IFRO Working Paper

Enabling reciprocity through blockchain design

Jens Gudmundsson
Jens Leth Hougaard

2020 / 14

See the full series IFRO Working Paper here:
www.ifro.ku.dk/english/publications/ifro_series/working_papers/

Department of Food and Resource Economics (IFRO)
University of Copenhagen
Rolighedsvej 23
DK 1958 Frederiksberg DENMARK
www.ifro.ku.dk/english/

Enabling reciprocity through blockchain design[*]

Jens Gudmundsson

Department of Food and Resource Economics, University of Copenhagen, Denmark

Jens Leth Hougaard

NYU-Shanghai, China
Department of Food and Resource Economics, University of Copenhagen, Denmark

Abstract

We introduce a *reciprocity protocol*, an innovative approach to coordinating and sharing rewards in blockchains. Inherently decentralized and easy to implement, it puts emphasis on incentives rather than forcing specific sharing rules from the outset. Analyzing the non-cooperative game the protocol induces, we identify a robust, strict, and Pareto-dominant symmetric equilibrium. In it, even self-centered participants show extensive reciprocity to one another. Thus, despite a setting that is generally unfavorable to reciprocal behavior, the protocol manages to build trust between the users by taking on a role akin to a social contract.

Keywords: Blockchain, reciprocity, protocol design, Nash equilibrium
JEL: C62, C72, D02, D63, D91

1. Introduction

The prevalence of centralized mining pools in the Bitcoin community is in stark contrast to the decentralized structure of the Bitcoin network itself and is viewed by some authors as a threat to system security (see e.g. Böhme et al., 2015; Arnosti and Weinberg, 2019; Cong et al., 2020; Leshno and Strack, 2020). Yet, pooled mining is a natural means of risk sharing among risk-averse miners and we therefore ask whether pooled mining can be successfully organized by a decentralized mechanism. It is intrinsically more challenging to align incentives in decentralized pools: one needs to overcome that miners prefer others to be generous to them while they themselves prefer to be selfish when sharing the pool's joint rewards. While there is substantial experimental evidence of people exhibiting kindness based on reciprocity (see e.g. Sobel, 2005), a decentralized pool—with an

ever-changing population of anonymous miners—resembles a traditional market. These have been shown to erode moral values (Falk and Szech, 2013; Bartling et al., 2015), pushing participants towards selfish behavior (compare Kranton, 1996; Bowles, 1998; Leider et al., 2009). We propose to solve this problem by designing a *reciprocity protocol*, which works as a coordination device, encouraging miners to reciprocate the generosity of fellow miners. We analyze the repeated game induced by the protocol and our main finding identifies a Pareto-dominant equilibrium in which risk-averse miners minimize payoff variance by sharing rewards generously with others. The strategies most "obvious" to coordinate on in practice, namely the equal split, turn out to be optimal.[1]

To explain our contribution more precisely, some background information on blockchains and mining pools is needed.[2] A blockchain stores information split in "blocks" that are cryptographically "chained".[3] The technology can be used in many contexts; the most prominent example is in cryptocurrencies, for which the blocks contain transactions between the network users.[4] Because these systems are decentralized and asynchronous—users may hold conflicting versions of the blockchain, ordering transactions differently—there is a mechanism in place to reach consensus on the state of the blockchain. Specifically, users compete to solve a "cryptographic puzzle" and the winner gets the right to order the most recent transactions. The harder the puzzle is to solve, the more secure the system and thus the more valuable the currency; to increase miner participation and indirectly the puzzle difficulty, winners are also rewarded ("mine") new coins. This is represented through the block's first transaction, its *coinbase* transaction, which distributes some new coins as specified by the miner.

While our contribution is applicable far beyond Bitcoin, we use Bitcoin throughout as the main example due to its dominant position in the space of cryptocurrencies.[5] Originally intended to be mined by individual users, its huge success has lead to a sharp increase in the computational power exerted, pushing the mining difficulty out of reach for individual users. Instead, miners pool

---

[1]Compare Roughgarden's (2020, p. 20) discussion on user experience and symmetric Nash equilibria.

[2]Many excellent sources cover these topics in greater detail; we refer the interested reader to Nakamoto (2008), Ferguson et al. (2010), Katz and Lindell (2014), Damgård et al. (2020), and http://bitcoin.org.

[3]The data of a block is summarized in its "header". Applying the blockchain-specific *hash function* on the header yields the "block hash". A cryptographic hash function H maps inputs of any size to outputs of a fixed size; it is designed so that computing $y = H(x)$ is easy while reverse-engineering an input x from an output $H(x)$ is hard. The *previous* block hash is included in the *current* block header to chain blocks together. A block is valid only if its hash falls below a hard-coded threshold; this is the "cryptographic puzzle" referred to later. Thus, tampering with data registered in an earlier block changes the hash of that and all subsequent blocks and thereby likely invalidates the resulting new chain (i.e., at least one of the hashes exceeds the threshold).

[4]A transaction consists of inputs, output addresses, and output-associated amounts. An input points to an unspent output from an earlier transaction while the output addresses specify the intended recipients. In this way, no "balances" are kept on this type of blockchain: to compute the balance of an address, one iterates through the entire blockchain, adding the output amounts and subtracting the inputs linked to the address.

[5]Our study pertains mainly to blockchains based on the *proof of work* (Dwork and Naor, 1993; Jakobsson and Juels, 1999) consensus protocol that underlies Nakamoto's (2008) Bitcoin blockchain, but it may be used in other settings as well. Moreover, our focus is on mining pools for existing currencies, but one could in principle set up a new cryptocurrency for which our "sidechain" is the "main chain". Other applications are give in Section 4.

their computational resources. While this development, a form of risk sharing, was predictable for risk-averse miners, it can have severe drawbacks. Most mining pools are *centralized* in that there is a group of individuals coordinating its operations. This is clearly in conflict with the intended decentralized nature of the cryptocurrency as it puts block creation in the control of a handful of organizations;[6] the argument is not too different from that pertaining to antitrust laws and market concentration (see e.g. Hirschman, 1964; Tirole, 1988). Moreover, miners must invest a lot of trust in the mining pool. The pool operators can be anonymous and may choose to keep (parts of) the rewards for themselves.[7] Lastly, the pools charge the miners a variety of fees, which aid in creating avoidable transaction costs (compare Williamson, 1979).

To circumvent these issues, we study *decentralized* mining pools. These replace the "middle men" (the pool organizers) for instance by smart contracts (e.g. Christidis and Devetsikiotis, 2016), providing a more direct connection between the miners and the blockchain.[8] This eliminates the need for trust, reduces delays, and improves miner welfare by not having any centralized running costs to finance through fees. Specifically, we employ a "sidechain" to keep track of the "work" done by the miners for the pool's benefit. While a block is valid on the "main chain" (Bitcoin's blockchain, for instance) if its hash falls below a certain threshold—this is the "cryptographic puzzle" referred to earlier—the sidechain can operate under its own rules, in particular, with a more lenient threshold.[9] The design is sketched out in Figure 1. Our guiding principle is to let miners choose their blocks themselves—so not enforcing a particular structure—and design the protocol to encourage reciprocity among the miners. Plenty of evidence speaks in favor of using such a flexible protocol that maintains miner autonomy over fixing the rules from the outset; see for instance Bartling et al. (2014) and Falk and Kosfeld (2006).[10]

Our *reciprocity protocol* operates as follows. It assigns a value v_b to miner m's block that captures how generous m is towards their fellow pool miners: the more of the block reward that m awards to other miners, the larger v_b and the more m can expect future miners to award m in return. Hence, the protocol encourages reciprocal behavior by awarding more generous miners a higher value, which will be met with higher rewards by future miners. The approach is familiar from *nudging theory*, which pertains to easy interventions, cheap to avoid without restricting choices (Thaler and Sunstein, 2008). The protocol is not forcing, as miners can choose to ignore it, yet it plays a

---

[6]As of November, 2020, five pools contributed roughly two thirds of new blocks (`https://btc.com/stats/pool`). See also the empirical analysis by Romiti et al. (2019).

[7]For some recent controversies, see for instance the case against Bitclub (`https://www.justice.gov/usao-nj/bitclub`) and the Blockseer Mining Pool (`https://cointelegraph.com/news/slippery-slope-as-new-bitcoin-mining-pool-censors-transactions`).

[8]For examples of technical implementations, see P2Pool (`http://p2pool.in`) and SMARTPOOL (Luu et al., 2017).

[9]All parts of a block influence its hash and thus its validity on either chain. This includes the coinbase transaction. That is to say, if a miner finds a full solution in which the new coins are directed to the pool, then the miner cannot "steal the block" by redirecting the new coins to herself as this would change the block's hash, likely invalidating it. Indeed, pools typically only credit work done that addresses the new funds to the pool (see e.g. Rosenfeld, 2011).

[10]In psychology, an extensive literature is devoted to the *self-determination theory* (see e.g. Deci and Ryan, 2012).

Figure 1: Left: main chain; right: pool's sidechain. The header and transactions of sidechain blocks follow the format of the main chain, while the footer chains the sidechain blocks together. The shaded parts coincide and represent Bitcoin blocks mined by the pool. White blocks are (left) Bitcoin blocks mined outside the pool and (right) blocks valid on the sidechain but not on the main chain.

valuable role, namely as a coordination device. Without the protocol, it is unlikely that there would be cooperation; with it, miners have a simple, objective measure to help them coordinate.

