The Political Economy of Blockchain Governance*

Barton E. Lee,[†] Daniel J. Moroz,[‡] David C. Parkes.[§]

Latest version here

February 8, 2020

Abstract

We develop a theory of blockchain governance using tools from formal political theory. The software underlying blockchain projects is frequently updated (forked) to implement new policies, and these forks are the subject of our inquiry. We investigate the ways in which the decentralized governance structure and preferences of users influence which policies are implemented, considering network effects as well as user preferences for different policies. We describe several types of forks and identify the strategic conditions necessary for each. We show that network-effects can motivate less moderate policy proposals, and highlight the role of market frictions created by cryptocurrency exchanges in promoting non-contentious forks that are adopted by all users. The model explains counter-intuitive phenomena that have been observed in the governance of blockchain systems.

Keywords: Blockchain, governance, forking, political processes, network effect

JEL Classification: D72, D73, D85

*We owe thanks to Davide Cipullo, Tadge Dryja, Joshua Gans, Richard Holden, Anton Kolotilin, Neha Narula, and seminar participants at the MIT Digital Currency Initiative (DCI) and UNSW. This work is supported in part by two generous gifts to the Center for Research on Computation and Society at Harvard, both to support research on applied cryptography and society. Daniel J. Moroz was also partially supported by the Ethereum Foundation. Barton E. Lee is supported by a UNSW Scientia PhD fellowship and a Data61 scholarship.

†School of Economics, UNSW Sydney, Australia, and Data61. Email: barton.e.lee@gmail.com

‡School of Engineering and Applied Sciences, Harvard University, USA. Email: dmoroz@g.harvard.edu

§School of Engineering and Applied Sciences, Harvard University, USA. Email: parkes@eecs.harvard.edu

1 Introduction

'You will not find a solution to political problems in cryptography.'

Satoshi Nakamoto

A blockchain is a software project that enables individuals to securely maintain a sequence of data entries without a centralized authority. The Bitcoin system was the first such blockchain; in the Bitcoin system, the data entries are used to maintain a ledger that records transfers of a native currency from one account to another. The Bitcoin currency has become the largest cryptocurrency by market capitalization. Although blockchains may have far reaching applications, they have so far been used primarily as the backbone of cryptocurrency systems. Our focus is on blockchains that are associated with a native currency. Since 2009, thousands of cryptocurrencies have emerged, and they maintain a combined market capitalization of over $260 billion (USD).

Blockchain governance is the process by which the rules (that is, the software) of these systems are managed.[1] In the context of cryptocurrencies, a change may affect the speed of transactions, the security of the currency, or the size of transaction fees. Debates around proposed changes to the rules can lead to battles based on ideological preferences in addition to financial considerations, and blockchains have given rise to novel governance structures. The focus of the present paper is to understand how this governance structure influences cryptocurrencies, the development of policies, and the behavior of users of these systems.[2]

Blockchain governance is unique. Unlike a sovereign state, policy development is decentralized and different users can choose to adopt different policies. A proposal to change the status quo can be made by any individual, there may be no authority or government who can restrict the agenda, and users can reject policies they dislike. Further-

[1]Industry commentator Kai Sedgwick has referred to blockchain governance as one of the greatest challenges facing blockchain applications (Sedgwick, 2018).

[2]These questions directly address the research agenda promoted and discussed in Beck et al. (2018).

more, because rules are represented through code, individuals can make credible policy proposals at relatively low cost.

The creation of a new set of rules, and together with this a new blockchain, is referred to as *forking*. To initiate a fork, an individual may copy the original blockchain's code and implement the desired policy changes into the code. For blockchains associated with native currencies, such as those considered here, a fork is also associated with a new currency. In a typical fork, every user of the current chain is endowed with an amount of currency on the new chain that is equal to the amount of original currency they own on the original chain.[3] Once created, it is up to users to decide whether to use the forked chain. In this way, individuals can accept and reject policy changes at will and also implement any change they desire at any time.

Thus, blockchain governance is a decentralized process both in terms of how proposals for change are made and also in regard to individual users' freedom to accept or reject policy changes. Governance is also complicated by the importance of network effects. All else being equal, users prefer to transact with a currency that is also used by a large number of other users. And yet, the forking of a chain can split the user population between the original chain and the new chain. A second complication for governance comes from the power held by cryptocurrency exchanges.[4] These exchanges facilitate the trade of cryptocurrencies for other assets, such as the U.S. dollar. Thousands of cryptocurrencies exist; however, a typical exchange lists only a few hundred. If a fork causes a currency to have a user base that is too small this may cause exchanges not to list the currency, rendering it much less useful and unlikely to survive. Thus, exchanges induce a market friction that prefers to avoid any one chain with a very small user base.

Contribution: To the best of our knowledge, this is the first paper to introduce a model of blockchain governance via the tools of formal political theory. Our model captures a

[3]See Alsindi (2018), who provides a detailed description of the forking process and an in depth analysis of prominent forks (see also Alsindi (2019a,b)).

[4]The three largest cryptocurrency exchanges are Binance, Huobi, and OKEX.

number of intuitive features of the blockchain ecosystem: a user's decision to adopt a proposed fork involves consideration of their preferences for the policy content of the fork and the (expected) network effects; a policy proposal can lead to a contentious, non-contentious or a null fork; and the distribution of user preferences has a consequential effect on the range of proposals that can survive as (non-null) forks. Our model can also be used to understand some of the counterintuitive governance events that have occurred in the blockchain ecosystem. For example, a fork that proposes a more extreme policy change may attract more supporters than a proposal with a moderate policy,[5] an individual may deliberately compromise the performance of the current chain in order to promote the adoption of a new chain,[6] and individuals may benefit from proposing a fork even when they know it will fail to be adopted by any user.[7]

1.1 Our results

We present a positive theory that seeks to describe and explain observed forking phenomena, such as those described in the previous paragraph, and answer questions such as 'when and why do forks occur?' These questions are important because forks have the potential to split a network and, in some cases, even destroy a currency via its replacement.

We introduce a model that captures the decentralized process of blockchain governance, with users who care about network effects and who have heterogeneous prefer-

[5]Since at least 2010, the block-size limit on Bitcoin has been a contentious policy issue. Forks to change the 1MB block-size limit to 2MB, 4MB, and even 7MB, were attempted between 2010 and 2017. It was only in 2017, with the fork to Bitcoin Cash and a block-size limit of 8MB, that there was a successful fork (see Figure 7 in Appendix B for how the fork affected Bitcoin price and hashrate, a measure of network popularity).

[6]As of early 2020, the Ethereum cryptocurrency is the number two cryptocurrency by market capitalization. In 2015, the Ethereum Foundation, a non-profit institution created to support the development of the Ethereum, implemented a "difficulty bomb" that would, after twelve months time, make transactions increasingly slow and render the current chain useless (Tual, 2015). This process of slowing down transactions became known as the Ice Age (further discussion to be provided in Section 5.1).

[7]On the Bitcoin network, a number of forks have been proposed and failed. Some discussion suggests that some of these forks were so-called *troll-blockchains*: fork proposals that are not intended to be taken seriously (Redman, 2017).

ences over policies, and market friction as induced by cryptocurrency exchanges.

Following the classical spatial models of politics (Hotelling, 1929; Downs, 1957), we model the preferences of users via a left-right political spectrum of ideologies.[8] Each user has an ideal policy, represented by some point on the real line. A unit mass of users transact with an original currency that operates with some status quo policy. In our benchmark model, we consider a single user, the *proposer*, with the power to propose a fork. Each user chooses whether to adopt the new chain or stick with the status quo.[9] In making this decision, a user considers the proposed policy as well as the network effects. A user's policy preference is modeled through a quadratic distance between a policy and a user's ideal policy. The network effect is modeled through a component of utility that increases with the number of users who adopt the same chain. In addition, the power of cryptocurrency exchanges is modeled through the requirement of a minimum threshold of users for a chain and its associated currency to be viable.

The proposal of a new policy can induce three different types of forks: *contentious forks, non-contentious forks*, and *null forks*. A contentious fork occurs when both the original and new policy attract enough support for their chains to be viable, in which case the user base is split. A non-contentious fork occurs when the new policy attracts so much support that the original chain is not viable, and the new chain replaces the original chain and attracts the entire user base. A null fork occurs when the new policy attracts insufficient support for the new chain to be viable, and the original chain retains the entire user base.

Given the parallels between our model and spatial models of politics, one may expect that moderate policies (close to the status quo) will receive more support than more extreme policy changes— as per the Median Voter Theorem (Downs, 1957; Hotelling, 1929). Indeed, when focusing on the policy dimension of user utility, the portion of users who prefer an extreme policy change is smaller than those who prefer a more moderate

[8] This follows Kurokawa's (2018) observation.

[9] Following Barrera and Hurder (2018), our model focuses on users that actively transact on a single chain. Thus, the users in our model most closely resemble currency holders that also run full nodes.

change. However, the influence of network effects and the ability of users to make individual decisions about adoption distorts this effect. In particular, a user with an extreme policy preference may be unwilling to adopt a new chain that proposes only a moderate policy change if this change is insufficient to compensate for a smaller network. In this way, a more extreme policy change can have more support than a moderate change. This provides a possible explanation for why the Bitcoin Cash fork was successful, while earlier and more moderate attempts failed.

The best outcome for a proposer is a non-contentious fork that features the proposer's ideal policy. More typically, the proposer faces a choice: maintain the network size and settle for the status quo, or create a contentious fork that implements a policy closer to the ideal. The best option depends on the proposer's preferences and the distribution of user preferences. For some distributions, any non-null fork will be contentious. For other distributions, non-contentious forks can arise when there exists policy changes that are sufficiently popular that the users who prefer the status quo are too small in number for the original chain to remain viable. In this case, we say that the distribution of preferences expresses a *consensus for change*. Intuitively, this occurs when user preferences are skewed away from the status quo. According to our model, for a proposer with an ideal policy that is in the opposite direction to the consensus for change, the best choice is to maintain the status quo. This limit to the proposer's power is a consequence of the decentralized model of governance. For a proposer with an ideal policy that is in the direction of the consensus for change, there exists policies that the proposer prefers to the status quo and that induce non-contentious forks. If one of these policies is the ideal policy, then this will be proposed. Otherwise, there is a tradeoff between proposing a moderate policy and inducing a non-contentious fork and proposing a more extreme policy and inducing a contentious fork. A proposer with more extreme policy preferences or who places little value on network size will prefer the contentious fork.

The market friction induced by cryptocurrency exchanges helps to promote non-contentious forks by causing a discontinuity where it no longer becomes useful for users to stick with the original chain. In this way, the presence of exchanges allow a majority to coerce a minority into adopting a new policy that they would otherwise reject. In particular, a proposer may introduce a more moderate fork to exploit this market friction and induce a non-contentious fork, choosing to harm a minority that is insufficient in size to maintain the original currency.

