of  $i$  that  $\bar{v}_C$  maximizes over is restricted to values in which the distributor is still willing to serve the consumer. By Assumption 6.5, a consumer that is more timely for the remainder of the exchange expects its expected timeliness will be no worse, implying that the consumer expects that its distributor's perception of  $r_i$  will be no worse. By Lemma 6.3, this implies that a distributor is no less likely to continue serving a consumer that is timely, i.e., the values of  $i$  over which  $\bar{v}_C$  is maximized does not shrink as a result of the consumer being timely. It follows then that  $d\bar{v}_C/dt \leq 0$  remains true. □

LEMMA 6.7. In the context of the current exchange, a distributor expects to be better off being timely for the remainder of an exchange than not.

*Proof.* As in the proof of Lemma 6.6, we first assume that a consumer never cuts off an exchange early. Then a distributor, at any instant  $r_0$  with an acknowledgment for  $i_0$  blocks, expects to cut off an exchange as specified in Lemma 6.3, resulting in an expected payoff of

$$\bar{v}_{\mathcal{D}} = \max_{i \geq 0} v_{\mathcal{D}}(i_0 + i, r_i)$$

As before, since  $dv_{\mathcal{D}}/dt \leq 0$ , it follows that  $d\bar{v}_{\mathcal{D}}/dt \leq 0$ , i.e., a distributor expects to be no worse off being timely.

Consider now if a consumer may cut off an exchange early if the exchange is expected to be sufficiently untimely. As before, this constrains the values of  $i$  over which  $\bar{v}_{\mathcal{D}}$  maximizes. By Assumption 6.5 and Lemma 6.2, a distributor does not expect that being more timely should worsen its perceived timeliness (and thus increase  $r_i$ ). As a result, the values of  $i$  over which  $\bar{v}_{\mathcal{D}}$  is maximizing remains the same, so  $d\bar{v}_{\mathcal{D}}/dt \leq 0$  still remains true.  $\square$

### From one channel to all channels

Through Assumption 6.5, we can argue that a client prefers to be as timely as possible in all its exchanges that it chooses to participate in within a particular channel. In reality, however, a client is likely to be participating in simultaneous exchanges in multiple channels, and how a client allocates its bandwidth among these channels has an obvious impact on its payoff.

For instance, suppose that a node is involved in two exchanges, one of which is timely and one of which is not. Even if it is in a client's best interest to be timely in its exchanges, a client may actually be better off allocating bandwidth from its timely exchange to its untimely exchange to improve its overall utility.

Fortunately, for situations where a client is involved in an exchange with a peer that it is not colluding with, the capped-bandwidth channel saves the day:

LEMMA 6.8. With any peer that a client is not colluding with, the client never expects that the bandwidth experienced in the remainder of a particular exchange will ever exceed the channel bandwidth.

This lemma simply follows from the properties of the capped-bandwidth channel. Simply put, a node does not believe that sending or receiving at a higher speed will result in higher payoff, and so a node (weakly) prefers sending at the channel bandwidth. If a client does not expect to receive above channel bandwidth for any exchange, the best a client can do is transfer at the channel bandwidth on all channels. This encourages clients to continue to be on-time in all exchanges that they choose to participate in.

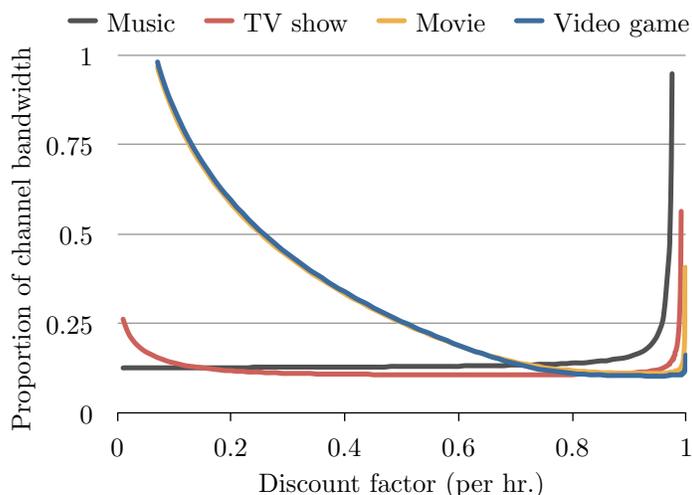
Note that, unlike our previous lemmas, Lemma 6.8 only holds for non-colluding peers that are involved in an exchange (in fact, it is because of Lemma 6.8 that Theorem 6.1 only holds for non-colluding peers). If peers are colluding, then we cannot guarantee that they will necessarily abide to the channel bandwidth. §6.4.5 describes this issue in more detail.

We now prove Theorem 6.1 using all our previous results.

*Proof (of Theorem 6.1).* A node's payoff is effectively determined by the exchanges that it takes part in. Being timely does not make a node's reputation worse (Assumption 6.5); consequently, Lemmas 6.6 and 6.7 show that a node is no worse off being timely in the context of every individual exchange. Finally, by Lemma 6.8, the timeliness of one exchange cannot affect the timeliness of other simultaneous exchanges, making being timely a node's expected best response. □

Content	Song	TV show	Movie	Game
Benefit (\$)	1	3	15	60
Size (MB)	8	1024	6144	8192
Block size (MB)	0.5	8	32	64

**Table 6.1:** Parameters used for various pieces of content (§6.4.3).

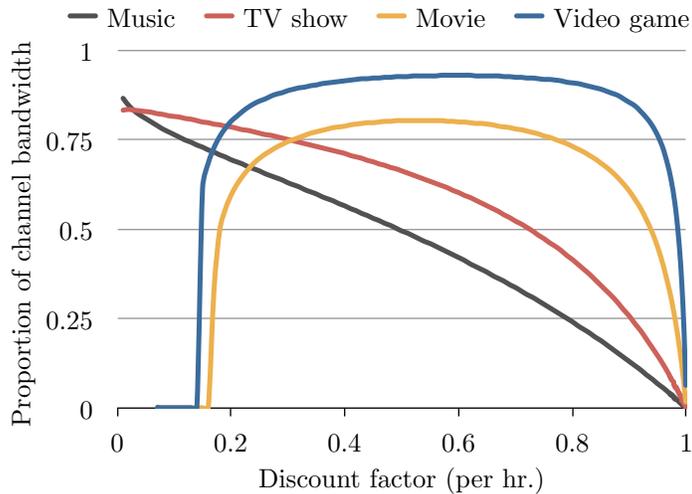


**Figure 6.5:** The minimum average effective channel bandwidth, expressed as a fraction of the channel bandwidth, that a client must experience to prefer using P2P over the free server tier. For values of the discount factor ( $x$ -axis) that do not have corresponding points, the client never prefers using P2P.

### 6.4.3 When do clients choose to use P2P?

While §6.4.2 showed that clients that chose to use P2P would prefer to be timely, recall that clients still have a choice regarding which service level they choose to use. Under what conditions do clients prefer using P2P over simply downloading from the server itself?

The answer to this question is complicated, depending on many factors, including how patient a client is, how much benefit a client gets from the content, how much the client must pay in order to get said content via the various methods, how quickly it believes it will get said content via the various



**Figure 6.6:** The minimum average effective channel bandwidth, expressed as a fraction of the channel bandwidth, that a client must experience to prefer using P2P over the paid server tier. For values of the discount factor ( $x$ -axis) that do not have corresponding points, the client never prefers using P2P.

methods, and what a client will do in the future. Moreover, this may also depend on how much credit a client has and how it values its credit.

To get an idea of what a client will choose, we consider the following scenario. Suppose we consider a client that is solely deciding which tier of service it will use exclusively for some piece of content (listed in Table 6.1), ignoring what it will do in the future. We conservatively assume that the client pays all the costs of acquiring the credits needed for the associated level of service right at the beginning of the exchange. We set the channel bandwidth to 1024 Kbps, 10 downstream channels, \$0.05/GB to upload content, and we set the cost of using paid download via the server to be 5 times more than via P2P.

Figures 6.5 and 6.6 shows, for various values of  $\delta$ , how much effective downstream bandwidth a client must experience in order to prefer using P2P over either the free or paid server tier of service using various pieces of content. As expected, very patient clients—those with values of  $\delta$  approaching 1—will

simply prefer getting its content via the cheapest method possible; for instance, for  $\delta = 0.9999$ , a client is always better off downloading via the free tier given the example parameters above. Clients who are extremely impatient ( $\delta < 0.1$ ) also prefer the cheapest option possible for large pieces of content: by the time these clients receive the content via any paid means, the content has been devalued beyond the amount the client had to pay to use the paid service in the first place.

