Rational Consensus: Extended Abstract

Joseph Y. Halpern
Cornell University
halpern@cs.cornell.edu

Xavier Vilaça
INESC-ID Instituto Superior Técnico
Universidade de Lisboa
xvilaca@gsd.inesc-id.pt

ABSTRACT

We provide a game-theoretic analysis of consensus, assuming that processes are controlled by rational agents and may fail by crashing. We consider agents that *care only about consensus*: that is, (a) an agent's utility depends only on the consensus value achieved (and not, for example, on the number of messages the agent sends) and (b) agents strictly prefer reaching consensus to not reaching consensus. We show that, under these assumptions, there is no *ex post Nash Equilibrium*, even with only one failure. Roughly speaking, this means that there must always exist a *failure pattern* (a description of who fails, when they fail, and which agents they do not send messages to in the round that they fail) and initial preferences for which an agent can gain by deviating. On the other hand, if we assume that there is a distribution π on the failure patterns and initial preferences, then under minimal assumptions on π , there is a Nash equilibrium that tolerates f failures (i.e., π puts probability 1 on there being at most f failures) if $f + 1 < n$ (where n is the total number of agents). Moreover, we show that a slight extension of the Nash equilibrium strategy is also a *sequential* equilibrium (under the same assumptions about the distribution π).

Keywords

Consensus; Game Theory; Crash Failures

1. Introduction

Consensus is a fundamental problem in distributed computing; it plays a key role in state machine replication, transaction commitment, and many other tasks where agreement among processes is required. It is well known that consensus cannot be deterministically achieved in asynchronous systems [9], but can be achieved in synchronous systems even if we allow Byzantine failures (see, e.g., [16]). The assumption in all these solutions is that the reason that processes do not follow the protocol is that they have been taken over by some adversary.

There has been a great deal of interest recently in viewing at least some of the processes as being under the control of rational agents, who try to influence outcomes in a way that promotes their self interest. Halpern and Teague [13] were perhaps the first to do this. Their focus was on secret sharing and multiparty computation. Following [2, 3, 6, 12], we are interested in applying these ideas to standard problems in game theory. And like [1, 4, 6], we are interested in what happens when there is a mix of rational and faulty agents. For the purposes of this paper, we restrict to crash failures. As we shall see, a number of subtle issues arise even in this relatively simple setting.

We focus on the *fair* consensus problem, where fairness means that the input of every agent is selected with equal probability. Fairness seems critical in applications where we do not want agents to be able to influence an outcome unduly. For instance, when agents must decide whether to commit or abort a transaction, it is useful to ensure that the outcome reflects the preferences of the agents, so that if a majority of agents prefers a particular outcome, it is selected with higher probability. Abraham, Dolev and Halpern [2] present a protocol for fair leader election that even tolerates coalitions of rational agents. That is, in equilibrium, a leader is elected, and each agent is elected with equal probability. Fair leader election can be used to solve fair consensus (for example, once a leader is elected, the consensus value can be taken to be the leader's value). However, the protocol of [2] assumes that there are no faulty agents. Groce et al. [12] directly provide protocols for consensus with rational agents, but again, they do not consider faulty agents and do not require fairness. Afek et al. [3] and Bei, Chen, and Zhang [6] provide protocols for consensus with crash failures and rational agents. However, Afek et al.'s protocol works only under strong assumptions about agents' preferences, such as an agent having a strict preference for outcomes where it learns the input of other agents, while Bei, Chen, and Zhang require that their protocol be robust to deviations (that is, it achieves agreement even if rational agents deviate), a requirement that we view as unreasonably strong (see Section 3). Neither of these protocols satisfy the fairness requirement. Moreover, the protocol proposed by Afek et al. is not even an equilibrium if some agent knows the input of other agents. As we show, this is not an accident.

To explain our result, we need to briefly recall the standard notion of *ex post Nash equilibrium*. In a setting where we have an adversary, a protocol is an *ex post* equilibrium if no agent has any incentive to deviate no matter what the adversary does. Formally, "no matter what the adversary

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODC'16 July 25–28, 2016, Chicago, IL, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-3964-3/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2933057.2933088>

does” is captured by saying that even if we fix the adversary’s choice (so that the agents essentially know what the adversary does), agents have no incentive to deviate. Abraham, Dolev, and Halpern [2] provide protocols for leader election (and hence consensus) that achieve ex post Nash equilibrium if there are no failures. Here, we show that even in synchronous systems, there is no consensus protocol that is an ex post Nash equilibrium if there can be even one crash failure.

In the case of crash failures, the adversary can be viewed as choosing two things: the *failure pattern*—which agents fail, when they fail, and which other agents they send a message to in the round that they fail, and the *initial configuration*—what the initial preference of each of the agents is. Roughly speaking, the reason that we cannot obtain an ex post Nash equilibrium is that if the failure pattern and initial configuration have a specific form, a rational agent i can take advantage of knowing this to increase the probability of obtaining consensus on its preferred value.

There might seem to be an inconsistency here. It is well known that we can achieve consensus in synchronous systems with crash failures, so it seems that we shouldn’t have any difficulty dealing with one possibly faulty agent and one rational agent who does not follow the protocol. After all, we can view a rational agent who deviates from the protocol as a faulty agent. But there is no contradiction. When the agent deviates from the purported equilibrium, consensus is still reached, just on a different value. That is, a rational agent may want to deviate so as to bias the decision, although a consensus is still reached.

To get around our impossibility result, rather than trying to achieve ex post Nash equilibrium, we assume that there is some distribution π on *contexts*: pairs $(F, \vec{v})$ consisting of a failure pattern F and an initial configuration $\vec{v}$. We show that under appropriate assumptions about π , if agents *care only about consensus*—specifically, if (a) an agent’s utility depends only on the consensus value achieved (and not, for example, on the number of messages the agent sends) and (b) agents strictly prefer reaching consensus to not reaching consensus—then there is a Nash equilibrium that tolerates up to f failures, as long as $f + 1 < n$, where n is the total number of agents. Specifically, we make two assumptions about π , namely, we assume that π *supports reachability* and is *uniform*. Roughly speaking, we say that π supports reachability if it attributes small probability to particular failure patterns that prevent information from one agent reaching an agent that has not crashed by the end of the protocol; we say that π is uniform if it attributes equal probability to equivalent failures of different agents. We believe that these assumptions apply in many practical systems; we discuss this further in Section 4.

Our Nash equilibrium strategy relies on “threats”; the threat that there will be no consensus if an agent deviates (and is caught). There might be some concern that these are empty threats, which will never be carried out. The notion of *sequential equilibrium* [15] is intended to deal with empty threats. Roughly speaking, a strategy is a sequential equilibrium if all agents are best responding to what the others are doing even off the equilibrium path. We generalize sequential equilibrium to our setting, where there might be failures, and show that the strategy that gives a Nash equilibrium can be slightly extended to give a sequential equilibrium that tolerates up to f failures.

The rest of this paper is organized as follows. In Section 2 we discuss the model that we are using. Our main technical results on Nash equilibrium and sequential equilibrium are given in Section 3. We conclude with some discussion of the assumptions in Section 4.

2. Model

We consider a synchronous message-passing system with n agents and reliable communication channels between each pair of agents. Time is divided into synchronous rounds. Each round is divided into a *send phase*, where agents send messages to other agents, a *receive phase*, where agents receive messages sent by other agents in the send phase of that round, and an *update phase*, where agents update the value of variables based on what they have sent and received. We denote by N the set of agents and assume that they have commonly-known identifiers in $\{0, \dots, n - 1\}$. Round m takes place between time m and time $m + 1$.

We now formalize the notion of run. We take a *round- m history* for agent i to be a sequence of form $(v, t_1, \dots, t_{m-1})$, where v is agent i ’s initial preference and t_j has the form (s_j, r_j, d_j) , where s_j is the set of messages that i sent in round j tagged by who they were sent to, r_j is the set of messages that i received in round j tagged by who they were sent by, and $d_j \in \{\lambda\} \cup V$ is i ’s decision (where λ denotes that no decision has been made yet and V is the set of decision values). A *global (round- m) history* has the form $(h_1, \dots, h_n)$ where h_i is a round- m history, if j receives a message $\mathbf{m}$ from i in round m' of h_j , then i sends $\mathbf{m}$ to j in round m' in history h_i . A *run* r is a function from time (which ranges over the natural numbers) to global histories such that (a) $r(m)$ is a global round- m history and (b) if $m < m'$, then for each agent i , i ’s history in $r(m)$ is a prefix of i ’s history in $r(m')$.