When the proposer wishes to change the status quo, they are more likely to make a proposal if there is consensus for change toward the proposer's ideal policy. This creates an incentive for a proposer to first modify the original chain by making the status quo less desirable.[10]

We extend the benchmark model to allow for the possibility of multiple proposers. In the case of multiple proposers, we model this as a lottery, where one proposal will be *recognized* by users and adopted or not as per the benchmark model. This model more closely resembles the power structure of Bitcoin, where there is no single user with disproportionate power and reveals the *spoiler effect* (Besley and Coate, 1997; Osborne and Slivinski, 1996): where a proposer may make a proposal knowing that it will fail, this proposal being used to reduce the likelihood of other proposals being recognized and taking effect. This effect persists even if the proposer experiences a fixed cost from making proposals.[11] In contrast to a single proposer model, the multiple proposer model shows that null forks can exist in equilibrium via this spoiler effect, and even when making proposals is costly. Indeed, the Bitcoin chain has experienced a plethora of null forks.

[10]We return to this in Section 5.1, where we relate our model to Ethereum, where the Ethereum Foundation has reduced the value of the status quo by initiating what has become known as the Ice Age.

[11]Anecdotal evidence that proposals in the Bitcoin network could be used as spoilers has been captured by cryptocurrency commentator Greenspan (2017), who writes "*these incessant unsupported forks are nothing more than a distraction from progress.*"

1.2 Related work

This paper is an application of formal political theory to blockchain governance. We incorporate features from spatial models of politics (Hotelling, 1929; Downs, 1957), citizen-candidate models (Besley and Coate, 1997; Osborne and Slivinski, 1996), and the literature on legislative bargaining (Baron and Ferejohn, 1989). We deviate from these classical models in two key ways. First, our setting does not require a single election outcome, or winning policy, to be enforced upon all users. Instead, an election outcome may result in two distinct policy outcomes that are experienced by different groups of users. Second, we model network effects (Akerlof et al., 2018; Becker, 1991; Katz and Shapiro, 1985), which play a larger role in blockchain applications than in standard political settings. Introducing network effects leads to new predictions; for example, the Median Value Theorem need not hold. There is a small literature that studies network effects and electoral competition.[12,13] Lee (2011), for example, considers a two-party election in the presence of a type of network effect called bandwagoning. Similar to us, Lee finds that extreme policies may attract more support than moderate proposals. However, in equilibrium, he shows that the Median Voter Theorem prevails. Faravelli et al. (2015) consider a model of voting in elections where voters benefit from the size of their party's mandate, i.e., the proportion of votes their party receives. However, unlike the present paper, Faravelli et al. do not explicitly consider the distinct policies of the different political parties.

Whereas we present a positive theory, the current literature on blockchain governance has focused on proposing alternative governance processes (Calcaterra, 2018), or analyzed the normative properties of the current structure (Barrera and Hurder, 2018), who consider a game-theoretic model and analyze the coordination issues surrounding

[12] In a related setting of party formation, Osborne and Tourky (2008) consider network effects. Their main result shows that, in equilibrium, network effects limit the number of political parties to two.

[13] There is a more expansive literature on the role of network effects in Hotelling models of price competition and product differentiation; see, for example, Navon et al. (1995); Conrad (2006); Griva and Vettas (2011); Peters et al. (2018).

blockchain governance. Berg et al. (2018) provide a discussion on how blockchain technology allows for policy experimentation and the role of governance in this process.

Industry white papers have proposed new governance structures and discussed their implications. For example, Calcaterra (2018) analyzes different ways of structuring decentralized companies (DAOs), and Goodman (2014) proposes so-called *on-chain governance*, in which the rules of the system could be amended within the system itself.[14] Our model applies equally well to systems with on-chain governance, where the off-chain forks studied in the present paper are still possible. On-chain governance has led to significant debate into the benefits and costs of various kinds of governance structures by prominent industry figures (Zamfir, 2017; Ehrsam, 2017; Buterin, 2019). Meanwhile, industry commentators frequently question the future governance of private blockchains such as Libra (e.g., Telford (2019)). In these private blockchains, only an approved set of participants, the companies in the Libra Association, for example, can engage in validation and governance of the chain. Our results apply equally well to this case, if we consider only the preferences of these approved participants, instead of any user, as in the general case.

1.3 Outline

Section 2 presents the benchmark model and Section 3 establishes some preliminary results. Section 4 focuses on the predictions of the model when there is consensus for change. Section 5 considers a perverse incentive that facing the proposer in our model and discusses an application to the Ethereum blockchain. Section 6 introduces an extension to the benchmark model, where there are multiple proposers, and discusses the models applicability to the Bitcoin blockchain. Section 7 concludes the paper.

[14]In an effort to create non-contentious forks, on-chain governance provides a default set of rules for how upgrades should be voted on and when they should be accepted.

2 A Benchmark Model

There is a unidimensional space of *policy alternatives*, $\Theta = \mathbb{R}$, and a unit mass of users, with i denoting a user. The *status quo policy* is denoted by θ_0, and is associated with the original currency.

Each user i has an *ideal policy*, θ_i, independently and identically distributed according to continuous distribution F; we denote the density function by f. Without loss of generality, we assume that $\theta_0 = 0$.[15]

When a new chain is proposed, it features a policy, $\theta \neq \theta_0$, i.e., a change to the status quo, and is associated with a new currency. Each user chooses whether to adopt the new chain and new currency, or to continue to adopt the original chain and original currency. A user who switches is said to *adopt* policy θ, and a user who does not is said to *maintain* the status quo.[16]

A user's utility from adopting a new chain depends on the policy and the decisions of other users. Given a mass of users, $q \in [0, 1]$, who adopt the chain, the utility user i attains is

$$u(\theta, q; \theta_i, \alpha) := -(\theta - \theta_i)^2 + \alpha\, G(q), \tag{1}$$

where $\alpha > 0$ is a measure of *network importance*, and G is a continuous, strictly increasing function, with $G(0) = 0$ and $G(1) = 1$.[17] The mass of users across both chains sums to one; thus, if the new chain attracts q users, then the original chain maintains $(1 - q)$ users.

We model the influence of cryptocurrency exchanges through a market friction, or

[15]Our model is agnostic about the location of the status quo policy. As will be discussed in the conclusion, the status quo policy may not be related to the preferences of current users, and understanding this relationship is an interesting direction for future research.

[16]Recall from Footnote 9, our model focuses on users that choose a single chain. Nonetheless, in equilibrium, it will be shown that users will not have an incentive to adopt more than one chain (except for a measure zero set of users).

[17]Our results hold for a more general family of utility functions, $u(\theta, q; \theta_i, \alpha) = V(\theta, \theta_i) + \alpha G(q)$, when continuity, concavity, single-peakedness and a positive mixed-derivative requirement is applied to V. However, for simplicity we take $V(\theta, \theta_i) = -(\theta - \theta_i)^2$ which ensures that $V(\theta, \cdot) - V(0, \cdot)$ has a convenient closed-form, partial inverse.

feasibility constraint, which requires a currency or chain to have a *minimum threshold* of adopters, $\bar{q} \in [0, 1/2)$, to be considered viable. We model users as only being able to adopt a viable currency. This captures the discontinuous effect of an exchange's choice of whether to list a currency.[18]

In the benchmark model, there is a single user, the proposer, denoted P, with the power to fork the original chain and propose a new chain and an associated new currency. Similar to other users, the proposer has preferences over policies and the size of the network, given by the utility function $u(\theta, q; \theta_P, \alpha_P) = -(\theta - \theta_P)^2 + \alpha_P G(q)$, where $\theta_P \in \Theta$ and $\alpha_P > 0$. We allow the proposer to have a distinct preference over network size with $\alpha_P \neq \alpha$.

Timing: We model the political process as proceeding in two stages. At stage one, P proposes a new chain with policy, θ, to replace the status quo, θ_0, or chooses to maintain the original chain. If making a proposal, there is a second stage, and each user forms a belief about the mass of adopters and decides whether to adopt, conditional on believing the new currency will be viable. If only one currency is viable then all users necessarily choose this currency.

Equilibrium Concept: We assume a complete information setting, with a rational proposer and rational users. As such, given any proposal, the proposer and users hold a common belief about the mass of adopters, denoted by $\hat{q}_\theta$, and this belief is *self-fulfilling*: if q_θ denotes the mass of adopters, then, in equilibrium, $\hat{q}_\theta = q_\theta$. This concept of self-fulling beliefs can be found in Katz and Shapiro (1985); Myerson and Weber (1993). We will also assume a parameter restriction (Assumption 1) that ensures the equilibrium is unique and, hence, avoids coordination issues. Below, we provide more details and a discussion of our equilibrium concept.

[18]With some care, relating to equilibrium selection, our results can be attained in an alternative model that embeds this viability effect directly into a user's utility function. Our approach is chosen for expositional simplicity.

We define $\tilde{q}_\theta$ to be a solution of the following fixed point equation

$$x = \int_{\{\theta_i \,:\, u(\theta,x;\theta_i,\alpha)\geq u(\theta_0,1-x;\theta_i,\alpha)\}} dF(\theta_i). \tag{2}$$

That is, $\tilde{q}_\theta = x$ such that x satisfies (2). The value $\tilde{q}_\theta$ corresponds to a mass of adopters of the new currency such that if all users held the belief $\tilde{q}_\theta$, and there were no feasibility constraints, the belief would be self-fulfilling.[19] Unlike the value $\tilde{q}_\theta$, the proposer and users' equilibrium beliefs, $\hat{q}_\theta$, will account for the feasibility constraint $\bar{q}$.

We now define an equilibrium of our model. An equilibrium occurs whenever the following three conditions hold:

(C1) Each user adopts the policy that maximizes their utility.

(C2) Each user and the proposer have beliefs, $\hat{q}_\theta$, that are self-fulfilling and account for the feasibility constraint. Formally, if $\tilde{q}_\theta \in (\bar{q}, 1-\bar{q})$ then $\hat{q}_\theta = q_\theta$, if $\tilde{q}_\theta \leq \bar{q}$ then $\hat{q}_\theta = 0$, and otherwise $\hat{q}_\theta = 1$.

(C3) The proposer acts to maximize their own expected utility, proposing a new policy, θ, or choosing to maintain the original chain.

Condition (C1) and (C3) are standard, and imply sequentially rational behavior.

Condition (C2) requires that users first form beliefs that would be self-fulfilling absent the feasibility constraint, and update their beliefs to account for the feasibility constraint. This removes coordination problems arising from users who hold different beliefs. All agents will have the same, correct belief about whether each chain will be viable. Hence, no user will adopt a currency that they expect to be viable but then later find that it attracted insufficient support.[20]

[19] In general, Equation (2) can have multiple solutions, however, our parameter restrictions, to be described below, will ensure that there exists a unique solution.

[20] Indeed, features within blockchain technology can be utilized to avoid these coordination issues. For example, as described in Lombrozo (2017), the Bitcoin Core developers implemented an activation rule for proposed forks that required 95% affirmative signalling among network validators (miners) before the fork was activated. If that supermajority was not reached, the fork would not be put into effect even though the code was widely distributed.

Parameter restriction: We restrict our attention to proposals that are not too close to the status quo. Combined with conditions (C1)-(C3), this suffices to ensure a unique self-fulfilling belief and hence a unique equilibrium.

