

Chapter I

Achieving Consensus on Blockchains

Achieving Consensus on Blockchains*

Zahra Ebrahimi[†]

Maxi Guennewig[‡]

Bryan Routledge[†]

Ariel Zetlin-Jones[†]

November 7, 2025

Abstract

Blockchains enable self-interested users to maintain a distributed ledger without relying on a trusted third party. We develop a dynamic game-theoretic framework in which users strategically decide how to update the ledger. Consensus—agreement on a single, correct ledger—is desirable but arises (or fails) as an equilibrium outcome. Within this framework, we analyze Nakamoto’s (2008) *longest chain rule*, the core consensus protocol underlying Bitcoin. We show that it fails to achieve consensus when users are sufficiently heterogeneous, encompassing previously identified double-spending attacks and revealing new incentive-based failures. We then present modified equilibrium strategies that relax double spending incentives and support consensus as an equilibrium outcome. Finally, we show that no strategy that fully incorporates all available information can achieve consensus, establishing a blockchain analog of the Grossman–Stiglitz paradox.

Keywords: Blockchain, consensus, double-spending, information.

*Previously circulated as "Getting Blockchain Incentives Right." First version: February 2020. We thank Yackolley Amoussou-Guenou, Catherine Casamatta, Jacob Leshno, Dmitry Orlov, Shengxing Zhang, and seminar participants at the the EFA 2025, the 2025 Summer Workshop on Money, Banking, Payments and Finance; the Virtual Finance Theory Seminar, The Joint Renmin University, Hong Kong Baptist University, and National Taiwan University Virtual Seminar, the 2nd Tokenomics Conference, the Madison Money Workshop (2021), SED Meetings (Minneapolis), Michigan State University, and the University of Calgary for insightful comments and discussions. Zahra Ebrahimi, Bryan Routledge and Ariel Zetlin-Jones are grateful for the financial support of the PNC Center for Financial Services Innovation at Carnegie Mellon University and The Ripple Foundation. Maxi Guennewig gratefully acknowledges support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through the CRC TR 224 (Project C03).

[†]Tepper School of Business, Carnegie Mellon University.

[‡]Department of Economics, University of Bonn.

1 Introduction

Blockchains are decentralized, distributed ledgers. They are records of (sequenced) data, possibly transactions as with bank ledgers, maintained by a disperse group of self-interested individuals or users. Unlike ledgers maintained by banks, governments, or other parties, in a blockchain setting, there is no single party responsible for maintenance and security of the ledger, nor is there a single party to resolve conflicts in the ledgers held by distinct individuals. For such decentralized, distributed ledgers to be socially valuable, the dispersed group of individuals who each maintain their own record of the blockchain’s data must agree on the “correct” version of the ledger. In other words, establishing and maintaining *consensus* among the individuals who maintain the data is paramount for blockchain systems.

A long tradition from theoretical computer science studies various consensus protocols and their security properties (see Lynch (1996) for a textbook treatment of such protocols). The analysis of consensus protocols inspired by this literature typically proceeds by positing strong assumptions on the nature of agents’ strategies, effectively imposing certain behavioral types. So-called honest agents have singleton strategy sets: they are required to follow the consensus protocol as proposed by the protocol designer. In the face of these honest types, models typically also include malicious or Byzantine agents who may adopt arbitrary strategies (with potentially unlimited costs) to disrupt the nature of consensus among the honest types. Classical results on consensus protocols explore the extent to which these protocols can achieve specific security properties under varying assumptions on the relative mix of honest and malicious agents.¹

Nakamoto (2008) famously proposed a novel protocol that powers the Bitcoin blockchain. This protocol makes no underlying assumptions about the mix of honest or malicious types. Indeed, Nakamoto effectively treats protocol participants as “rational” and free to choose any record validation process they find individually optimal. In practice, the Bitcoin consensus protocol has been remarkably secure since its inception.

However, almost two decades after Nakamoto first published the Bitcoin protocol, we have a limited understanding of why the Bitcoin protocol is secure in the absence of guarantees on the number of honest miners in the system. What is well understood are the limitations of the specific strategy proposed in Nakamoto (2008). In Bitcoin, the most widely studied blockchain, the proposed consensus protocol (or strategy) is for users to agree that the longest

¹A leading example is Fischer et al. (1985) who show that a single faulty node makes it impossible to reach consensus among deterministic asynchronous processes.

chain—technically, the chain that represents the most computational work—is the correct chain. For example, Biais et al. (2019) and Budish (2025) have both shown that if the value of modifying the data on the blockchain is sufficiently large or the users’ ability to write data to the blockchain is not evenly distributed, then the longest chain consensus protocol is not sufficient to prevent users from modifying past data. Such critiques call into question the economic viability and security of blockchain-based ledgers.²

Our paper’s main contribution is to develop a new theoretical model of blockchain consensus that permits us to study the robustness of protocols—strategies for appending data to blockchain ledgers—that generate consensus. The key innovation in our model relative to the nascent literature on consensus with rational agents is that agents in our model have direct preferences over information in different versions of the blockchain ledger. This innovation allows us to derive the relevant incentive constraints for any candidate protocol as opposed to only that proposed in Nakamoto (2008). We use this new model to highlight how small modifications to agents’ behavior may significantly enhance the robustness of consensus as an equilibrium outcome even when an individual agent may have an outside ability to append data to the blockchain (e.g. concentrated mining power under Proof-of-Work or concentrated stake under Proof-of-Stake). In particular, we derive modifications of Nakamoto’s longest chain strategy that resolve the double-spend attacks raised as a critical technological fault underlying the Bitcoin protocol by Budish (2025).