#### 6.4.4 The effects of Byzantine failure

A client's beliefs regarding how Byzantine failures occur and their prevalence have an impact a client's incentives. However, because we abstract away these beliefs in the three parameters mentioned in §6.4.1, many of our results and their affiliated conditions in §6.4.2 and 6.4.3 continue to hold, albeit under more limited circumstances (e.g., a Byzantine peer may make it more costly to get an additional fragment or increase the expected amount of time needed to get fragments).

However, we do have to make one additional restriction on a rational client's beliefs regarding Byzantine failures to ensure that Theorem 6.1 still holds:

ASSUMPTION 6.9. Rational clients expect to be strictly worse off interacting (downloading as a consumer, serving as a distributor) with a known Byzantine peer than not.

The reason this assumption is needed is because if rational clients interacted freely with Byzantine peers (or possibly even preferred interacting with them), rational participants may be willing to be labeled as Byzantine if a deviation is expected to be sufficiently profitable. If rational and Byzantine clients alike deviate, then any stigma (and Byzantine labeling) behind and any disincentive for deviating is removed, making cooperation difficult to achieve.

In particular, without this assumption, it is possible that a consumer

would be able to profit in expectation by being receptive to receiving at rates higher than channel bandwidth and having a random Byzantine distributor send at a higher rate. Given Assumption 6.9, Lemma 6.8 still holds since a client believes no client will ever send at a rate higher than channel bandwidth except possibly Byzantine peers, which the client does not want to associate with anyway. As a result, even in the presence of Byzantine participants, Theorem 6.1 holds.

### 6.4.5 The effects of collusion

In previous sections, we considered incentives with respect to individual consumers and distributors. In this section, we explore how incentives are affected if consumers or distributors can form coalitions.

Much like §6.4.4, the existence of coalitions have an impact on the parameters from §6.4.1. Although we cannot ensure that coalitions cannot profitably deviate from Seer (in particular, Lemma 6.8 no longer holds), despite these profitable deviations, Theorem 6.1 continues to hold between clients that are not colluding, thus guaranteeing that clients that are not colluding still benefit from the cooperation of non-colluding clients in Seer.

To see why, we briefly consider a few ways colluders can potentially game Seer. We refer to a peer whom the client is colluding with as an insider (with respect to the client); all other peers are outsiders.

*A coalition could work together and use the free channels to get content.* A coalition could simply work together and use the limited free channels that each individual has at the free tier to download content. While possible, this greatly restricts the aggregate downstream bandwidth that the coalition possesses, which could prove to be useful if multiple participants are interested in downloading content from Seer. Note that this does not affect the actions of non-colluding clients.

*A coalition could exchange content among its members at a rate faster than the channel bandwidth.* A coalition may request a peer to download blocks,

but then choose to exchange content between coalition members at a rate exceeding the channel bandwidth. Such a deviation could benefit the coalition in that members get content more quickly, especially since they may also share the block key.

Ultimately, there is nothing in Seer that can disincentivize a coalition from performing this deviation. Note that even if coalition members finish data transfers early with one another, the server does not issue credit until the expected time ends, so finishing transfers early provides no benefit with respect to credit. If a coalition prefers to deviate even at the risk of being late for exchanges with outsiders, then coalition members may appear to outsiders as clients that transfer more slowly on average, which may adversely affect how outsiders view the timeliness of these members.

*A coalition could exchange content among its members outside of Seer.* A coalition may choose to exchange content among themselves using their own optimized protocol. As before, there is little we can do.

*A coalition could register to serve blocks they do not have.* A colluder could register for a block that it has previously downloaded but no longer has. Consumers, however, check to ensure that content is sent from the expected source address. Thus, a colluder here must acquire the content from a fellow insider before sending it to the consumer, which is effectively the same as transferring content outside of Seer.

*A coalition could send reports for exchanges that are not actually completed.* Two colluders that are paired together could feign the exchange and report to the server even though such an exchange never occurred. The result of such a deviation is one coalition member transfers credit to a fellow insider, which in itself does not enable the coalition to gain credit as a whole. Moreover, such a deviation uses up one of the distributor's channels that could have been used to earn credit.

*A coalition could start spurious exchanges with outsiders.* A member of a coalition could start spurious exchanges with outsiders in order to waste their resources. While this could slow down outsiders from obtaining content or

serving content (thereby allowing other coalition members to serve content in the meantime), recall that clients ultimately have free will as to whether they believe continuing this exchange is in their best interest. Thus, the outsider can always abort the exchange and receive pro-rated credit from the server (via the coalition), and distributors can always reject this coalition member in the future. Moreover, performing these spurious exchanges as a consumer wastes credit, since the consumer has to pay (potentially the server rate) for the block; performing this deviation as a distributor wastes a channel that could have been used to serve other consumers.

#### **6.4.6 Discussion**

In this section, we discuss some other considerations and issues.

**Macroeconomic issues** By providing a credit system and a bank, Seer effectively induces an economy among its participants. As with any economy, there are macroeconomic issues that need to be addressed. For instance, what if a particular client hoards all the credits it has? Alternatively, what if clients who earn credit spend it to download from the server, thus draining the economy of credit?

Many of the policies and decisions to shape the economy are outside the the scope of this paper. That said, various policies and decisions could be adopted by content providers to provide different guarantees. For instance, one way to introduce additional credit in the service is by rewarding credit to clients that have contributed more in the past. While coalition members may then have some incentive to generate spurious exchanges, one way to combat this is by, once again, leveraging the impatience of clients: charge a tax on distributors (making the exchanges negative sum) and then refund some multiple of the tax later on to the distributor. As a result of discounting, the coalition is actually worse off: while the tax refund is higher than the tax charged, the refund is heavily discounted by being in the future.

Alternatively, there could simply be periodic credit grants to all clients or those that have little credit and have not downloaded much recently. We must ensure that clients do not sit around for a “free lunch”; this can be achieved by simply making the amount low and the frequency rare to ensure that a client is unlikely to significantly benefit from this grant.

**Bootstrapping the service.** Until this point, we have ignored how clients gets credits in the first place when joining the service. Ultimately, this is a policy decision that depends on the content provider and its willingness to serve clients directly with the server. For instance, clients could be issued a certain number of credits to start, and if the participant purchases any content, additional credits could be issued. A client that wants to earn extra credit could potentially ask the server to assign it blocks to serve. In such cases, the server encrypts the block with a special symmetric key that prevents the client from peeking into the block’s contents. When a service first starts, the tracker itself could serve to help bootstrap the service by acting as a consumer downloading content from clients. While this would waste the tracker’s resources, this would also ensure that clients do their fair share to earn credit. Finally, a content provider could provide a method for clients to purchase credit with cash.

**Relaxing the reliance on the capped-bandwidth channel.** Lemma 6.8 relies on a combination of the capped-bandwidth channel and clients weakly preferring to continue using exchanging at channel bandwidth.

However, even if the channel did not have the bandwidth cap, note that an exchange’s bandwidth cannot be unilaterally decided by one node. The bandwidth is determined by both nodes, and to send at a higher bandwidth, a node must attempt sending at a faster rate, and its peer must also be willing to transfer at this rate. How much bandwidth a peer is willing or able to allocate to an exchange is unknown to a node, since it depends on a peer’s bandwidth and the other transfers that the peer is simultaneously undertaking, including how timely these other exchanges are, this peer’s perceived timeliness to its

own peers, how likely this peer’s own peers will themselves transfer at a faster rate, and so on.

Ultimately, to best respond, a node needs to form expectations over how all other nodes in the system are behaving in the system: whether they are acting as consumers and distributors, for which blocks they are processing, with whom they are exchanging, when the exchange began/is expected to end, etc. Such knowledge is unlikely to be known with any meaningful accuracy by any single node in the system. Thus, it may be realistic to simply assume that clients do not possess sufficient knowledge to assess whether it can expect to send more quickly with a particular non-Byzantine peer.

