Working Paper presented at the

Peer-to-Peer Financial Systems 2018 Workshop

June, 2018

Blockchain Upgrade as a Coordination Game

Cathy Barrera
Prysm Group

Stephanie Hurder
Prysm Group

P2P Financial Systems

Powered by
Exponential Science

Blockchain Upgrade as a Coordination Game

Cathy Barrera, Prysm Group
Stephanie Hurder, Prysm Group*

This Version: June 7, 2018

Abstract

The choice by a blockchain community of whether to implement a proposed upgrade constitutes a coordination game, for which governance rules can help community members to select an equilibrium. Unlike typical models of governance, in which the entire population is subject to the outcome of the governance process, hard forks are possible and may be desirable for a subset of the community. We develop a simple model of a coordination game in which coordinating on a single chain is always an equilibrium for the community, whereas initiating a hard fork can be an equilibrium depending on the composition of the community and the upgrade being proposed. We explore the performance of two common voting schemes – majority rule and quadratic voting – in assisting equilibrium selection and maximizing the total surplus of the community. We show that, given the constant threat of a hard fork, both the governance rules adopted and restrictions on the policy proposal space are important for maximizing community value.

*Corresponding author: stephanie@prysmgroup.io

1 Introduction

As more decentralized institutions leverage blockchain to improve function and transparency, the question of how to govern an open blockchain is becoming increasingly important. The rules governing the consensus protocol, whether to rollback the chain after a hack, or the token's target exchange rate with other currencies are just some examples of things that may need to be adjusted after the platform launches. Economists have studied voting and governance systems for decades, and their research can provide guidance for organizations seeking to design governance systems to make these decisions (Mueller, 2003).

Blockchain has two distinct features that make its governance different than other situations in which governance has traditionally been studied. First, if some participants in a blockchain don't like the output of a governance process, they can choose to hard fork, or to break off from the existing chain and start their own independent chain. Second, most blockchain platforms exhibit network effects: having more individuals on a chain, all other things equal, increases the value of the chain to the other users. As a result, when choosing among policies in a governance setting, an individual's preferred policy depends not only on what the individual wants, but what the other individuals on the chain want. This complexity of preferences cannot be captured by commonly advocated voting procedures, leading to problematic outcomes when those voting procedures are used.

In a blockchain community, a group that feels sufficiently strongly about the outcome of a governance process can choose to secede from the chain. Such a scenario is generally assumed away in contexts such as government or corporate governance because the equivalent choice is extremely costly in those other contexts. Inefficiencies in blockchain governance arise not because an undesirable outcome is imposed on some subset of the population, but because network effects inflict large costs on the community when that subset of the community cannot be forced to

stay. Under this constraint, choosing policy proposals that everyone can live with becomes more important than the method for selecting from some arbitrary set of proposals.

In this paper, we illustrate some of the complexities that arise when using standard governance techniques to govern a blockchain. We first build a game-theoretic model of a simple coordination game. Users of the blockchain must choose between two chains: either stay on the main chain with a status quo policy, or join a new chain with an upgraded policy. There are two types of individuals on the blockchain, and the value of each chain to that individual depends on the individual's type and the number of other individuals choosing that policy.

We characterize the pure-strategy Nash equilibria of this game. It is always an equilibrium for everyone to choose to stay on the main chain or for everyone to choose the upgraded chain. Depending on the composition of the user base and the relative benefit of the two policies to each type of user, a hard fork may be an equilibrium, Pareto dominant, or total surplus maximizing.

We then investigate the ability of two standard governance mechanisms, majority rule and quadratic voting, to choose the socially optimal outcome for the entire blockchain community. We first discuss the various criteria that can be used by the social planner to measure the welfare of the community. We then show that the type of governance process selected is critical for preventing hard forks and non-equilibrium outcomes. Governance systems that work well in other contexts (such as state or national elections) will produce complicated and potentially suboptimal outcomes in this setting.

We end with a discussion of the implications of our model and some areas for further investigation. One critical observation is that, as blockchain governance processes are developed, blockchain developers may wish to encourage frequent, incremental proposals rather than infrequent, radical proposals in order to maintain a single community.

All proofs appear in the appendix[1].

2 Model of Individual Chain Participation

We first study a model where there is no governance process. A large blockchain community is considering a proposed upgrade to its system. Each community member can choose to stay on the original chain, or to move to a new chain with an upgraded policy to conduct future transactions.

2.1 Setup and Definitions

The value, v, to an individual community member, i, of transacting on a particular chain, j, is a function of the fraction of the overall community also transacting on that chain, x. Assume this function, $v_{ij}(x)$, is continuously increasing, $v'_{ij}(x) > 0$ (network effects), twice differentiable, and that the returns to scale are not increasing, $v''_{ij}(x) \leq 0$. Assume that the value of transacting on a chain on which no other community members are transacting is 0, $v_{ij}(0) = 0$ for all i, j. Denote $\bar{v}_{ij} = v_{ij}(1)$, the maximum value individual i can receive from transacting on chain j.

Let $j = 0$ be the status quo chain, and chain $j = 1$ be the chain containing the proposed upgrade.