Assumption 1 *The proposed policy θ is never too close to the status quo policy. Formally, we assume*

$$|\theta| \geq \max\{\sqrt{\alpha},\, \alpha \sup_{x\in\Theta} f(x) \max_{q\in[0,1]} G'(q)\}, \tag{3}$$

and we denote the subset of policies θ satisfying (3) by $\tilde{\Theta}$.

Without Assumption 1, multiple equilibria can arise. For example, if a proposal equal to the status quo $\theta_0 = 0$ is proposed, then three equilibria exist: one where all users adopt the proposal, one where all users maintain the status quo, and one where the users split 50-50 between the two (identical) policies. Assumption 1 rules out the possibility of the proposer making such proposals. This assumption can be justified in a variety of ways. First, if whenever a policy θ violating Assumption 1 is proposed the equilibrium outcome is chosen at random from the set of possible equilibria, then a proposer that places a sufficiently high value on network effect or is sufficiently risk averse would never prefer to propose such a policy. Alternatively, if users are assumed to exhibit status quo bias (Samuelson and Zeckhauser, 1988), whereby they only adopt a new chain if a sufficiently large policy change is proposed, then the proposer will never have a strict incentive to propose a policy that violates Assumption 1.

Lemma 1 *Under Assumption 1, there is a unique value $x \in [0,1]$ satisfying Equation (2). Thus, there is a unique self-fulfilling belief $\hat{q}_\theta \in (\bar{q}, 1-\bar{q}) \cup \{0,1\}$ and a unique equilibrium.*

Proof. In Appendix A. ■

Given Lemma 1, we will no longer distinguish between the belief that users hold about the mass of other adopters, $\hat{q}_\theta$, and the actual mass of adopters, q_θ, since the two coincide.

3 Preliminaries

3.1 User decisions to adopt

We begin by focusing on a user's decision to adopt a new policy. Proposition 1 says that a user's decision depends on the user's preferences, the mass of other adopters, and how much the user values network size relative to the policy dimension. Furthermore, whether a user will adopt the new policy can be reduced to determining whether the user has an ideal policy to the right or to the left of the *marginal user's ideal policy*, denoted by θ_{marg}. Thus, the decision to adopt the new policy or maintain the original policy partitions the policy space of users into two contiguous regions.

Proposition 1 (User adoption decision rule) *Suppose both the currency, with policy $\theta \in \tilde{\Theta}$, and the original currency are viable. Let $q_\theta \in (\bar{q}, 1-\bar{q})$ denote the mass of users adopting the new policy. A user i adopts the new policy if and only if*

$$\underbrace{-(\theta-\theta_i)^2 + (\theta_0-\theta_i)^2}_{\text{Policy benefit}} \geq -\alpha \underbrace{\left[G(q_\theta) - G(1-q_\theta)\right]}_{\text{Network externalities}}. \tag{4}$$

Furthermore, there exists a marginal user's ideal policy,

$$\theta_{marg} := \frac{\theta}{2} - \alpha \frac{G(q_\theta) - G(1-q_\theta)}{2\theta}, \tag{5}$$

such that: if $\theta > 0$ ($\theta < 0$) user i adopts the new policy if and only if $\theta_i \geq \theta_{marg}$ ($\theta_i \leq \theta_{marg}$). Thus, the mass of users adopting the new policy, q_θ, is equal to $1 - F[\theta_{marg}]$ if $\theta > 0$ and is equal to $F[\theta_{marg}]$ otherwise.

Proof. In Appendix A. ■

Absent network preferences, that is, with $\alpha = 0$, a user's decision is purely policy-driven. The user adopts the chain operating with a policy closest to their ideal policy, θ_i. That is, the user adopts the new chain, operating with θ, if and only if $\theta \in (0, 2\theta_i)$.

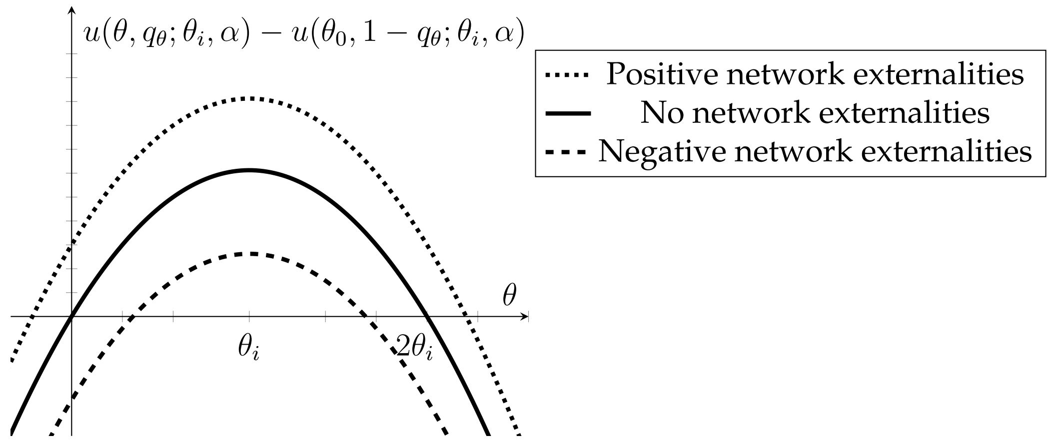

Figure 1: User i's adoption decision rule: adopt if and only if $u(\theta, q_\theta; \theta_i, \alpha) - u(\theta_0, 1 - q_\theta; \theta_i, \alpha) > 0$.

However, in the presence of network preferences, $\alpha > 0$, the mass of other adopters becomes a factor. If more than a majority of others adopt, then $[G(q_\theta) - G(1 - q_\theta)] > 0$ and there are *positive network externalities* on the new chain. In this case, user i may be willing to adopt the new chain even if the policy θ is further away from the user's ideal policy than the status quo. When less than a majority of the users adopt, then $[G(q_\theta) - G(1 - q_\theta)] < 0$ and there are *negative network externalities*; in this case, the reverse holds. Figure 1 illustrates a user's adoption decision rule.

3.2 The proposer's problem

We now turn our attention toward the proposer's decision as to which policy to propose, if any. We begin by defining the three different types of forks that the proposer can induce.

Definition 1 (Types of forks) A new proposal is said to lead to a

(i) **contentious fork** if the currencies on both chains are viable,

(ii) **non-contentious fork** if only the currency on the new chain is viable,

(iii) **null fork** if only the currency on the original chain is viable.

A non-contentious fork occurs when the new currency attracts a sufficiently large mass of adopters that the original currency is no longer viable. A null fork occurs when the new currency attracts so few adopters that it is not viable.

The type of fork occurs depends on the distribution of user preferences, F, as well as the proposer's preferences. In equilibrium, the proposer chooses a policy $\theta \in \tilde{\Theta}$ to maximize their utility

$$\begin{cases} u(0, 1; \theta_P, \alpha_P) & \text{if } \hat{q}_\theta \leq \bar{q}, \text{ or} \\ u(\theta, \hat{q}_\theta; \theta_P, \alpha_P) & \text{if } \hat{q}_\theta \in (\bar{q}, 1 - \bar{q}), \text{ or} \\ u(\theta, 1; \theta_P, \alpha_P) & \text{if } \hat{q}_\theta \geq 1 - \bar{q}, \end{cases} \tag{6}$$

or chooses to maintain the status quo policy, which leads to utility $u(\theta_0, 1; \theta_P, \alpha_P)$.

Each case in (6) corresponds to one of the types of forks identified in Definition 1. The first corresponds to a null fork, the second to a contentious fork, and the third, a non-contentious fork.

3.3 Incentives for extreme proposals

The mass of adopters q_θ need not change in a monotonic way as the proposed policy moves away from the status quo. In fact, a more extreme policy may attract more users than a less extreme policy. This result is in contrast to standard political economy models, where individuals are typically assumed not to have network preferences (i.e., $\alpha = 0$), and where implemented policies are enforced upon all individuals. In these standard models, the prediction is that as the policy moves further away from the status quo (becoming more extreme) the mass of users who would choose to adopt the policy decreases. The non-monotonicity in the benchmark model is driven by users also having preferences over network size.

Proposition 2 (Non-monotonic adoption rate) *For certain parameter ranges, the policy adoption rate q_θ is non-monotonic as the proposal, $\theta \in \tilde{\Theta}$, moves away from the status quo.*

Proof. In Appendix A. ■

Proposition 2 can be intuitively understood. For simplicity, suppose the feasibility constraint $\bar{q} = 0$. Consider the most extreme user, with ideal policy denoted θ_E. When the status quo policy is maintained with a unit mass of adopters, this user attains utility $u(\theta_0, 1; \theta_E, \alpha) = -\theta_E^2 + \alpha$. Now suppose that θ_E is sufficiently large such that $u(\theta_0, 1; \theta_E, \alpha) < 0$. Then, as per Figure 2, there must exist some intermediate policy $\theta^* \in (0, \theta_E)$ such that user E is indifferent between the status quo policy with a unit mass of users, and the policy $\underline{\theta}$ where no users adopt the policy. Now if the policy $\underline{\theta}$ (or any policy to the left) is proposed, then the mass of users that adopt is zero in equilibrium. However, for a policy $\underline{\theta} + \varepsilon$ with $\varepsilon > 0$, user E strictly benefits from adopting the policy and so too does a non-zero mass of users with ideal policies close to θ_E. Thus, in equilibrium there will be a non-zero mass of adopters. That is, the policy adoption rate increases. As θ further increases, the additional adopters have less extreme preferences and hence adopt at a slower rate, leading to a diminishing adoption rate. Eventually, an increase in θ will lead to the marginal adopter being to the left of the policy proposal, and any further increases in θ lead to a decline in the mass of adopters. This result is illustrated in Figure 3.

As a corollary of Proposition 2, network preferences can lead the proposer to propose a policy more extreme than theur ideal policy. Suppose θ^* is the policy that maximizes the mass of adopters— as per Figure 3 —and consider a proposer with ideal policy θ_P equal to $\theta^* - \varepsilon$, for sufficiently small $\varepsilon > 0$ and $\alpha_P > 0$ large. The proposer is close to indifferent between the policies θ_P and θ^*; however, due to network preferences, the proposer will strictly prefer to propose $\theta^* > \theta_P$, which entails a larger mass of adopters.