Another way to relax the capped-bandwidth channel is by having nodes believe that Byzantine peers transfer above channel bandwidth with positive probability. Because non-Byzantine nodes (at least) initially run the specified protocol and transfer at channel bandwidth, then any node observed to be deviating by transferring above channel bandwidth will be labeled as Byzantine by its peer in the exchange. If interacting with Byzantine peers is expected to have an adverse impact on a rational node’s payoff, a rational node will neither interact with any Byzantine peer nor transfer above channel bandwidth, as this may result in that node being labeled Byzantine and cutoff from a non-Byzantine peer.

Another, related method is to leverage acquiescent nodes, which are willing to follow whatever strategy they are given. For instance, suppose that, following the protocol, a node that observes its peer exceeding channel bandwidth in an exchange terminates the exchange—which makes the exchange no less timely or costly for a rational node than finishing at channel bandwidth—and blacklists the peer—which is no better than simply being labeled as untimely for this exchange. A rational node will then never deviate in this manner with a peer if a rational node believes there are a sufficiently large risk of encountering a peer who will enforce this punishment.

Both of these prior approaches assume that a node never appears to be transferring at a faster bandwidth than it actually is. We believe this is a

reasonable assumption, given strong identities of nodes and the reality that network loss makes nodes appear to be slower, not faster.

Finally, note that this lemma is regarding a node’s expectation, and reality does not necessarily have to match this expectation.

## 6.5 Implementation and evaluation

Seer is implemented using around 11,800 lines of heavily-multithreaded Java code and 430 lines of Apache Thrift [3] for implementing the control channels (RPC and serialization).

Data communication is done through an implementation of the capped-bandwidth channel, which is implemented by writing a wrapper around the Java’s socket library and rate-limiting the bandwidth by checking the throughput every second and idling the channel if the channel is sending or receiving too quickly. The channel attempts to maintain an average transfer bandwidth equal to the channel bandwidth *within* a particular fragment, i.e., it will try to catch up if it falls behind while transferring the fragment. *Between* fragments, the channel will delay if it is running ahead of schedule. This implies that if the channel falls behind on a particular fragment, it will not attempt to make up for this tardiness by transferring future fragments at a bandwidth higher than channel bandwidth (as specified in §6.3.2).

Our implementation separates a membership server, where clients sign in and log into the service, from the tracker server, which effectively handles tracking the state of all clients in the system. To allow users to access the tracker after logging in through the membership server, the membership server provides a server-signed access token that is effectively a lease on when they can access the service. This access token is presented with every message sent to the tracker. We separate out the server-backed distributor, but in our experiments, we co-locate the two together.

We use SSL channels to secure all control communication (i.e., any communication outside of the data transfers themselves). In our implementation,

when the consumer initiates a block exchange with a particular distributor (via a secure control channel), the consumer provides a nonce. The distributor, when initiating a data connection to the consumer, must present this nonce as the first piece of data. The consumer only accepts an incoming connection if the nonce received is the one expected from this IP address. We do not secure the data channel otherwise.

The Seer client, upon receiving the block ids that make up a piece of content, download the block ids in completely random order. For simplicity, we do not implement the bandwidth test that the server performs. The implementation assumes a single fixed block and fragment size. Our implementation of the client also currently does not perform any batching; thus, for a given workload, our experiments will put more load on the server.

We implemented the client policy based on the conditions described in §6.4.2. We set the client to expect the best-case scenario for continuation utility: full credit (and, for the consumer, the block) in a timely fashion. As a result, a client believes it can do really well in future exchanges, making it less patient with the current exchange it is part of. We fix the expected marginal net payoff per fragment to 1, and we calculate the expected time to get more fragments based on the observed bandwidth in the current exchange starting 1 second into the exchange. Clients re-evaluate their policies every time they have an opportunity to send a `FRAGMENT` or `ACK` message or after 1 second after the last message, whichever is sooner.

**Experimental setup.** We run all server-related code (membership, tracker, and server-backed distributor) on a single server with a two 8-core, 3 GHz AMD Opteron 4284 processors and 64 GB of RAM. The peers are running on about 180 department cluster machines with varying configurations, with one machine with a dual-core, single-threaded, 2.33 GHz Intel Core 2 Duo processor and 4 GB of RAM to a large number of machines (about half) with quad-core, dual-threaded 3.5 GHz Intel Xeon-powered machines with 16 GB of RAM. Overall, we have a total of about 670 cores with 1260 hardware threads

Block size (KB)	Throughput (Kbits/sec)
512	8721.6
1024	9606.36
2048	9714.63
4096	9510.08
8192	8870.26

**Table 6.2:** Average throughput per client vs. block size.

with approximately 14.6 GB of RAM on average per machine. We distribute peers on machines in a weighted fashion; the number of clients we put on each machine is linear to the number of cores the machine has.

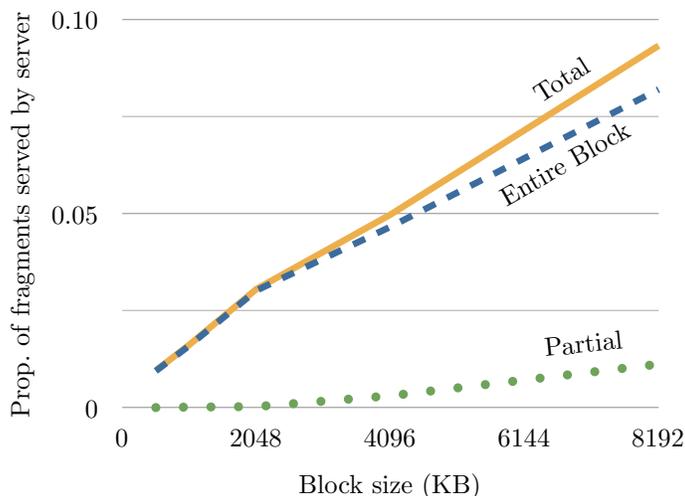
All of our machines run Ubuntu Linux 12.04. While our machines are connected through a department network via 1 Gbps links, we set the channel bandwidth to 1 Mbps and restrict the network bandwidth of peers to 10 capped-bandwidth downstream channels and 10 capped-bandwidth upstream channels (10 Mbps full-duplex). We do not restrict the number of channels for the server.

### 6.5.1 Selecting block size

We first explored the tradeoffs involved in selecting different block sizes. Intuitively, bigger block sizes incur lower overhead on the tracker because clients check in less frequently, but larger blocks result both in less entropy (due to fewer available blocks) and more involvement of the server if nodes fail. Moreover, because each block is served by a single peer and each channel has a capped bandwidth, larger blocks reduce the number of blocks that can be downloaded in parallel.

In this experiment, we deploy a 100-peer flash crowd which start downloading a 500 MB file simultaneously. To simulate more realistic network conditions, we introduce an artificial network latency of 100 ms at the application level on all network communication (control and data).

Table 6.2 shows the average client throughput observed by clients for



**Figure 6.7:** Proportion of data served by the server to a flash crowd with varying block sizes.

varying block sizes. As expected, a small block size can cause lower throughput due to the increased overhead imposed to the server. On the other hand, a large block size may hurt the throughput due to the reduced content entropy in the P2P network. In our experiment, a block size of 1-4 MB achieves about the same performance: about 11.4% higher than a block size of 512 KB. We believe that it is possible that with an implementation of Seer that uses finer-grained locks, smaller block sizes may perform better.

Figure 6.7 shows the proportion of fragments that Seer is serving using the server when block sizes are varied. We counted the number of fragments that clients downloaded from the server as a result of not using the assigned distributor at all (“Entire Block”) and aborting its exchange early (“Partial”). We calculated the proportion over the total number of fragments that all clients downloaded from the server. As block sizes increase, the decrease in entropy and increase in proportion of possibly aborting an exchange early results in more fragments being served from the server.