The reciprocity protocol induces an infinitely repeated game played among a set of infinitely many miners. Preferences are private information and restricted only to the extent that miners are risk averse, an uncontroversial restriction as risk-averse miners are exactly those who benefit from partaking in pooled mining. Through well-justified assumptions to limit the strategy space from the outset, for instance assuming that miners do not play dominated strategies, we are able to analyze a one-shot game. In essence, a miner's strategy is a level of reciprocity shown towards the others. We focus on symmetric equilibria in which all miners adopt the same strategy, but note still that they may be supported by *asymmetric* preference profiles. Our first result, Proposition 1, shows that each miner has a unique best response when the others universally choose the same strategy. This implies that equilibria will be strict. Next, Proposition 2 establishes bounds on the best response function: typically, a miner does not benefit from being more selfish than the others, but there is also an upper bound on how much kinder one should be. Our main result, Theorem 1, shows that the game has a strict and Pareto-dominant symmetric Nash-equilibrium x^* in which every risk-averse miner minimizes her payoff variance by splitting the block reward equally between herself and the n previous miners on the sidechain, where n is a parameter of the reciprocity protocol.

While these findings are robust under all forms of risk aversion, we turn to a well-known family of preferences for choice under uncertainty—exponential utility functions—to get a deeper understanding of the structure of the best response function (these exhibit *constant absolute risk aversion*; see e.g. Markowitz, 1952; Arrow, 1965; Pratt, 1964; Savage, 1971). Proposition 3 gives a precise description of the best response function, showing in particular that it is non-decreasing in the strategy universally chosen by the others. Proposition 4 pertains to the extremes of the class of exponential utility functions. In particular, the bounds established in Proposition 2 are "tight" even when restricting to such functions. Within the class, the more risk averse the miners, the fewer

4

equilibria there are: in the limit, all are essentially equal to x^* identified above. And, to reiterate, even when there are multiple equilibria, x^* maintains a focal position as it Pareto-dominates all others.

Our work relates to several strands of literature. A large literature pertains to cooperation and trust in repeated games (see e.g. Mertens et al., 2015), for instance with findings that speak in favor of our Pareto-dominant equilibrium at x^*: Bo and Fréchette (2011) experimentally conclude that "cooperation does prevail under some treatments—namely when the probability of continuation and the payoff from cooperation are high enough" (see also Bo and Fréchette, 2018). Our continuation probability, say the probability that the Bitcoin protocol is not completely revised, is arguably close to one. Provided players are sufficiently patient, cooperation through the "contagion strategy" (see also Kandori, 1992) is easier to maintain the more of the players of a fixed pool participate (Duffy and Xie, 2016). As our reciprocity protocol is "pool-hopping proof" (see e.g. Rosenfeld, 2011), miner incentives to leave the pool are unchanged over time: once they make the choice to join the pool, they have no reason to leave. Thus, this can encourage more cooperation.

We contribute also to the literature on blockchain-based games as recently surveyed by Liu et al. (2019). The strategic concerns pertain to "forking" the blockchain (Biais et al., 2019), witholding blocks (Kiayias et al., 2016; Koutsoupias et al., 2019), investing in computational power (Ma et al., 2018), and between-pool attacks (Kim and Hahn, 2019, using evolutionary game theory). Put succinctly, the scope for strategic behavior is often found to be small for "small" miners, while "large" miners may take advantage of their strong position. Fisch et al. (2017) examine optimal reward sharing in mining pools and find evidence in favor of so called "fixed-rule pools", which include the "Pay-Per-Last-N-Shares"-rule that our reciprocity protocol induces in the equilibrium at x^*. Budish (2018) analyzes the costs involved in keeping cryptocurrencies secure from sabotage. While he identifies limitations of using blockchains based on proof-of-work, it should be noted that these, some years later, remain prevalent.

Lastly, we relate to the literature on behavioral economics. Geanakoplos et al. (1989) introduce "psychological games" in which preferences depend on players' actions and their beliefs (before and during play) to capture aspects such as surprise and gratitude (see also Gilboa and Schmeidler, 1988; Segal and Sobel, 2007, 2008; Battigalli and Dufwenberg, 2009). Rabin (1993) develops on their approach, deriving psychological games from basic "material games" and introducing "fairness equilibria" through a "kindness function" that evaluates how kind one player is to another given actions and beliefs. This naturally captures a preference for reciprocity by having players desire to be kind to others they perceive as kind. Rabin's (1993) work is extended to extensive-form games by Dufwenberg and Kirchsteiger (2004), who also explore a different kindness function (see also Dufwenberg and Kirchsteiger, 2019), and by Sebald (2010) to situations in which outcomes also depend on chance. While Bolton and Ockenfels (2000) model the source of reciprocity as a desire to maintain equity, the above approach stresses punishment of hostile intentions and rewards to kind

intentions (Fehr and Gächter, 2000). Falk and Fischbascher (2006) develop an alternative theory of reciprocity in which kindness is not evaluated only on how well off one player makes another, but also on interpersonal comparisons. There is also an extensive experimental literature (see e.g. Fehr et al., 1997; Fehr and Schmidt, 2006). Compared to this type of approach, our reciprocity protocol offers a direct way to motivate reciprocal behavior through the rules of the game itself—preferences remain rational in the conventional sense of Nash (1950a). In this way, we also enable reciprocity among computational (non-human) agents for which "psychological" preferences are less evident.[11]

The paper is structured as follows. Our approach to decentralized pooled mining and the reciprocity protocol is introduced in Section 2. In Section 3, we analyze and solve for the equilibria of the induced game. We conclude in Section 4.

2. Decentralized mining pools and reciprocity

In this section, we introduce our design proposal of decentralized mining pools. Recall, these pools are run without a central, governing authority organizing and coordinating the actions within the pool. This has many benefits: rewards go directly to the miners, bypassing and eliminating the need for any form of trusted third party; as there are minimal operating costs, miner payoffs increase as there is no need for miner fees. The pool is made operational through a "sidechain" that runs parallel to the main chain of the cryptocurrency. As on the main chain, a block is valid on the sidechain only if its hash falls below a particular threshold, but the sidechain can be made more lenient. A block that is valid on the sidechain but possibly not on the main chain is a *partial solution*; a block valid on the main chain is a *full solution*. We let p and q denote the respective probabilities of finding a block valid on the main chain and the sidechain, so $0 < p \leq q \leq 1$.[12]

We take the current state of the two chains as given. Each block b holds a possibly large list of transactions, but for our purposes only the coinbase transaction is relevant. In turn, the coinbase transaction has an ordered list of outputs, recipients and amounts, restricted only to the extent that the amounts add up to the block reward. Our analysis will pertain to the creation of a new block, "block 0"; its reward will be shared with the most recent miners of the sidechain to influence the reward distribution of future blocks. Figure 2 describes how blocks are labelled. Hence, the miner of the present block awards an amount α_1 to the most recent miner in the sidechain, α_2 to the second-most, and so on, expecting to be rewarded some amounts $\beta_1, \beta_2, \ldots$ in the upcoming blocks.

---

[11]See e.g. Zhong et al. (2002) on artificial agents evolving when playing games. Behavioral biases can be affected both by delegating decisions (compare Hamman et al., 2010), say to machines, and by precommitment (e.g. Ariely and Wertenbroch, 2002; Augenblick et al., 2015). Compare also Kahneman's (2011) dichotomy of systems of though.

[12]Compare footnote 3, the probabilities correspond to the fraction of valid hashes (threshold over maximum output). These parameters are hard-coded, exogenous to the miners, with p set at the network level and q at the pool level.

Figure 2: Past blocks, $-1, -2, \ldots$, are taken as given; the analysis pertains mainly to the present block, block 0, for which the miner has to form expectations about the future, blocks $1, 2, \ldots$. Specifically, the miner awards amounts α_i to prior miners to obtain awards β_j of future block rewards.

2.1. Risk aversion and inducing kind and reciprocal behavior

The imperative rationale for pooled mining is that it allows miners to obtain small but frequent rewards, reducing payoff variance with little effect on payoff expectation (e.g. Rosenfeld, 2011). The fact that individual mining is completely dominated by pooled mining makes it evident that miners are risk averse, seeking to reduce payoff variance. In this way, every miner in the pool benefits from everyone seeking blocks in which the new coins are shared within the pool compared to purely selfish mining, but in a typical *tragedy of the commons*, the individual incentives may push in another direction: every miner wants the other miners to be generous while the miner herself prefers to remain selfish.

In what follows, we will show that the blockchain technology enables the design of protocols that induce kind and reciprocal behavior even from those purely self-centered, and thereby can help to overcome the misalignment of incentives. For this purpose, we will evaluate how "generous" a miners is towards other members of the pool and set up the protocol to induce miners to be more generous towards more generous miners; the rationale then for being generous in your blocks is that it will be reciprocated by others in subsequent blocks. Example 1 illustrates how this can induce self-centered miners to prefer being generous to being selfish.

Example 1. Suppose every miner in the pool seeks blocks in which the new coins are shared equally among the $n \in \mathbb{N}$ latest "generous" block finders and the miner herself. For now, we leave open what exactly "generous" means, but it suffices to label a miner acting as just specified as generous, while a purely selfish miner is not generous.