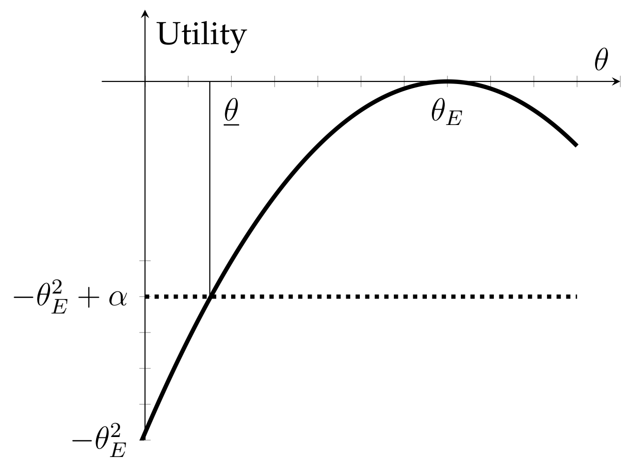

Figure 2: The most extreme agent's points of indifference

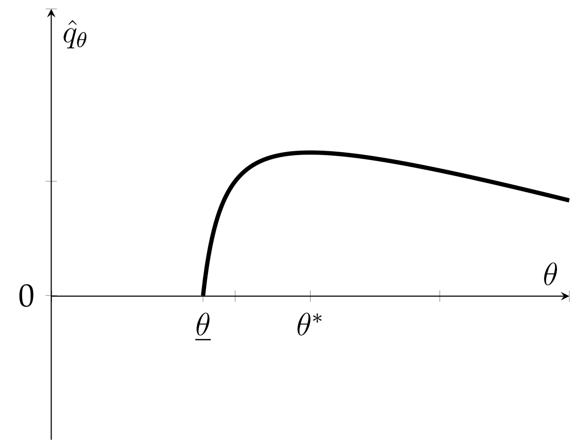

Figure 3: An endogenous mass of adoption

4 Consensus for Change, and Forking Conditions

We now explore conditions under which we expect forks to occur. First, we identify a sufficient condition for a non-contentious fork to occur, we call this condition *consensus for change*. Consensus for change requires that, holding all else equal, a mass of at least $(1-\bar{q})$ users prefer some policy in $\tilde{\Theta}$ compared to the status quo. Intuitively, this occurs when the distribution of user ideal policies are skewed away from the status quo policy. Definition 2 formally states the consensus for change condition and Proposition 3 states its sufficiency in generating non-contentious forks.

Definition 2 (Consensus for Change) *Let* $\tilde{\theta}_{min} := \arg\min_{\theta \in \tilde{\Theta}}\{|\theta|\}$. *There is* consensus for change towards the right *if* $F(\tilde{\theta}_{min}/2) \leq \bar{q}$. *There is* consensus for change towards the left *if* $1 - F(-\tilde{\theta}_{min}/2) \leq \bar{q}$.

Proposition 3 (Consensus condition) *Suppose there is consensus for change, then for some values of ideal policy,* θ_P*, the proposer will optimally induce a non-contentious fork.*

Proof. In Appendix A. ■

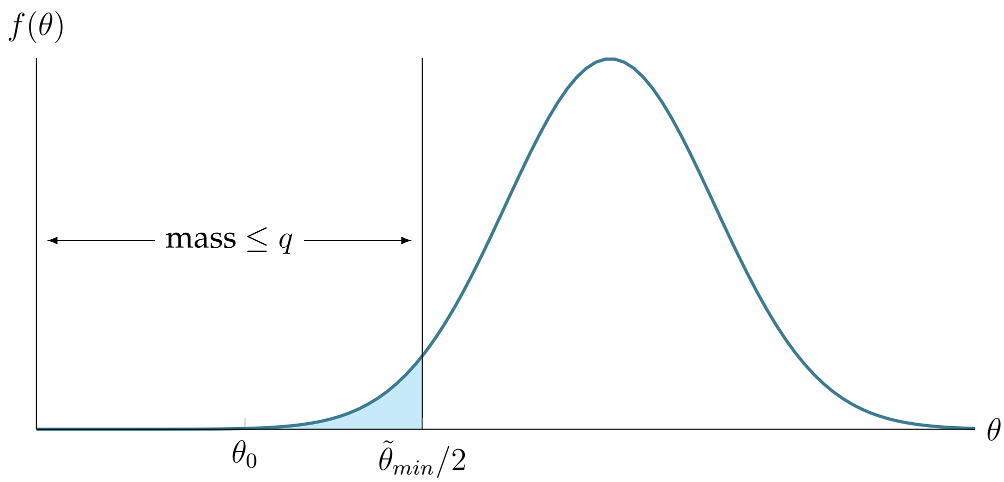

Figure 4: Consensus for change

It is intuitive why the consensus for change condition is sufficient for generating non-contentious forks. When the distribution of user ideal policies is skewed away from the status quo policy, there exists some new policy θ that is preferred and will be adopted by a large mass of users. If the distribution is sufficiently skewed, then the mass of adopters will be large enough (exceeding $1 - \bar{q}$) to make the original chain not viable. This leads to a non-contentious fork whenever a new currency with policy θ is proposed. Thus, if the proposer has ideal policy θ_P equal to this policy, then the proposer will always induce a non-contentious fork, since the proposer is able to implement the ideal policy without any effect on network size.

We now characterize the behavior of the proposer when there is consensus for change. Let $\theta_{max} > 0$ and $\theta_{min} > 0$ denote the maximum and minimum policies in $\tilde{\Theta}$, respectively, that ensure a non-contentious fork. Such a policy exists whenever there is consensus for change. Note that θ_{min} is equal to $\tilde{\theta}_{min}$ as defined in Definition 2. Proposition 4 partitions the proposer's ideal policy into different regions where a null fork, a non-contentious fork, and a contentious fork will be induced. This is schematically illustrated in Figure 5.

Proposition 4 (Proposer's forking strategy) *Suppose there is consensus for change toward the right. There exists a δ, with $\delta > 0$, such that*

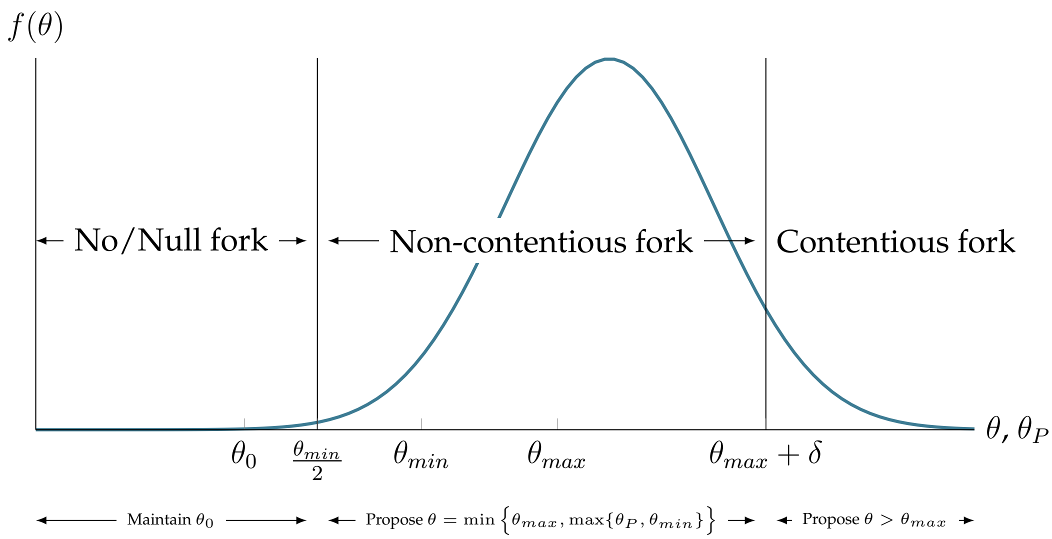

Figure 5: Forking strategies with consensus for change to the right

- *If $\theta_P \in (-\infty, \theta_{min}/2)$, the proposer will maintain the status quo policy (or propose a null fork),*
- *if $\theta_P \in (\theta_{min}/2, \theta_M + \delta)$, the proposer will propose $\theta = \min\left\{\theta_{max}, \max\{\theta_P, \theta_{min}\}\right\}$, and there will be a non-contentious fork, and*
- *if $\theta_P > \theta_M + \delta$, then the proposer will propose some $\theta > \theta_M$, and a contentious fork occurs.*

This is schematically illustrated in Figure 5.

Proof. In Appendix A. ■

When the proposer has an ideal policy to the left of the status quo and there is consensus for change towards the right, the proposer can do no better than maintain the status quo (a null fork). Any proposal that would be adopted would only make the proposer worse off. This also holds true for a small region of policies to the right of the status quo, i.e., $\theta_P \in (0, \theta_{min}/2)$.

A proposer with an ideal policy to the right of the status quo, and beyond $\theta_{min}/2$, will induce a fork. If the proposer's ideal policy θ_P belongs to the interval $[\theta_{min}, \theta_{max}]$, then the proposer can propose a new currency with ideal policy θ_P and this will induce a non-contentious fork. If the proposer has an ideal policy slightly outside of interval $[\theta_{min}, \theta_{max}]$, they will propose a policy that is closest to the proposer's ideal while ensuring a non-contentious fork. Finally, if the proposer's ideal policy is sufficiently far from θ_{max}, they will propose a contentious fork. This occurs because the ideal policy is sufficiently far from the status quo and θ_{max} that the proposer is willing to sacrifice at least $\bar{q}$ of network size to induce a chain with a more preferred policy.

Proposition 4 leads immediately to the following corollary. Corollary 1 says that, when there is consensus for change, the proposer's ideal policy influences the types of forks that will occur. In particular, as the proposer's ideal policy increases, the type of fork will transition from a null fork to a non-contentious fork to a contentious fork. Furthermore, the proposed policy proposed is weakly increasing to the right as the proposer's ideal policy moves to the right.

Corollary 1 (Comparative statics I) *Suppose there is consensus for change towards the right. As the proposer's ideal policy, θ_P, increases,*

- *the fork induced by the proposer will shift from a null fork to a non-contentious fork to a contentious fork, and*
- *the proposed policy, θ, weakly increases.*

We omit the proof of this corollary.

The types of forks induced by the proposer are also influenced by how much the proposer values network size, α_P, and the amount of market friction generated by exchanges, $\bar{q}$. However, the influence of each of these parameters also depends on the proposer's ideal policy.

Proposition 5 considers two cases. First, if the proposer's ideal policy is close or to the left of the status quo yet there is consensus for change towards policies to the right, there is no effect of increasing α_P or increasing $\bar{q}$. In this case, the proposer maintains the status quo, since the proposer can do no better, and increasing these parameters does not change this. Second, if the proposer has an ideal policy further to the right than $\theta_{min}/2$, they will induce either a non-contentious or a contentious fork. Increasing either α_P or $\bar{q}$ will never cause the proposer to avoid a fork but it will affect the proposed policy, and in some cases the type of fork.

Proposition 5 (Comparative statics II) *Suppose there is consensus for change towards the right. If $\theta_P \leq \theta_{min}/2$, then there is no effect of a marginal increase in α_P or $\bar{q}$ on the proposal. If $\theta_P \geq \theta_{min}/2$, then*

- *a marginal increase in the proposer's ideal policy, α_P, can cause contentious forks to become non-contentious, but never the reverse, and leads to more moderate forks, i.e., proposals closer to the status quo.*
- *a marginal increase in market friction, $\bar{q}$, can cause contentious forks to become non-contentious, but never the reverse, and has an ambiguous effect on the positioning of forks relative to the status quo.*

Proof. In Appendix A. ■

A proposer who is inducing contentious forks is willing to sacrifice network size for a more preferred policy. Because of this, increasing the value of the network, α_P, will motivate the proposer to maintain a larger network, and hence in some cases a proposer inducing a contentious fork will prefer to induce a non-contentious fork.

Importantly, the maximum policy θ_{max} that can be proposed can still ensure a non-contentious fork is unaffected by changes in α_P, and hence a proposer who transitions from a contentious fork to a non-contentious fork will do so while proposing more moderate policies.