Definition. *If $v_{ij}(x) \geq v_{ik}(x)$ for $k \neq j$ for all x, then chain j is individual i's **preferred chain**.*

Suppose the community contains two types of members $\theta \in \{A, B\}$, where type A prefers the status quo chain, and type B prefers the upgraded chain. Each individual of type θ shares the same value functions $v_{\theta j}(x)$. Let β be the fraction of community members who are type B, or prefer the upgraded chain.

Observation. *The maximum value a type A individual can receive from transacting on the status quo chain is larger than the maximum value a type A individual can receive from transacting on*

[1] Appendix in progress in this draft.

the upgraded chain. The maximum value a type B individual can receive from transacting on the status quo chain is less than the maximum value a type B individual can receive from transacting on the upgraded chain.

$$\bar{v}_{A0} \geq \bar{v}_{A1}$$

$$\bar{v}_{B1} \geq \bar{v}_{B0}$$

Let q denote the fraction of the population that chooses to transact on the upgraded chain, $j = 1$. Then $(1 - q)$ is the fraction of the population that chooses to transact on the status quo chain, $j = 0$.

2.2 Independent Strategic Chain Choice

In order to determine what chain each individual will choose in equilibrium, the value function of the two chain options must be compared at each q, or each relevant combination of actions of all other community members. We can find the level of q for which each type is indifferent between the two chains.

Lemma 1. *For each $\theta \in \{A, B\}$, $\exists\ q_\theta$ such that $v_{\theta j}(q_\theta) = v_{\theta k}(q_\theta - 1)$, where j is type θ's preferred chain and k is the alternative chain. For each type, there is a fraction of the population transacting on their preferred chain, q_θ, for which that type is indifferent between choosing their preferred chain and choosing the alternative chain.*

1. $0 < q_B \leq .5 \leq (1 - q_A) < 1$

2. *A type θ individual will choose the status quo chain whenever $q \leq q_\theta$.*

3. *A type θ individual will choose the upgraded chain whenever $q > q_\theta$.*

As more community members choose the upgraded chain instead of the status quo chain (q increases), the value each type can receive increases for the upgraded chain and decreases for the status quo chain. As such, each type is more likely to choose the upgraded chain when q is higher. Because type A prefers the status quo chain, the level of adoption required to make individuals of this type indifferent between choosing the upgraded chain and their preferred chain is higher than that of type B individuals.

We can now describe the full set of pure strategy equilibria of this individual choice game, based on ranges of β, the fraction of the community who are type B, in relation to the indifference points of both types, q_A and q_B.

Proposition 1. *The pure strategy Nash equilibria of this game are as follows:*

a) For all $\beta \in [0, 1]$, it is an equilibrium for all community members to use the same chain. Specifically, there are at least two pure strategy equilibria, one in which all community members use chain $j = 0$ the status quo chain, and one in which all community members use chain $j = 1$, the upgraded chain.

b) For $\beta < q_B$ or $\beta > (1 - q_A)$, there are no other pure strategy equilibria.

c) For $\beta \in [q_B, (1 - q_A)]$, there is one additional pure strategy equilibrium, in which all type A community members use the status quo chain and all type B community members choose the upgraded chain.

In other words, it is always an equilibrium for everyone to stay on the original chain, or for everyone to move to the upgraded chain. For certain population compositions, it is also an equilibrium for the individuals who prefer the upgrade policy to hard fork and start their own chain.

For brevity, in the remainder of the paper, we will refer to the equilibria described in Proposition 1 a) as single-chain equilibria (upgraded and status quo when it is necessary to specify), and to the equilibrium described in Proposition 1 c) as the hard fork equilibrium.

3 The Social Planner's Objective

The equilibria discussed above arise when there is no governance process, and each individual is choosing the best chain for him or herself. A blockchain developer may want to institute governance processes to ensure that certain equilibria occur rather than others. The governance process serves as a decentralized method for the community to achieve these equilibria.

Before choosing a governance process, the social planner needs to know what he or she is trying to optimize. We'll consider two criteria for the optimality of equilibria. One criterion is total surplus maximization, or if the equilibrium maximizes the total surplus of the community. Another criterion is Pareto Optimality, or if there is no other equilibrium that can make one party better off, without making the other party worse off.[2]

The social choice (governance) literature has long debated which criteria should be used by the social planner to select equilibria. Majority or plurality rule, while not generally total surplus maximizing, have benefited from their practicality and simplicity of implementation (Bowen, 1943; Bergstrom, 1979; Rothkopf, 2007). Recent research by Posner and Weyl (2015) has argued that the total surplus maximizing policy can be achieved with the relatively uncomplicated quadratic voting mechanism. The blockchain social planner has the additional concern of whether the desired policy will be applied to all community members.

[2]If utility is transferable, Pareto optimal and total surplus maximizing equilibria coincide, since individuals can make side payments to ensure that the surplus-maximizing equilibria are also Pareto Optimal. For now, we assume non-transferable utility. We discuss the validity of this assumption later in Section 5.3.

3.1 Optimality of Single-Chain Equilibria

Both single-chain equilibria are Pareto Optimal as long as there are some users of each type in the population. Depending on the community composition, either the status quo or upgraded chain can be total surplus maximizing.