Agents are either correct or faulty in a run. An agent fails only by crashing. If it crashes in round m of run r , then it may send a message to some subset of agents in round m , but from then on, it sends no further messages. We assume that all messages sent are received in the round in which they are sent. Thus, we take a *failure* $\mathbf{f}$ of agent i to be a tuple (i, m, A) , where m is a round number (intuitively, the round at which i crashes) and A is a set of agents (intuitively, the set of agents j to whom i can send a message before it fails). We assume that if $m > 1$, then A is non-empty, so that i sends a message to at least one agent in round m if i fails in round m . (Intuitively, if $m > 1$, we are identifying the failure pattern where i crashes in round m and sends no message with the failure pattern where i crashes in round $m - 1$ and sends messages to all the agents.) A *failure pattern* F is a set of failures of distinct agents i . A run r has *context* $(F, \vec{v})$ if (a) $\vec{v}$ describes the initial preferences of the agents in r , (b) if $(i, m, A) \in F$, then i sends all messages according to its protocol in each round $m' < m$, sends no messages in each round $m' > m$, and sends messages according to its protocol only to the agents in A in round m , and (c) all messages sent in r are received in the round that they are sent. Let $\mathcal{R}(F, \vec{v})$ consist of all runs r that have context $(F, \vec{v})$. Let $\mathcal{R}(F)$ consist of all runs that have F as the set of failures.

In the consensus problem, we assume that each agent i has an initial preference v_i in some set V . For ease of exposition, we take $V = \{0, 1\}$. (Our results can easily be extended to deal with larger sets of possible values.) A protocol achieves

consensus if it satisfies the following properties [9]:

- **Agreement:** No two correct agents decide different values.
- **Termination:** Every correct agent eventually decides.
- **Integrity:** All agents decide at most once.
- **Validity:** If an agent decides v , then v was the initial preference of some agent.

We are interested in one other property: fairness. Note that, once we fix a context, a protocol for the agents generates a probability on *runs*, and hence on outcomes, in the obvious way. Fairness just says that each agent has probability at least $1/n$ of having its value be the consensus value, no matter what the context. More precisely, we have the following condition:

- **Fairness:** For each context $(F, \vec{v})$, if c of the nonfaulty agents in F have initial preference v , then the probability of v being the consensus decision conditional on $\mathcal{R}(F, \vec{v})$ is at least c/n .

It is straightforward to view a consensus problem as a game once we associate a utility function u_i with each agent i , where u_i maps each outcome to a utility for i . Technically, it is an *extensive-form Bayesian game*. In a Bayesian game, agents have *types*, which encode private information. In consensus, an agent’s type is its initial preference. A strategy for agent i in this game is just a protocol: a function from information sets to actions. As usual, we view an extensive-form game as being defined by a game tree, with the nodes where an agent i moves into *information sets* where, intuitively, two nodes are in the same information set of agent i if i has the same information at both. In our setting, the nodes in a game tree correspond to global histories, and agent i ’s information set at a global history is determined by i ’s history in that global history; that is, we can take i ’s information set at a global history h to consist of all global histories where i ’s history is the same as it is at h . Thus, we identify an information set I_i for agent i with a history h_i for agent i . If I_i is the information set associated with history h_i , we denote by $\mathcal{R}(I_i)$ the set of runs r where i has history h_i in $r(m)$.

In game theory, a *strategy* for agent i is a function that associates with each information set I_i for agent i a distribution over the actions that i can take at I_i . In distributed computing, a protocol for agent i is a function that associates with each history h_i for agent i a distribution over the actions that i can take at h_i . Since we are identifying histories for agent i with information sets, it is clear that a protocol for agent i can be identified with a strategy for agent i . In consensus, the actions involve sending messages and deciding on values. We assume that there is a special value $\perp$ that an agent can decide on. By deciding on $\perp$, an agent guarantees that there is no consensus. If we assume that an agent prefers to reach consensus on some value to not reaching consensus at all, in the language of Ben Porath [7], this means that each agent has a *punishment strategy*.

We next want to define an appropriate solution concept for our setting. The standard approach is to say that an equilibrium is a *strategy profile* (i.e., a tuple of strategies, one for each agent) where no agent can do better by deviating. “Doing better” is typically taken to mean “gets a higher

expected utility”. However, if we do not have a probability on contexts, we cannot compute an agent’s expected utility. We thus consider two families of solution concepts. In the first, we take “doing better” to mean that, for each fixed context, no agent can do better by deviating. Once we fix the context, the strategy profile generates a probability distribution on runs, and we can compute the expected utility. In the second approach we assume a distribution on contexts.

A strategy profile $\vec{\sigma}$ is an ϵ - f -Nash equilibrium if, for each fixed context $(F, \vec{v})$ where there are at most f faulty agents in F , and all agents i , there is no strategy σ'_i for agent i such that i can improve its expected utility by more than ϵ . Formally, if $u_i(\vec{\tau} \mid \mathcal{R}(F, \vec{v}))$ denotes i ’s expected utility if strategy profile $\vec{\tau}$ is played, conditional on the run being in $\mathcal{R}(F, \vec{v})$, we require that for all strategies σ'_i for i , $u_i((\sigma'_i, \vec{\sigma}_{-i}) \mid \mathcal{R}(F, \vec{v})) \leq u_i(\vec{\sigma} \mid \mathcal{R}(F, \vec{v})) + \epsilon$. An f -Nash equilibrium is a 0- f -Nash equilibrium. The notion of f -Nash equilibrium extends the notion of ex post Nash equilibrium by allowing up to f faulty agents; a 0-Nash equilibrium is an ex post Nash equilibrium.¹

Given a distribution π on contexts and a strategy profile $\vec{\sigma}$, π and $\vec{\sigma}$ determine a probability on runs denoted $\pi_{\vec{\sigma}}$ in the obvious way. We say that $\vec{\sigma}$ is an ϵ - π -Nash equilibrium if, for all agents i and all strategies σ'_i for i , we have $u_i(\sigma'_i, \vec{\sigma}_{-i}) \leq u_i(\vec{\sigma}) + \epsilon$, where now the expectation is taken with respect to the probability $\pi_{\vec{\sigma}}$. A π -Nash equilibrium is a 0- π -Nash equilibrium. If π puts probability 1 on there being no failures, then we get the standard notion of (ϵ) -Nash equilibrium.

3. Possibility and Impossibility Results for Consensus

In this section, we consider the consensus problem from a game-theoretic viewpoint. We focus on the case where agents care only about consensus, since this type of utility function seems to capture many situations of interest. For the rest of this section, let β_{0i} be i ’s utility if its initial preference is decided, let β_{1i} be i ’s utility if there is consensus, but not on i ’s initial preference, and let β_{2i} be i ’s utility if there is no consensus. The assumption that agents care only about consensus means that, for all i , $\beta_{0i} > \beta_{1i} > \beta_{2i}$.

Note that although we assume that agents prefer consensus to no consensus, unlike Bei, Chen, and Zhang. [6], we do not require that our algorithms guarantee consensus when rational agents deviate. Our algorithm does guarantee that there will be consensus if there are no deviations. On the other hand, we allow for the possibility that a deviation by a rational agent will result in there being no consensus. For example, suppose that a rational agent pretends to fail in a setting where there is a bound f on the number of crash failures. That means that if f other agents actually do crash, then some agent will detect that $f + 1$ agents seem to have crashed. Our algorithm requires that if an agent detects such an inconsistency, then it aborts. If the probability that f agents actually crash is low, in our framework, a rational agent may decide that it is worth the risk of pretending to

¹This definition is in the spirit of the notion of (k, t) -robustness as defined by Abraham et al. [1], where coalitions of size k are allowed in addition to t “faulty” agents, but here we restrict the behavior of the faulty agents to crash failures rather than allowing the faulty agents to follow an arbitrary protocol, and take $k = 1$. We also allow the deviating agents to fail (this assumption has no impact on our results).

crash if the potential gain is sufficiently large. Bei, Chen, and Zhang would not permit this, since they require consensus even if rational agents deviate from the algorithm. This requirement thus severely limits the possible deviations.