Likewise, increasing the market friction $\bar{q}$ can induce a proposer who would propose a contentious fork to instead propose a non-contentious fork. This occurs because the largest policy θ_{max} that ensures a non-contentious fork increases as $\bar{q}$ increases. Let the new value of θ_{max} be denoted by θ'_{max}. For this reason, the maximum utility the proposer can attain from a non-contentious fork increases, while the maximum utility from a contentious fork is unchanged. In some cases, a proposer who previously proposed a contentious fork will now propose a more moderate policy θ'_{max}, and induce a non-contentious fork. On the other hand, a proposer who is in the region $(\theta_{max}, \theta_{max} + \delta)$ will propose a non-contentious fork, and policy θ_{max}. An increase in $\bar{q}$ will increase θ_{max} to θ'_{max}, and any proposer in this region will benefit from proposing a more extreme policy, such as θ'_{max}, over θ_{max}.

Finally, we remark on instances where there is no consensus for change. If the condition in Definition 2 does not hold, then it may still be the case that there exists policies in $\tilde{\Theta}$ that induce non-contentious forks. In this case, Proposition 4 continues to hold. If there is no such policy in $\tilde{\Theta}$, then any new proposal leads to either a contentious fork or a null fork. It is immediate that only a proposer with θ_P far from the status quo, or a proposer who places a low enough value on network size, will engage in a non-null fork, since any fork will come at the cost of a reduced network.

5 The Value of Being the Proposer

We now consider how the value that comes from being able to make a proposal depends on parameters of the model. We define the *value of being the proposer* as the difference between the proposer's maximum attainable utility when making a proposal and the utility from the status quo. Formally, the value of being the proposer is

$$W(\theta_P, \alpha_P) := \max_{\theta \in \tilde{\Theta}} \{u(\theta, q_\theta; \theta_P, \alpha_P)\} - u(\theta_0, 1; \theta_P, \alpha_P). \tag{7}$$

Proposition 6 says that the value of the proposer is increasing as θ_P moves away from the status quo and the proposer's ideal policy become more extreme, as α_P decreases and the proposer values the network less, and as user preferences *shift* in the direction of the proposer's ideal policy. Formally, suppose $\theta_P > 0$ ($\theta_P < 0$) and recall that F is the distribution of user policy preferences, we say that users' preferences shift in the direction of the proposer's ideal policy if users' preferences have distribution $F_{shift} : F_{shift}(\theta) = F(\theta - a)$ for all θ and for some $a > 0$ ($a < 0$).

Proposition 6 (Value of proposal power) *The value of being the proposer is weakly increasing as*

(i) θ_P *diverges from* θ_0*,*

(ii) α_P *decreases, and*

(iii) users' preferences shift in the direction of the proposer's ideal policy.

Proof. In Appendix A. ■

The results in Proposition 6 are intuitive. A proposer expects more utility from having proposal power whenever the proposer will exercise the power and propose a fork. From Propositions 4 and 5, and Corollary 1, this is more likely to occur when the proposer's ideal policy is far from the status quo, the proposer places little value on network size, or there is consensus for change in the direction of the proposer's ideal policy.

We conclude this subsection with a brief remark. So far our model has assumed that the proposer can make proposals at no cost. It is straightforward to relax this simplifying assumption. If the proposer experiences a fixed cost, say $C > 0$, from making a proposal, then our main insights continue to hold with only slight changes to the theoretical results. For example, Proposition 4 will be modified in two ways. First, the proposer will never propose a null fork since if proposals are costly, then they would be strictly better off maintaining the status quo. Second, the proposer will only propose a contentious or non-contentious fork if, in addition to the conditions specified in Proposition 4, the proposer's value (7) exceeds the cost of making a proposal C.

Remark 1 (Costly proposals) *If proposals are costly, then the proposer will never propose a null fork.*

5.1 The Ethereum network's Ice Age phenomenon

Among cryptocurrencies, a highly polarizing policy-issue relates to how *consensus* is achieved in regard to what is the correct sequence of blocks that comprise a blockchain. The most commonly applied method to achieve consensus is the *Proof of Work* (PoW) protocol, but many within the cryptocurrency community have pushed for a policy reform and, in particular, the adoption of an alternate protocol called *Proof of Stake* (PoS) (see Schumann (2018) for a discussion of the differences between these two consensus protocols).

The Ethereum network makes uses of the PoW protocol. As of early 2020, the Ethereum cryptocurrency is the number two cryptocurrency by market capitalization at $19 Billion (USD). However, in September 2015, the Ethereum Foundation, a non-profit institution created to support the development of Ethereum, announced their intention to move the network to the PoS protocol. Perhaps to show commitment, the Ethereum Foundation proposed at the same time a fork that introduced a "difficulty bomb" into the code, justifying it as an update that would "elegantly guarantees a hard-fork point in the next 16 months." After twelve months, the policy change would start to make transactions increasingly slow, and eventually render the chain useless (Tual, 2015). The anticipated effect of this change became known as the "Ice Age." The difficulty bomb fork was universally accepted as a non-contentious fork.

Proposition 6 provides an insight into the Ice Age phenomenon, and how it may have enabled the Ethereum network to implement six non-contentious forks between 2015 and 2019 (Moos, 2020; BokkyPooBah, 2019). In particular, the Ice Age has been repeatedly delayed via the proposal and adoption of new (non-contentious) forks. The first delay occurred in October 2017, delaying the difficulty bomb by 18 months, and then again in February 2019, delaying the bomb by another 12 months. With each non-contentious fork that occurred, not only was the bomb delayed, but a bundle of other policy changes were also introduced.

How does the behavior of the Ethereum Foundation connect with our results? We posit that a reasonable approximation for the Ethereum network is the single proposer model, where the Ethereum Foundation takes the role of proposer. Our results highlight that the proposer has limited ability to implement reform when there is no consensus for change. This is especially true for a proposer who places a high value on the network size, or a proposer who does not have extreme policy preferences. In this light, a possible motivation of the Ice Age may be to generate consensus for change, which leads to increased utility for the proposer (Proposition 6). This occurs because the difficulty bomb makes the original currency increasingly undesirable for users, and so when faced with a new proposal this can lead to a non-contentious fork.

6 An Extension to Handle Multiple Proposers

We now consider an extension of the benchmark model where there may be multiple users who simultaneously propose a new policy.

There is a set of m proposers, denoted by $\{P_j\}_{j=1}^{m}$, perhaps with $m > 1$.[21] Not every proposer needs to make a proposal. Let $S \subseteq \{1, 2, \ldots, m\}$ be the subset of proposers who actually make proposals. We say that one of these proposals will be *recognized*. The probability that a proposal is recognized is a function of the set of proposers that make proposals.[22] We assume that the probability of any proposer's proposal being recognized is positive and decreasing as the set of proposers making proposals expands. In particular, proposer j's proposal is recognized with probability $r_j(S)$, where $r_j(S) \in (0, 1)$, $r_j(S)$ is decreasing with respect to set-inclusion ordering, and $\sum_{k \in S} r_k(S) = 1$. That is, we model this recognition process as a contest.

Given the policy that is recognized, the game in terms of proposer utilities and the adoption decisions of users then proceeds as per the benchmark model.

[21]The set of proposers is finite and, hence, of measure zero. Thus, as with the users in our model, the proposers do not contribute to network effects themselves.

[22]We do not allow the probability that a proposal is recognized to depend on the proposal itself. This may be an interesting avenue for future work.

6.1 Two proposer model

We now consider the case where there is more than one proposer. In particular, we focus on the $m = 2$ case. The main insight that comes from this is the presence of a spoiler effect, whereby a proposer may propose a new currency that they know will not be viable simply to reduce the probability that another proposal will be adopted.

Proposition 7 (The spoiler effect) *There can exists a spoiler effect in equilibrium. That is, a proposer may make a proposal that, if recognized, would lead to a null fork. The incentives for a given proposer to act as a spoiler are increasing as*

(i) the other proposer has more extreme views, and

(ii) users' policy preferences shift towards the other proposer's ideal policy.

Proof. In Appendix A. ∎

To illustrate Proposition 7, we consider a situation where there is consensus for change toward a policy $\theta > 0$, and the proposers have ideal policies $\theta_{P_2} < \theta_0 = 0 < \theta_{P_1}$. This is schematically illustrated in Figure 6.

Proposer P_1 is able to propose a new currency with policy $\theta = \theta_{P_1}$ and induce a non-contentious fork. If recognized, then all users (including proposer P_2) will move to this new chain. Proposer P_2 has an ideal policy to the left of the status quo, and hence if recognized cannot induce a fork and would prefer to the status quo. Despite this, the existence of proposer P_1 motivates proposer P_2 to propose a null fork, in order to reduce P_1's chance at forking the network and despite there being consensus for change toward P_1. The incentive for P_2 to engage in this spoiling behavior increases as P_1's ideal policy becomes more extreme or as users' preferences shift towards P_1's ideal policy.

It is straightforward to see that the spoiler effect persists even if proposals are costly. So long as the cost of a proposals is sufficiently small, P_2 will have a strict incentive to propose a spoiler policy that will fail to be adopted even if it is recognized.

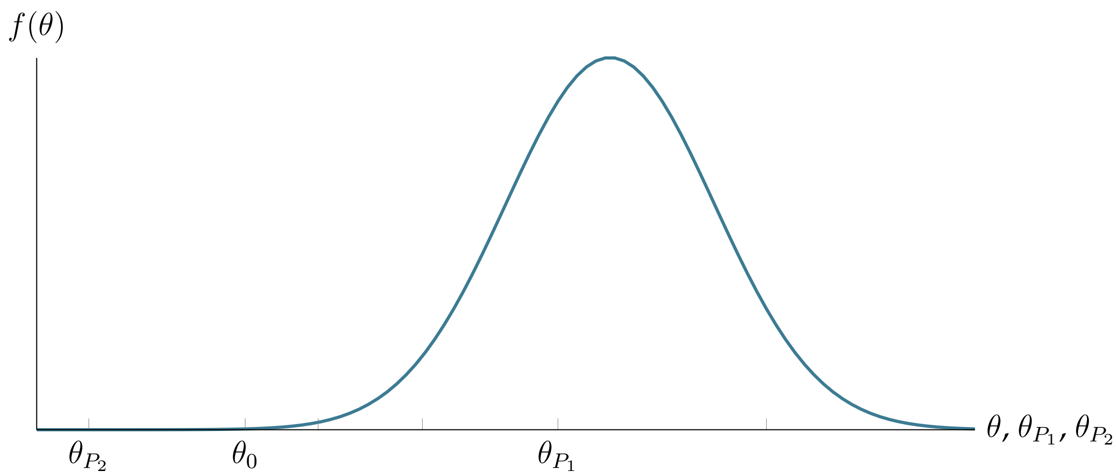

Figure 6: The spoiler effect.