Observation. *There exists $\hat{s}$ such that for all $\beta > \hat{s}$ the total surplus of the upgrade equilibrium is greater than the total surplus of the status quo equilibrium, and for all $\beta < \hat{s}$ the total surplus of the status quo equilibrium is greater than the total surplus of the upgrade equilibrium.*

The upgraded chain delivers more total surplus than the status quo chain if the community contains sufficiently many individuals who prefer the upgrade. The reverse holds for the status quo chain. Whether the surplus-maximizing chain is the total surplus maximizing equilibrium depends on whether a hard fork is an equilibrium and, if so, what its total surplus is. We'll solve that in Section 3.3.2.

The total surplus maximizing chain is the same as the chain that is preferred by the majority when the relative benefit to the type B individuals from the upgrade is the same as the relative benefit to the type A individuals from the status quo chain.

Proposition 2. *The single-chain equilibrium with higher total surplus is equivalent to the majority's preferred chain if and only if the value difference between the preferred chain and the alternative chain is the same for both types of community member: $\hat{s} = \frac{1}{2}$ if and only if $\bar{v}_{A0} - \bar{v}_{A1} = \bar{v}_{B1} - \bar{v}_{B0}$.*

If the two types are symmetric in the difference between their preferred and not preferred chains—when all other community members adopt the same chain—then the majority's preferred chain *is* the total surplus maximizing chain. However, if there are asymmetries between the two types, then this is not the case.

Because each type's payoff is increasing in the number of people on a chain, the highest payoff for each type is realized when everyone joins their preferred chain. Therefore, both single-chain equilibria are Pareto Optimal.

Observation. *Each single-chain equilibrium is Pareto Optimal for $\beta \in (0,1)$. Since one type is receiving its maximum possible payoff, that type will be worse off in any other configuration.*

3.2 Blocking Coalitions

Even if a single chain equilibrium is Pareto Optimal or even total surplus maximizing, it may still be the case that a coalition of community members would like to block that outcome from occurring by hard forking. The following proposition describes the threshold of adoption above which type θ individuals would rather have their own chain than to participate with the other type in the chain they do not prefer.

Proposition 3. *For each $\theta \in \{A, B\}$, there exists p_θ such that $v_{\theta j}(p_\theta) = \bar{v}_{\theta k}$, where j is the preferred chain of type θ and k is the other chain.*

a) $q_\theta < p_\theta$

b) For $\beta > p_B$ each type B individual would prefer all type Bs to choose the upgraded chain than to coordinate with the type A's on the status quo chain. And for $\beta \leq p_B$ each type B individual would prefer the single-chain status quo equilibrium to participating with β type Bs on the upgraded chain.

c) For $\beta < (1 - p_A)$, each type A individual would prefer all type As to choose the status quo chain than to coordinate with the type B's on the upgraded chain. And for $\beta \geq (1 - p_A)$ each type A individual would prefer the single-chain upgrade equilibrium to participating with $(1 - \beta)$ type As on the status quo chain.

In the presence of blocking coalitions, the social planner's job becomes much more difficult. The social planner needs to take in to account that a dissatisfied group of users always has the option to leave the chain and fork. This can happen even if a single chain equilibrium is total surplus maximizing.

3.3 Optimality of Hard Forks

As we discussed earlier, blockchain situations are unique because governance decisions cannot be forced on the entire community. A dissatisfied coalition is always free to leave. An important question for the social planner is when a hard fork is optimal for the community – either Pareto Optimal or total surplus maximizing. And if so, under what conditions should the social planner attempt to deter or impede a hard fork?

We begin by solving for when a hard fork is Pareto Optimal, and then the conditions for a hard fork to be total surplus maximizing.

3.3.1 Pareto Optimality of Hard Forks

A hard fork is Pareto Optimal if and only if both of the types A and B are blocking coalitions.

Proposition 4 describes the ranges of β, the fraction of the community who are type B, for which both types will prefer the hard fork equilibrium to the single-chain equilibrium coordinated on the chain that type does not prefer. In this case, coordination on a single-chain equilibrium implies that one type will be worse off than if that type initiated a hard fork, so the hard fork itself is Pareto optimal.

However, there are ranges for which one or both types does at least as well in the single-chain equilibrium coordinated on the chain they do not prefer as they do in a hard fork equilibrium. In that case, there is a single-chain equilibrium that makes both types better off than the hard fork

equilibrium. This implies that there are always ranges of β for which a hard fork is an equilibrium but not Pareto Optimal.

Proposition 4. *There exists a range of β for which a hard fork is Pareto Optimal if and only if $p_B < (1 - p_A)$.*

a) When this condition is satisfied, a hard fork is Pareto Optimal if and only if $\beta \in (p_B, (1 - p_A))$.

b) If this range exists, it is inside the range for which a hard fork is an equilibrium, $(p_B, (1 - p_A)) \in [q_B, (1 - q_A)]$.

c) If this range exists, it need not include $\frac{1}{2}$.

If this range exists and β is inside it, then no matter which coordinated equilibrium is chosen, the 'losing' side constitutes a blocking coalition.

3.3.2 Total Surplus of Hard Forks

It is also possible for the hard fork to be total surplus maximizing for some values of β.