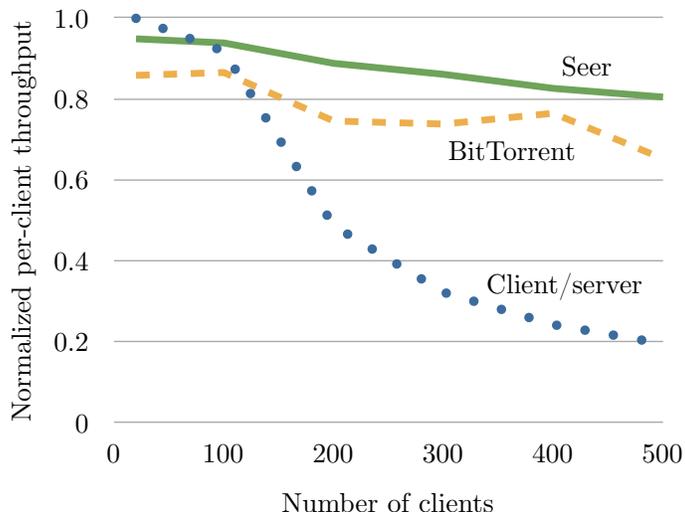
## 6.5.2 Scalability

To investigate the scalability of Seer, we ran experiments with 20 to 500 peers. For comparison, we ran the same experiments with BitTorrent and with a conventional client/server system. For BitTorrent, we use cTorrent 3.3.2 [6] to provide the clients and initial seeder; we use opentracker [13] to provide the tracker. We put no network bandwidth constraint on the seeder but limit all other clients to 10 Mbps upstream/downstream bandwidth. We leave all other parameters set to their defaults (e.g., 4 slots for serving peers, one of which is an unchoke slot).

For client/server, we use our own simple implementation of a server and client that uses multithreading to serve requests from clients directly. We use our capped-bandwidth data channels in this implementation to limit channel bandwidth to 10 Mbps/10 Mbps upstream/downstream bandwidth but use no other mechanisms from Seer. In all experiments, clients download 500MB of content and log out 10 seconds after completing their download. The block size used for Seer is 4MB with 64KB fragment. The size of a piece—BitTorrent’s version of the block—is 512KB (the recommended size for 500 MB of content).

**Flash crowd.** We first tested Seer with all clients joining the system at the same time. Figure 6.8 shows how Seer compares to BitTorrent and a traditional client/server system. As expected, in Seer, the average download throughput of client decreases as the number of clients increase due to increased load on the server having to coordinate block exchanges and make up for aborted exchanges. With 20 clients, Seer achieves 9.7 Mbps; at 200 clients, 9.1 Mbps; and by 500 clients, 8.2 Mbps. Not surprisingly, while client/server system works better for small numbers of clients, its throughput significantly drops with the larger number of clients as the bandwidth of the server becomes the bottleneck of system performance.

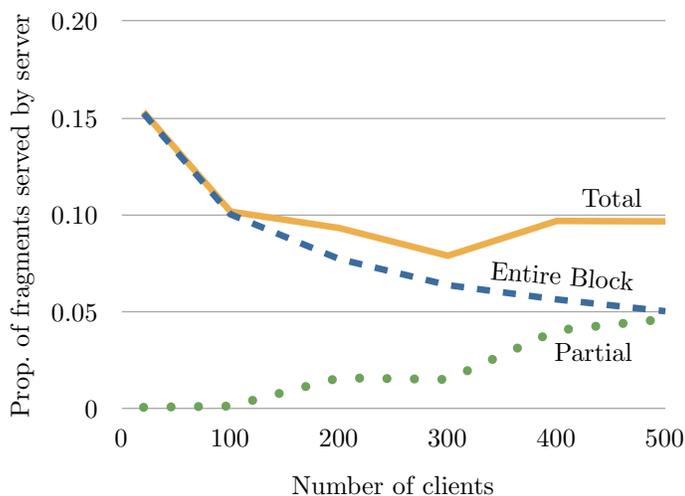
Surprisingly, Seer outperforms BitTorrent for the number of clients we evaluated. We believe this benefit comes from the server in Seer ensuring that consumers are always matched with distributors that have the desired block,



**Figure 6.8:** Comparison of throughput, normalized to the theoretical maximum of 10 Mbps, between Seer, BitTorrent, and client/server with a varying number of clients in a flash crowd.

have bandwidth, and are willing to send said block to the requesting consumer, whereas in BitTorrent, clients are given a random subset of peers to download from and upload to, some of which may not have bandwidth or may not be willing to serve the client. Note that another benefit that Seer’s capped-bandwidth channels would provide in more heterogeneous environments than the one in our evaluation is that nodes are effectively bandwidth-matched and have a clear understanding of what they are expected to do. On the other hand, BitTorrent clients may end up uploading content to peers in an vain attempt to try to curry their favor by unchoking them, which may occur when nodes have vastly different connections [71]. However, we expect that the performance of Seer would drop below BitTorrent with larger number of clients since the server is the performance bottleneck.

To get an idea of how many fragments the Seer server is serving, we measured what proportion of fragments Seer is serving using the server with varying number of clients. Figure 6.9 displays the results of this experiment.

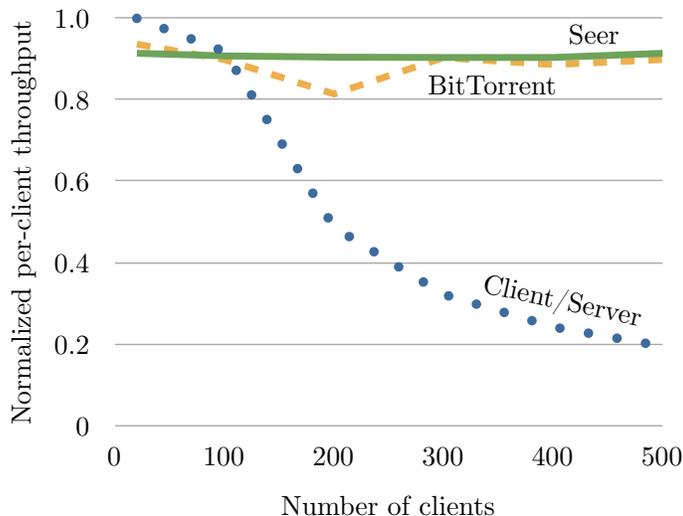


**Figure 6.9:** Proportion of received data served by the server with a varying number of clients in a flash crowd.

Unsurprisingly, for a small number of clients, the clients receive more data from the server. As we increase the number of clients, less than 10% of data is served directly by the server. On the other hand, the portion of partial block transfer served by the server increases as the number of clients increase: in our tests, Seer server serves only approximately 10% of data when there are at least 100 clients.

**Randomly-distributed downloads.** We then tested how well Seer performed versus other systems when clients logged in and downloaded the content at random times. We spawned clients at random delays, where the random delay was an exponentially-distributed random variable with  $\lambda = 1$  (per second, i.e., one client per second on average).

Given this, Figure 6.10 shows how Seer compares to BitTorrent and client/server systems. Both Seer and BitTorrent perform better, roughly providing stable throughput throughout the number of clients we tested, even with a relatively high rate of clients joining the system. The reason that throughput

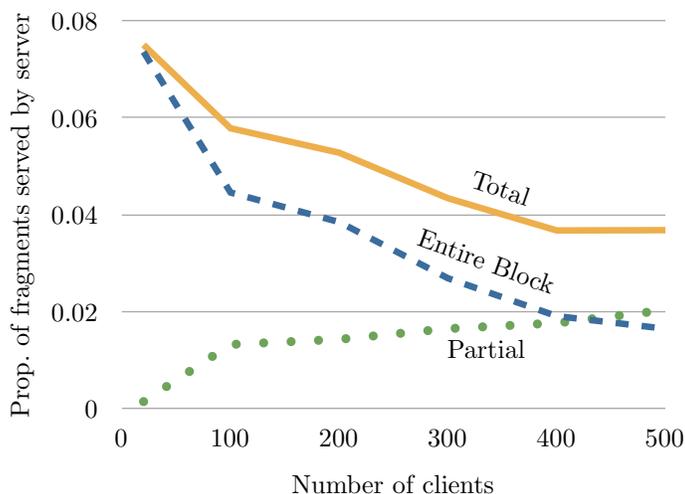


**Figure 6.10:** Comparison of throughput, normalized to the theoretical maximum of 10 Mbps, between Seer, BitTorrent, and client/server with a varying number of clients that download at random times.

remains more stable is because there are more opportunities for clients to serve as distributors to help other clients. This is particularly true for BitTorrent; since clients are given a random set of peers to work with, having clients join at random times ensures that the random set of peers is more likely to have the content than in the flash crowd scenario. For Seer, staggering when clients join the system ensures that fewer peers have to go through the server to get the content, reducing the load on the server itself. Unsurprisingly, client/server still performs poorly even in this case, as bandwidth continues to be a bottleneck in the system for most clients downloading the content.