First, let miner m be selfish, seeking blocks in which the entire reward (normalized to 1) is kept to m herself. Thus, with probability p she finds a full solution and obtains the block reward of 1. If she instead merely finds a partial solution, then she does not expect any future rewards from the pool. In this way, m's expected payoff of selfish mining is $\mathbb{E}^s = p$ and her payoff variance is

$$\mathbb{V}^s = p \cdot (1-p)^2 + (1-p) \cdot (0-p)^2 = p(1-p).$$

Suppose instead m employs the same strategy as the others. With probability q, she finds a

partial solution. Conditional on this, it is a full solution with probability p/q and she obtains $1/(n+1)$ out of the reward. Regardless of whether m's block is a full or a partial solution, the same holds for the n blocks that follow: each is a full solution with probability p/q and then awards $1/(n+1)$. Thus, m's expected payoff is unchanged,

$$\mathbb{E}^g = q \cdot \frac{p}{q} \cdot \frac{1}{n+1} \cdot (n+1) = p = \mathbb{E}^s,$$

but her payoff variance is now considerably smaller:[13]

$$\mathbb{V}^g = q \cdot \sum_{i=0}^{n+1} \binom{n+1}{i} \left(\frac{p}{q}\right)^i \left(1 - \frac{p}{q}\right)^{n+1-i} \left(\frac{i}{n+1} - p\right)^2 + (1-q) \cdot (0-p)^2$$

$$= p \left(\frac{p}{q} - \frac{p}{(n+1)q} + \frac{1}{n+1} - p\right).$$

A reasonable approximation is that $\mathbb{V}^s$ is of the order of $\min\{q/p, n+1\} \cdot \mathbb{V}^g$. Intuitively, m's payoff variance decreases in her probability q of finding a partial solution and in the number n of miners who share the rewards. ○

Next, we make precise what we mean by a "generous" miner in the context of our *reciprocity protocol*.

2.2. The reciprocity protocol

The protocol is designed to encourage rather than enforce reciprocal behavior in the sense that miners should find it in their interest to reciprocate kindness with kindness. The intuition behind it is as follows. At each point in time, the miners seek blocks to put on the blockchain. As part of this, they specify the coinbase transaction, describing whereto the potential block rewards go. In particular, the miner may choose to share them with others in the pool. If the block turns out to be a full solution, the payments get realized; if not, the block may still be a partial solution, which could signal the miner's good intentions. Hence, this emphasizes the trade-off between leaving a lot to yourself—in the event that your block is a full solution—and sharing generously with others, which may be reciprocated through someone else's block. The protocol assigns a value to the block based on how much the miner has signaled she will give to the others in the pool. This value increases with the total amount shared. Still, recall that the miner remains free to divide the reward in any way she likes; the role of the protocol is to serve as a simple coordination device to facilitate collaboration.

For the protocol to be operational in practice, it cannot be obviously exploitable. For this reason, we take steps to minimize some particular strategic opportunities. In practice, we would

---

[13]As $p < q$ and $n \geq 1$, $np < nq \iff (n+1)p - p + q < (n+1)q \iff \frac{p}{q} - \frac{p}{(n+1)q} + \frac{1}{n+1} < 1$.

expect miners to hold several blocks on the sidechain. When that is the case, the miner may attempt to appear generous by rewarding her "previous self". This cannot be detected as miners are able to create new "identities" freely: it is impossible to distinguish one miner, who is using two addresses, from two separate miners. For this reason, it is inadequate just to see how much the miner awards *someone* in the sidechain. We take two steps to mitigate this issue. First, the protocol does not allow "skips": to count as generous towards miner m_j, you need to fully compensate every miner m_i, $i < j$, as well. Second, the protocol does not take excessive amounts into account: even if you "overcompensate" m_j, so $\alpha_j > v_j$, only the smaller of the two, $\min\{\alpha_j, v_j\} = v_j$, is counted towards your "level of generosity". Thus, if miner m holds the block that sits 100 blocks deep in the sidechain, we do allow m to count as "generous" to herself, but only if she adequately compensates the 99 more recent blocks that were added thereafter.

Put formally, the protocol defines a value $v_b \geq 0$ of how "generous" a block b is on the basis of the block's coinbase transaction. Recall from Figure 2 that the block's coinbase transaction specifies the amounts $\alpha_1, \alpha_2, \ldots$ awarded to the most recent miners $m_1, m_2, \ldots$, respectively, such that the α's sum to the total block reward. Thus, take as given the most recent blocks of the sidechain, found by miners $m_1, m_2, \ldots$ and valued $v_1, v_2, \ldots$, and consider the construction of a new block, block 0. The more of the block's potential reward that is awarded to recent miners, the higher the value v_0 of block 0. That is, we check whether block 0 matches up with the current state of the sidechain: put simply, we iterate through the coinbase transaction of block 0 as long as the ith most recent miner m_i is awarded at least v_i in the ith output, adding $\min\{\alpha_i, v_i\}$ to the "level of generosity" x in every step. Thereafter, the protocol computes the block value as $v_0 = x/n$ for some predetermined parameter $n \in \mathbb{N}$.[14] The protocol parameter n can be chosen freely. The larger n, the smaller the value of miner m's own block—so the smaller the amount that m expects from future blocks—but also the smaller the value of *other* miner's blocks, so the "longer" m's block stays relevant and the more times m gets rewarded. In equilibrium, we will find that payoff expectation is constant while payoff variance decreases in n.

Next, Definition 1 formalizes the procedure; thereafter, Example 2 illustrates the computation in a couple of numerical examples.

Definition 1 (Reciprocity protocol with parameter $n \in \mathbb{N}$)**.** Initialize $x = 0$ and a counter $i = 1$ that will iterate through the existing blocks.

1. Label the recipient and the amount of the ith output of block 0's coinbase transaction r_i and α_i, respectively.
2. Check if the recipient is the ith recent miner, that is, if $r_i = m_i$. If so, then add $\min\{\alpha_i, v_i\}$ to x; otherwise, terminate.

---

[14]The approach can readily be generalized to $v_0 = v(x)$ for some arbitrary increasing function v. Much of the intuition pertaining to x^* as defined later would then extend to x for which $v(x) = 1 - x$.

3. Check if a sufficient amount is transferred, that is, if $\alpha_i \geq v_i$. If so, then increase i by 1 (to continue to the next block) and return to step 1; otherwise, terminate.

Once the above terminates, compute the value of block 0 through $v_0 = x/n$. ○

Example 2. We showcase the protocol through the example in Table 1. On the left, we display the four most recent blocks: for instance, Alice found the most recent one and its coinbase transaction is valued to $v_1 = 4$. To the right, we list the coinbase transaction of two new blocks 0 and $0'$ through their recipients and amounts. For instance, the second output of block 0 awards 3 to Eve.

Sidechain		New block 0			New block $0'$		
Miner	Value	Recipient	Amount	Generosity	Recipient	Amount	Generosity
Alice	4	Alice	4	4	Alice	4	4
Bob	2	Eve	3	0	Bob	3	2
Charlie	2	Charlie	1	0	Charlie	1	1
Dave	2	Dave	2	0	Dave	2	0
		Total		4			7

Table 1: Left: the four most recent blocks on the sidechain; middle and right: two new blocks 0 and $0'$. "Generosity" is the amount added in step 2 of the reciprocity protocol for each output.

The protocol computes the "generosity" of each output, that is, the increment to x. For block 0, the second output has the wrong recipient (Eve, not Bob), so we do not account for this or later outputs. For block $0'$, the second output overcompensates Bob by one unit which does not count as generosity, and the third output does not fully compensate Charlie, so we only partially account for this but no later outputs. In this way, blocks 0 and $0'$ are valued $4/n$ and $7/n$, respectively. ○

In essence, the protocol gives the miners a simple coordination device that shines light on the trade-off between keeping the potential reward to yourself versus sharing it with others (hoping to later be reciprocated for this). The miner hedges, on the one hand, the event that their own block is a full solution with, on the other, the event that a future block will be a full solution. Taking a closer look at this trade-off, we can make some observations.

First, "overcompensating" someone through $\alpha_i > v_i$ is wasteful: the excess $\alpha_i - v_i$ does not get accounted for when evaluating the block (compare $0'$ in Example 2). In the same way, setting $\alpha_i < v_i$ and $\alpha_j > 0$ for some $i < j$ is wasteful (compare 0 is Example 2). Hence, normalizing the reward to 1, if the miner intends to keep $1 - x$ to herself, then she should share x by "matching" the sidechain: set $\alpha_1 = v_1$, $\alpha_2 = v_2$, and so on, until $\alpha_j < v_j$ with $\alpha_1 + \cdots + \alpha_j = x$. At its one extreme, the miner can set $x = 0$ by not sharing the reward with anyone; at the other, $x = 1$, the miner matches the sidechain with the entire reward, maximizing the block's value. Optimally, we expect the miner to find a compromise between these two extremes, balancing the value of the block, $v(x)$, with what they leave to themselves, $1 - x$.

3. Game-theoretic analysis

In this section, we analyze the non-cooperative game induced by the reciprocity protocol. When a miner decides how generous to be, she takes into account that her expected payoff depends also on the decisions of the other pool members through the reciprocity protocol. To derive a prediction of optimal, rational behavior, we next define the game and thereafter solve for its equilibria.

3.1. Formalizing the game

We consider a set of infinitely many miners engaged in an infinitely repeated game. All stage games are identical and correspond to the creation of a single block by a randomly drawn miner. In what follows, we examine the decision of an arbitrary miner in an arbitrary stage game.

In each stage game, every miner chooses a block (strategy spaces are made precise below) and nature draws an "active" miner. The chosen block is then "evaluated": with probability q, it is a partial solution that gets added to the sidechain. Conditional on this, with probability p/q, it is a full solution and its reward is distributed as specified in the block's coinbase transaction. While q is pertinent to the payoff variance, as shown in Example 1, it essentially is a multiplicative factor in the miner's preference and thus inconsequential for the miner's optimal choice of strategy. For this reason, it is without loss to set q to 1; after all, blocks that fail the "evaluation", that do not qualify even as partial solutions, do not add anything to the analysis. Moreover, we normalize the block reward to 1 and fix the protocol parameter $n \in \mathbb{N}$ used to compute the block's value $v(x) = x/n$. Focusing on a single block in each stage game allows us to abstract from some technicalities such as network delays and attempts to "fork" the blockchain; that is to say, every miner has complete information on the state of the blockchain when choosing their block and always mines "on top of" the most recent block. Miners turn "inactive" once chosen and thus appear only once. While this enables a one-shot analysis of the game, it likely lessens the incentives to be generous towards the pool:[15] if miners could be rechosen, then they may benefit from being even more generous towards the pool (in particular, towards their "old self"). Thus, any findings of reciprocal equilibrium behavior likely underestimates the effect.