We begin by proposing a dynamic model of decentralized, distributed record-keeping. We interpret the records being kept in our model as records of transactions involving agents’ balances of a unit of account—as in the Bitcoin blockchain. However, we may also interpret this data as computational code intended to be conducted by the network of validators as well—as with smart contracts in the Ethereum blockchain. In each period, rational, self-interested agents or validators, append a new block of transactions to a particular existing block, or location, in the blockchain. We refer to validators as miners as they are typically called in the Bitcoin protocol. Given this “locational” choice, the likelihood a miner’s block is added to the blockchain depends on her (exogenous) mining power, which we model as a probability. If a miner’s block of data is added, the block not only includes transactions but also a mining reward for that miner—the reward is an increment to the miner’s balance of a unit of account on the blockchain. This process of mining implies that in any period, the blockchain resembles a graph (or tree) of blocks. Each possible path or *chain* of blocks from

²See also Eyal and Sirer (2014) and more recently Bahrani and Weinberg (2024) for similar critiques from the computer science literature on consensus protocols.

the first, “genesis” block to any other block in the tree of blocks, which may be thought of as a fork from any chain in the tree, represents a distinct ledger with a possibly distinct value for each miners’ aggregate balances.

We assume that miners value these ledger balances because positive net balances are in theory spendable for physical goods, services, or other currencies. However, these balances are only spendable if other miners agree that the balances are valid. To formalize a notion of agreement, we assume miners “vote with their feet” or, more aptly, their mining power. That is, we assume a given miner values balances on ledgers being mined by other miners more than balances on ledgers that are not being mined by other miners. Miners then have a direct preference for agreement or consensus. Whether consensus is achieved, however, is an equilibrium outcome.

To capture the idea that balances are spendable, we consider transactions involving negative balances, or spend transactions, as corresponding to some form of consumption off the blockchain. A spend transaction, once added to the blockchain graph, persists in the graph in perpetuity. We assume that the miner is compensated for this spend by the one-time receipt of “goods” from an external party. To the extent settlement of goods is delayed (as it is in practice with individuals typically waiting for at least six additional confirmed blocks before settlement in the case of the Bitcoin protocol), the settlement transaction may also involve a premium paid to the miner who endures the cost of the negative transaction once it is mined before settlement occurs.

In this environment, we evaluate different strategies that may generate consensus on ledgers. We first show that Nakamoto’s proposed longest chain strategy features the same flaws in our model as found in earlier work. First, miners may have incentives to remove “spend” transactions. Once a miner spends balances in order to obtain physical goods or other services, they have incentives to work on chains that omit this spend transaction. To the extent other miners are assumed to follow the longest chain protocol, it is straightforward to develop conditions where so-called double-spending attacks are profitable for a rational miner. The reason, as identified in both Biais et al. (2019) and Budish (2025), is that the costs of a double-spend attack do not scale with the potential benefits.

Our more general model also reveals additional reasons miners may deviate from Nakamoto’s proposal: miners may have incentives to work on ledgers that are not the longest chain but involve large transactions that increase their balances in the hopes that these ledgers become the longest chain. We describe these incentives as consensus re-direction attacks. Additionally, miners may have incentives to work on ledgers that are not the longest chain but involve

large transactions that decrease their balances in the hopes that these ledgers do not immediately become the consensus chain but do so eventually. The incentive to induce a conflicting ledger to control when spend transactions become part of the longest chain we describe as saving on consensus attacks.

Our second contribution is to use our model to identify modifications to proposed strategies that relax miners’ incentives to deviate in these situations. To address the first concern—removing spend transactions—a modified version of the longest-chain strategy featuring “checkpoints” is an equilibrium for a much larger range of concentrated mining power and spend transaction sizes. The checkpoint rule embedded in miners’ strategies, introduces history dependence. Miners’ strategies call on them to ignore any blockchain forks that omit too much old information—that is, that omit blocks behind the current checkpoint. These checkpoints ensure that if any miner attempts to omit blocks behind the current consensus checkpoint, then no miners will treat this new deviation chain as the correct chain regardless of its length. We show that for some checkpoint rules as miners’ discount factors tend to 1 that consensus is as an equilibrium outcome even when a single miner may have substantial mining power and the chain contains arbitrarily large transactions.