## 6.6 Summary

In this paper, we designed and implemented a new hybrid P2P service—Seer—that we believe represents a significant advancement in the design and implementation of provably robust dependable systems. However, we believe Seer is



**Figure 6.11:** Proportion of received data served by the server with a varying number of clients that download at random times.

just the first of many. We are interested in seeing if our approach to dealing with collusion can be used to deal with Sybil identities—which complicates the granting of credit to participants—as well as trying to enable more flexibility in various parts of the protocol without weakening the robust guarantees provided by the service. For example, being able to provide consumers the ability to blacklist distributors may be useful in dealing with faulty (or even colluding) entities, but it is unclear how this may affect our incentives. We also believe that a Seer-like approach can be adapted to streaming media since the incentives in Seer are based on timing and node impatience. With the principles we learned in designing and building Seer, we believe that Seer provides us a useful blueprint for building new dependable robust systems.

# Chapter 7

## Related Work

There is extensive related work in the areas that this thesis covers. We describe some of the more relevant work that we are aware of here.

### 7.1 Incentives in the presence of failure

There has been much work in providing incentives when nodes may fail. The two primary examples that this thesis focuses on are the BAR model [24] and  $(k, t)$ -robustness [19, 21]. This thesis formalizes and extends the work of the BAR model and shows that, while appealing, the guarantees that  $(k, t)$ -robustness provides with respect to failure and coalition are hard to achieve in practice. As this thesis has covered these two models in detail, we omit further details here.

The  $t$ -maximin-like approach (Definition 3.14) has been used by Moscibroda et al. [86] to consider worst-case Byzantine behavior in the context of computer virus propagation. Furthermore, several systems have been built using the BAR model [78, 79].

Eliaz [49] has described a solution concept which is effectively  $(1, t)$ -robustness. Gradwohl [56] explored regret-free equilibria with  $t$  arbitrary or colluding nodes in leader election and random sampling games. As shown in §3.2, these notions do not admit equilibria in many real-world cooperative

services.

Our results are similar in spirit to previous work in mechanism design [47, 55, 66, 88, 97] which have explored the feasibility of using mechanisms to incentivize nodes to reveal their true types. Much of this work has found that mechanisms that incentivize nodes to reveal their true preferences or types for every possible realization of types are found to be often impossible or heavily restricted. Others [47, 88] achieved positive results by using Bayesian solution concepts instead of dominant ones. Mookherjee et al. [84] define conditions in which Bayesian incentive-compatible mechanisms can be replaced by equivalent dominant-strategy mechanisms.

Maximin strategies have been previously explored in conjunction with adversarial or possibly irrational agents. Alon et al. [25] quantify how, in a two-player zero-sum game, the payoff of playing a mixed maximin strategy is affected by an adversary who can choose its actions based on some information about its peer’s realized strategy. Tennenholtz [101], extending the work of Aumann et al. [30, 32], explores how maximin strategies can approximate the payoff of a Nash equilibrium when a rational node may not want to rely on the rationality of its peers.

## 7.2 Leveraging acquiescent participants

The existence of acquiescent or altruistic participants extend far beyond cooperative services into the real world (e.g., [27]). Consequently, there has been extensive work in game theory that has covered imperfect knowledge, private signaling, and the use of “irrationally” correct nodes. The use of acquiescence to achieve cooperation in the finitely-repeated prisoner’s dilemma game was first proposed by Kreps et al. [73]. It was shown that reputations could be maintained even when there was imperfect observation of actions [53]. Cripps et al. later showed that, under certain conditions, reputations cannot be maintained forever unless the action played by the irrational node was part of a rational node’s equilibrium strategy [45, 46].

The scenario considered in this thesis differs from much of this work considering Byzantine participants and network loss. None of the previous work explicitly consider the possibility of Byzantine and acquiescent players. Much of the previous work has also assumed that actions or their corresponding signals can either be observed at least publicly [45, 53], if not perfectly [73], or that any signal can occur as a result of any action with positive probability [46]. More importantly, previous work has focused on the existence (or nonexistence) of equilibrium under general conditions, whereas we are interested in applying this theory to a specific problem and a realistic model that we believe to be applicable to many distributed protocols.

Martin [83] introduced a notion of equilibrium in which rational nodes do not deviate regardless of Byzantine or acquiescent nodes' actions. Our work differs from this work by showing the need for acquiescence to address a key problem in cooperative services and considering real-world issues such as network costs and lossy links.

Vassilakis et al. [103] study how acquiescence affects content sharing in P2P services at the application level. Their approach, which does not address Byzantine participants, complements our own; we focus on network-level incentives and issues (such as lossy links) that motivate participants to actually send the content they share at the application level.

Note that the pester mechanism that we implement using acquiescent nodes in §4 can be applied to many systems, e.g., BAR Gossip [79], FOX [77], and PropShare [76] can use acquiescence to incentivize key exchange. Our technique may provide insight into the larger fair exchange problem [72, 90]. Finally, rational secret sharing [60] faces a similar problem to the last-exchange problem we considered in §4. However, without a pestering mechanism, our work is not directly applicable.

### 7.3 Incentives in the presence of collusion

Coalitions have been studied in depth in the game theory literature. We have seen how hard it is to achieve useful equilibria in cooperative services using strong Nash equilibrium [31], which is effectively  $n$ -resilience and thus requires a strategy profile be Pareto optimal, and  $k$ -resilience [19, 21], which has weaker but similar requirements. Green et al. [57] has similarly demonstrated the difficulty of dealing with coalitions in mechanism design by showing that, in the presence of coalitions, revealing one’s private information truthfully is not a dominant strategy.

Bernheim et al. [33] describe coalition-proof Nash equilibria, which weaken strong Nash equilibria by requiring that the equilibrium be preferable only to self-enforcing deviations, i.e., a deviation by a coalition in which no sub-coalition of this coalition can further deviate and profit. Considering only self-enforcing deviations provides little benefit when coalitions have exogenous means to ensure that coalition members deviate together, which we argue is often the case in cooperative services. For instance, friends may avoid hurting each other because of social repercussions (which can be formally modeled using notions of binding commitments or multimarket contact [34]). Finally, there has been work in defining correlated versions of strong Nash and coalition-proof equilibria (e.g., [35, 48, 85]); like their non-correlated counterparts, these equilibria require a best response despite how nodes collude and thus have similar shortcomings.

As previously mentioned in §5.1, the theory of the core and cooperative game theory studies how nodes can cooperate in order to maximize their own payoffs, but the focus on cooperative game theory—how nodes choose to collude—differs from the focus of our (and traditional, non-cooperative game theory) work—the interactions between non-cooperative nodes and coalitions.

In the context of mechanism design and auctions, Chen et al. [40] describe rationally-robust implementation, an interesting non-equilibrium-based solution concept that primarily aims to ensure that even if every individual

or coalition is given no initial hint of what to play, the underlying mechanism induces individuals or coalitions to choose strategies that ultimately preserve some desired system property. As a result, players may play multiple strategies, as in our equilibria; unlike equilibrium-based approaches, rationally-robust implementation does not predict the exact strategies that will be used, which ultimately may be any undominated strategy. It is unclear whether rationally-robust implementation's notion of dominance can remove enough strategies in games based on cooperative services to enable the existence of useful properties that hold for all surviving strategies.

Another way to deal with collusion is by aiming for an approximate best response or  $\epsilon$ -equilibrium (e.g., [77, 78]), which guarantees that deviations only provide minimum benefit. This approach could be used to disincentivize coalitions if colluding provides limited benefit (which, as seen in Section 5.1, may not be the case) and is largely complementary to our approach. Similarly, DCast [109] is an overlay multicast protocol that guarantees each node that follows the protocol some baseline payoff, even if others may collude. However, DCast does not aim to be an equilibrium and thus provides no guarantees that nodes will actually follow the protocol.

In some cases (e.g., in a multicast cost-sharing game [29]), mechanisms can be designed that are robust to coalitional deviations. However, since it is difficult to devise such mechanisms, many systems focus instead on detecting or reducing the effects of collusion. Several content distribution systems (e.g., [92, 94]) use incentives that attempt to reduce the benefits of collusion. Lian et al. [80] use a variety of techniques to detect collusion in a popular P2P service. Reiter et al. [95] design a reputation mechanism that require nodes to solve puzzles to prove they have the content in question. Tran et al. [102] develop a credit-based system in which a node's reputation is based on the number of distinct credit issuers it has received credit from and filters out those issuers that have issued excessive credits. EigenTrust [69] uses trusted peers to provide reputations that are robust against limited misbehavior (due to coalitions or failure). Similarly, Feldman et al. [52] and Marti et al. [82]

describe reputation systems that place more trust and weight in certain peers' opinions. Finally, Zhang et al. [110] describe a heuristic for preventing colluding administrators from using links to increase the ranks of their pages in Google's PageRank algorithm. These systems can only ameliorate, not eliminate, the effects of collusion and provide no rigorous assurance that rational nodes will not deviate.