3.1 An Impossibility Result

We start by showing that there is no fair consensus protocol that is an f -Nash equilibrium.

THEOREM 1. *If $\vec{\sigma}$ solves fair consensus, agents care only about consensus, and $f \geq 1$, then $\vec{\sigma}$ is not an f -Nash equilibrium*

PROOF. Consider the initial configuration $\vec{v}$ where all agents but i have initial preference 0 and i has initial preference 1. If F^1 is the failure pattern where no agent fails, by Fairness, the agents must decide 1 with positive probability in context $(\vec{v}, F^1)$. It follows that there must be a failure pattern F^2 where only agent i fails but the agents decide 1 with positive probability in context $(\vec{v}, F^2)$. (In F^2 , i fails only after a decision has been made in F^1 .) If F^0 is the failure pattern where only i fails, and i fails immediately, before sending any messages, then it is clear that no agents can distinguish this context from one where all agents have initial preference 0, so all agents must decide 0, by the Validity requirement.

Put a partial order $\leq$ on failure patterns where only i crashes by taking $F \leq F'$ if either i crashes in an earlier round in F than in F' , or i crashes in the same round m in both F and F' , but the set of agents to whom i sends a message in F is a subset of the set of agents to whom i sends a message in F' . Clearly $F^0 < F^2$. Thus, there exists a minimal failure pattern F^* such that $F^0 < F^* \leq F^2$, only i fails in F^* , the consensus is on 1 with positive probability in context $(\vec{v}, F^*)$, the consensus is 0 with probability 1 in all contexts $(F, \vec{v})$ where only agent i fails in F and $F < F^*$. We can assume without loss of generality that i sends a message to some agent j in the round m in which i fails. To see this, note that if i crashes in the first round then i must send a message to some agent (otherwise $F^* = F^0$ and the decision is 0 with probability 1). And if i crashes in round $m > 1$, we have assumed that i sends at least one message before crashing (recall that we identify an agent crashing at round $m > 1$ and sending no messages with the agent crashing at round $m - 1$ and sending to all agents).

Now suppose that an agent j that receives a message from i in round m pretends not to receive that message. This makes the situation indistinguishable from the context $(F, \vec{v})$ where F is just like F^* except that i does not send a message to j in round m . Since $F^0 \leq F < F^*$, the decision must be 0 with probability 1 in context $(\vec{v}, F)$. Since j has initial preference 0 in $\vec{v}$, j can increase its expected utility by this pretense, so $\vec{\sigma}$ is not an f -Nash equilibrium. $\square$

3.2 Obtaining a π -Nash equilibrium

We now prove a positive result. If we are willing to assume that there is a distribution π on contexts with some reasonable properties, then we can get a fair π -Nash equilibrium. But, as we show below, there are some subtle problems in doing this.

Before discussing these problems, it is useful to recall some results from social choice theory. Consider a setting with n agents where each has a preference order (i.e., a total order) over some set O of outcomes. A *social-choice function* is a (possibly randomized) function that maps a profile of preference orders to an outcome. For example, we can consider

agents trying to elect a leader, where each agent has a preference order over the candidates; the social-choice function chooses a leader as a function of the expressed preferences. A social-choice function is *incentive compatible* if no agent can do better by lying about its preferences. The well-known *Gibbard-Satterthwaite* theorem [10, 18] says that if there are at least three possible outcomes, then the only incentive-compatible deterministic social-choice function f is a *dictatorship*; i.e., the function f just chooses a player i and takes the outcome to be i 's most-preferred candidate, ignoring all other agents' preferences. Gibbard [11] extends this result to show that if there are at least three outcomes, then the only randomized incentive-compatible social-choice function is a *random dictatorship*, which essentially amounts to choosing some player i according to some probability distribution and then choosing i 's value.

Bei, Chen, and Zehang [6] point out that a strategy profile that solves consensus can be viewed as a social-choice function: agents have preferences over three outcomes, 0, 1, and $\perp$, and the consensus value (or $\perp$, if there is no consensus) can be viewed as the outcome chosen by the function. A strategy profile that is a Nash equilibrium is clearly incentive-compatible; no agent has an incentive to lie about its preferences. Thus, it follows from Gibbard's [11] result that a solution to rational consensus must be a randomized dictatorship. And, indeed, our protocols can be viewed as implementing a randomized dictatorship: one agent is chosen at random, and its value becomes the consensus value. However, implementing such a randomized dictatorship in our setting is nontrivial because of the possibility of failures.²

A naive protocol.

We start with a protocol that, while not solving the problem, has many of the essential features of our solution, and also helps to point out the subtleties. Consider the following slight variant of one of the early protocols for consensus [8]: In round 1, each agent i broadcasts a tuple $(i, v_i, x_{i0}, \dots, x_{if})$, where v_i is i 's initial preference, and x_{it} is a random element in $\{0, \dots, n - t\}$. For round 2, $\dots$, $f + 1$, each agent i broadcasts all the tuples $(j, v_j, \vec{x}_j)$ that i received and did not already forward in earlier rounds. At the end of round $f + 1$, each agent checks for consistency; specifically, it checks that it has received tuples from at least $n - f$ agents and that it has not received distinct tuples claimed to have been sent by some agent j . If i detects an inconsistency, then i decides $\perp$. Otherwise, suppose that i received tuples from $n - t$ agents. Then i computes the sum mod $n - t$ of the values x_{jt} for each agent j from which it received a tuple. If the sum is S , then i decides on the value of the agent with the $(S + 1)$ st highest id among the $n - t$ agents from which it received tuples. (Here is where we are implementing the random dictatorship.) Note that the random value x_{jt} is used by i in computing the consensus value if

²We remark that Theorem 1 of Bei, Chen, and Zhang [6] claims that, given a fixed failure pattern, a strategy profile for consensus that is a Nash equilibrium must implement a dictatorship, rather than randomized dictatorship. While this is true if we restrict to deterministic strategies, neither we nor Bei, Chen, and Zhang do so. We have not checked carefully whether results of Bei, Chen, and Zhang that depend on their Theorem 1 continue to hold once we allow for randomized dictatorships.

exactly t faulty agents are discovered; the remaining random values sent by agent j in the first round are discarded.

It is straightforward to check that if all nonfaulty agents follow this protocol, then they will all agree on the set of tuples received (see the proof of Theorem 2 for an argument similar in spirit), and so will choose the same decision value, and each agent whose value is considered has an equal chance of having their value determine the outcome. But this will not be in general a π -Nash equilibrium if π allows up to f failures, that is, π puts probability 0 on all failure patterns that have more than f failures and $f \geq 2$.

Consider a distribution π that puts positive probability on all contexts with at most f failures, and an initial configuration where agent 1 prefers 1, but all other agents prefer 0. Agent 1 follows the protocol in the first round, and receives a message from all the other agents. We claim that agent 1 may have an incentive to pretend to fail (without sending any messages) at this point. Agent 1 can gain by doing this if one of the other agents, say agent 2, crashed in the first round and sent a message only to agent 1. In this case, if 1 pretends to crash, no other agent will learn 2's initial preference, so 1's initial preference will have a somewhat higher probability (at least $\frac{1}{n-1} - \frac{1}{n}$) of becoming the consensus decision. Of course, there is a risk in pretending to crash: if f agents really do crash, then an inconsistency will be detected, and the decision will be $\perp$. Let $\alpha_{<f}$ be the probability of there being fewer than f failures and at least one agent crashing in the first round who does not send to any agent other than 1 (this is the probability that 1 gains some utility by its action); let $\alpha_{=f}$ be the probability of there being f crashes other than 1 (this is an upper bound on the probability that 1 loses utility by its action). Then 1's expected gain by deviating is at least

$$(\beta_{0i} - \beta_{1i}) \left(\frac{1}{n-1} - \frac{1}{n} \right) \alpha_{<f} - (\beta_{0i} - \beta_{2i}) \alpha_{=f}.$$

This is a small quantity. However, if f is reasonably large and failures are unlikely, we would expect $\alpha_{=f}$ to be much smaller than $\alpha_{<f}$, so as the number f of failures that the protocol is designed to handle increases, deviating becomes more and more likely to produce a (small) gain.