Lemma 2. *For each θ there exists an s_θ such that the amount gained by type θ from initiating a hard fork is equal to the amount lost by the other type when type θ leaves their preferred chain. $q_\theta < p_\theta < s_\theta$. s_θ solves:*

$$s_\theta \bar{v}_{\theta k} + (1 - s_\theta)\bar{v}_{-\theta k} = s_\theta \bar{v}_{\theta j}(s_\theta) + (1 - s_\theta)\bar{v}_{-\theta k}(1 - s_\theta)$$

Where j is type θ's preferred chain and k is type $-\theta$'s preferred chain; $-\theta$ is the other type.

Proposition 5. *There is a range of β for which a hard fork is total surplus maximizing if and only if $s_B < (1 - s_A)$. If this condition is satisfied:*

a) A hard fork is total surplus maximizing if and only if $\beta \in [s_B, (1 - s_A)]$.

b) $\hat{s} \in [s_B, (1-s_A)]$ *so that for* β *above the range the upgrade equilibrium is total surplus maximizing, and for* β *below the range the status quo is total surplus maximizing, and in the range the hard fork is total surplus maximizing.*

c) $[s_B, (1 - s_A)] \subset [p_B, (1 - p_A)]$

d) It may or may not be that $\frac{1}{2} \in [s_B, (1 - s_A)]$.

Corollary 1. *If* $p_B < (1-p_A)$*, then there are ranges of* β *for which the hard fork is Pareto optimal but not total surplus maximizing.*

Corollary 2. *There is no range for which a hard fork is Pareto optimal and/or total surplus maximizing if* $p_B + p_A \geq 1$.

The set of results above show two things. First, there are a very limited range of values for which a hard fork is total surplus maximizing. However, there are a wider range of values for which the hard fork is Pareto optimal but not total surplus maximizing.

This presents a particularly difficult situation for the social planner. If he or she tries to implement the surplus-maximizing outcome, there will be a blocking coalition that will choose to fork, and the social planner cannot prevent that sub-optimal outcome. This is where the assumption of non-transferable utility is particularly crucial, as if we had transferable utility, the surplus-maximizing outcome would also be Pareto optimal.

4 Implications for Governance Systems

The blockchain social planner—who by nature of the technology is operating under network effects and costless defection—faces a complicated set of issues.

First, as we discussed previously, governance in the context of blockchain is better thought of as a coordination mechanism. The results of a governance decision process are a suggestion that

the community is free to accept and implement or to form coalitions to block. A good governance system will facilitate productive coordination, but the designer of such a system cannot enforce how the system is used.

Second, solving for the optimal policies requires detailed knowledge of the community composition. Without this information, the indifference points and total surplus thresholds cannot be determined, and therefore it is unclear what policy the governance process should implement. While theoretical solutions to this problem, such as the Vickrey-Clarke-Groves mechanism, have been developed, such solutions are exceedingly complicated and none has ever been implemented (Vickrey, 1961; Clarke, 1971; Groves, 1973).

It is for these reasons, among others, that blockchain will probably require new methods of governance that will be developed over time. In this section, we focus on the ambiguities and strategic issues that can arise using two frequently advocated governance methods. The two methods are:

- Majority rule: The policy preferred by more than 50% of the population is implemented[3].
- Quadratic voting: The policy that maximizes total surplus is implemented[4].

4.1 Majority Rule Governance

Majority rule informs the community what fraction of community members prefer each option. A policy decision using majority rule will choose to upgrade if $\beta > .5$ and otherwise will choose the status quo.

[3]Many blockchain developers have noted the problem of implementing a system such as majority rule under an anonymous system. Our focus here is on the desirability of specific governance processes. Implementation of a governance process with desirable properties is a topic for further research.

[4]Posner and Weyl (2015) proves that the quadratic voting mechanism proposed in that paper always results in the total surplus maximizing policy. Like majority rule, quadratic voting faces a variety of implementation challenges, which should be studied further.

The potential existence of a blocking coalition complicates the practicality of majority rule governance in the context of blockchain. In addition to providing a mechanism for policy choice, a majority rule governance structure also provides information that can facilitate coordination for a blocking coalition. If a coalition's preferred policy is not chosen to be implemented by the majority, they may still learn through the voting process that their coalition is large enough that they would prefer to hard fork.

In the following proposition, we solve for the conditions under which there is no blocking coalition that would prefer to secede and start a new chain.

Proposition 6. *The majority rule outcome has no blocking coalition for any value of* β *if and only if* $p_B \geq \frac{1}{2}$ *and* $(1 - p_A) \leq \frac{1}{2}$.

If this condition is satisfied then majority rule will always implement a single-chain equilibrium, specifically the one preferred by a majority of the community.

4.2 Quadratic Voting

A common critique of majority rule voting is that, while it measures the fraction of the population that prefers each option, it does not measure the intensity of those relative preferences. For that reason, Posner and Weyl (2015) developed the concept of quadratic voting. In that paper the authors demonstrate that under quadratic voting, a voter will optimize his votes according to the value a policy change will induce for him, and the community will select the equilibrium that is total surplus maximizing. Quadratic voting is potentially useful to a blockchain community because it will capture if the minority coalition has particularly intense preferences.

However, like majority rule, quadratic voting is designed to apply to contexts where blocking coalitions and hard forks are not present. The additional information provided by this mechanism, regarding the strength of relative preferences, is not enough to prevent a blocking coalition from

forming. In order to account for these conditions, the voting mechanism would need to elicit the full value functions for each individual for each policy for each level of q, not just at $q \in \{0, 1\}$[5].