## 7.4 Other P2P and hybrid P2P systems

There have been a plethora of hybrid P2P services used in commercial applications (e.g., Spotify [17], Blizzard [5], PPTV [14, 64], LiveSky [108], and Akamai [2]). As previously mentioned, these services largely assume that clients will contribute their fair share, an assumption that may not hold in practice [22, 65, 96].

Hybrid P2P systems have also been proposed in previous research. Dandelion [100] is a hybrid P2P service that shares many similarities with Seer. Where Dandelion and Seer differs is in its incentives: while Dandelion mentions malicious, altruistic, and colluding nodes, there is little in the way of argument that shows that following Dandelion is in rational node's best interest despite these "irrational" agents, especially since altruistic nodes give away content for free.

Antfarm [92] is a hybrid P2P content distribution system that uses a BitTorrent-like protocol for disseminating content. Unlike BitTorrent and like Seer, a coordinator issues credits to clients, which use them to pay for content. Unlike Seer, Antfarm's main purpose for these tokens is for resource allocation: by forcing clients to cash in credits, the coordinator can track which swarms are actively downloading, and the coordinator can allocate trusted seeders to those swarms that would benefit the most from additional seeders. Although these tokens could play a part in incentivizing cooperation, Antfarm does not provide rigorous incentives for encouraging nodes to disseminate content or participate faithfully in bandwidth resource allocation, leaving BitTorrent to

deal with the former issue (which itself is known to be susceptible to gaming [68, 76, 81, 93, 99]). Moreover, Antfarm, like BitTorrent, does not handle malicious or colluding nodes.

Floodgate [89] is a P2P system that uses a credit system based on a trusted server. Floodgate uses tokens that do not have to be checked against the server that allow nodes to download some piece of content. Unlike Seer, Floodgate only uses the server for credits, does not leverage the server for anything else, and provides few theoretical guarantees that nodes want to participate in the service.

The benefits of P2P techniques in video-on-demand services have been studied in much previous work (e.g., [62, 63]). Zebra [41] and Zebroid [42] are two systems that implement a hybrid P2P-assisted video-on-demand service. Huang et al. [64] showed the benefits of using P2P techniques in PPLive [14], which has since become a hybrid architecture known as PPTV similar to the one proposed for Seer. Only a few of these services consider rational behavior at all, and none of this work provide guarantees that rational participants will faithfully participate (even ignoring collusion or Byzantine behavior, which are also not considered). Contracts [94] attempts to improve PPLive’s incentives for rational cooperation and has some facilities for handling collusion, but it provides no formal guarantees and does not consider the possibility of Byzantine participants.

Reliable Client Accounting (RCA) [23] describes a way of performing reliable accounting of all client interactions in hybrid CDNs. Clients maintain logs of interactions, report these logs to a centralized infrastructure, which audits the logs to find anomalous behavior. Such behavior is then quarantined and restricted to downloading from infrastructure (not P2P) nodes. RCA is largely complementary to our approach and could potentially be used to catch misbehaving nodes in the system, thus making the P2P component more attractive to clients.

Aperjis et al. [28] propose PACE, a mechanism that matches up nodes for multilateral exchanges—which Aperjis et al. argue are more efficient and

robust than bilateral exchanges such as those used by BitTorrent—and enables nodes to set their own price for serving content—which enable efficient use of network resources. While Seer also uses a multilateral exchange, Seer effectively provides only one single global price that is dictated by the tracker. However, the pricing mechanism provided by PACE is largely orthogonal to the mechanisms used by Seer, and a PACE-like pricing mechanism could be adopted by Seer.

Beyond hybrid P2P systems, there have been a variety of P2P systems used in practice (e.g., [7, 8, 11, 16]) which rely on users faithfully serving their fair share of content. Such assumptions are unlikely to hold in practice [22, 65, 96]. Kazaa [9] was formerly a popular P2P service that provided some incentives in the form of a reputation, which was later gamed by free-riding clients [10]. BitTorrent [4, 43] is, at the time of writing, a popular P2P service for distributing content, which leverages a tit-for-tat scheme to give preference to peers that have uploaded to a node in the past. There has been much subsequent work that has shown that the incentives provided by the protocol are susceptible to strategic manipulation [68, 76, 81, 93, 99].

EquiCast [70] and FOX [77] are a few other examples of research systems that have been developed to deal with rational incentives. Of these systems, only AntFarm deals with Byzantine behavior; none of them are able to leverage acquiescent nodes to encourage rational cooperation nor deal with rational collusion. Finally, as previously mentioned, there has been several systems that, like Seer, have been built upon the BAR model [24, 79, 78]. Unlike these systems, Seer leverages the server to tolerate the same types of nodes and deal with collusion without requiring some of the strong assumptions (e.g., risk-aversion) that these works relied on.

# Chapter 8

## Conclusion

This thesis spans both theory and systems by describing contributions to the design, reasoning, and implementation of incentives in cooperative services where nodes may be faulty, selfish, or correct. More generally, this thesis demonstrates the importance of considering the applicability of theory to real systems; after all, as we showed, theoretically-rigorous solution concepts that require a rational node to always prefer the equilibrium strategy, despite how Byzantine failures may occur or how nodes may collude, are unlikely to yield useful results in many practical cooperative services. Furthermore, while it may be desirable to aim for guarantees that hold in the absence of acquiescent nodes, such goals may sacrifice the ability to achieve cooperation at the end of the service. This loss may, in turn, result in the inability to guarantee cooperation at all in the entire service.

Fortunately, this dissertation is not all bad news. We showed that, by taking expectation into account, solution concepts can both provide rigorous guarantees regarding rational behavior in the presence of Byzantine failure and yield useful equilibria in applications of interest. We demonstrate that acquiescence, which often enables selfish free-riding, can be wielded to encourage rational cooperation. We introduce a novel approach for reasoning about the incentives when nodes may collude and solution concepts that allow us to provide strong guarantees about the behavior of rational nodes that may collude.

Finally, we describe the design and implementation of Seer, a system that provides robust, reliable, and scalable content dissemination and demonstrates that our theoretical insights can be applied to real systems.

# Bibliography

- [1] “1 in 5 Macs has malware on it. Does yours?”. <http://nakedsecurity.sophos.com/2012/04/24/mac-malware-study>.
- [2] Akamai. <http://www.akamai.com>.
- [3] Apache Thrift. <http://thrift.apache.org>.
- [4] BitTorrent. <http://bittorrent.com>.
- [5] Blizzard Entertainment. <http://www.blizzard.com>.
- [6] CTorrent. <http://http://ctorrent.sourceforge.net>.
- [7] Gnutella. <http://en.wikipedia.org/wiki/Gnutella>.
- [8] Incriptus. <http://www.incriptus.com>.
- [9] Kazaa. <http://en.wikipedia.org/wiki/Kazaa>.
- [10] Kazaa Lite. [http://en.wikipedia.org/wiki/Kazaa\\_Lite](http://en.wikipedia.org/wiki/Kazaa_Lite).
- [11] Napster. <http://en.wikipedia.org/wiki/Napster>.
- [12] OpenGarden. <http://opengarden.com>.
- [13] opentracker. <http://erdgeist.org/arts/software/opentracker>.
- [14] PPTV. <http://www.pptv.com>.
- [15] SETI@home. <http://setiathome.ssl.berkeley.edu>.