A π -Nash equilibrium.

There are three problems with the preceding protocol. The first is that, even if 1 pretends to fail, 1's value will be considered a potential consensus value, since everyone received the value before 1 failed. This means that there is little downside in pretending to fail. Roughly speaking, we deal with this problem by taking into consideration only the values of nonfaulty agents when deciding on a consensus value. The second problem is that since agents learn the random values $(x_{i0}, \dots, x_{if})$ that will be used in determining the consensus value in round 1, they may be able to guess with high probability the value that will be decided on at a point when they can still influence the outcome. To address this problem, agents do not send these random values in the first round; instead, they use secret sharing [19], so as to allow the nonfaulty agents to reconstruct these random values when they need to decide on the consensus value. This prevents agents from being able to guess with high probability what the decision will be too early. The third problem is that in some cases agents can safely lie about the messages sent by other agents (e.g., i can pretend that another agent

did not crash). We could solve this by assuming that messages can be signed using unforgeable signatures. We do not need this or any other cryptographic assumption. Instead, we use some randomization to ensure that if an agent lies about a message that was sent, it will be caught with high probability.

Thus, in our algorithm, an agent i generates random numbers for two reasons. The first is that it generates $f+1$ random numbers $(x_{i0}, \dots, x_{if})$, where x_{it} is used in choosing the consensus value if there are exactly t faulty agents discovered, and then, as we suggested above, shares them using secret sharing, so that the numbers can be reconstructed at the appropriate time (see below). The second is that it generates $n-1$ additional random numbers, denoted $z_{ij}^m[i]$, one for each agent $j \neq i$, in each round m , and sends them to j in round m . Then if agent j claims that it got a message in round m from i , it will have to also provide $z_{ij}^m[i]$ as proof.

In more detail, we proceed as follows. Initially, each agent i generates a random tuple $(x_{i0}, \dots, x_{if})$, where x_{it} is in $\{0, \dots, n-t\}$. It then computes $f+1$ random polynomials $q_{i0}, \dots, q_{if}$, each of degree 1, such that $q_{it}(0) = x_{it}$. It then sends $(q_{i0}(j), \dots, q_{if}(j))$ to agent j . The upshot of this is that no agent will be able to compute x_{it} given this information (since one point on a degree-1 polynomial q_{it} gives no information regarding $q_{it}(0)$). In addition, in round 1, each agent i sends v_i to each agent j , just as in the naive algorithm; it also generates the random number $z_{ij}^1[i]$ and a special random number z , and sends each agent the vector $z_{ij}^1 = (z_{ij}^1[1], \dots, z_{ij}^1[n])$, where $z_{ij}^1[j'] = 0$ for $j' \neq i, j$ and $z_{ij}^1[j] = z$. (As we said, these random numbers form a "signature"; their role will become clearer in the proof.) Finally, in round 1, agent i sends a status report SR_i^1 ; we discuss this in more detail below. In the receive phase of round 1, agent i adds all the values received from other agents to the set ST_i .

In round m with $2 \leq m \leq f$, i again sends a status report SR_i^m and a vector z_{ij}^m . For each agent j , $SR_i^m[j]$ is a tuple of the form (m, x) , where m is the first round that i knows that j crashed ($m = \infty$ if i believes that j has not yet crashed), and x is either the vector z_{ji}^{m-1} of random values sent by j in $m-1$ (if i believes that j has not yet crashed) or an agent that told i that j crashed in round m . The tuple z_{ij}^m is computed by setting $z_{ij}^m[l]$ for $l \neq i, j$ to be $z_{li}^{m-1}[l]$, the random number sent by l in the previous round (this will be used to prove that i really got a message from l in the previous round—it is our replacement for unforgeable signatures); again, $z_{ij}^m[i]$ is a random value generated by i . In round $f+1$, i also sends j the secret shares y_{ji}^t it received in round 1 from each agent l (i.e., the value $q_l^t(i)$ that it received from l , assuming that l did not lie). This enables j to compute the polynomials q_{it} , and hence the secret $q_{it}(0) = x_{it}$ for $0 \leq t \leq f$.

If i detects an inconsistency in round $m \leq f+1$, then i decides $\perp$, where i detects an inconsistency in round m if the messages received by i are inconsistent with all agents following the protocol except that up to f agents may crash. (In the full paper [14], we give an exhaustive list of all the ways that i can detect an inconsistency.)

If agent i does not detect an inconsistency at some round $m \leq f+1$, i proceeds as follows in round $f+1$. For each round $1 \leq m \leq f+1$ in a run r , agent i computes $NC_m(r)$, the set of agents that it believes did not crash up to and including round m . Take $NC_0(r) = N$ (the set

of all agents). Say that round m in run r *seems clean* if $NC_{m-1}(r) = NC_m(r)$. As we show (Theorem 2), if no inconsistency is detected in run r , then there must be a round in r that seems clean. Moreover, we show that if m^* is the first round in r that seems clean to a nonfaulty agent i , then all the nonfaulty agents agree that m^* is the first round that seems clean in r , and they agree on the initial preference of all agents in $NC_{m^*}(r)$, and the random numbers sent by these agents in round 1 messages in run r . The agents then use these random numbers to choose an agent j among the agents in $NC_{m^*}(r)$ and take v_j to be the consensus value.

The pseudocode for the strategy (protocol) $\bar{\sigma}^{cons}$ that implements this idea is given in Figure 1. We discuss the algorithm in more detail in the full paper.

We now prove that $\bar{\sigma}^{cons}$ gives a π -Nash equilibrium, under reasonable assumptions about π . We first prove that the protocol satisfies all the properties of fair consensus without making any assumptions about π .

THEOREM 2. $\bar{\sigma}^{cons}$ solves fair consensus if at most f agents crash, $f+1 < n$, and all the remaining agents follow the protocol.

The proof of this theorem (and all others) is in the full paper.

It remains to show that $\bar{\sigma}^{cons}$ is a π -Nash equilibrium. We show that $\bar{\sigma}^{cons}$ is a π -Nash equilibrium under appropriate assumptions about π . Specifically, we assume that π *supports reachability* and is *uniform*, notions that we now define. The reachability assumption has three parts. The first two parts consider how likely it is that some information that an agent j has will reach an agent that will decide on a value; the third part is quite similar, and considers how likely it is that a nonfaulty agent becomes aware that an agent j failed in round m . Of course, the answer to these questions depends in part on whether agents are supposed to send messages in every round (as is the case with $\bar{\sigma}^{cons}$). In the formal definition, we implicitly assume that this is the case. (So, effectively, the reachability assumption is appropriate only for protocols where agents send messages in every round.) Given agents i and $j \neq i$, a round- m information set I_i for i , a failure pattern F compatible with I_i , in that $\mathcal{R}(F) \cap \mathcal{R}(I_i) \neq \emptyset$, and $m' \geq m$, say that a nonfaulty agent $l \neq i$ is *reachable from j without i between rounds m' and $f+1$ given F* if there is a sequence $j_{m'}, \dots, j_{f+1}$ of agents different from i such that $j = j_{m'}$, for $m'' = \{m', \dots, f\}$, $j_{m''}$ has not failed prior to round m'' according to F , and either does not fail in round m'' or, if $m'' < f+1$, $j_{m''}$ fails in round m'' but sends a message to $j_{m''+1}$ before failing (i.e., if $(j_{m''}, m'', A) \in F$, then $j_{m''+1} \in A$), and $l = j_{f+1}$.