Miners choose blocks, specifically coinbase transactions, to maximize expected utility for an increasing and strictly concave utility function. Preferences are private information and may differ between miners. Given the design of the protocol and the one-shot assumption, miner m has to account for two aspects when designing her block: first, if m's block is a full solution, then m is better off the more she has left to herself in her own block; second, if another miner ℓ in a future round finds a full solution, then some of this reward may be assigned to m, the amount depending indirectly on how generous m now is. Given the anonymous nature of the miners, m does not condition her choice on the current state of the sidechain but rather optimally trades off

---

[15]The assumption is reminiscent of Bolton and Ockenfels's (2000) justification to focus on a one-shot game.

the amount kept, in the event her own block is a full solution, with the potential rewards she may be granted by others through her block's value. In this sense, miner choices are "stationary"—the block design is the same in each stage game. At some point, m's block has fallen so far down the chain that it is no longer rewarded in future blocks.[16] In this way, m only takes a fairly limited period of time into account. For this reason, we assume that miners do not discount future payments.

Given the way that "overcompensations" get truncated in the reciprocity protocol, there is no reason to award previous miners more than the value of their block, $\alpha_i > v_i$, as the miner is better off keeping the excess, $\alpha_i - v_i$, to herself. In this way, the undominated choices the miner can make are of a particular form. Specifically, these are given by a number $x \in (0, 1]$ corresponding to the following coinbase transaction (compare x in Definition 1). Let the most recent blocks in the sidechain be valued $v_1, v_2, \ldots$ and let $j \in \mathbb{N}$ be such that

$$v_1 + \cdots + v_j \leq x < v_1 + \cdots + v_j + v_{j+1}.$$

That is, x can be used to fully compensate the latest j blocks. The block design associated to x sets $\alpha_i = v_i$ for $i \leq j$ and $\alpha_{j+1} = x - (\alpha_1 + \cdots + \alpha_j)$.[17] The rest of the block reward, $1 - x$, is kept to the miner herself and the block is valued x/n. Any other way of redistributing x among the previous miners is "dominated" in the sense that it reduces the block's value without increasing the amount kept. To summarize, the restriction to undominated and stationary strategies—that choices are driven by how they affect the miner herself, not the other, for m anonymous, miners—reduces each miner's strategy space to the interval $(0, 1]$.[18]

3.2. Best responses, symmetric equilibria, and expectations of future rewards

Our focus throughout is on symmetric equilibria.[19] For this reason, consider miner m and assume that all but m adopt strategy $z \in (0, 1]$. We seek m's best responses $BR(z) \subseteq (0, 1]$ to z. In particular, if $z \in BR(z)$, then z is a symmetric equilibrium. While equilibrium existence is immediate (e.g. Nash, 1950a), existence of *symmetric* equilibria despite preference asymmetry is not obvious.

Fix the strategy x chosen by miner m and the strategy z universally chosen by the others. Thus far, focus has been on the past blocks, specifically on their values v_i and the amounts α_i that

---

[16] In the equilibrium of our analysis, this will occur after n additional blocks have been found.

[17] We may also have $x < v_1$ if the miner is particularly selfish. Then $\alpha_1 = x$.

[18] We exclude purely selfish miners, $x = 0$, as these add nothing to the pool. The sidechain can be set up to ignore such blocks. When the sidechain is initialized, the first miners should be treated as fully generous.

[19] Asymmetric strategy profiles, where some miners are more generous than others, can of course also be relevant. In equilibrium, a miner would need to form beliefs about which miners will be chosen in future rounds and on how generous they will be. This could be addressed by being more explicit about the distribution under which nature selects miners, but is left for future research. Still, we contend that the equilibrium identified in Theorem 1 likely would hold its ground also in such a setting: we conjecture that it would remain an equilibrium and that it would Pareto dominate also any asymmetric equilibria.

m awards the blocks' miners. From now on, we shift to the shares β_j that m expects to acquire from future blocks. By design, m leaves $\beta_0(x, z) = 1 - x$ to herself in her own block. Imagine next that m's block is added to the sidechain. If m is relatively selfish, $x \leq z$, then m is awarded $\beta_i(x, z) = v(x) = x/n$ in each block $i \in \{1, \ldots, n\}$ that follows her own; we say that m is *fully compensated* in block i when $\beta_i(x, z) = v(x)$. For yet later blocks, the n blocks that follow m's "cost" $n \cdot v(z) = z$ to fully compensate, leaving nothing to m (compare footnote 16). If instead $x > z$, then m may not be fully compensated in some of the n subsequent blocks. To be compensated in the jth block following her own, each block $i < j$ valued $v(z) = z/n$ needs to be fully compensated. There are $j - 1$ such blocks, leaving at most $z - (j - 1) \cdot z/n$ to award m. Put succinctly, for blocks $i \in \{1, \ldots, n\}$, we have

$$\beta_i(x, z) = \min\left\{\frac{x}{n}, \frac{n + 1 - i}{n} \cdot z\right\}.$$

Let $u: \mathbb{R}_{\geq 0} \to \mathbb{R}$ denote miner m's increasing and strictly concave utility function. The probability that all blocks in $S \subseteq \{0, \ldots, n\}$ but no blocks in its complement $\{0, \ldots, n\} \setminus S$ are full solutions is $p^{|S|}(1 - p)^{n + 1 - |S|}$. When that is the case, the payoff to m is made out of the awards $\beta_i(x, z)$ from the blocks $i \in S$. In this way, m's expected utility $\mathbb{E}u(x, z)$ of choosing x "against" z chosen by the others is as follows:

$$\mathbb{E}u(x, z) = \sum_{S \subseteq \{0, \ldots, n\}} p^{|S|}(1 - p)^{n + 1 - |S|} u\left(\sum_{i \in S} \beta_i(x, z)\right).$$

Next, Example 3 illustrates these concepts for a particular utility function.

Example 3. Let miner m's preference be represented by $u(w) = \sqrt{w}$, everywhere increasing and strictly concave. Moreover, let $n = 2$, so that there only are three relevant blocks to consider: m's own block ("0") and the two blocks that follow thereafter ("1" and "2"). Each of the $2^{n+1} = 8$ subsets S of $\{0, 1, 2\}$ represents a different situation in which the blocks in S are full solutions while those outside S are not. To simplify the algebra, let $p = 1 - p = 1/2$. In this way, each subset occurs with equal probability, $1/8$. Using $\beta_i \equiv \beta_i(x, z)$, the expected utility is as follows:

$$\mathbb{E}u(x, z) = \frac{1}{8}\left(\sqrt{\beta_0} + \sqrt{\beta_1} + \sqrt{\beta_2} + \sqrt{\beta_0 + \beta_1} + \sqrt{\beta_0 + \beta_2} + \sqrt{\beta_1 + \beta_2} + \sqrt{\beta_0 + \beta_1 + \beta_2}\right).$$

We drop the multiplicative constant $1/8$ henceforth. Table 2 shows how the awards β_i depend on x and z. As $x > 2z$ is dominated by $x = 2z$, it suffices to consider $0 < x \leq 2z$.

The best response first follows the line $x = 2z$ and then "turns" towards the line $x = z$. To see this, consider first the second row of Table 2, namely $z \leq x \leq 2z$. In this case,

$$\mathbb{E}u(x, z) = \sqrt{1 - x} + \sqrt{x/2} + \sqrt{z/2} + \sqrt{1 - x/2} + \sqrt{1 - x + z/2} + \sqrt{x/2 + z/2} + \sqrt{1 - x/2 + z/2}.$$

x	$\beta_0(x, z)$	$\beta_1(x, z)$	$\beta_2(x, z)$
$0 < x \leq z$	$1 - x$	$x/2$	$x/2$
$z \leq x \leq 2z$	$1 - x$	$x/2$	$z/2$
$2z \leq x \leq 1$	$1 - x$	z	$z/2$

Table 2: The amounts that agent m expects to be awarded in her own and the two future blocks. Recall that $n = 2$, so the amount "$2z$" corresponds to nz while "$x/2$" and "$z/2$" pertain to $v(x) = x/n$ and $v(z) = z/n$.

Differentiating with respect to x, we find that the derivative is positive when evaluated at $x = 2z$ for small z, say up to $z \approx 0.08$. That is to say, for small z, the best response is $BR(z) = 2z$. If we instead evaluate the derivative at $x = z$, we find it to be negative for slightly larger z, say beyond $z \approx 0.15$. Still, provided z is not too large, m's best response is $BR(z) = z$. Once z is large, m's best response is constant and derived from Table 2's first row. Figure 3 shows $x = BR(z)$.

Figure 3: Illustration of the best response function $BR(z)$ for $n = 2$ and $u(w) = \sqrt{w}$. For $0 < z \leq n/(n + 1) = 2/3$, the function is bounded by the lines $x = z$ and $x = nz = 2z$: for small z, it follows the line $x = 2z$, for larger z, it follows $x = z$. For $n/(n + 1) \leq z \leq 1$, the best response is fixed at $BR(z) = n/(n+1)$. There is an interval of symmetric equilibria (that is, z such that $BR(z) = z$), starting at $z \approx 0.15$ and ending at $z = n/(n + 1) = 2/3$.