- [16] SpaceMonkey. <http://spacemonkey.com>.
- [17] Spotify. <http://www.spotify.com>.
- [18] Ittai Abraham, Lorenzo Alvisi, and Joseph Y. Halpern. Distributed computing meets game theory: Combining insights from two fields. *SIGACT News*, 42(2):69–76, 2011.
- [19] Ittai Abraham, Danny Dolev, Rica Gonen, and Joe Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *Proceedings of the 25th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, July 2006.
- [20] Ittai Abraham, Danny Dolev, and Joseph Y. Halpern. Private communication.
- [21] Ittai Abraham, Danny Dolev, and Joseph Y. Halpern. Lower bounds on implementing robust and resilient mediators. In *Proceedings of the 5th IACR Theory of Cryptography Conference*, March 2008.
- [22] Eytan Adar and Bernardo A. Huberman. Free riding on Gnutella. *First Monday*, 5(10):2–13, October 2000.
- [23] Paarijaat Aditya, Mingchen Zhao, Yin Lin, Andreas Haeberlen, Peter Druschel, Bruce Maggs, and Bill Wishon. Reliable client accounting for P2P-infrastructure hybrids. In *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation*, April 2012.
- [24] Amitanand S. Aiyer, Lorenzo Alvisi, Allen Clement, Michael Dahlin, Jean-Philippe Martin, and Carl Porth. BAR fault tolerance for cooperative services. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles*, October 2005.

- [25] Noga Alon, Yuval Emek, Michal Feldman, and Moshe Tennenholtz. Adversarial leakage in games. In *Proceedings of the 1st Symposium on Innovations in Computer Science*, January 2010.
- [26] Danielle Alyias, Dennis Batchelder, Joe Blackbird, Joe Faulhaber, David Felstead, Paul Henry, Jeff Jones, Jimmy Kuo, Marc Lauricella, Le Li, Nam Ng, Tim Rains, Vidya Sekhar, Holly Stewart, Matt Thomlinson, and Terry Zink. Microsoft Security Intelligence Report, volume 14, July through December, 2012.
- [27] James A. Andreoni and John H. Miller. Rational cooperation in the finitely repeated prisoner’s dilemma: Experimental evidence. *Economic Journal*, 103(418):570–585, May 1993.
- [28] Christina Aperjis, Michael J. Freedman, and Ramesh Johari. Peer-assisted content distribution with prices. In *Proceedings of the 4th ACM International Conference on emerging Networking EXperiments and Technologies*, December 2008.
- [29] Aaron Archer, Joan Feigenbaum, Arvind Krishnamurthy, Rahul Sami, and Scott Shenker. Approximation and collusion in multicast cost sharing. *Games and Economic Behavior*, 47(1):36–71, April 2004.
- [30] R. J. Aumann and M. Maschler. Some thoughts on the minimax principle. *Management Science*, 18(5):54–63, January 1972.
- [31] Robert J. Aumann. Acceptable points in general cooperative  $n$ -person games. *Annals of Mathematics Study* 40, 4:287–324, 1959.
- [32] Robert J. Aumman. On the non-transferable utility value: A comment on the Roth-Shaper examples. *Econometrica*, 53(3):667–677, May 1985.
- [33] B. Douglas Bernheim, Bezalel Peleg, and Michael D. Whinston. Coalition-proof Nash equilibria, I. Concepts. *Journal of Economic Theory*, 42(1):1–12, 1987.

- [34] B. Douglas Bernheim and Michael D. Whinston. Multimarket contact and collusive behavior. *The RAND Journal of Economics*, 21(1):1–26, Spring 1990.
- [35] Francis Bloch and Bhaskar Dutta. Correlated equilibria, incomplete information and coalitional deviations. *Games and Economic Behavior*, 66(2):721–728, July 2009.
- [36] Paul Bond. What Hollywood execs privately say about Netflix. <http://www.hollywoodreporter.com/news/hollywood-execs-privately-netflix-71957>.
- [37] Robert Boyd. Mistakes allow evolutionary stability in the repeated prisoner’s dilemma game. *Journal of Theoretical Biology*, 136(1):47–56, January 1989.
- [38] Jonathan Browning and Omar R. Valdimarsson. Iceland, data-center hub? <http://www.businessweek.com/articles/2012-03-28/iceland-data-center-hub>.
- [39] Brian X. Chen. Apple plans new data center in Oregon. <http://bits.blogs.nytimes.com/2012/02/22/apple-data-center-prineville/>.
- [40] Jing Chen, Silvio Micali, and Paul Valiant. Robustly leveraging collusion in combinatorial auctions. In *Proceedings of the 1st Symposium on Innovations in Computer Science*, January 2010.
- [41] Yih-Farn Chen, Yennun Huang, Rittwik Jana, Hongbo Jiang, Michael Rabinovich, Jeremy Rahe, Bin Wei, and Zhen Xiao. Towards capacity and profit optimization of video-on-demand services in a peer-assisted IPTV platform. *ACM Multimedia Systems*, 15(1):19–32, February 2009.
- [42] Yih-Farn Robin Chen, Rittwik Jana, Daniel Stern, Bin Wei, Mike Yang, and Hailong Sun. Zebroid: Using IPTV data to support peer-assisted VoD content delivery. In *Proceedings of the 19th International Workshop*

*on Network and Operating Systems Support for Digital Audio and Video*, 2009.

- [43] Bram Cohen. Incentives build robustness in BitTorrent. In *Proceedings of the First Workshop on Economics of Peer-to-Peer Systems*, June 2003.
- [44] Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, November 1982.
- [45] Martin W. Cripps, George J. Mailath, and Larry Samuelson. Imperfect monitoring and impermanent reputations. *Econometrica*, 72(2):407–432, March 2004.
- [46] Martin W. Cripps, George J. Mailath, and Larry Samuelson. Disappearing private reputations in long-run relationships. *Journal of Economic Theory*, 127(1):287–316, May 2007.
- [47] Claude d’Aspremont and Louis-André Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11(1):25–45, February 1979.
- [48] Ezra Einy and Bezalel Peleg. Coalition-proof communication equilibria. In *Social Choice, Welfare, and Ethics: Proceedings of the Eighth International Symposium in Economic Theory and Econometrics*, June 1995.
- [49] Kfir Eliaz. Fault tolerant implementation. *Review of Economic Studies*, 69:589–610, August 2002.
- [50] Joseph Farrell and Matthew Rabin. Cheap talk. *Journal of Economic Perspectives*, 10(3):103–118, Summer 1996.
- [51] Joan Feigenbaum and Scott Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, September 2002.

- [52] Michal Feldman, Kevin Lai, Ion Stoica, and John Chuang. Robust incentive techniques for peer-to-peer networks. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, May 2004.
- [53] Drew Fudenberg and David K. Levine. Maintaining a reputation when strategies are imperfectly observed. *Review of Economic Studies*, 59(3):561–579, July 1992.
- [54] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, MA, August 1991.
- [55] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, July 1973.
- [56] Ronen Gradwohl. Rationality in the full-information model. In *Proceedings of the 7th IACR Theory of Cryptography Conference*, February 2010.
- [57] Jerry Green and Jean-Jacques Laffont. On coalition incentive compatibility. *The Review of Economic Studies*, 46(2):243–254, April 1979.
- [58] Vassos Hadzilacos and Sam Toueg. Fault-tolerant broadcasts and related problems. In *Distributed Systems*, pages 97–145. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 2nd edition, 1993.
- [59] Andreas Haeberlen, Petr Kouznetsov, and Peter Druschel. PeerReview: Practical accountability for distributed systems. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles*, October 2007.
- [60] Joseph Halpern and Vanessa Teague. Rational secret sharing and multiparty computation. In *Proceedings of the 36th ACM Symposium on Theory of Computing*, June 2004.
- [61] Joseph Y. Halpern and Rafael Pass. Algorithmic rationality: Adding cost of computation to game theory. *ACM SIGecom Exchanges*, 10(2):9–15, June 2011.

- [62] Cheng Huang, Jin Li, and Keith W. Ross. Can Internet video-on-demand be profitable? In *ACM SIGCOMM Computer Communication Review*, volume 37, pages 133–144, October 2007.
- [63] Cheng Huang, Angela Wang, Jin Li, and Keith W. Ross. Understanding hybrid CDN-P2P: Why Limelight needs its own Red Swoosh. In *Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, May 2008.
- [64] Yan Huang, Tom Z. J. Fu, Dah-Ming Chiu, John C. S. Lui, and Cheng Huang. Challenges, design and analysis of a large-scale P2P-VoD system. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, August 2008.
- [65] Daniel Hughes, Geoff Coulson, and James Walkerdine. Free riding on Gnutella revisited: The bell tolls? *IEEE Distributed Systems Online*, 6(6), June 2005.
- [66] Philippe Jehiel, Moritz Meyer ter Vehn, Benny Moldovanu, and William R. Zame. The limits of ex-post implementation. *Econometrica*, 74(3):585–610, May 2006.
- [67] J. Jonsson and B. Kaliski. Public-key cryptography standards (PKCS) #1: RSA cryptography specifications version 2.1. <http://tools.ietf.org/html/rfc3447>.
- [68] Seung Jun and Mustaque Ahamad. Incentives in BitTorrent induce free riding. In *Proceedings of the Third Workshop on Economics of Peer-to-Peer Systems*, August 2005.
- [69] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The EigenTrust algorithm for reputation management in P2P networks. In *Proceedings of the 12th International World Wide Web Conference*, 2003.