Note that if j is nonfaulty according to F , then a nonfaulty agent is certainly reachable from j without i between rounds m' and $f+1$; just take $j_{m'} = \dots = j_{f+1} = j$. But even if j fails in round m' according to F , as long as j can send a message to a nonfaulty agent other than i , or there is an appropriate chain of agents, then a nonfaulty agent is reachable from j without i by round $f+1$. The probability of there being a failure pattern for which a nonfaulty agent is reachable from j without i depends in part on how many agents are known to have failed in I_i ; the more agents are known not to have failed, the more likely we would expect a nonfaulty agent to be reachable from j without i .

We also want this condition to hold even conditional on a set of failure patterns, provided that the set of failure

patterns does not favor particular agents failing. To make this precise, we need a few more definitions. Say that an agent j is *known to be faulty in I_i* if j is faulty in all runs in $\mathcal{R}(I_i)$; thus, j is known to be faulty in I_i if j did not send a message to i at round $m-1$ according to I_i . Say that a set $\mathcal{F}$ of failure patterns *satisfies the permutation assumption with respect to a set F of failures and an information set I_i* if, for all permutations g of the agents that keep fixed the agents that fail in F or are known to be faulty in I_i , if $F' \in \mathcal{F}$, then so is $g(F')$, where $g(F')$ is the failure pattern that results by replacing each triple $(j, m'', A) \in F'$ by $(g(j), m'', g(A))$. $\mathcal{F}$ *satisfies the permutation assumption with respect to I_i* if $\mathcal{F}$ satisfies it with respect to the empty set of failures and I_i . Let $\mathcal{R}(F) = \cup_{F' \in \mathcal{F}} \mathcal{R}(F')$.

We say that π *supports reachability* if for all agents i , all time- m information sets I_i such that M agents are not known to be faulty in I_i , failure pattern F , and all sets $\mathcal{F}$ of failure patterns that satisfy the permutation assumption with respect to F and I_i , we have that

1. if $j \neq i$ is not known to be faulty in I_i and is not in F , then

$$\pi(\text{no nonfaulty agent } l \neq i \text{ is reachable from } j \text{ without } i \text{ between rounds } m \text{ and } f+1 \mid \mathcal{R}(I_i) \cap \mathcal{R}(F) \cap \mathcal{R}(\mathcal{F})) \leq \frac{1}{2M};$$
2. if $j \neq i$ is not known to be faulty in I_i and is not in F , then

$$\pi(\text{no nonfaulty agent } l \neq i \text{ is reachable from } j \text{ without } i \text{ between rounds } m-1 \text{ and } f+1 \mid \mathcal{R}(I_i) \cap \mathcal{R}(F) \cap \mathcal{R}(\mathcal{F})) \leq \frac{1}{2M};$$
3. if a message from some agent j not in F was received up to and including round $m-2$ but not in round $m-1$, then

$$\pi(\text{no nonfaulty agent } l \neq i \text{ is reachable from an agent } j' \neq i \text{ that did not receive a message from } j \text{ in round } m-1 \text{ without } i \text{ between rounds } m \text{ and } f+1 \mid \mathcal{R}(I_i) \cap \mathcal{R}(F) \cap \mathcal{R}(F)) \leq \frac{1}{2M}.$$

The first two requirements essentially say that if i hears from j in round $m-1$, then it is likely that other agents will hear from j as well in a way that affects the decision, even if i does not forward j 's information. That is, it is unlikely that j will fail right away, and do so in a way that prevents its information from having an effect. Similarly, the third requirement says that if i does not hear from j in round $m-1$ (as reflected in I_i), then it is likely that other agents will hear that j crashed at or before round $m-1$ even if i does not report this fact.

We next define the notion of uniformity. Given two failure patterns F^1 and F^2 , we say that F^1 and F^2 are *equivalent* if there is a permutation g of the agents such that $F^2 = g(F^1)$. We say that π is *uniform* if, for all equivalent failure patterns F^1 and F^2 and vectors $\vec{v}$ of initial preferences, we have $\pi(F^1, \vec{v}) = \pi(F^2, \vec{v})$. Intuitively, if π is uniform, then the probability of each failure pattern depends only on the number of messages omitted by each agent in each round; it does not depend on the identity of faulty agents.

THEOREM 3. If $f+1 < n$, π is a distribution that supports reachability, is uniform, and allows up to f failures, and agents care only about consensus, then $\bar{\sigma}^{cons}$ is a π -Nash equilibrium.

PROOF. Fix an agent i and a strategy σ_i . We must show that we have

$$u_i(\vec{\sigma}^{cons}) \geq u_i(\sigma_i, \vec{\sigma}_{-i}^{cons}). \quad (1)$$

Suppose, by way of contradiction, that (1) does not hold. Then i must deviate from σ_i^{cons} at some round m . Consider all the ways that i can deviate in round m that can affect the outcome (we discuss what it means to affect the outcome shortly):

1. i pretends to crash; it does not send messages to some subset of agents in round m (and then does not send messages from then on).
2. $m = 1$ and i sends $(i, 1 - v_i)$ to some agent $j \neq i$ (i.e., i lies about its initial preference to at least one agent).
3. i sends an incorrectly formatted message to $j \neq i$ (i.e., i sends a message that is different in format from that required by $\vec{\sigma}^{cons}$).
4. $m = 1$ and i sends values y_{ij}^t to an agent $j \neq i$ such that there is no polynomial q_i^t of degree 1 that interpolates them all or does not choose the polynomials q_i^t at random.
5. i does not choose some z_{ij}^m appropriately (as specified by $\vec{\sigma}^{cons}$).
6. $m < f + 1$ and i decides on a value in $\{0, 1\}$ in round m or $m = f + 1$ and i decides on an incorrect value on the equilibrium path.
7. $m = f + 1$ and i sends a value y_{ji}^t to $j' \neq i$ different from the value y_{ji}^t that i received from j in round 1.
8. i does not send a round $m' < m$ message to some agent j that i does not know at round m to have been faulty in round m' , and sends a round m message to $j' \neq i$.
9. i lies about j 's status to $j' \neq i$; that is, i sends j' a status report $\overline{SR}_i^m$ such that $\overline{SR}_i^m[j] \neq SR_i^m[j]$.

Note that in a deviation of type 8, we did not consider the case where i deviates by not sending a message to j in round m' and then sending a message to j' if i knows that j failed in round m' . In this case, i 's deviation is undetectable, and will not affect the outcome. Clearly if i performs only such undetectable deviations, then σ_i is equivalent to σ_i^{cons} , so we do not need to worry about these deviations.

We consider these deviations one by one, and show that none of them makes i better off. More precisely, we show that if σ_i involves only deviations 1– d on the list above for appropriate choices of d , then (1) holds. But even this “brute force” argument requires some care, using a somewhat delicate induction on the number of deviations that i is better off not deviating. We sketch the argument for deviations of type 1 (which turns out to be the hardest case) here, leaving the remaining details to the full paper.

So suppose that σ_i involves only i pretending to crash and that if I_i^* is a time- m^* information set for i , $\mathcal{F}$ is a set of failure patterns that satisfies the permutation assumption relative to I_i^* , $\pi_{\vec{\sigma}^{cons}}(\mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F})) > 0$, and either there are no deviations in runs in $\mathcal{R}(I_i^*)$ or the first deviation in a run in $\mathcal{R}(I_i^*)$ occurs at or after information set I_i^* , then

$$u_i(\vec{\sigma}^{cons} \mid \mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F})) \geq u_i((\sigma_i, \vec{\sigma}_{-i}^{cons}) \mid \mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F})). \quad (2)$$

(1) clearly follows from (2) by taking I_i^* to be the initial information set and letting $\mathcal{F}$ be the set of all failure patterns compatible with I_i^* .