Remark 2 (Costly proposals and the spoiler effect) *The spoiler effect persists even if proposals are costly.*

With a single proposer and costly proposals, we should never expect to see null forks in equilibrium (Remark 1). Indeed, in the Ethereum network, where the single proposer model is a reasonable fit, this is the case. In other networks the single proposer model is not a good fit. For example, in the Bitcoin network no single user has disproportionate influence and there are many proposers. Accordingly, many null forks arise in Bitcoin. For example, Bitcoin World and Bitcoin Ore are two null forks that occurred in the Bitcoin network; see also Balashova (2018). Proposition 7 and Remark 2 provide an insight into the presence of null forks even in situations where proposals are costly, and even when the proposer knows they will fail to be adopted if recognized.

7 Conclusion

In this paper, we have introduced a formal, political economy model of blockchain governance. Our model has abstracted away from many details, but maintained three core components that are salient to blockchains: the governance process is decentralized, users care about network effect and have heterogeneous preferences over policies, and the presence of cryptocurrency exchanges induces market frictions.

Our results highlight how the governance structure of blockchain applications and the cryptocurrency environment more broadly influences the incentives to fork chains, and the types of policy changes that we would expect to be proposed. We showed that, in equilibrium, a variety of forks can occur: contentious forks, non-contentious forks, and null forks. Non-contentious forks are more likely to occur when the proposer places a high importance on network effects, or has an ideal policy that is popular among users. Furthermore, we showed that cryptocurrency exchanges induce market frictions that facilitate non-contentious forks, and that this may disproportionately harm small groups of users. This occurs because exchanges list only a small number of currencies, and so a small minority, wanting to reject a proposed fork, may be unable to keep using the original currency and are instead coerced into accepting a fork that they would otherwise (in the absence of these frictions) reject. We also showed that blockchain governance produces an incentive for extremism. In standard models of political economy there exists incentives that push policy proposals toward the median voter's preferences (the Median Voter Theorem). In our benchmark model of blockchain governance, this incentive need not arise. Rather, a more extreme policy proposal may attract more support from users than a moderate policy change This occurs because adopting a new policy proposal can be costly for users in terms of network effects, and so small policy changes may fail to incentivize users that prefer more extreme policy changes. Yet, a large policy change may be able to incentivize users that desire extreme policy changes to adopt the policy, and even if it is costly in terms of network effects.

Our model provides a possible explanation for some of the counter-intuitive phenomena that have been observed within the cryptocurrency environment. For example, the success of the Bitcoin Cash fork after many more moderate forks had failed, and that a proposer may have an incentive to sabotage the status quo (e.g., the Ethereum Ice Age).

We also considered an extension to the benchmark model that incorporated multiple proposers. This extension explains some of the observed differences between governance of the Ethereum network and the Bitcoin network. In particular, it explains why users may have the incentive to propose forks that they know are doomed to fail (e.g., the Bitcoin community's tendency to frequently propose forks that fail).

Finally, we conclude by noting some directions for future research. First, our model assumes complete information. This simplifying assumption is appropriate for the purposes of this paper, and our results are robust to small amounts of uncertainty. However, it may be interesting to formally explore the role of incomplete information in blockchain governance. Second, our model is agnostic about the position of the status quo policy. Naturally, one may anticipate that the status quo policy should reflect the preferences of existing users (e.g., approximating the median of users' preferences). Yet, since a currency's user base is dynamic, the status quo may more realistically reflect the preferences of the initial users or a small group of developers that were involved in the currency's creation. Furthermore, how responsive a status quo policy is to changes in the user base may depend on the distribution of user preferences and whether the currency has adopted an on-chain governance protocol. Understanding the dynamics of a changing user base and status quo policy is left for future work.

A Omitted Proofs

Proof of Lemma 1. We solve for the case where $\theta > 0$, the case for $\theta < 0$ is proven similarly. First note that the value $\tilde{q}_\theta = x$, which solves (2), satisfies

$$\begin{aligned} & x = \int \mathbb{I}\Big(u(\theta, x; \theta_i, \alpha) \geq u(\theta_0, 1-x; \theta_i, \alpha) \Big)\, dF(\theta_i) \\ \iff & x = \int \mathbb{I}\Big(-(\theta - \theta_i)^2 + (\theta_0 - \theta_i)^2 \geq -\alpha[G(x) - G(1-x)] \Big)\, dF(\theta_i) \\ \iff & x = \int \mathbb{I}\Big(\theta_i \geq \frac{\theta}{2} - \alpha \frac{G(x) - G(1-x)}{2\theta} \Big)\, dF(\theta_i) \\ \iff & x = 1 - F\Big[\frac{\theta}{2} - \alpha \frac{G(x) - G(1-x)}{2\theta} \Big]. \end{aligned} \tag{8}$$

We now establish that the value $\tilde{q}_\theta \in [0, 1]$ exists and is unique. This is a fixed-point problem $T_\theta(x) = x$, where $T_\theta(x) := 1 - F[\theta/2 - \alpha \frac{G(x) - G(1-x)}{2\theta}]$. The mapping $T_\theta(x)$ is a contraction mapping whenever $|T'(x)| < 1$, and when this holds, by Banach's Fixed Point Theorem and an application of the Mean Value Theorem, there is a unique solution x solving (2). Thus, a sufficient condition for a unique solution x in $[0, 1]$ is

$$\Big| \frac{d\,T_\theta(x)}{dx} \Big| = \frac{\alpha}{2\theta} f\Big[\frac{\theta}{2} - \alpha \frac{G(x) - G(1-x)}{2\theta} \Big] [G'(x) + G'(1-x)] < 1,$$

which simplifies to

$$\theta > \frac{\alpha}{2} f[\frac{\theta}{2} - \alpha \frac{G(\hat{q}_\theta) - G(1-\hat{q}_\theta)}{2\theta}][G'(\hat{q}_\theta) + G'(1-\hat{q}_\theta)].$$

It is now immediate that this condition is satisfied whenever Assumption 1 holds. Thus, under Assumption 1, a unique solution x in $[0, 1]$ holds and, by (C2), the unique self-fulfilling belief will be such that $\hat{q}_\theta \in \{0, 1\} \cup (\bar{q}, 1 - \bar{q})$, as required. ■

Proof of Proposition 1. Let θ be the policy proposal and $q_\theta \in (\bar{q}, 1 - \bar{q})$ be the mass of adopters. The user's decision is between adopting the new policy θ with mass q_θ of users, or maintaining the status quo policy θ_0 with the remaining $(1-q_\theta)$ mass of users. The user adopts the proposal if and only if

$$
\begin{aligned}
& u(\theta, q_\theta; \theta_i, \alpha) \geq u(\theta_0, 1-q_\theta; \theta_i, \alpha) \\
\iff & -(\theta-\theta_i)^2 + \alpha G(q_\theta) \geq -(\theta_0-\theta_i)^2 + \alpha G(1-q_\theta) \\
\iff & -(\theta-\theta_i)^2 + (\theta_0-\theta_i)^2 \geq -\alpha[G(q_\theta) - G(1-q_\theta)].
\end{aligned}
$$

For $\theta > 0$, it is immediate that a user i adopts if and only if

$$
\begin{aligned}
& -\theta^2 + 2\theta\theta_i \geq -\alpha[G(q_\theta) - G(1-q_\theta)] \\
\iff & \theta_i \geq \frac{\theta}{2} - \alpha \frac{G(q_\theta) - G(1-q_\theta)}{2\theta}.
\end{aligned} \tag{9}
$$

Thus, the mass of adopters is simply

$$
\mathbb{P}[\theta_i \geq \theta_{marg}] = 1 - F[\theta_{marg}],
$$

where $\theta_{marg} := \frac{\theta}{2} - \alpha\frac{G(q_\theta)-G(1-q_\theta)}{2\theta}$. ■

Proof of Proposition 2. We prove the result in the case that $\bar{q} = 0$, it is easy to see that for $\bar{q} > 0$ not too large the result continues to hold. Furthermore, for simplicity we assume that $\alpha \in (0,1)$, and $F(1) = 1$, i.e., there is a most extreme agent with ideal policy $\theta_E = 1$. However, it will be immediate from the steps of the proof that the result holds more generally.

Consider the policy proposal $\underline{\theta} \in (0,1]$ such that the most extreme user E, with ideal policy $\theta_E = 1$, is indifferent between adopting the policy $\underline{\theta}$ when $q_{\underline{\theta}} = 0$, and maintaining the status quo with a unit mass of users. That is, $\underline{\theta} \in (0,1]$ solves

$$
-(1-\underline{\theta})^2 = -1 + \alpha,
$$

and so $\underline{\theta} := 1 - \sqrt{1-\alpha}$. In equilibrium, the mass of users adopting policy $\underline{\theta}$ is zero.

Now consider a policy $\theta_\varepsilon = \underline{\theta} + \varepsilon$ for $\varepsilon > 0$ sufficiently small. The most extreme user E strictly prefers to adopt the policy θ_ε for any $q_{\theta_\varepsilon} \geq 0$ rather than maintain the status quo policy (regardless of the mass of users maintaining the status quo). Furthermore, since F is continuous and strictly increasing in $(-1, 1)$, there is a mass of users with ideal policies in a neighborhood of $\theta_E = 1$ who also strictly prefer to adopt the policy θ_ε rather than maintain the status quo. It is now immediate that $q_{\theta_\varepsilon} > 0$; that is, the mass of users adopting the policy has increased.

It remains to prove that, at some point, the mass of adopters ceases to increase and begins to decrease; that is, the mass of adopters is non-monotonic. Consider the policy $\bar{\theta} > 2$ such that the most extreme user is indifferent between the status quo policy when no users adopt and the policy $\bar{\theta}$ when all users adopt, i.e., $\bar{\theta}$ solves

$$-(1-\bar{\theta})^2 + \alpha = -1,$$

and so $\bar{\theta} = 1 + \sqrt{1+\alpha}$. In equilibrium, the mass of users adopting policy $\bar{\theta}$ is zero. Since the mass of adopters is continuous it follows from the Intermediate Value Theorem that, at some point, the mass of adopters is decreasing. This completes the proof. ■

Proof of Proposition 3. Suppose there is consensus for change toward the right, i.e., $F(\tilde{\theta}_{min}/2) \leq \bar{q}$. When the policy $\tilde{\theta}_{min}$ is proposed the marginal adopter has ideal policy

$$\theta_{marg} = \frac{\tilde{\theta}_{min}}{2} - \alpha \frac{G(\tilde{q}_{min}) - G(1-\tilde{q}_{min})}{2\tilde{\theta}_{min}},$$

where $\tilde{q}_{min}$ is the mass of adopters of policy $\tilde{\theta}_{min}$. From Lemma 1, there is a unique solution $q = \tilde{q}_{min} \in [0, 1]$ such that

$$q = 1 - F[\frac{\tilde{\theta}_{min}}{2} - \alpha \frac{G(q) - G(1-q)}{2\tilde{\theta}_{min}}]. \tag{10}$$

Note that the RHS of (10) is increasing in q and at $q = 1$ the LHS is at least as large as the RHS. Thus, by the Intermediate Value Theorem, if

$$(1-\bar{q}) \leq 1 - F[\frac{\tilde{\theta}_{min}}{2} - \alpha\frac{G((1-\bar{q})) - G(1-(1-\bar{q}))}{2\tilde{\theta}_{min}}], \tag{11}$$

then the unique solution $q = \tilde{q}_{min} \in [0,1]$ must be such that $q \geq 1-\bar{q}$, and hence the actual mass of adopters will equal one. We conclude that proposing the policy $\tilde{\theta}_{min}$ leads to a non-contentious fork. It is immediate that a proposer with ideal policy $\tilde{\theta}_{min}$ will optimally propose this policy and induce a non-contentious fork.