Interesting to note is that $BR(0.08) = 0.16 > 0.15 = BR(0.15)$, so generosity need not be complementary in the strategic sense: the best response to more generous play may actually be to be more selfish. Most best responses pertain to the second case of Table 2; we return to this in Proposition 2. Moreover, there is a vast range of symmetric equilibria, which are Pareto ranked: equilibria with a higher z Pareto dominate those with a smaller z. Lastly, there is a Pareto-dominant symmetric equilibrium at $z = n/(n + 1) = 2/3$, a point which seemingly plays an important role; indeed, this will be confirmed in Theorem 1. ○

3.3. Main results

We now turn to our main results. First, Proposition 1 shows that each miner has a unique best response when the other miners make identical choices. In this way, the equilibria that we later find will be strict.

Proposition 1. For each $z \in (0, 1]$, there is a unique best response $BR(z)$.

Proof. The miner maximizes $\mathbb{E}u(\cdot, z)$. This is linear in u, which is strictly concave in β_i, which is concave in x. Hence, all in all, $\mathbb{E}u$ is strictly concave in x and thus has a unique maximizer. □

14

To prepare our second result, we define the strategy x^* as the solution $x \in (0, 1]$ to $\beta_0(x, z) = v(x)$, namely $x^* \equiv n/(n+1)$ familiar from Example 3. For convenience, let also $z^* \equiv n/(n+1)$. Proposition 2 identifies bounds on the best response function; Figure 4 illustrates the result.


Figure 4: Illustration of Proposition 2. Miner m's best response $x = BR(z)$ lies in the shaded area for each choice $z \leq z^*$ and depends on m's utility function. For choices $z \geq z^*$, the best response is always $x^* = z^*$.

Proposition 2. The best response function is bounded as follows:

A. For $0 < z \leq z^*$, $z \leq BR(z) \leq nz$ and $BR(z) \leq x^*$;

B. For $z^* \leq z \leq 1$, $BR(z) = x^*$.

Proof. We compute miner m's expected utility of choosing x when $z \in (0, 1]$ is universally chosen by the others. We proceed in three parts to establish the various bounds.

Part 1. Consider an arbitrary $x \leq z$. As noted in Subsection 3.2, m is fully compensated in the n blocks that follow her own, $\beta_1(x, z) = \cdots = \beta_n(x, z) = v(x) = x/n$, and excluded in all later blocks. In this way, m's expected utility depends on how many of the $n + 1$ blocks (m's and the n following) are full solutions. Conditional on there being $k \in \{0, \ldots, n\}$ full solutions, m's own block is one of these k blocks with probability $k/(n+1)$. Thus, conditional on k full solutions, m's payoff is $1 - x + (k - 1) \cdot x/n$ with probability $k/(n + 1)$; otherwise, with probability $1 - k/(n + 1)$, the payoff is $k \cdot x/n$. Therefore, the expected utility conditional on k full solutions is independent of x:

$$\mathbb{E}u(x, z; k) = \frac{k}{n+1}\left(1 - x + (k - 1) \cdot \frac{x}{n}\right) + \frac{n + 1 - k}{n + 1} \cdot \frac{kx}{n}$$
$$= \frac{k}{n(n+1)}\left(n \underbrace{- nx + (k - 1)x + (n + 1 - k)x}_{0}\right) = \frac{k}{n+1}.$$

As this applies for each k, the "unconditional" expected utility is also unaffected by x. Still, m's choice x influences her payoff distribution. In particular, the induced payoff variance is smaller the closer what she leaves to herself, $\beta_0(x, z) = 1 - x$, is to what she is awarded by others, $\beta_1(x, z) = \cdots = \beta_n(x, z) = x/n$. As m is risk averse, she prefers a smaller variance. Thus, under the condition

15

$x \leq z$, m prefers to minimize $|1 - x - x/n|$, or, equivalently, $|x^* - x|$. For $z \leq z^* = x^*$, this is decreasing in x; for $z \geq z^*$, it is minimized at x^*. Therefore, we obtain the the following bounds: for $0 < z \leq z^*$, m prefers $x = z$ to each $x' < z$, so $BR(z) \geq z$; for $z^* \leq z \leq 1$, m prefers $x = x^*$ to each $x \leq z$, $x \neq x^*$, so $BR(z) \geq x^*$.

Part 2. Next, we show that, for $z^* \leq z \leq 1$, $x = x^*$ is preferred to $x' > z$. By the arguments presented in Part 1, x^* is preferred to each $x \neq x^*$ for which m is fully compensated in each of the n blocks following m's own. When m is relatively generous, $x' > z$, the situation is even worse: m is for instance only awarded $z/n < x/n$ in the nth block. In this way, the payoff distribution is a "mean-reducing spread" of that induced by x^*: payoff expectation decreases while payoff variance increases. As m is risk averse, she prefers x^* to $x' > z$.

Part 3. Finally, consider $x = nz$ and $x' > x$. Then $\beta_0(x, z) = 1 - x > 1 - x' = \beta_0(x', z)$ and, for each $i \in \{1, \ldots, n\}$, $\beta_i(x, z) = \beta_i(x', z)$. Hence, $\mathbb{E}u(x, z) > \mathbb{E}u(x', z)$. Therefore, for each $z \in (0, 1]$, $BR(z) \leq nz$. $\qquad\square$

A consequence of Proposition 2 (actually of both its parts) is that the best response to z^* is $x^* = z^*$. This finding relies only on miner risk aversion and holds regardless the underlying utility function. Together with the reciprocity protocol itself, this is our main contribution. It confirms the intuition developed in Example 1, namely that our protocol induces reciprocal equilibrium behavior even from self-centered miners. To see that the equilibrium is Pareto dominant (compare end of Example 3), Proposition 2 implies that any other symmetric equilibrium is based on $x < x^*$. At such x, the miner leaves more to herself than the value of her block, increasing payoff variance.

Theorem 1. *For risk-averse miners, $x^* = z^* = n/(n + 1)$ is the Pareto-dominant symmetric equilibrium of the game induced by the reciprocity protocol.*

In this equilibrium, the miner leaves exactly as much to herself as the value of her block. In this way, the incentives of the miner is "aligned" with that of the pool: the miner wants to maximize the number of full solutions found by the pool, irrespective of who the finders of those blocks are. In contrast, for any $x < x^*$, the miner prefers finding a full solution herself to it being found by a fellow pool member (and vice versa for $x > x^*$). Moreover, each block is designed to share the reward equally among the miner herself and the n latest miners on the sidechain. Hence, the simple, natural solution of an equal split turns out to be optimal. In this way, equilibrium behavior resembles the "Pay-Per-Last-N-Shares"-scheme that is popular in practice (see e.g. Rosenfeld, 2011), applied for instance in AntPool (`https://antpool.com`). However, it is here derived through a decentralized implementation, alleviating any need for miners to put trust in others and increasing miner payoffs by removing pool-related fees.

Lastly, we repeat that payoff expectation is invariant while payoff variance decreases in n. In this way, the equal-split outcome that is obtained in equilibrium improves miner welfare, in the

Pareto sense, the more miners, n, that the rewards are split between. And on the other side of the coin, if we fix n, then the outcome Pareto improves the closer we get to the equal split. For instance, prioritizing "recently active" miners, in the sense of giving higher weights to more recent blocks than older ones, would increase payoff variance. A simple example of this, previously employed by P2Pool, is to have a "miner's bonus", say setting aside some of the reward to the miner and sharing the rest equally.

3.4. Exponential utility functions

While Propositions 1, 2, and Theorem 1 are findings that hold generally for risk-averse preferences, we now turn to a particular class of utility functions. A staple in the literature on choice under uncertainty is the class of *exponential utility functions* (see e.g. Arrow, 1965; Pratt, 1964),

$$u(w) = 1 - e^{-Aw}, \tag{$\star$}$$

which entail risk-averse preferences for $A > 0$. Their wealth-independent risk aversion make these functions particularly tractable to analyze and interpret (see e.g. Howard, 1971; Eliashberg and Winkler, 1978; Haubrich, 1994; Gerchak and Kilgour, 1999; Çanakoğlu and Özekici, 2009; Canbolat and Rothblum, 2019; Delong, 2019), yet the class is flexible enough to often match up well with real-world data (see e.g. Jullien and Salanié, 2000; Botti et al., 2008). In this subsection, we explore the structure of the best response functions under exponential utility functions. First, Proposition 3 shows that miner m never is less generous in response to more generous fellow miners.[20]

Proposition 3. For risk-averse miners with exponential utility functions, the best response function is non-decreasing in the other miners' universally chosen strategy z for each $0 < z \le 1$.

Proof. Let $k \in \{1, \ldots, n\}$ and consider the case in which miner m is fully compensated in blocks $1, \ldots, k$ but not in blocks $k+1, \ldots, n$. This is the case for x and $z \in (0,1]$ such that $x \in [(n-k)z, (n-k+1)z] \equiv \mathcal{X}^k(z)$, as then $\beta_0(x,z) = 1 - x$, $\beta_1(x,z) = \cdots = \beta_k(x,z) = v(x) = x/n$, and $\beta_i(x,z) = (n+1-i)z/n$ for $i = k+1, \ldots, n$; see Figure 5 (left). We extend this to the hypothetical situation in which m is fully compensated in precisely the first k blocks *regardless* her choice x, that is, also to $x \notin \mathcal{X}^k(z)$. To do so, define $\gamma_0(x,z) = 1 - x$, $\gamma_1(x,z) = \cdots = \gamma_k(x,z) = x/n$, and $\gamma_i(x,z) = (n+1-i)z/n$ for $i = k+1, \ldots, n$. Moreover, redefine the function $\mathbb{E}u$ on the basis of γ_i rather than β_i through $f_k \colon (0,1] \times (0,1] \to \mathbb{R}$:

$$f_k(x,z) = \sum_{S \subseteq \{0,\ldots,n\}} p^{|S|}(1-p)^{n+1-|S|} u\left(\sum_{i \in S} \gamma_i(x,z)\right).$$

---

[20]Similar results (available upon request) are obtained when preferences exhibit increasing absolute risk aversion.