In addition, the option of conducting a hard fork is not explicitly considered by this voting scheme. As demonstrated by Proposition 5 there are some conditions under which a hard fork itself delivers more surplus than either single-chain equilibrium does. Under those conditions, quadratic voting as it is currently specified would not result in the total surplus maximizing policy if it did succeed in coordinating the full community on a single chain.

In the next proposition, we solve for the community compositions where the policy resulting from quadratic voting will not have a blocking coalition.

Proposition 7. *The quadratic voting outcome has no blocking coalition for any value of β if and only if $p_B \geq \hat{s}$ and $(1 - p_A) \leq \hat{s}$.*

If these thresholds could be guaranteed, then quadratic voting would always result in a single chain equilibrium. Further, under these condition, a hard fork is not total surplus maximizing, so the resulting single-chain equilibrium does maximizes total surplus.

4.3 Optimal Governance

Figure 1 summarizes the findings from Section 3 so that two alternative governance systems—majority rule and quadratic voting—can be analyzed and compared.

The table has four columns:

1. In the first column, neither type wants to block the other type's preferred chain, and thus neither will hard fork. The governance process only decides which policy will be implemented on the single chain.

[5]This is exactly what the Vickrey-Clarke-Groves mechanism, which has been criticized for its impracticality (see Rothkopf (2007)), does.

2. In the second column, the minority wants to block the majority's preferred chain, but not vice versa.

3. In the third column, the majority wants to block the minority's preferred chain, but not vice versa.

4. Each type wants to block the other's preferred chain.

The rows indicate which equilibrium outcome is total surplus maximizing:

1. In the first row, the majority's preferred chain is total surplus maximizing.

2. In the second row, the minority's preferred chain is total surplus maximizing.

3. In the third row, the hard fork is total surplus maximizing.

The table highlights some of the cases in which the choice of governance process is particularly important in a blockchain setting.

If $\beta \in [(1-p_A), p_B]$ (column 1), then either single-chain equilibrium can be implemented without the risk of a hard fork occurring. A majority rule governance system will implement the policy preferred by the majority, while quadratic voting will implement the policy that maximizes surplus (even if it is preferred only by a minority). All community members will choose to coordinate on that chain.

Columns 2 and 3 are cases in which the choice of governance system have material implications for the future of the community. In column 2, where $\beta \in [\frac{1}{2}, (1-p_A)]$ or $\beta \in [p_B, \frac{1}{2}]$, the minority will want to block the majority's preferred chain, but not vice versa. Under majority rule, the minority will choose to fork from the policy that the governance process chooses. Under quadratic voting, however, the minority's preferred policy will be implemented if it is total surplus maximizing (row 2) in which case the community will stay in a single chain.