Given a strategy profile $\vec{\sigma}$, let $\mathcal{R}(\vec{\sigma})$ denote the possible runs of $\vec{\sigma}$. If there are no runs in $\mathcal{R}(\sigma_i, \vec{\sigma}_{-i}^{cons}) \cap \mathcal{R}(I_i^*)$ in which i pretends to fail, then conditional on $\mathcal{R}(I_i^*)$, σ_i and σ_i^{psc} agree, so (2) holds. If there are runs in $\mathcal{R}(\sigma_i, \vec{\sigma}_{-i}^{cons}) \cap \mathcal{R}(I_i^*)$ in which i pretends to fail, then we proceed by induction on the number of information sets I_i at or after I_i^* at which i first pretends to crash such that $\pi_{(\sigma_i, \vec{\sigma}_{-i}^{cons})}(\mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F}) \cap \mathcal{R}(I_i)) > 0$. Suppose that i first pretends to crash at some information set I_i that comes at or after I_i^* and $\pi_{(\sigma_i, \vec{\sigma}_{-i}^{cons})}(\mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F}) \cap \mathcal{R}(I_i)) > 0$. Thus, there are no runs in $\mathcal{R}(I_i)$ in which i pretends to crash prior to information set I_i . Let σ'_i be identical to σ_i except that i does not pretend to fail at or after I_i . By (2),

$$u_i(\vec{\sigma}^{cons} \mid \mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F})) \geq u_i((\sigma'_i, \vec{\sigma}_{-i}^{cons}) \mid \mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F})).$$

We now show that

$$\begin{aligned} &u_i((\sigma'_i, \vec{\sigma}_{-i}^{cons}) \mid \mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F})) \\ &\geq u_i((\sigma_i, \vec{\sigma}_{-i}^{cons}) \mid \mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F})). \end{aligned} \quad (3)$$

(2) follows immediately.

To prove (3), since $\mathcal{R}(I_i^*)$ is the union of all the time- m information sets for i that follow I_i^* , it suffices to prove that for all time- m information sets I'_i for i that follow I_i^* , we have

$$u_i((\sigma'_i, \vec{\sigma}_{-i}^{cons}) \mid \mathcal{R}(I'_i) \cap \mathcal{R}(\mathcal{F})) \geq u_i((\sigma_i, \vec{\sigma}_{-i}^{cons}) \mid \mathcal{R}(I'_i) \cap \mathcal{R}(\mathcal{F})) \quad (4)$$

(provided, of course, that $\pi_{\vec{\sigma}^{cons}}(\mathcal{R}(I'_i) \cap \mathcal{R}(\mathcal{F})) > 0$; in the future, we take it for granted that the relevant results apply only if we are conditioning on a set with positive measure). (3) clearly follows from (4), since the time- m information sets for i partition $\mathcal{R}(I_i^*) \cap \mathcal{R}(\mathcal{F})$.

If $I'_i \neq I_i$, then (4) holds trivially, since in that case σ'_i agrees with σ_i at I'_i and all subsequent information sets. Thus, it suffices to prove (4) in the case that $I'_i = I_i$. We can assume without loss of generality that i 's actions at and after I_i are deterministic. If i is better off by pretending to fail at I_i with some probability, then i is better off by pretending to fail at I_i with probability 1. Note that (a) whether or not there is a seemingly clean round, (b) which is the first seemingly clean round if there is one, and (c) which agents are considered nonfaulty at that round are completely determined by the failure pattern. Specifically, a particular failure pattern $F' \in \mathcal{F}$ determines the first seemingly clean round m^* . We partition the set $\mathcal{F}$ into four sets, $\mathcal{F}_1, \dots, \mathcal{F}_4$, and show that conditional on $\mathcal{R}(I_i) \cap \mathcal{R}(F_j)$, agent i does at least as well by using σ'_i as it does by using σ_i , for $j = 1, \dots, 4$.

- $\mathcal{F}_1$ consists of the failure patterns in $\mathcal{F}$ where with $(\sigma_i, \vec{\sigma}_{-i}^{cons})$ an inconsistency is detected (because $f + 1$ agents seem to fail).
- $\mathcal{F}_2$ consists of the failure patterns $F' \in \mathcal{F} - \mathcal{F}_1$ such that in all runs r' in $\mathcal{R}((\sigma'_i, \vec{\sigma}_{-i}^{cons})) \cap \mathcal{R}(I_i) \cap \mathcal{R}(F')$, the first seemingly clean round occurs at some round $m^* < m$.
- $\mathcal{F}_3$ consists of the failure patterns in $\mathcal{F} - \mathcal{F}_1$ that result in m being the first seemingly clean round with both $(\sigma_i, \vec{\sigma}_{-i}^{cons})$ and $(\sigma'_i, \vec{\sigma}_{-i}^{cons})$.

- $\mathcal{F}_4$ consists of the failure patterns in $\mathcal{F} - \mathcal{F}_1$ where the first seemingly clean round m^* with $(\sigma'_i, \bar{\sigma}_{-i}^{cons})$ comes at or after m while with $(\sigma_i, \bar{\sigma}_{-i}^{cons})$, the first seemingly clean round m^* comes strictly before m or strictly after m .

The arguments in the first three cases is relatively straightforward. To show that, conditional on $\mathcal{R}(I_i) \cap \mathcal{R}(\mathcal{F}_4)$, i 's utility is at least as large with $(\sigma'_i, \bar{\sigma}_{-i}^{cons})$ as with $(\sigma_i, \bar{\sigma}_{-i}^{cons})$ requires the reachability and uniformity assumptions. We leave details to the full paper. $\square$

3.3 A π -Sequential Equilibrium for Fair Consensus

Our π -Nash equilibrium requires an agent i to decide on $\perp$ whenever i detects a problem. While this punishes the agent that causes the problem, it also punishes i . Would a rational agent actually play such a punishment strategy? Note that the need to punish occurs only off the equilibrium path; if all agents follow $\bar{\sigma}^{cons}$, agents never decide $\perp$. But to get agents to play according to $\bar{\sigma}^{cons}$ requires the threat of playing $\perp$. There might be a concern that this is an empty threat; a rational agent might not be willing to play $\perp$ if it detects a deviation.

The solution concept of *sequential equilibrium* [15] is a refinement of Nash equilibrium that, roughly speaking, requires that agents also make best responses not only on the equilibrium path, but off the equilibrium path as well. We now define π -*sequential equilibrium*, a generalization of sequential equilibrium that allows for faulty agents (where, as before, π is a distribution on failure contexts). We then show that $\bar{\sigma}^{cons}$ is essentially a π -sequential equilibrium.

Defining π -sequential equilibrium.

Roughly speaking, a strategy profile $\bar{\sigma}$ is a sequential equilibrium if, for each agent i and information set I_i for agent i , σ_i is a best response to $\bar{\sigma}_{-i}$ conditional on reaching I_i (i.e. conditional on $\mathcal{R}(I_i)$). The problem is that the probability of $\mathcal{R}(I_i)$ is 0 if I_i is not on the equilibrium path, so we cannot condition on $\mathcal{R}(I_i)$.

Define a *belief system* μ to be a function that associates with each agent i and information set I_i for agent i a probability μ_{I_i} on histories in I_i . Say that a belief system μ is *consistent with $\bar{\sigma}$ and π* if there exists a sequence of *completely mixed* strategy profiles $\bar{\sigma}^1, \bar{\sigma}^2, \dots$ (where a strategy profile is completely mixed if it gives positive probability to every action at every information set) converging to $\bar{\sigma}$ such that

$$\mu_{I_i}(h) = \lim_{M \rightarrow \infty} \frac{\pi_{\bar{\sigma}^M}(h)}{\pi_{\bar{\sigma}^M}(I_i)}.$$

Note that μ_{I_i} , π , and $\bar{\sigma}$ together define a probability distribution over runs in $\mathcal{R}(I_i)$. Let $\mu_{I_i, \pi, \bar{\sigma}}$ denote this probability distribution.

A pair $(\bar{\sigma}, \mu)$ is a π -*sequential equilibrium* if μ is a belief system consistent with $\bar{\sigma}$ and π such that, for every agent i , information set I_i , and strategy σ'_i , $u_i((\sigma_i, \bar{\sigma}_{-i}) \mid \mathcal{R}(I_i)) \geq u_i((\sigma'_i, \bar{\sigma}_{-i}) \mid \mathcal{R}(I_i))$, where now the expected utility is taken with respect to $\mu_{I_i, \pi, \bar{\sigma}}$. (Kreps and Wilson's [15] definition of sequential equilibrium is identical, except that there is no distribution π on failure contexts.)

Extending $\bar{\sigma}^{cons}$ to a π -sequential equilibrium.