17

For $x \in \mathcal{X}^k(z)$, $\gamma_i(x, z) = \beta_i(x, z)$ and $f_k(x, z) = \mathbb{E}u(x, z)$; for $x \notin \mathcal{X}^k(z)$, $\gamma_i(x, z) \geq \beta_i(x, z)$ and $f_k(x, z) \geq \mathbb{E}u(x, z)$.[21] Therefore, if $f_k(\cdot, z)$ is maximized at $\hat{x} \in \mathcal{X}^k(z)$, then so is $\mathbb{E}u(\cdot, z)$, and $BR(z) = \hat{x}$. Below, Claim 1 asserts that we can separate the x's and z's in f_k:

Claim 1. There are functions $g_k \colon (0, 1] \to \mathbb{R}$ and $h_k \colon (0, 1] \to \mathbb{R}$ such that

$$f_k(x, z) = g_k(x) + h_k(z) - g_k(x)h_k(z),$$

where $h_k(z) < 1$.

Figure 5: Left: $\mathcal{X}^k(z)$ corresponds to the white cone. Middle: f_k is increasing in the direction of the arrows and maximized at $\hat{x}$ for each z; subject to $x \in \mathcal{X}^k(z)$, f_k is maximized along the solid line; in particular, for z such that $\hat{x} \in \mathcal{X}^k(z)$, $BR(z) = \hat{x}$ (horizontal segment). Right: two levels k and k'; the best response function again follows the thick line. Arrows point towards the maximizers $\hat{x}$ (black, associated to k) and $\hat{x}'$ (red, associated to k').

A consequence of Claim 1 is that if $\hat{x}$ maximizes g_k, then $\hat{x}$ also maximizes $f_k(\cdot, z)$ *for all z*. In particular, for each $\hat{z} \in (0, 1]$ such that $\hat{x} \in \mathcal{X}^k(\hat{z})$, we have $BR(\hat{z}) = \hat{x}$; see Figure 5 (middle). To complete the picture, we repeat the exercise for each $k \in \{1, \dots, n\}$. Figure 5 (right) illustrates this for two levels k and $k' > k$ for which the functions f_k and $f_{k'}$ are maximized at $\hat{x}$ and $\hat{x}' \geq \hat{x}$, respectively.[22] The "diagonal" segment at the cones' edges is part of the best response function: the arguments pertaining to k (k') imply that $\mathbb{E}u$ increases in the direction of the black (red) arrow. For the left-most and right-most areas ($k = 0$ and $k = n$), Proposition 2 asserts that the best response is $BR(z) = nz$ and $BR(z) = x^*$, respectively. Thus, in conclusion, BR is non-decreasing in z.

To prove Claim 1, we first show that equation $(\star)$ implies that $u(w+c) = u(w)+u(c)-u(w)u(c)$:

$$u(w + c) + u(w)u(c) = 1 - e^{-A(w+c)} + (1 - e^{-Aw})(1 - e^{-Ac})$$

$$= 1 - e^{-A(w+c)} + e^{-A(w+c)} - e^{-Aw} + 1 - e^{-Ac}$$

$$= 1 - e^{Aw} + 1 - e^{Ac} = u(w) + u(c).$$

---

[21] This can be seen as $\beta_i(x, z)$ is the minimum of two terms, one which equals $\gamma_i(x, z)$. The inequality may be strict: for instance, $\gamma_n(0, z) = z/n > 0 = \beta_n(0, z)$.

[22] We cannot have $\hat{x} > \hat{x}'$: there then is $z \in (0, 1]$ such that $\hat{x} \in \mathcal{X}^k(z)$, so $BR(z) = \hat{x}$, and $\hat{x}' \in \mathcal{X}^{k'}(z)$, so $BR(z) = \hat{x}'$. This would contradict Proposition 1.

Next, we partition each $S \subseteq \{0, \ldots, n\}$ into $S_x \equiv S \cap \{0, \ldots, k\}$ and $S_z \equiv S \cap \{k+1, \ldots, n\}$; then, for $i \in S_x$, $\gamma_i(x, z)$ depends only on x; for $i \in S_z$, $\gamma_i(x, z)$ depends only on z. Let $|S|$ denote the size of the set S; as $|S_x| + |S_z| = |S|$,

$$\left(p^{|S_x|}(1-p)^{k+1-|S_x|} \right) \left(p^{|S_z|}(1-p)^{n+1-k-|S_z|} \right) = p^{|S|}(1-p)^{n+1-|S|}.$$

Therefore, using $u(w + c) = u(w) + u(c) - u(w)u(c)$, we can reformulate $f_k(x, z)$ as follows:

$$\sum_{S \subseteq \{0, \ldots, n\}} p^{|S|}(1-p)^{n+1-|S|} u \left(\sum_{i \in S} \gamma_i(x, z) \right)$$

$$= \sum_{S_x \subseteq \{0, \ldots, k\}} p^{|S_x|}(1-p)^{k+1-|S_x|} \sum_{S_z \subseteq \{k+1, \ldots, n\}} p^{|S_z|}(1-p)^{n-k-|S_z|} u \left(\sum_{i \in S_x} \gamma_i(x, z) + \sum_{i \in S_z} \gamma_i(x, z) \right)$$

$$= \underbrace{\sum_{S_x \subseteq \{0, \ldots, k\}} p^{|S_x|}(1-p)^{k+1-|S_x|} u \left(\sum_{i \in S_x} \gamma_i(x, z) \right)}_{g_k(x)} \underbrace{\sum_{S_z \subseteq \{k+1, \ldots, n\}} p^{|S_z|}(1-p)^{n-k-|S_z|}}_{1}$$

$$+ \underbrace{\sum_{S_x \subseteq \{0, \ldots, k\}} p^{|S_x|}(1-p)^{k+1-|S_x|}}_{1} \underbrace{\sum_{S_z \subseteq \{k+1, \ldots, n\}} p^{|S_z|}(1-p)^{n-k-|S_z|} u \left(\sum_{i \in S_z} \gamma_i(x, z) \right)}_{h_k(z)}$$

$$- \underbrace{\sum_{S_x \subseteq \{0, \ldots, k\}} p^{|S_x|}(1-p)^{k+1-|S_x|} u \left(\sum_{i \in S_x} \gamma_i(x, z) \right)}_{g_k(x)} \underbrace{\sum_{S_z \subseteq \{k+1, \ldots, n\}} p^{|S_z|}(1-p)^{n-k-|S_z|} u \left(\sum_{i \in S_z} \gamma_i(x, z) \right)}_{h_k(z)}$$

$$= g_k(x) + h_k(z) - g_k(x)h_k(z). \qquad \square$$

Proposition 3 shows not only that the best response function is non-decreasing (in contrast to that of Example 3), but gives a precise description of its shape: it consists of horizontal segments connected by segments along the lines $x = kz$ for $k \in \mathbb{N}$. As we vary the parameter A—the larger A, the more risk averse the miner—we also vary the miner's best response function. Figure 6 (left) illustrates its effect on the best response function.

Next, Proposition 4 shows that every point in the shaded area of Figure 4 corresponds to a best response under some exponential utility function; Figure 6 (right) shows the "extreme" best responses for $A \to 0$ and $A \to \infty$. This implies that the bounds established in Proposition 2 are minimal: it is not possible to derive "tigher" bounds without restricting to a particular class of utility functions. Taken to its first extreme, $A \to \infty$, we see that every symmetric equilibria may be arbitrarily close to x^*, the equilibrium derived in Theorem 1. On the other hand, Proposition 4 also shows that there may be many symmetric equilibria—in its other extreme, $A \to 0$, every $x \in (0, x^*]$ can be supported as a symmetric equilibrium for a particular utility profile. However, all these

Figure 6: Left: best response functions for parameters $A \in \{1, 3, 9\}$ of the exponential utility function with $n = 3$. Right: the full scope of best response functions for parameters $A > 0$ and the limiting cases.

equilibria are Pareto ordered: they all yield expected payoff p but the payoff variance decreases with x. Again, this further strengthens the focality and appeal of the symmetric equilibrium at x^*: not only is it strict and robust to all risk-averse preferences—it also Pareto dominates all other symmetric equilibria.

Proposition 4. For each $0 < z < z^*$ and $z \le x \le nz$ such that $x < z^*$, there is an exponential utility function with parameter $A > 0$ for which $BR(z) = x$.