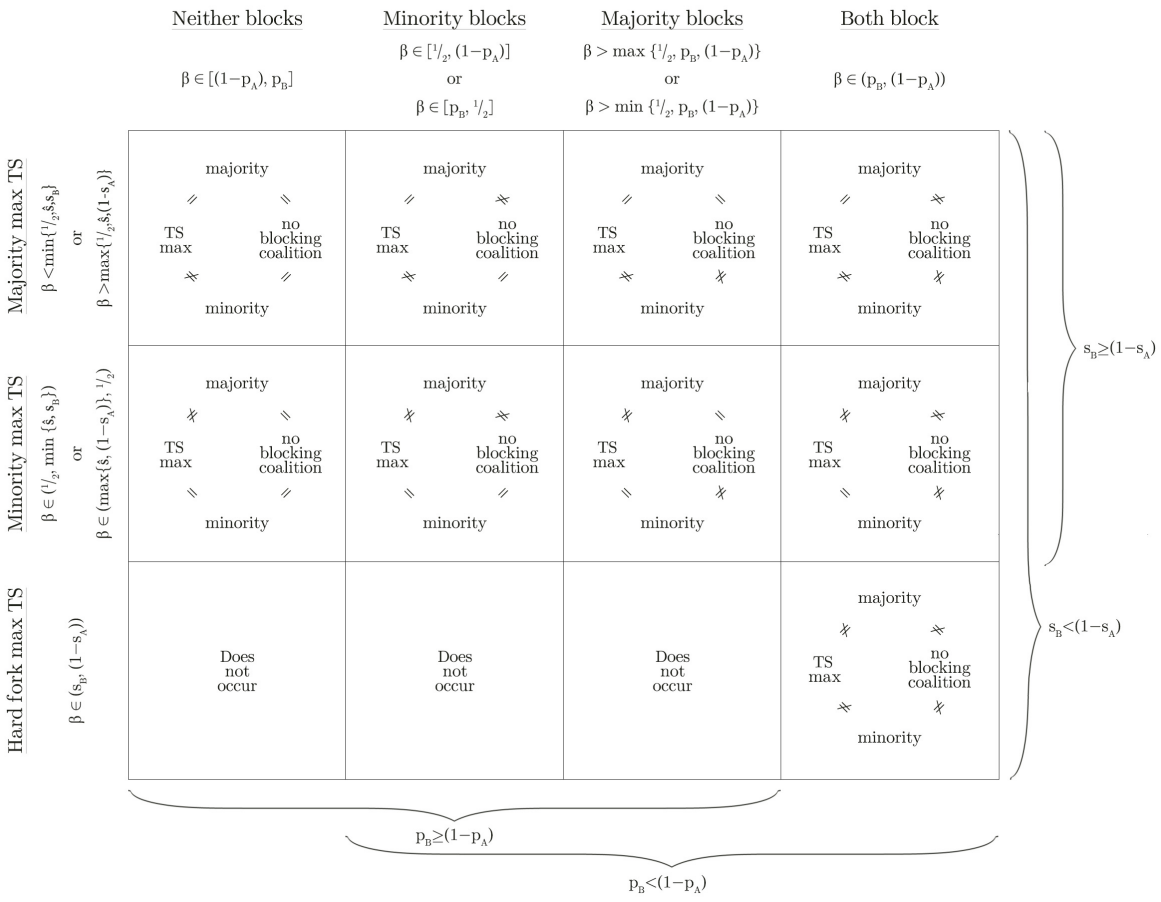

Figure 1: Single Chain Outcomes and Their Blocking Coalitions for Ranges of β

If $\beta > \max\{\frac{1}{2}, p_B, (1-p_A)\}$ or $\beta < \min\{\frac{1}{2}, p_B, (1-p_A)\}$ (column 3), the reverse occurs. The majority wants to block the minority's preferred chain, but not vice versa. This means that while majority rule will result in outcomes that are not subject to blocking, a governance system—such as quadratic voting—intended to select the total surplus maximizing outcome can result in a blocking coalition wanting to initiate a fork.

Both majority rule voting and quadratic voting suffer in column 4, where any single-chain equilibrium has a blocking coalition. The voting process used here is irrelevant, because the fork will always occur.

The above analysis assumes that, despite the knowledge that a hard fork is possible, voters would utilize the same voting strategies that they would in the absence of this possibility. Of course, because these voting systems provide limited information that can be utilized by blocking coalitions to change the outcomes of the governance process, strategic voting would have to be re-analyzed in this context. Strategic voting would further complicate the basic problems that arise due to blocking coalitions.

5 Discussion

The model in this paper explores some of the complications that can arise when designing a governance process for a blockchain. In this section, we discuss some extensions and paths for future research.

5.1 Optimal Policy Proposal Processes

The problems discussed in the previous section arise when one or both types values the two policy options very differently—loving one option and hating the other—and when the two types are asymmetric regarding their preferences.

In this context, many of these problems can be solved if the possible set of thresholds is limited. By reducing the scope of policies that can be considered, the social planner can reduce the probability that a proposed policy is polarizing enough to provoke a hard fork, or to stretch the limits of the governance process.

While these proposal processes are not often explicitly modeled, the analysis above indicates that in the context of blockchain, this aspect of the governance system is vital to facilitating community decisions. For example, under majority rule, if proposals are moderate enough to guarantee that each type would rather participate in a single chain with the other type than use its preferred chain with only half of the community, then hard forks and blocking coalitions can be avoided entirely. The solution for quadratic voting is more complex because the threshold depends on $\hat{s}$, which itself is a function of the policy proposed.

Understanding the interplay between the governance system and the policy proposal process is a fertile area for future research.

5.2 The Social Planner's Objective

The idea of 'the tyranny of the majority' cannot apply in a context where blocking coalitions are possible. As such, it is unclear that total surplus maximization should be the goal of a blockchain social planner. The primary questions for those designing blockchain governance is when should hard forks be deterred and when should they be facilitated.

In this paper, any differences in continuation values from maintaining a single-chain versus initiating a hard fork are not explicitly modeled. Further research regarding the tradeoffs between foregoing a desired policy today and enabling future outcomes with a larger community should be explored.

Further, if people are complex, and have sets of preferences that align with different groups across different dimensions, then the majority is not a monolith and the boundaries between potential coalitions are not well defined. Instead, different subsets of the community are 'winning' or getting their way on different issues over time. In this case, the majority looks less like a tyrant and more like an amorphous entity, of which any community member could be a part at any given time. In this case the goal of the social planner may be to facilitate a set of intertemporal compromises through which the governance process serves all community members.

5.3 Transferable Utility

We have assumed in this analysis that utility is not transferable across community members, so that it is impossible to compensate the 'losing' type in order to achieve the total surplus maximizing outcome.

Existing governance theories do not consider whether utility is transferable, because transferability is irrelevant when outside options are assumed away—those who lose out do not have to be compensated when they do not have a choice in the matter. For blockchain, on the other hand, since a single policy choice cannot be enforced, it becomes a question whether community members will abide by the decision that results from a governance process. While a social planner may want to implement the total surplus maximizing policy, the inability to transfer utility together with the inability to enforce a policy choice on the whole community may imply that such an implementation is infeasible.

This is an issue that needs to be explored further. We anticipate that there would be a variety of challenges to creating a system that would allow the losing side to be compensated. Free-riding and adverse selection are just two. Exactly how a system can be designed to prevent this type of problem should be the subject of future research.

6 Conclusion

In this paper, we have explored the performance of classic governance mechanisms in governing a blockchain community. Both majority rule and quadratic voting can provide effective governance, or lead to suboptimal outcomes, depending on the policies proposed and the composition of the community.