We now show that the protocol $\bar{\sigma}^{cons}$ can be extended to a π -sequential equilibrium with minimal changes. In the proof of Theorem 3, we showed that i could not gain by deviating at an information set I_i where there were no deviations of type 1–9 prior to I_i . We did not show that i does not gain from deviating at I_i if an inconsistency is detected at I_i , so that i is expected to decide $\perp$. In fact, if i believes that the inconsistency may go unnoticed by other agents due to crashes and consensus may still be reached on some value in $\{0, 1\}$, then i always gains by not deciding $\perp$. However, suppose that μ^{se} is a belief system such that at an information set I_i for i that is off the equilibrium path due to a deviation (or multiple deviations) from $\bar{\sigma}^{cons}$ by agents other than i , i believes that these agents decided $\perp$ when they deviated. (Intuitively, i believes that if the agents were crazy enough to deviate in the first place, then they were also crazy enough to decide $\perp$.) In that case, deciding $\perp$ is also a best response for i .

The belief system μ^{se} is not enough to deal with information sets I_i off the equilibrium path due to i himself having deviated. Agent i cannot believe that it played $\perp$ when it in fact did not. To get a sequential equilibrium, we modify σ_i^{cons} at information sets off the equilibrium path that are reached due only to agent i 's deviations. Define the strategy σ_i^{se} so that it agrees with σ_i^{cons} at every information set I_i where agent i has not deviated in the past. Thus, in particular, i decides $\perp$ with σ_i^{se} if i detects an inconsistency at one of these information sets. More generally, say that an information set I_i is *unsalvageable* if i knows at I_i that another agent j deviated or detected an inconsistency at a point when j had not crashed, and thus decided $\perp$. I_i is certainly unsalvageable if reaching I_i requires deviations by agents other than i (for then the agent that performed that deviation decided $\perp$). But even if i is the only agent who deviates at I_i , I_i may be unsalvageable. For example, i does not send a message to j in round m_1 , i sends a message to j in round $m_2 > m_1$, and then j sent a message to i in round $m_2 + 1$, the round- $(m_2 + 2)$ information set where i receives j 's message is also unsalvageable. If I_i is unsalvageable, i decides $\perp$. Finally, if I_i is salvageable, then at I_i agent i acts in a way that is most likely to have the other agents think that there has been no inconsistency. In general, there may be more than one failure pattern that will prevent a non-faulty agent from realizing that there is an inconsistency. For example, if $f = 1$, $n = 3$, and agent 1 did not send a message to agent 2 in round m , but did send a message to agent 3, then i can either not send a message to any agent in round $m + 1$, or it can send a message to agent 3. If it is more likely that neither 2 nor 3 failed in round m than agent 2 failed before telling agent 3 that it did not hear from 1, then it would be better for i not to send a message to 2 or 3 in round $m + 1$. If there is more than one best response, then i chooses a fixed one according to some ordering on actions. (Note that this means that, unlike $\bar{\sigma}^{cons}$, the behavior of $\bar{\sigma}^{se}$ may depend on π .)

Having defined $\bar{\sigma}^{se}$, we can now define μ^{se} formally. We assume that there are only finitely many actions that i can play at each of its information sets I_i : it can send one of K_{I_i} possible messages and/or decide one of $\perp$, 0, or 1 if it has not yet made a decision, or do nothing. Given an integer $M > 0$, let $\bar{\sigma}^M$ be the strategy profile where at each information set I_i , agent i plays $\sigma_i^{se}(I_i)$ with probability $1 - 1/M$, and di-

vides the remaining probability $1/M$ over all the actions that can be played at I_i as follows: if i has already decided before, then i sends each of the K_{I_i} possible messages with equal probability $\frac{1}{M(K_{I_i}+1)}$ and does nothing with probability $\frac{1}{M(K_{I_i}+1)}$; if i has not yet decided at I_i , then for each of the K_{I_i} messages $\mathbf{m}$ that it can send, it decides $\perp$ and sends $\mathbf{m}$ with probability $\frac{1}{M(K_{I_i}+1)} - \frac{1}{M^2(K_{I_i}+1)}$, decides $\perp$ and sends no message with probability $\frac{1}{M(K_{I_i}+1)} - \frac{1}{M^2(K_{I_i}+1)}$, and performs each of the remaining $3(K_{I_i}+1)$ possible actions with equal probability $\frac{1}{3M^2(K_{I_i}+1)}$. Clearly $\bar{\sigma}^M$ is completely mixed and the sequence $\bar{\sigma}^M$ converges to $\bar{\sigma}^{se}$. Given a round- m information set I_i and global history $h \in I_i$, let

$$\mu_{I_i}^{se}(h) = \lim_{M \rightarrow \infty} \frac{\pi_{\bar{\sigma}^M}(h)}{\pi_{\bar{\sigma}^M}(I_i)}.$$

The effect of this definition of $\mu_{I_i}^{se}$ beliefs is that if I_i is off the equilibrium path as a result of some other agent j 's deviation, then i believes that j played $\perp$. Moreover, i believes that other agents j have similar beliefs.

THEOREM 4. *If $f+1 < n$, π is a distribution that supports reachability, is uniform, and allows up to f failures, and agents care only about consensus, then $(\bar{\sigma}^{se}, \mu^{se})$ is a π -sequential equilibrium.*

4. Discussion

We have provided a strategy for consensus that is a π -Nash equilibrium and can be extended to a π -sequential equilibrium, where π is a distribution on contexts that allows up to f failures and satisfies minimal conditions, as long as $n > f+1$. Although our argument is surprisingly complicated, we have considered only the simplest possible case: synchronous systems, crash failures, and only one player deviating (i.e., no coalitions). A small variant of our strategy also gives a Nash and sequential equilibrium even if coalitions are allowed, but proving this seems significantly more complicated. We are currently writing up the details carefully. Of course, things will get even worse once we allow more general types of failures, such as omission failures and Byzantine failures. But such failure types, combined with rational agents, are certainly of interest if we want to apply consensus in, for example, financial settings of the type considered by Mazières [17]. Consensus is known to be impossible in an asynchronous setting, even with just one failure [9], but algorithms that attain consensus with high probability are well known (e.g., [5]). We may thus hope to get an ϵ - π -Nash equilibrium in the asynchronous setting if we also allow rational agents. We believe that the techniques developed in this paper will be applicable to these more difficult problems.

It is also worth examining our assumptions regarding distributions in more detail. The uniformity assumption implies that no agent is more likely to fail than any other. If all agents can be identified with identical computers, then this seems quite reasonable. But if one agent can be identified with a computer that is known to be more prone to failure, then the uniformity assumption no longer holds. Note that the uniformity assumption does allow for correlated failures, just as long as the permutation of a correlated failure is just as likely as the unpermuted version.

Now consider the assumption that π supports reachability. If we are considering Nash equilibrium (where there

is only one deviating agent), the assumption says that the probability, conditional on an information set I_i (and some assumptions about failures), that some information (about a message sent by an agent that crashes or about the fact that an agent crashed in a particular round) is quite high, where “quite high” is a function of the number of agents M that are nonfaulty according to I_i . Since the more nonfaulty agents there are, the more likely it is that an agent $l \neq i$ is reachable from j without i .

Our final comment concerns the fairness assumption. While this assumption distinguishes our work from some of the other related work (e.g., [3, 6]), since, as we observed above, a consensus protocol must essentially implement a randomized dictatorship, achieving fairness once we get consensus in the presence of rational and faulty agents is not that difficult; we must simply ensure that the rational agents cannot affect the probability of a particular agent being selected as dictator. We enforce this using appropriate randomization in our protocol. The requirement in [6] that consensus must be achieved no matter what the deviating agents do turns out to have far more impact on the technical results than the fairness requirement.

In any case, we believe that the need for dealing with both rational and faulty agents in consensus protocols is compelling. There is clearly much more to be done on this problem.

Acknowledgments: Joe Halpern's work supported in part by NSF grants IIS-0911036 and CCF-1214844, and by AFOSR grant FA9550-12-1-0040, and ARO grant W911NF-14-1-0017. Xavier Vilça has been partially supported by Fundação para a Ciência e Tecnologia (FCT) through projects with references PTDC/EEI-SCR/1741/2014 (Abyss), UID/CEC/-50021/2013, and ERC-2012-StG-307732, and through the PhD grant SFRH/BD/79822/2011.