It remains to prove that the consensus for change condition is sufficient to ensure that (11) holds. Rearranging (11) gives

$$\bar{q} \geq F[\frac{\tilde{\theta}_{min}}{2} - \alpha\frac{G((1-\bar{q})) - G(1-(1-\bar{q}))}{2\tilde{\theta}_{min}}], \tag{12}$$

and since $\bar{q} < 1/2$ we have $G((1-\bar{q})) > G(1-(1-\bar{q}))$ and so

$$F[\frac{\tilde{\theta}_{min}}{2} - \alpha\frac{G((1-\bar{q})) - G(1-(1-\bar{q}))}{2\tilde{\theta}_{min}}] < F[\frac{\tilde{\theta}_{min}}{2}]$$

The RHS is no larger than $\bar{q}$ by the consensus for change condition. This completes the proof. ■

Proof of Proposition 4. Suppose there is consensus for change toward the right; the proof for when consensus for change is toward the left is similar.

If a policy $\theta \in \tilde{\Theta}$ such that $\theta < 0$ is proposed, then it will result in a null fork. To see this, recall that the marginal adopter is the user with ideal policy

$$\frac{\theta}{2} - \alpha\frac{G(q) - G(1-q)}{2\theta}.$$

An upper bound for this ideal policy is

$$\frac{\theta}{2} - \alpha\frac{G(1) - G(0)}{2\theta} = \frac{\theta}{2} - \alpha\frac{1}{2\theta}.$$

However, since $|\theta| > \sqrt{\alpha}$ it follows that

$$\frac{\theta}{2} - \alpha\frac{G(q) - G(1-q)}{2\theta} < \frac{-\sqrt{\alpha}}{2} + \alpha\frac{1}{2\sqrt{\alpha}} = 0.$$

Thus, the marginal adopter has ideal policy to the left of the status quo and, hence, also to the left of $\theta_{min}/2$. The consensus for change condition then implies that the mass of adopters of any policy proposal $\theta < 0$ attract support from less than $\bar{q}$ users and hence leads to an adoption rate of zero. A proposer with ideal policy $\theta_P < 0$ will optimally maintain the status quo since there is no policy proposal that the they could benefit from. Furthermore, any proposer with $\theta_P < \theta_{min}/2$ will also come to the same conclusion, since the smallest policy that could be proposed and not lead to a null fork is θ_{min} (recall that Assumption 1 restricts the policy space).

Now, consider a proposer with $\theta_P \in (\theta_{min}/2, \theta_{min}]$. Proposing any policy $\theta \in \tilde{\Theta}$ such that $\theta < \theta_{min}$ leads to a null fork; however, proposing θ_{min} leads to a non-contentious fork. Clearly, a proposer with $\theta_P \in (\theta_{min}/2, \theta_{min}]$ will optimally propose θ_{min}.

Recall that θ_{max} is defined as the maximum policy $\theta \in \tilde{\Theta}$ such that proposing θ_{max} leads to a non-contentious fork. Also note that the sufficient condition (11) from within the proof of Proposition 3 shows that for any policy $\theta \in [\theta_{min}, \theta_{max}]$, proposing θ will lead to a non-contentious fork. Thus, any proposer with ideal policy $\theta_P \in [\theta_{min}, \theta_{max}]$ will propose $\theta = \theta_P$, which will lead to all users adopting their ideal policy.

Now consider a proposer with $\theta_P > \theta_{max}$. The proposer can attain utility

$$-(\theta_{max} - \theta_P)^2 + \alpha$$

by proposing policy $\theta = \theta_{max}$. This provides maximal utility among all policies in $\tilde{\Theta}$ and no larger than θ_{max}. In contrast, by proposing a policy $\theta > \theta_{max}$ the maximal utility is strictly bounded above by

$$-(\theta - \theta_P)^2 + \alpha\, G(1 - \bar{q}).$$

When θ_P approaches θ_{max} from above, it is straightforward to see that proposing θ_{max} is optimal. While for larger values of θ_P this need not be the case and a contentious fork may arise. Let $\theta_{max} + \delta$ for $\delta > 0$ denote the largest proposer value such that the proposer induces a non-contentious fork. This value exists (i.e. is well-defined) since if a proposer with θ_P engages in a contentious fork, then so too will a proposer with a more extreme ideal policy. This completes the proof. ■

Proof of Proposition 5. First, consider a proposer with $\theta_P < \frac{\theta_{min}}{2}$. For a given $\alpha_P, \bar{q}$, such a proposer maintains the status quo (Proposition 4). A marginal increase in α_P or $\bar{q}$ does not change the optimal behavior. This follows since the closest policy $\theta \neq \theta_0$ that the proposer could propose without inducing a null fork is the policy θ_{min}. This value is independent of both $\bar{q}$ and α_P. Thus, such a proposer would still prefer to maintain the status quo policy, and there is no effect of a marginal increase in α_P nor $\bar{q}$ on the proposer's behavior.

Now, consider a proposer with $\theta_P \geq \theta_{min}/2$. We focus on the two cases (an increase in α_P and an increase in $\bar{q}$) separately.

Case 1: A marginal increase in α_P to α'_P.

Recall that θ_{max} denotes the maximum proposal that induces a non-contentious fork. This value is unchanged by an increase in α_P to α'_P, since each user's utility is unchanged. Thus, a proposer with $\theta_P \in (\theta_{min}/2, \theta_{max}]$ will induce a non-contentious fork by proposing $\theta = \max\{\theta_{min}, \theta_P\}$, as per Proposition 4. Similarly, and again by Proposition 4, a proposer with $\theta_P > \theta_{max}$ but close to θ_{max} will also propose θ_{max} and induce a non-contentious fork. Recall that for a given value of α_P, the interval of proposer ideal policies where this behavior occurs is denoted by $(\theta_{max}, \theta_{max} + \delta)$. A marginal increase in α_P leads to a larger δ value, which we denote by δ'. To see this, consider a proposer with $\theta_P = \theta_{max} + \delta$ who is indifferent between proposing θ_{max} and proposing some other policies $\theta > \theta_{max}$ when the value of the network importance is α_P. That is,

$$-(\theta_P - \theta_{max})^2 + \alpha_P = -(\theta_P - \theta)^2 + \alpha_P\, G(q)$$
$$-\delta^2 + \alpha_P = -(\theta_P - \theta)^2 + \alpha_P\, G(q),$$

where q is the mass of adopters when proposing θ. An increase in α_P to α'_P leads to the LHS being strictly larger than the RHS. Thus, the proposer with $\theta_P = \theta_{max} + \delta$ strictly prefers to propose $\theta = \theta_{max}$ and induce a non-contentious fork, as opposed to inducing a contentious fork. By continuity this also hold for a neighborhood of proposer ideal policies about $\theta_{max} + \delta$. Thus, a marginal increase in α_P causes some proposer (those in the region of $(\theta_{max} + \delta, \theta_{max} + \delta')$) to propose more moderate policies and non-contentious forks, rather than contentious forks.

Case 2: A marginal increase in $\bar{q}$ to $\bar{q}'$.

In this case, the maximum policy proposal that will induce a non-contentious fork θ_{max} increase to some value θ'_{max}. To see this, note that by definition θ_{max} has a mass of adopters equal to $1 - \bar{q}$, and the mass of adopters, as a function of θ, $q_\theta \in [0, 1]$ is continuous. Furthermore, since θ_{max} is the maximum policy proposal that leads to a non-contentious fork (with feasibility constraint $\bar{q}$) it must be that $(1 - q_\theta) < (1 - \bar{q})$ for all $\theta > \theta_{max}$. However, an increase in $\bar{q}$ to $\bar{q}'$ leads to θ_{max} still inducing a non-contentious, since

$$1 - \bar{q}' < 1 - \bar{q},$$

but, in addition, a policy $\theta_{max} + \varepsilon$ for $\varepsilon > 0$ sufficiently small will also lead to a non-contentious fork. That is, the maximum policy proposal that induces a non-contentious fork increases with $\bar{q}$. Let θ'_{max} denote this new value. A proposer with ideal policy

$\theta_P = \theta_{max} + \varepsilon$ will now optimally propose the more extreme policy θ_P and induce a non-contentious fork rather than proposing θ_{max} (which also leads to a non-contentious fork). Thus, in this instance the increase in $\bar{q}$ leads some non-contentious forks to become more extreme.

In contrast, consider a proposer with ideal policy $\theta_P = \theta_{max} + \delta$. Under $\bar{q}$, this proposer is indifferent between inducing a contentious fork with some policy $\theta > \theta_{max}$ and proposing $\theta = \theta_{max}$, i.e.,

$$-(\theta_P - \theta_{max})^2 + \alpha_P = -(\theta_P - \theta)^2 + \alpha_P\, G(q)$$
$$-\delta^2 + \alpha_P = -(\theta_P - \theta)^2 + \alpha_P\, G(q),$$

where q is the mass of adopters when proposing θ. An increase in $\bar{q}$ leads to the possibility of proposing policy $\theta'_{max} > \theta_{max}$ and maintaining the entire network size. That is, the proposer is able to attain utility

$$-(\theta_P - \theta'_{max})^2 + \alpha_P > -(\theta_P - \theta_{max})^2 + \alpha_P.$$

Thus, a proposer with $\theta_P = \theta_{max} + \delta$ will strictly prefer to propose a more moderate policy $\theta = \theta'_{max}$ and induce a non-contentious fork rather than propose a policy θ that leads to a contentious fork (and is necessarily larger than θ_{max} and θ'_{max}). By continuity, this also holds for a neighborhood of proposer ideals around θ_P. Thus, in this instance the proposer is motivated to induce a non-contentious fork with a more moderate policy.

■

Proof of Proposition 6. Suppose $0 < \theta_P$ – the proof follows similarly for the opposite case. Let θ^* be an optimal policy proposal when the proposer's ideal policy is θ_P. Thus, θ^* maximizes

$$-(\theta_P - \theta)^2 + \alpha_P\, G(q_{\theta^*}) - \Big[-\theta_P^2 + \alpha_P \Big],$$

let W denote this maximal value. Note that θ^* is necessarily positive.

Proof of (i) and (ii): Notice that the above expression (for fixed $\theta \geq 0$) is increasing in θ_P and is decreasing in α_P. Thus, a proposer with ideal policy $\theta'_P > \theta_P$ or network-importance value of $\alpha'_P < \alpha_P$ has a value of being proposer at least as large as W.