Given the complexity of governing a blockchain, we anticipate that blockchain developers will end up designing new and unique governance mechanisms. Some open questions to consider include what to use as the social planner's objective, whether and when to try to prevent hard forks, and how to design a policy proposal process that works in harmony with the chosen governance.

References

Ted C. Bergstrom. When does majority rule supply public goods efficiently? *The Scandanavian Journal of Economics*, 81(2):216–226, 1979.

Howard R. Bowen. The interpretation of voting in the allocation of economic resources. *The Quarterly Journal of Economics*, 58(1):27–48, 1943.

Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 8:19–33, 1971.

Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

Dennis C. Mueller. *Public Choice III*. Cambridge University Press, 2003.

Eric A. Posner and E. Glen Weyl. Voting squared: Quadratic voting in democratic politics. *Vanderbilt Law Review*, 68(2), 2015.

Michael H. Rothkopf. Thirteen reasons why the vickrey-clarke-groves process is not practical. *Operations Research*, 55(2):191–197, 2007.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

A Proofs

Lemma 1. *For each $\theta \in \{A, B\}$, $\exists\ q_\theta$ such that $v_{\theta j}(q_\theta) = v_{\theta k}(q_\theta - 1)$, where j is type θ's preferred chain and k is the alternative chain. For each type, there is a fraction of the population transacting on their preferred chain, q_θ, for which that type is indifferent between choosing their preferred chain and choosing the alternative chain.*

1. $0 < q_B \leq .5 \leq (1 - q_A) < 1$

2. *A type θ individual will choose the status quo chain whenever $q \leq q_\theta$.*

3. *A type θ individual will choose the upgraded chain whenever $q > q_\theta$.*

Proof. Existence holds due to the intermediate value theorem. 2) and 3) hold because $v'_{\theta j} > 0$. 1) Holds because by definition, if j is type θ's preferred chain, it must be that $v_{\theta j}(\frac{1}{2}) \geq v_{\theta k}(\frac{1}{2})$. Thus, $q_B \in (0, \frac{1}{2}]$ and $(1 - q_A) \in [\frac{1}{2}, 1)$. □

Proposition 1. *The pure strategy Nash equilibria of this game are as follows:*

a) *For all $\beta \in [0, 1]$, it is an equilibrium for all community members to use the same chain. Specifically, there are at least two pure strategy equilibria, one in which all community members use chain $j = 0$ the status quo chain, and one in which all community members use chain $j = 1$, the upgraded chain.*

b) *For $\beta < q_B$ or $\beta > (1 - q_A)$, there are no other pure strategy equilibria.*

c) *For $\beta \in [q_B, (1 - q_A)]$, there is one additional pure strategy equilibrium, in which all type A community members use the status quo chain and all type B community members choose the upgraded chain.*

Proof. If neither type is randomizing then $q \in \{0, \beta, 1\}$.

a) By Lemma 1, $q_B, (1-q_A) \in (0,1)$, thus when $q = 0$ individuals of both types choose to participate in the status quo chain, and when $q = 1$ individuals of both types choose to participate in the upgraded chain.

b) By Lemma 1 if $\beta < q_B$ then $\beta < (1 - q_A)$, therefore it cannot be an equilibrium for β type Bs to use the upgraded chain; either the type As or the type Bs would prefer to switch to the other chain in that circumstance.

c) By Lemma 1, if $\beta \in [q_B, (1 - q_A)]$, then type A individuals prefer to use the status quo chain if the other type As are also using the status quo chain—because $(1 - \beta) > q_A$—and type B individuals prefer to use the upgraded chain the other type Bs are also using the upgraded chain—because $\beta > q_B$.

□

Proposition 2. *The single-chain equilibrium with higher total surplus is equivalent to the majority's preferred chain if and only if the value difference between the preferred chain and the alternative chain is the same for both types of community member:* $\hat{s} = \frac{1}{2}$ *if and only if* $\bar{v}_{A0} - \bar{v}_{A1} = \bar{v}_{B1} - \bar{v}_{B0}$.

Proof. The total surplus of the status quo chain is $(1 - \beta)\bar{v}_{A0} + \beta\bar{v}_{B0}$. The total surplus of the upgraded chain is $(1 - \beta)\bar{v}_{A1} + \beta\bar{v}_{B1}$. $\hat{s}$ satisfies:

$$\hat{s}(\bar{v}_{B1} - \bar{v}_{B0}) = (1 - \hat{s})(\bar{v}_{A0} - \bar{v}_{A1})$$

If $\hat{s} = \frac{1}{2}$, then

$$(\bar{v}_{B1} - \bar{v}_{B0}) = (\bar{v}_{A0} - \bar{v}_{A1})$$

and if $(\bar{v}_{B1} - \bar{v}_{B0}) = (\bar{v}_{A0} - \bar{v}_{A1})$, then,

$$\hat{s} = (1 - \hat{s})$$

or $\hat{s} = \frac{1}{2}$. □

Proposition 3. *For each $\theta \in \{A, B\}$, there exists p_θ such that $v_{\theta j}(p_\theta) = \bar{v}_{\theta k}$, where j is the preferred chain of type θ and k is the other chain.*

a) $q_\theta < p_\theta$

b) For $\beta > p_B$ each type B individual would prefer all type Bs to choose the upgraded chain than to coordinate with the type A's on the status quo chain. And for $\beta \leq p_B$ each type B individual would prefer the single-chain status quo equilibrium to participating with β type Bs on the upgraded chain.

c) For $\beta < (1 - p_A)$, each type A individual would prefer all type As to choose the status quo chain than to coordinate with the type B's on the upgraded chain. And for $\beta \geq (1 - p_A)$ each type A individual would prefer the single-chain upgrade equilibrium to participating with $(1 - \beta)$ type As on the status quo chain.