Proof. We first recall the structure of the best response function identified in Proposition 3: the space $(0, 1] \times (0, 1]$ can be partitioned in $n + 1$ cones in which the miner is fully compensated in the $k \in \{0, \ldots, n\}$ blocks following her own. In the interior of each cone, the best response x is independent of the choice of the others, z. In particular, by Claim 1, for $k = 1$ it suffices to maximize the following function:

$$g(x) = p(1 - p) \cdot (u(1 - x) + u(x/n)) + p^2 \cdot u(1 - x + x/n).$$

Label $\hat{x}(A)$ the maximizer of g for the exponential utility function with parameter $A > 0$. Suppose, for contradiction, that there exists $\bar{x} < x^*$ such that, for each $A > 0$, $\hat{x}(A) \le \bar{x}$. First, we construct an upper bound on $g(\bar{x})$,

$$g(\bar{x}) < p(1 - p) \cdot (u(1) + u(\bar{x}/n)) + p^2 \cdot u(1) = p(1 - p) \cdot u(\bar{x}/n) + p \cdot u(1).$$

and, second, a lower bound on $g(x^*)$ simplified using $1 - x^* = x^*/n$:

$$g(x^*) > p(1 - p) \cdot (u(1 - x^*) + u(x^*/n)) + p^2 \cdot u(x^*/n) = p(2 - p) \cdot u(x^*/n).$$

Divide both bounds by $p(2 - p) > 0$ and define $\pi \equiv p/(2 - p) \in (0, 1)$; then, $g(x^*) > g(\hat{x})$ if

$$u(x^*/n) \ge \frac{p \cdot u(1) + (1 - p) \cdot u(\bar{x}/n)}{2 - p} = \pi \cdot u(1) + (1 - \pi) \cdot u(\bar{x}/n).$$

20

Put in words, the miner prefers x^* to $\bar{x}$ whenever obtaining x^*/n with certainty is better than a gamble between $\bar{x}/n$ and 1. The *certainty equivalent* C of this latter lottery is as follows (see e.g. Raiffa, 1968; Kirkwood, 1997):

$$C = -\frac{1}{A} \ln\left(\pi \cdot e^{-A} + (1 - \pi) \cdot e^{-A\bar{x}/n}\right).$$

As $A \to \infty$, C approaches $\bar{x}/n < x^*/n$. Hence, there exists a large-enough parameter A for which the miner prefers x^* to $\bar{x}$. Indeed, this analysis applies to all $x \leq \bar{x}$. This is a contradiction as there then exists $A > 0$ for which $\hat{x}(A)$ exceeds $\bar{x}$. In terms of Figure 6 (right), this shows that there are exponential utility functions for which the best response function follows (arbitrarily close to) the solid line labeled "$A \to \infty$".

The second part of the proof pertains to the other extreme, namely the dashed line marked "$A \to 0$". Again, for contradiction, suppose that there exists $\bar{z} > 0$ and $\bar{x}$ such that, for each $A > 0$, $BR(\bar{z}) \geq \bar{x} > \bar{z}$. We claim that there are small-enough parameters $A > 0$ for which $\bar{z}$ is a better response to $\bar{z}$ than all $x \geq \bar{x}$, which implies the desired contradiction. As $A \to 0$, the miner is approximately risk neutral: the certainty equivalent of a gamble reduces to the gamble's expected payoff (again, see e.g. Kirkwood, 1997). In general, when playing z against z, the expected payoff is p; playing $x > z$ against z, on the other hand, reduces the expected payoff as the miner is not fully compensated in every block, for instance not in the nth: $\beta_n(x, z) = z/n < x/n$. □

To emphasize further, a desirable property of the equilibrium at x^* is its robustness: it holds universally for risk-averse preference profiles, prescribing the same "obvious" strategy for every miner; yet despite them all possibly having different preferences, it turns out to be optimal for everyone. While Proposition 2 showed that any candidate for a symmetric equilibrium must be based on $x < x^*$, Proposition 4 shows that there are preference profiles (even symmetric ones) for which such x is *not* an equilibrium. Hence, x^* is unique in this regard: no other level is an equilibrium for every preference profile.

4. Concluding remarks

We have suggested a specific design of decentralized mining pools operationalized through a sidechain that runs parallel to the main chain, for instance to the Bitcoin blockchain. A key element is to share the pool's block rewards among its members. In line with experimental evidence showing positive intrinsic value of decision rights and hidden costs of control, we employ a reciprocity protocol that incentivizes risk-averse miners to reciprocate kindness with kindness. The more generous miner m is towards past successful miners, the more generous future successful miners will be towards m herself. The need for such a protocol is driven by the fact that the interaction resembles a conventional market, with an ever-changing population of anonymous participants, leaving little reason to expect miners to exhibit "psychological preferences". This is reinforced by

21

the fact that miners may employ computational agents to act on their behalf. Analyzing miners' decisions through a repeated game, we show that the protocol enables reciprocity even from purely self-optimizing miners. In particular, the game induced by the protocol has a Pareto-dominant symmetric Nash equilibrium in which the miners show considerable generosity.

As noted in the introduction, this paper relates also to an extensive experimental literature. An interesting complement to this theoretical study would be to empirically investigate the reciprocity protocol. This can be done either by coding the software and using it in practice or, as a second best, through a lab experiment.

The winner-determination problem of blockchains relates also to contests as introduced by Tullock (1980) and recently surveyed by Corchón and Serena (2018). There, players typically exert effort, which gets mapped through a "contest success function" to determine the winner (compare e.g. Skaperdas, 1996; Rai and Sarin, 2009; Nitzan and Ueda, 2011; Vázquez-Sedano, 2017). Our reciprocity protocol could be used in this context, for instance by letting the right to publish a new block be determined through the contest success function. That is to say, players exert effort, which rewards them the opportunity to extend the sidechain; with some probability, their block is a full solution which becomes the outcome of the contest.

We end on an alternative and potentially promising application for the reciprocity protocol, namely in dispute settlement. The rich theory on bargaining, dating back to Nash (1950b), typically uses veto rights among the participants to reach consensus on a solution. In its simplest two-participant case, alternating offers may be made (see e.g. Ståhl, 1972; Rubinstein, 1982), and the risk of having an offer rejected encourages the proposer to be generous to the other participant (see also Rubinstein and Wolinsky, 1985; Binmore et al., 1992). In richer environments with many participants, proposers may be drawn randomly (compare e.g. Okada, 1996; Compte and Jehiel, 2010; Eraslan and McLennan, 2013) and unanimous agreement might, for instance, be required to settle the dispute. Here, the reciprocity protocol entails a different approach. Our mining pool members are the bargaining participants; once a "full solution" has been found, that is the outcome of the negotiations; a "partial solution" is a proposal that although not implemented is still observed by all participants, showcasing, perhaps, that a player has suggested a generous resolution. Hence, participants would be pushed towards an equitable agreement not because anything else would be vetoed, but rather because of the later-round costs associated with the reputation of being selfish. Exploring this aspect of the reciprocity protocol further is left for future research.

References

Ariely, D., Wertenbroch, K., 2002. Procrastination, Deadlines, and Performance: Self-Control by Precommitment. Psychological Science 13, 219–224.

Arnosti, N., Weinberg, S.M., 2019. Bitcoin: A Natural Monopoly. ITCS—Innovations in Theoretical Computer Science.

Arrow, K.J., 1965. Aspects of the Theory of Risk Bearing. Yrjo Jahnssonin Saatio, Helsinki.

Augenblick, N., Niederle, M., Sprenger, C., 2015. Working over Time: Dynamic Inconsistency in Real Effort Tasks. Quarterly Journal of Economics 130, 1067–1115.

Bartling, B., Fehr, E., Herz, H., 2014. The Intrinsic Value of Decision Rights. Econometrica 82, 2005–2039.

Bartling, B., Weber, R.A., Yao, L., 2015. Do Markets Erode Social Responsibility? Quarterly Journal of Economics 130, 219–266.

Battigalli, P., Dufwenberg, M., 2009. Dynamic psychological games. Journal of Economic Theory 144, 1–35.

Biais, B., Bisiere, C., Bouvard, M., Casamatta, C., 2019. The blockchain folk theorem. Review of Financial Studies 32, 1662–1715.

Binmore, K., Osborne, M.J., Rubinstein, A., 1992. Noncooperative models of bargaining, in: Aumann, R., Hart, S. (Eds.), Handbook of Game Theory with Economic Applications. volume 1. chapter 7, pp. 179–225.

Bo, P.D., Fréchette, G.R., 2011. The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence. American Economic Review 101, 411–429.

Bo, P.D., Fréchette, G.R., 2018. On the Determinants of Cooperation in Infinitely Repeated Games: A Survey. Journal of Economic Literature 56, 60–114.

Bolton, G.E., Ockenfels, A., 2000. ERC: A Theory of Equity, Reciprocity, and Competition. American Economic Review 90, 166–193.

Botti, F., Conte, A., Di Cagno, D., D'Ippoliti, C., 2008. Risk Attitude in Real Decision Problems. The B.E. Journal of Economic Analysis & Policy 8, 1–32.

Bowles, S., 1998. Endogenous Preferences: The Cultural Consequences of Markets and other Economic Institutions. Journal of Economic Literature 36, 75–111.

Budish, E., 2018. The Economic Limits of Bitcoin and the Blockchain. Working Paper 24717. NBER.

Böhme, R., Christin, N., Edelman, B., Moore, T., 2015. Bitcoin: Economics, Technology, and Governance. Journal of Economic Perspectives 29, 213–238.

Canbolat, P.G., Rothblum, U.G., 2019. Constant Risk Aversion in Stochastic Contests with Exponential Completion Times. Naval Research Logistics 66, 4–14.

Christidis, K., Devetsikiotis, M., 2016. Blockchains and Smart Contracts for the Internet of Things. IEEE Access 4, 2292–2303.

Compte, O., Jehiel, P., 2010. The Coalitional Nash Bargaining Solution. Econometrica 78, 1593–1623.

Cong, L.W., He, Z., Li, J., 2020. Decentralized Mining in Centralized Pools. The Review of Financial Studies , 1–40.

Corchón, L.C., Serena, M., 2018. Contest theory, in: Corchón, L.C., Marini, M.A. (Eds.), Handbook of Game Theory and Industrial Organization. volume 2. chapter 6, pp. 125–146.

Damgård, I., Nielsen, J.B., Orlandi, C., 2020. Distributed Systems and Security. `https://cs.au.dk/~orlandi/dsikdist`. Accessed 2020-11-04.

Deci, E.L., Ryan, R.M., 2012. Motivation, personality, and development within embedded social contexts: An overview of self-determination theory. Oxford University Press.

Delong, L., 2019. Optimal investment for insurance company with exponential utility and wealth-dependent risk aversion coefficient. Mathematical Methods of Operations Research 89, 73–113.