5. References

- [1] I. Abraham, D. Dolev, R. Gonen, and J. Y. Halpern. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *Proc. 25th ACM Symposium on Principles of Distributed Computing*, pages 53–62, 2006.
- [2] I. Abraham, D. Dolev, and J. Y. Halpern. Distributed protocols for leader election: a game-theoretic perspective. In *Proc. 27th International Symposium on Distributed Computing*, pages 61–75, 2013.
- [3] Y. Afek, Y. Ginzberg, S. Landau, and M. Sulamy. Distributed computing building blocks for rational agents. In *Proc. 33rd ACM Symposium on Principles of Distributed Computing*, pages 406–415, 2014.
- [4] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J. P. Martin, and C. Porth. BAR fault tolerance for cooperative services. In *Proc. 20th ACM Symposium on Operating Systems Principles (SOSP 2005)*, pages 45–58, 2005.
- [5] J. Aspnes. Randomized protocols for distributed consensus. *Distributed Computing*, 16(2–3):165–176, 2003.
- [6] X. Bei, W. Chen, and J. Zhang. Distributed consensus resilient to both crash failures and strategic

manipulations. Available at

<http://arxiv.org/abs/1203.4324>; version 3, 2012.

- [7] E. Ben-Porath. Cheap talk in games with incomplete information. *Journal of Economic Theory*, 108(1):45–71, 2003.
- [8] D. Dolev and H. R. Strong. Polynomial algorithms for multiple processor agreement. In *Proc. 14th ACM Symposium on Theory of Computing*, pages 401–407, 1982.
- [9] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty processor. *Journal of the ACM*, 32(2):374–382, 1985.
- [10] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.
- [11] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- [12] A. Groce, J. Katz, A. Thiruvengadam, and V. Zikas. Byzantine agreement with a rational adversary. In *Proc. 39th International Colloquium on Automata, Languages, and Programming, Part II*, pages 561–572, 2012.
- [13] J. Y. Halpern and V. Teague. Rational secret sharing and multiparty computation: extended abstract. In *Proc. 36th ACM Symposium on Theory of Computing*, pages 623–632, 2004.
- [14] J. Y. Halpern and X. Vilaca. Rational consensus. available at www.cs.cornell.edu/home/halpern/papers/ratconsensus.pdf, 2016.
- [15] D. M. Kreps and R. B. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.
- [16] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, San Francisco, 1997.
- [17] D. Mazières. The Stellar consensus protocol: a federated model for internet-level consensus. Available at www.stellar.org/papers/stellar-consensus-protocol.pdf, 2015.
- [18] M. Satterthwaite. Strategy-proofness and Arrow’s conditions: existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [19] A. Shamir. How to share a secret. *Communications of the ACM*, 22:612–613, 1979.

Algorithm 1 $\sigma_i^{cons}(v_i)$: i ’s consensus protocol with initial value v_i

```

1: decided  $\leftarrow$  false
2:  $ST_i \leftarrow \{(i, v_i)\}$ 
3:  $z \leftarrow$  random in  $\{0 \dots n-1\}$ 
4: for all  $j \neq i$  do
5:    $SR_i^1[j] \leftarrow (\infty, \perp)$   $\triangleright$  All agents are initially active
6:   for all  $l \neq j, i$  do
7:      $z_{ij}^1[l] \leftarrow 0$ 
8:    $z_{ij}^1[i] \leftarrow$  random in  $\{0 \dots n-1\}$   $\triangleright$  Random number to be
      used by  $j$  in round 2
9:    $z_{ij}^1[j] \leftarrow z$   $\triangleright$  Proves that  $i$  sends round 1 message to  $j$ 
10: for all  $0 \leq t \leq f$  do
11:    $x_i[t] \leftarrow$  random in  $\{0, \dots, n-t-1\}$   $\triangleright$  A random number for
      each possible value of  $t$ 
12:    $q_i^t \leftarrow$  random polynomial of degree 1 with  $q_i^t(0) = x_i[t]$ 
13:   for all  $j \neq i$  do
14:      $y_{ij}[t] \leftarrow q_i^t(j)$ 

15: for all round  $1 \leq m \leq f+1$  such that  $\neg$ decided do
16:   Phase 1: send phase
17:   for all  $j \neq i$  do
18:     if  $m = 1$  then Send  $\langle v_i, SR_i^m, (y_{ij}^0, \dots, y_{ij}^f), z_{ij}^m \rangle$  to  $j$ 
19:     if  $2 \leq m \leq f$  then Send  $\langle SR_i^m, z_{ij}^m \rangle$  to  $j$ 
20:     if  $m = f+1$  then Send  $\langle SR_i^m, (y_{li}^0, \dots, y_{li}^f)_{l \neq j} \rangle$  to  $j$ 
21:   EndPhase

22:   Phase 2: receive phase
23:    $SR_i^{m+1} \leftarrow SR_i^m$ 
24:   for all  $j \neq i$  do
25:     if receive valid message from  $j$  then
26:       if  $m = 1$  then  $ST_i \leftarrow ST_i \cup \{(j, v_j)\}$   $\triangleright ST_i$ 
          contains all the values that  $i$  has seen
27:        $SR_i^{m+1}[j] \leftarrow (\infty, z_{ji}^m)$   $\triangleright$  Note that  $j$  is still active
28:       for all  $l \neq i, j$  do
29:         if  $SR_j^m[l] = (m', j')$  and  $SR_i^m[l] = (m'', j'')$ 
          and  $m' < m''$  then
30:            $SR_i^{m+1}[l] \leftarrow (m', j)$   $\triangleright l$  crashed earlier
          than previously thought
31:            $z_{il}^{m+1}[j] \leftarrow \perp$ 
32:           else if  $SR_j^m[l] = (\infty, z_{ji}^m)$  then
33:              $z_{il}^{m+1}[j] = z_{ji}^m[j]$ 
34:           else if  $SR_i^{m+1}[j] = (\infty, z')$  for some  $z'$  then
35:              $SR_i^{m+1}[j] \leftarrow (m, i)$   $\triangleright i$  detects a crash of  $j$ 
36:           for all  $l \neq i$  do
37:              $z_{il}^{m+1}[j] \leftarrow \perp$ 
38:   EndPhase

39:   Phase 3: update phase
40:   if an inconsistency is detected then
41:     Decide( $\perp$ )  $\triangleright$  Punishment
42:     decided  $\leftarrow$  true
43:   else if  $m \leq f$  then
44:     for all  $j \neq i$  do
45:        $z_{ij}^{m+1}[i] \leftarrow$  random in  $\{0, \dots, n-1\}$ 
46:   else if decided = false then
47:      $NC_0 = N$ 
48:     for all  $1 \leq m' \leq f+1$  do
49:        $NC_{m'} \leftarrow \{j \in N - \{i\} \mid \forall m'' \leq m', l(SR_i^{f+2}[j] \neq$ 
           $(m'', l)) \cup \{i\}\}$   $\triangleright$  Agents that did not crash up to round  $m'$ 
50:        $m^* \leftarrow$  first round  $m'$  such that  $NC_{m'} = NC_{m'-1}$   $\triangleright$ 
          First round that seems clean
51:        $t \leftarrow n - |NC_{m^*}|$   $\triangleright$  Number of crashes prior to  $m^*$ 
52:       for all  $j \in NC_{m^*}$  do
53:          $q_j^t \leftarrow$  unique polynomial interpolating the values
           $y_{ji}^t$  received  $\triangleright$  otherwise, an inconsistency was detected
54:          $x_j[t] \leftarrow q_j^t(0)$ 
55:          $S \leftarrow \sum_{j \in NC_{m^*}} x_j[t] \bmod (n-t)$   $\triangleright$  Calculate a
          random number in  $0, \dots, n-t-1$ 
56:       Decide( $v_j$ ), where  $j$  is the  $(S+1)$ st highest id in
           $NC_{m^*}$ 
57:   EndPhase

```