Proof. Existence is due to the intermediate value theorem.

a) $\bar{v}_{\theta k} > v_{\theta k}(1 - q_\theta) = v_{\theta k}(q_\theta)$: the first inequality holds due to $v'_{\theta k} > 0$, the second by Lemma 1.

b) Holds because $v_{B1}(q)$ is increasing in β.

c) Holds because $v_{B1}(q)$ is decreasing in β (increasing in $(1 - \beta)$).

□

Proposition 4. *There exists a range of β for which a hard fork is Pareto Optimal if and only if $p_B < (1 - p_A)$.*

a) When this condition is satisfied, a hard fork is Pareto Optimal if and only if $\beta \in (p_B, (1 - p_A))$.

b) If this range exists, it is inside the range for which a hard fork is an equilibrium, $(p_B, (1 - p_A)) \in [q_B, (1 - q_A)]$.

c) If this range exists, it need not include $\frac{1}{2}$.

Lemma 2. *For each θ there exists an s_θ such that the amount gained by type θ from initiating a hard fork is equal to the amount lost by the other type when type θ leaves their preferred chain. $q_\theta < p_\theta < s_\theta$. s_θ solves:*

$$s_\theta \bar{v}_{\theta k} + (1 - s_\theta)\bar{v}_{-\theta k} = s_\theta \bar{v}_{\theta j}(s_\theta) + (1 - s_\theta)\bar{v}_{-\theta k}(1 - s_\theta)$$

Where j is type θ's preferred chain and k is type $-\theta$'s preferred chain; $-\theta$ is the other type.

Proposition 5. *There is a range of β for which a hard fork is total surplus maximizing if and only if $s_B < (1 - s_A)$. If this condition is satisfied:*

a) A hard fork is total surplus maximizing if and only if $\beta \in [s_B, (1 - s_A)]$.

b) $\hat{s} \in [s_B, (1 - s_A)]$ so that for β above the range the upgrade equilibrium is total surplus maximizing, and for β below the range the status quo is total surplus maximizing, and in the range the hard fork is total surplus maximizing.

c) $[s_B, (1 - s_A)] \subset [p_B, (1 - p_A)]$

d) It may or may not be that $\frac{1}{2} \in [s_B, (1 - s_A)]$.

Corollary 3. *If $p_B < (1-p_A)$, then there are ranges of β for which the hard fork is Pareto optimal but not total surplus maximizing.*

Proof. By Proposition 5 c) the range in which a hard fork is total surplus maximizing is strictly contained in $[p_B, (1-p_A)]$, when it exists. Then, if $p_B < (1-p_A)$, there is an $x \in [p_B, (1-p_A)]$ such that $x \notin [s_B, (1-s_A)]$. □

Corollary 4. *There is no range for which a hard fork is Pareto optimal and/or total surplus maximizing if $p_B + p_A \geq 1$.*

Proof. By Proposition 4, if $p_B + p_A \geq 1$, then there is no range for which the hard for is Pareto Optimal. By Lemma 2, $p_\theta < s_\theta$. Therefore, if $p_B + p_A \geq 1$, then $s_B > (1-s_A)$. By Proposition 5, there is not range for which the hard for is total surplus maximizing. □

Proposition 6. *The majority rule outcome has no blocking coalition for any value of β if and only if $p_B \geq \frac{1}{2}$ and $(1-p_A) \leq \frac{1}{2}$.*

Proof. By Proposition 3, type B blocks when $\beta > p_B$ and type A blocks when $\beta < (1-p_A)$. Type B (A) is in the majority when $\beta > \frac{1}{2}$ ($\beta < \frac{1}{2}$). Therefore, if $p_B \geq \frac{1}{2}$ ($(1-p_A) \leq \frac{1}{2}$), then B (A) $\beta > p_B$ ($\beta < (1-p_A)$) implies that B (A) is in the majority. □

Proposition 7. *The quadratic voting outcome has no blocking coalition for any value of β if and only if $p_B \geq \hat{s}$ and $(1-p_A) \leq \hat{s}$.*

Proof. By Proposition 3, type B blocks when $\beta > p_B$ and type A blocks when $\beta < (1-p_A)$. If $p_B \geq \hat{s}$ and $(1-p_A) \leq \hat{s}$, then $p_B > (1-p_A)$; therefore by Proposition 5. Then, by Proposition 2, Type B's (A's) preferred chain is total surplus maximizing $\beta > \hat{s}$ ($\beta < \hat{s}$). Therefore, if $p_B \geq \frac{1}{2}$ ($(1-p_A) \leq \frac{1}{2}$), then B (A) $\beta > p_B$ ($\beta < (1-p_A)$) implies that B's (A's) preferred chain is total surplus maximizing. □