Duffy, J., Xie, H., 2016. Group size and cooperation among strangers. Journal of Economic Behavior & Organization 126, 55–74.

Dufwenberg, M., Kirchsteiger, G., 2004. A theory of sequential reciprocity. Games and Economic Behavior 47, 268–298.

Dufwenberg, M., Kirchsteiger, G., 2019. Modelling kindness. Journal of Economic Behavior & Organization 167, 228–234.

Dwork, C., Naor, M., 1993. Pricing via Processing, Or, Combatting Junk Mail, Advances in Cryptology. CRYPTO'92: Lecture Notes in Computer Science 740, 139–147.

Eliashberg, J., Winkler, R.L., 1978. The Role of Attitude toward Risk in Strictly Competitive Decision-Making Situations. Management Science 24, 1231–1241.

Eraslan, H., McLennan, A., 2013. Uniqueness of stationary equilibrium payoffs in coalitional bargaining. Journal of Economic Theory 148, 2195–2222.

Falk, A., Fischbascher, U., 2006. A Theory of Reciprocity. Games and Economic Behavior 54, 293–315.

Falk, A., Kosfeld, M., 2006. The Hidden Costs of Control. American Economic Review 96, 1611–1630.

Falk, A., Szech, N., 2013. Moral and Markets. Science 340, 707–711.

Fehr, E., Gächter, S., 2000. Fairness and Retaliation: The Economics of Reciprocity. Journal of Economic Perspectives 14, 159–181.

Fehr, E., Gächter, S., Kirchsteiger, G., 1997. Reciprocity as a Contract Enforcement Device: Experimental Evidence. Econometrica 65, 833–860.

Fehr, E., Schmidt, K.S., 2006. The Economics of Fairness, Reciprocity and Altruism—Experimental Evidence and New Theories, in: Kolm, S.C., Ythier, J.M. (Eds.), Handbook of the Economics of Giving, Altruism and Reciprocity. Elsevier. volume 1. chapter 8, pp. 615–691.

Ferguson, N., Schneier, B., Kohno, T., 2010. Cryptography Engineering: Design Principles and Practical Applications. John Wiley & Sons.

Fisch, B., Pass, R., Shelat, A., 2017. Socially Optimal Mining Pools, in: WINE 2017: Lecture Notes in Computer Science, Springer. pp. 205–218.

Geanakoplos, J., Pearce, D., Stacchetti, E., 1989. Psychological games and sequential rationality. Games and Economic Behavior 1, 60–79.

Gerchak, Y., Kilgour, D.M., 1999. Optimal parallel funding of research and development projects. IIE Transactions 31, 145–152.

Gilboa, I., Schmeidler, D., 1988. Information dependent games: Can common sense be common knowledge? Economics Letters 27, 215–221.

Hamman, J.R., Loewenstein, G., Weber, R.A., 2010. Self-interest through Delegation: An Additional Rationale for the Principal-Agent Relationship. American Economic Review 100, 1826–1846.

Haubrich, J.G., 1994. Risk Aversion, Performance Pay, and the Principal-Agent Problem. Journal of Political Economy 102, 258–276.

Hirschman, A.O., 1964. The Paternity of an Index. American Economic Review 54, 761–762.

Howard, R.A., 1971. Proximal Decision Analysis. Management Science 17, 507–541.

Jakobsson, M., Juels, A., 1999. Proofs of Work and Bread Pudding Protocols. Communications and Multimedia Security , 258–272.

Jullien, B., Salanié, B., 2000. Estimating Preferences under Risk: The Case of Racetrack Bettors. Journal of Political Economy 108, 503–530.

Kahneman, D., 2011. Thinking, Fast and Slow. Farrar, Straus and Giroux.

Kandori, M., 1992. Social Norms and Community Enforcement. Review of Economic Studies 59, 63–80.

Katz, J., Lindell, Y., 2014. Introduction to Modern Cryptography. 2 ed., Chapman & Hall.

Kiayias, A., Koutsoupias, E., Kyropoulou, M., Tselekounis, Y., 2016. Blockchain Mining Games, in: EC '16: Proceedings of the 2016 ACM Conference on Economics and Computation, pp. 365–382.

Kim, S., Hahn, S.G., 2019. Mining Pool Manipulation in Blockchain Network Over Evolutionary Block Withholding Attack. IEEE Access 7, 144230–144244.

Kirkwood, C.W., 1997. Notes on Attitude Toward Risk Taking and the Exponential Utility Function. `http://www.public.asu.edu/~kirkwood/DAStuff/refs/risk.pdf`. Accessed 2020-11-04.

Koutsoupias, E., Lazos, P., Ogunlana, F., Serafino, P., 2019. Blockchain Mining Games with Pay Forward, in: WWW '19: The World Wide Web Conference, Association for Computing Machinery. pp. 917–927.

Kranton, R.E., 1996. Reciprocal Exchange: A Self-Sustaining System. American Economic Review 86, 830–851.

Leider, S., Möbius, M.M., Rosenblat, T., Do, Q.A., 2009. Directed Altruism and Enforced Reciprocity in Social Networks. Quarterly Journal of Economics 124, 1815–1851.

Leshno, J.D., Strack, P., 2020. Bitcoin: An Axiomatic Approach and an Impossibility Theorem. American Economic Review: Insights 2, 269–286.

Liu, Z., Luong, N.C., Wang, W., Niyato, D., Wang, P., Liang, Y.C., Kim, D.I., 2019. A Survey on Blockchain: A Game Theoretical Perspective. IEEE Access 7, 47615–47643.

Luu, L., Velner, Y., Teutsch, J., Saxena, P., 2017. SmartPool: Practical Decentralized Pooled Mining. `http://smartpool.io/docs/smartpool.pdf`. Accessed 2020-11-04.

Ma, J., Gans, J.S., Tourky, R., 2018. Market Structure in Bitcoin Mining. Working Paper 24242. NBER.

Markowitz, H., 1952. Portfolio Selection. Journal of Finance 7, 77–91.

Mertens, J.F., Sorin, S., Zamir, S., 2015. Repeated Games. Cambridge University Press.

Nakamoto, S., 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. `https://bitcoin.org/bitcoin.pdf`. Accessed 2020-11-04.

Nash, J.F., 1950a. Equilibrium points in n-person games. Proceedings of the National Academy of Sciences 36, 48–49.

Nash, J.F., 1950b. The Bargaining Problem. Econometrica 18, 155–162.

Nitzan, S., Ueda, K., 2011. Prize sharing in collective contests. European Economic Review 55, 678–687.

Okada, A., 1996. A Noncooperative Coalitional Bargaining Game with Random Proposers. Games and Economic Behavior 16, 97–108.

Pratt, J.W., 1964. Risk Aversion in the Small and in the Large. Econometrica 32, 122–136.

Rabin, M., 1993. Incorporating Fairness into Game Theory and Economics. American Economic Review 83, 1281–1302.

Rai, B.K., Sarin, R., 2009. Generalized contest success functions. Economic Theory 40, 139–149.

Raiffa, H., 1968. Decision Analysis: Introductory Lectures on Choices Under Uncertainty. Addison-Wesley.

Romiti, M., Judmayer, A., Zamyatin, A., Haslhofer, B., 2019. A Deep Dive into Bitcoin Mining Pools: An Empirical Analysis of Mining Shares. `https://arxiv.org/abs/1905.05999`.

Rosenfeld, M., 2011. Analysis of Bitcoin Pooled Mining Reward Systems. `https://arxiv.org/abs/1112.4980`. Accessed 2020-11-04.

Roughgarden, T., 2020. Transaction Fee Mechanism Design for the Ethereum Blockchain: An Economic Analysis of EIP-1559. `https://timroughgarden.org/papers/eip1559.pdf`. Accessed 2021-02-01.

Rubinstein, A., 1982. Perfect Equilibrium in a Bargaining Model. Econometrica 50, 97–109.

Rubinstein, A., Wolinsky, A., 1985. Equilibrium in a Market with Sequential Bargaining. Econometrica 53, 1133–1150.

Savage, L.J., 1971. Elicitation of Personal Probabilities and Expectations. Journal of the American Statistical Association 66, 783–801.

Sebald, A., 2010. Attribution and reciprocity. Games and Economic Behavior 68, 339–352.

Segal, U., Sobel, J., 2007. Tit for Tat: Foundations of Preferences for Reciprocity in Strategic Settings. Journal of Economic Theory 136, 197–216.

Segal, U., Sobel, J., 2008. A characterization of intrinsic reciprocity. International Journal of Game Theory 36, 571–585.

Skaperdas, S., 1996. Contest success functions. Economic Theory 7, 283–290.

Sobel, J., 2005. Interdependent Preferences and Reciprocity. Journal of Economic Literature 43, 392–436.

Ståhl, I., 1972. Bargaining Theory. The Economic Research Institute at Stockholm School of Economics.

Thaler, R., Sunstein, C., 2008. Nudge: Improving Decisions about Health, Wealth, and Happiness. Yale University Press.

Tirole, J., 1988. The Theory of Industrial Organization. MIT Press.

Tullock, G., 1980. Efficient rent seeking. A&M University Press.

Vázquez-Sedano, A., 2017. Sharing the Effort Costs in Group Contests. The B.E. Journal of Theoretical Economics 18.

Williamson, O.E., 1979. Transaction-cost economics: The governance of contractual relations. Journal of Law and Economics 22, 233—261.

Zhong, F., Kimbrough, S.O., Wu, D.J., 2002. Cooperative Agent Systems: Artificial Agents Play the Ultimatum Game. Group Decision and Negotiation 11, 433–447.

Çanakoğlu, E., Özekici, S., 2009. Portfolio selection in stochastic markets with exponential utility functions. Annals of Operations Research 166.