- [70] Idit Keidar, Roie Melamed, and Ariel Orda. EquiCast: Scalable multicast with selfish users. In *Proceedings of the 25th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, July 2006.
- [71] Umair Waheed Khan and Umar Saif. Bittorrent for the less privileged. November 2011.
- [72] Steve Kremer, Olivier Markowitch, and Jianying Zhou. An intensive survey of fair non-repudiation protocols. *Computer Communications*, 25:1606–1621, November 2002.
- [73] David Kreps, Paul Milgrom, John Roberts, and Robert Wilson. Rational cooperation in the finitely repeated prisoners’ dilemma. *Journal of Economic Theory*, 27(2):245–252, August 1982.
- [74] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.
- [75] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.
- [76] Dave Levin, Katrina LaCurts, Neil Spring, and Bobby Bhattacharjee. BitTorrent is an auction: Analyzing and improving BitTorrent’s incentives. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, August 2008.
- [77] Dave Levin, Rob Sherwood, and Bobby Bhattacharjee. Fair file swarming with FOX. In *Proceedings of the 5th International Workshop on Peer-to-Peer Systems*, February 2006.
- [78] Harry C. Li, Allen Clement, Mirco Marchetti, Manos Kapritsos, Luke Robison, Lorenzo Alvisi, and Mike Dahlin. FlightPath: Obedience vs. choice in cooperative services. In *Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation*, December 2008.

- [79] Harry C. Li, Allen Clement, Edmund L. Wong, Jeff Napper, Indrajit Roy, Lorenzo Alvisi, and Michael Dahlin. BAR Gossip. In *Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation*, November 2006.
- [80] Qiao Lian, Zheng Zhang, Mao Yang, Ben Y. Zhao, Yafei Dai, and Xiaoming Li. An empirical study of collusion behavior in the Maze P2P file-sharing system. In *Proceedings of the 27th International Conference on Distributed Computing Systems*, June 2007.
- [81] Thomas Locher, Patrick Moor, Stefan Schmid, and Roger Wattenhofer. Free riding in BitTorrent is cheap. In *Proceedings of the 5th Workshop on Hot Topics in Networks*, November 2006.
- [82] Sergio Marti and Hector Garcia-Molina. Limited reputation sharing in P2P systems. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, May 2004.
- [83] Jean-Philippe Martin. Leveraging altruism in cooperative services. Technical Report MSR-TR-2007-76, Microsoft Research.
- [84] Dilip Mookherjee and Stefan Reichelstein. Dominant strategy implementation of Bayesian incentive compatible allocation rules. *Journal of Economic Theory*, 56(2):378–399, April 1992.
- [85] Diego Moreno and John Wooders. Coalition-proof equilibrium. *Games and Economic Behavior*, 17(1):80–112, November 1996.
- [86] Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proceedings of the 25th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, July 2006.
- [87] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA, 1991.

- [88] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, April 1983.
- [89] Srijiith K. Nair, Erik Zentveld, Bruno Crispo, and Andrew S. Tanenbaum. Floodgate: A micropayment incentivized P2P content delivery network. In *Proceedings of the 17th International Conference on Computer Communications and Networks*, August 2008.
- [90] Henning Pagnia and Felix C. Gärtner. On the impossibility of fair exchange without a trusted third party. Technical Report TUD-BS-1999-02, Darmstadt University of Technology Department of Computer Science.
- [91] Anthony Palazzo. Netflix CEO Hastings uses Facebook to announce viewership. <http://www.bloomberg.com/news/2013-04-11/netflix-ceo-hastings-uses-facebook-to-announce-viewership.html>.
- [92] Ryan S. Peterson and Emin Gün Sirer. Antfarm: Efficient content distribution with managed swarms. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, April 2009.
- [93] Michael Piatek, Tomas Isdal, Thomas E. Anderson, Arvind Krishnamurthy, and Arun Venkataramani. Do incentives build robustness in BitTorrent? In *Proceedings of the 4th USENIX Symposium on Networked Systems Design and Implementation*, April 2007.
- [94] Michael Piatek, Arvind Krishnamurthy, Arun Venkataramani, Richard Yang, David Zhang, and Alexander Jaffe. Contracts: Practical contribution incentives for P2P live streaming. In *Proceedings of the 7th USENIX Symposium on Networked Systems Design and Implementation*, April 2010.

- [95] Michael K. Reiter, Vyas Sekar, Chad Spensky, and Zhenghao Zhang. Making peer-assisted content distribution robust to collusion using bandwidth puzzles. In *Proceedings of the 5th International Conference on Information Systems Security*, December 2009.
- [96] Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble. A measurement study of peer-to-peer file sharing systems. In *Proceedings of the 9th Annual ACM/SPIE Multimedia Computing and Networking*, January 2002.
- [97] Mark Allen Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, April 1975.
- [98] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, November 1979.
- [99] Michael Sirivianos, Jong Han Park, Rex Chen, and Xiaowei Yang. Free-riding in BitTorrent networks with the large view exploit. In *Proceedings of the 6th International Workshop on Peer-to-Peer Systems*, February 2007.
- [100] Michael Sirivianos, Jong Han Park, Xiaowei Yang, and Stanislaw Jarecki. Dandelion: Cooperative content distribution with robust incentives. In *Proceedings of the 2007 USENIX Annual Technical Conference*, June 2007.
- [101] Moshe Tennenholtz. Competitive safety analysis: Robust decision-making in multi-agent systems. *Journal of Artificial Intelligence Research*, 17:363–378, November 2002.
- [102] Nguyen Tran, Jinyang Li, and Lakshminarayanan Subramanian. Collusion-resilient credit-based reputations for peer-to-peer content dis-

- tribution. In *Proceedings of the 2010 Workshop on the Economics of Networks, Systems, and Computation*, October 2010.
- [103] Dimitrios K. Vassilakis and Vasilis Vassalos. An analysis of peer-to-peer networks with altruistic peers. *Peer-to-Peer Networking and Applications*, 2(2):109–127, June 2009.
- [104] Edmund L. Wong and Lorenzo Alvisi. What’s a little collusion between friends? In *Proceedings of the 32nd Annual ACM Symposium on Principles of Distributed Computing*, July 2013.
- [105] Edmund L. Wong, Ji Hong, Sangmin Lee, and Lorenzo Alvisi. Seer: Glancing into the future of dependable MAD systems. Technical Report TR-13-08, The University of Texas at Austin Department of Computer Science.
- [106] Edmund L. Wong, Joshua B. Leners, and Lorenzo Alvisi. It’s on me! The benefit of altruism in BAR environments. In *Proceedings of the 24th International Symposium on Distributed Computing*, September 2010.
- [107] Edmund L. Wong, Isaac Levy, Lorenzo Alvisi, Allen Clement, and Mike Dahlin. Regret freedom isn’t free. In *Proceedings of the 15th International Conference on Principles of Distributed Systems*, December 2011.
- [108] Hao Yin, Xuening Liu, Tongyu Zhan, Vyas Sekar, Feng Qiu, Chuang Lin, Hui Zhang, and Bo Li. Design and deployment of a hybrid CDN-P2P system for live video streaming: Experiences with LiveSky. In *Proceedings of the 17th ACM International Conference on Multimedia*, October 2009.
- [109] Haifeng Yu, Phillip B. Gibbons, and Chenwei Shi. DCast: Sustaining collaboration in overlay multicast despite rational collusion. In *Proceedings of the 19th ACM Conference on Computer and Communications Security*, October 2012.

- [110] Hui Zhang, Ashish Goel, Ramesh Govindan, Kahn Mason, and Benjamin Van Roy. Making eigenvector-based reputation systems against collusions. In *Proceedings of the 3rd International Workshop on Algorithms and Models for the Web-Graph*, October 2004.