Proof of (iii): Recall that q_θ solves

$$x = 1 - F\Big[\theta/2 - \alpha\frac{G(x) - G(1-x)}{2\theta}\Big].$$

If users' policy prefences shift toward $\theta_P > 0$ (as defined in the text preceding Proposition 6) then for any $x \leq q_\theta$ the LHS is less than the RHS. Further, since the fixed point equation has a unique solution it must be that the new fixed point solution $x' > x$. That is, when the user preferences shift, the mass of adopters q_θ is larger under F_{shift} than under F. Thus, a proposer facing distribution F_{shift} has a value of being proposer at least as large as W. ■

Proof of Proposition 7. Suppose there is consensus for change to the right, and the two potential proposers have ideal policy $\theta_{P_1} > 0$ and $\theta_{P_2} < 0$, respectively. In a single proposer model, it is immediate that proposer 2 is unable to propose any policy that would improve P_2's utility; while, proposer 1 will always propose a policy leading to either a contentious or a non-contentious fork. In the two proposer model, if proposer 2 makes a proposal, then it will be a null fork and will lead the status quo policy being maintained with the entire network mass.

Now, suppose that proposer 1 makes a proposal, this will be beneficial whenever the cost of proposals is sufficiently small. Proposer 2 experiences a reduction in utility if they do not make a proposal. If they make a proposal, then with probability $r \in (0, 1)$ their proposal will be recognized. This means that proposer 2 will avoid proposer 1's proposal and will be able to maintain the network size. Since proposals are costless, proposer 2 will make this proposal even though, if recognized, it will result in a null fork.[23] Note that proposer 2's choice to propose a fork, scales proposer 1's net expected utility from making a proposal down by a factor of $1 - r$, but it will still be optimal for proposer 1 to make a proposal. Thus, in equilibrium, proposer 2 makes a proposal simply to avoid proposer 1's proposal, and despite knowing that the proposer's own proposal will result in a null fork if it is recognized.

In such instances, the incentive for proposer 2 to engage in this spoiling behavior in increasing as their utility decreases from proposer 1's proposal. From earlier results, we know that this is the case when proposer 1 has a more extreme ideal policy or the users' preferences shift toward proposer 1. ■

[23]Of course, if the cost of making a proposal is positive but sufficiently small, then proposer 2 will still have a strict incentive to make a proposal.

B Blockchain Data

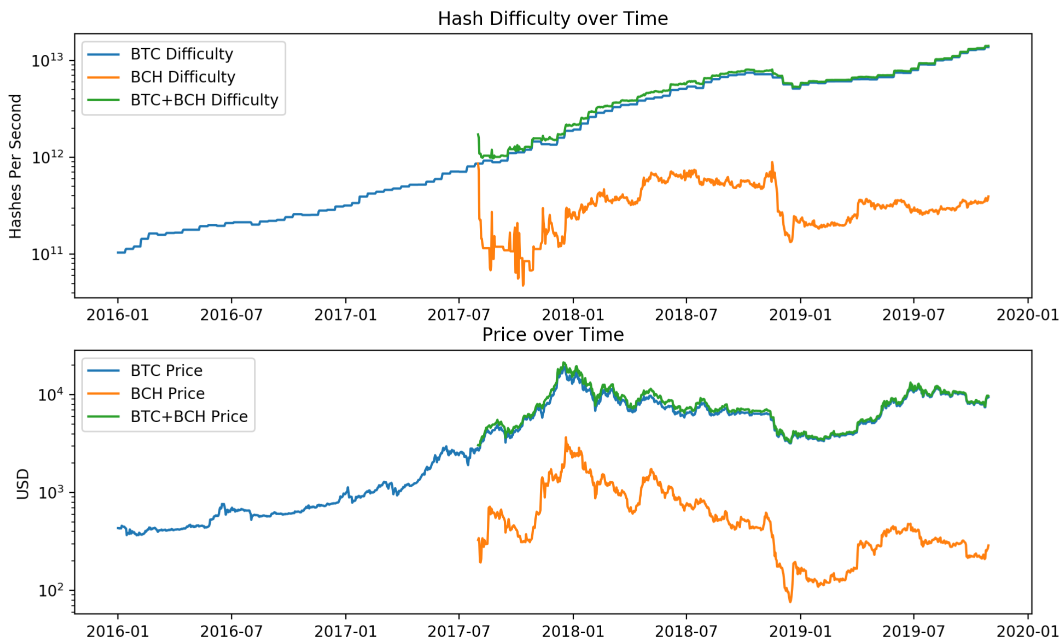

Figure 7: The BTC-BCH Fork. This graph shows the impact that the Bitcoin Cash (BCH) fork (July 2017) had on Bitcoin (BTC) hashrate difficulty and price. The hashrate (measured in hashes per second) of a proof-of-work network measures the quantity of computational work done to secure the underlying blockchain. The bottom graph gives the price of each asset in nominal USD terms.

References

Akerlof, Robert, Richard Holden, and Luis Rayo, "Network Externalities and Market Dominance," *Working paper*, 2018.

Alsindi, Wassim Z., "Towards an Analytical discipline of Forkonomy," `http://www.pllel.com/industries/wp-content/uploads/180810-forkonomy-v3.pdf` August 2018.

Alsindi, Wassim Z., "Forkonomy Revisited: Where Are They Now?," `https://hackernoon.com/forkonomy-revisited-where-are-they-now-73fbfbec6b4d` January 2019.

Alsindi, Wassim Z., "What the forkgov! Does staying together tear you apart even more?," `http://www.pllel.com/industries/what-the-forkgov-does-staying-together-tear-you-apart-even-more/` May 2019.

Balashova, Ana, "5 Failed Bitcoin Forks That Never Made It," `https://blockexplorer.com/news/5-failed-bitcoin-forks/` Oct 2018.

Baron, David P. and John A. Ferejohn, "Bargaining in Legislatures," *The American Political Science Review*, 1989, *83* (4), 1181–1206.

Barrera, Cathy and Stephanie Hurder, "Blockchain Upgrade as a Coordination Game," *SSRN Working Paper*, 2018.

Beck, Roman, Christoph Müller-Bloch, and John Leslie King, "Governance in the Blockchain Economy: A Framework and Research Agenda," *Journal of the Association for Information Systems*, 2018, *19* (10), 1020–1034.

Becker, Gary S., "A note on restaurant pricing and other examples of social influences on price," *Journal of Political Economy*, 1991, *99* (5), 1109–1116.

Berg, Chris, Sinclair Davidson, and Jason Potts, "Institutional discovery and competition in the evolution of blockchain technology," *Working Paper*, 2018.

Besley, Timothy and Stephen Coate, "An Economic Model of Representative Democracy," *The Quarterly Journal of Economics*, 1997, *112*, 85–114.

BokkyPooBah, The Officious, "Please provide a summary of the Ethereum hard forks," `https://ethereum.stackexchange.com/questions/13014/please-provide-a-summary-of-the-ethereum-hard-forks` Dec 2019.

Buterin, Vitalik, "Notes on Blockchain Governance," `https://hackernoon.com/notes-on-blockchain-governance-ob65o3pod` May 2019.

Calcaterra, Craig, "On-chain Governance of Decentralized Autonomous Organizations," *SSRN Working Paper*, 2018.

Conrad, Klaus, "Price competition and product differentiation when goods have network effects," *German Economic Review*, 2006, *7* (3), 339–361.

Downs, Anthony, *An Economic Theory of Democracy*, New York: Harper and Row, 1957.

Ehrsam, Fred, "Blockchain Governance: Programming Our Future," `https://medium.com/@FEhrsam/blockchain-governance-programming-our-future-c3bfe30f2d74` November 2017.

Faravelli, Marco, Priscilla Man, and Randall Walsh, "Mandate and paternalism: A theory of large elections," *Games and Economic Behavior*, 2015, *93*, 1–23.

Goodman, L.M., "Tezos: A Self-Amending Crypto-Ledger Position Paper," August 2014. `https://tezos.com/static/position_paper-841a0a56b573afb28da16f6650152fb4.pdf`.

Greenspan, Mati, "New Bitcoin Abomination," `https://hacked.com/new-bitcoin-abomination/` December 2017.

Griva, Krina and Nikolaos Vettas, "Price competition in a differentiated products duopoloy under network effects," *Information Economics and Policy*, 2011, *23*, 85–97.

Hotelling, Harold, "Stability in Competition," *Economic Journal*, 1929, *39*, 41–57.

Katz, Michael L. and Carl Shapiro, "Network externalities, competition, and compatibility," *The American Economic Review*, 1985, *75* (3), 424–440.

Kurokawa, Kay, "Left/Right Political Ideologies of Hodlers: Part 1," `https://hackernoon.com/left-right-political-ideologies-of-hodlers-part-1-faa71ed689e4` February 2018.

Lee, Woojin, "Bandwagon, underdog, and politcal competition: the uni-dimensional case," *Social Choice and Welfare*, 2011, *36* (3), 423–449.

Lombrozo, Eric, "Forks, Signaling, and Activation," `https://medium.com/@elombrozo/forks-signaling-and-activation-d60b6abda49a` June 2017.

Moos, Mitchell, "Ethereum Muir Glacier Hard Fork Complete, Difficulty Bomb Delayed," `https://cryptobriefing.com/ethereum-muir-glacier-hard-fork-complete-difficulty-bomb-delayed/` Jan 2020.

Myerson, Roger B. and Robert J. Weber, "A theory of Voting Equilibria," *The American Political Science Review*, 1993, *87* (1), 102–114.

Navon, Ami, Oz Shy, and Jacques-Francois Thisse, "Product Differentiation in the Presence of Positive and Negative Network Effects," *CEPR Discussion Paper: 1306*, 1995.

Osborne, Martin J. and Al Slivinski, "A Model of Political Competition with Citizen-Candidates," *The Quarterly Journal of Economics*, 1996, *111*, 65–96.

Osborne, Martin J. and Rabee Tourky, "Party Formation in Single-Issue Politics," *Journal of the European Economic Association*, 2008, *6* (5), 974–1005.

Peters, Hans, Marc Schröder, and Dries Vermeulen, "Hotelling's location model with negative network externalities," *International Journal of Game Theory*, 2018, *47*, 811–837.

Redman, Jamie, "A True Network or Troll? A Look at the "Bitcoin Clashic" Project," `https://news.bitcoin.com/a-true-network-or-troll-a-look-at-the-bitcoin-clashic-project/` November 2017.

Samuelson, William and Richard Zeckhauser, "Status quo bias in decision making," *Journal of Risk and Uncertainty*, 1988, *1*, 7–59.

Schumann, Turner, "Consensus Mechanisms Explained: PoW vs. PoS," `https://hackernoon.com/consensus-mechanisms-explained-pow-vs-pos-89951c66ae10` April 2018.

Sedgwick, Kai, "Why Governance is the Greatest Problem for Blockchains To Solve," `https://news.bitcoin.com/why-governance-is-the-greatest-problem-that-blockchains-must-solve/` July 2018.

Telford, Taylor, "Why governments around the world are afraid of Libra, Facebook's cryptocurrency," `https://www.washingtonpost.com/business/2019/07/12/why-governments-around-world-are-afraid-libra-facebooks-cryptocurrency` July 2019.

Tual, Stephan, "Ethereum Protocol Update 1," `https://blog.ethereum.org/2015/08/04/ethereum-protocol-update-1/` Aug 2015.

Zamfir, Vlad, "Against On-Chain Governance," `https://medium.com/@Vlad_Zamfir/against-on-chain-governance-a4ceacd040ca` December 2017.