Checkpoints may then arise as a strategic solution to double-spending attacks, the central problem identified in Budish (2025), suggesting that the concerns raised in previous work are features of particular equilibrium strategies rather than limitations of blockchain technology itself. The result that small changes in behavior makes consensus a much more robust equilibrium outcome also offers a novel explanation for why attacks previously identified in the literature as profitable when honest agents use Nakamoto’s proposed strategy have not been observed. Moreover, variants of the checkpoint rule are already implemented in practice—for example, on the Ethereum and Polkadot blockchains (Buterin and Griffith, 2017; Buterin et al., 2020; Stewart and Kokoris-Kogia, 2020) as well as more broadly across other systems (Xu et al., 2023).

An important consideration associated with introducing history dependence via checkpoints in blockchain strategies is the risk of network latency. Latency creates the potential for disagreement on histories. Furthermore, since the entire network of miners does not see new blocks at the same time it is likely that accidental forks will occur. If latency causes disagreement on the history of the chain, then latency may cause disagreement on the checkpoint as well causing risk to equilibrium consensus.

We study the interaction between history-dependent checkpoint strategies and latency risk by allowing for unexpected shocks to communication among the set of miners. Specif-

ically, we assume that in some period, miners may be unexpectedly partitioned into two disjoint sets and one miner in each group will be selected to append a block of new data. This notion of latency resolves at the end of the subsequent period when miners observe all blocks added by each group of miners and the histories they observed.

We show how a particularly stark checkpoint rule that resolves all double-spend attacks in our model without latency—a strategy where in each period, miners select the last block added to the previous checkpoint as the new checkpoint—is not robust to our notion of latency. Specifically, once latency resolves, miners have incentives to deviate from a rule that calls on them to ignore any blockchain forks that begin behind *their own* current checkpoint since rewards from newly mined blocks are more valuable on forks with more miners (or mining power). Ex-post disagreement on checkpoints creates incentives for miners to abandon their own checkpoint.

Instead, we consider an alternative checkpoint strategy where the checkpoint lags behind the most recently appended block. If this lag is sufficient so that the checkpoint block is not a block added during the latency period when miners’ communication sets are partitioned, then miners agree on the checkpoint once latency resolves. In this sense, despite the fact that checkpoints introduce history dependence, lagged checkpoints are robust to latency windows (of finite, known length). These results suggest that delayed settlement, as with the six block confirmation lag in Bitcoin, and checkpoints should be attuned to technological features such as latency windows rather than to resolve potential incentive constraints.

However, when checkpoint lags are long enough to absorb severe latency shocks, opportunities emerge for miners to redirect or save on consensus. We demonstrate that any equilibrium strategy that eliminates these attacks on consensus introduces a fundamental conflict: a tradeoff between achieving consensus and fully reflecting the most recent information appended to the blockchain.

We explain the trade-off using a general class of strategies that we call “full-information checkpoint rules.” Recall that miners value ledgers based on the balances in those ledgers and the extent of consensus (how many other miners are building on a given ledger). Full-information checkpoint rules are maximally informationally sensitive: they select the mining location as a function of the balances based on all blocks on the graph beyond the current checkpoint, including the latest block appended. For example, one strategy from this class selects the block associated with the highest aggregate ledger balances across miners, which has a natural interpretation as maximizing welfare. For these strategies, however, even a small change in the latest block appended would typically induce a large, sudden shift in the

consensus location and, therefore, potentially induce a large shift in an individual miner’s payoff. These shifts create strong incentives for miners to deviate, and we show that no such strategy is a perfect equilibrium.

In contrast, strategies that do not fully reflect the most recent information appended to the blockchain can successfully achieve consensus (for any history) and do prevent redirection and saving on consensus attacks. These “partial-information checkpoint rules” work by introducing a form of strategic inertia—they do not change mining locations in response to each new block appended to the chain. While the partial-information checkpoint rule successfully achieves consensus in equilibrium, the resulting allocations often incur welfare losses compared to maximally informationally sensitive strategies. In some cases, the ledger that becomes consensus is Pareto dominated by another ledger.

Our result therefore shows that a perfectly information-sensitive consensus mechanism destroys the incentive for miners to maintain such consensus. We view this necessary trade-off between conditioning consensus on all information and achieving consensus itself as a blockchain analog of the Grossman and Stiglitz (1980) Paradox.

Literature review. We develop a dynamic, game-theoretic model of blockchain consensus in which rational, self-interested agents have well-defined preferences over the data recorded on the blockchain and strategically decide how to update it. Our paper is therefore most closely related to papers on the economics of blockchain consensus and security.

Biais et al. (2019) and Budish (2025) also present game-theoretic models of blockchain environments, showing that blockchains are susceptible to double-spending attacks if mining power is concentrated and transaction values are large. In their analyses, honest miners follow the longest chain rule and therefore switch to an attacker’s fork once it becomes the longest chain. The expected cost of a successful attack is given by the expected cost of creating a fork and extending it to surpass the original chain—it is therefore fixed conditional on the honest miners’ computing power. The expected benefit of an attack is increasing in the size of the spend transaction. Budish (2025) argues that, since the honest miners’ computing power is increasing in the cost of using the blockchain (by a free entry logic), mining rewards must scale with transaction size on the blockchain.^{3,4}

³Biais et al. (2019) further highlight that consensus can be fragile: miners can coordinate on creating forks if mining strategies exhibit strategic complementarities.

⁴Other papers take honest miners’ strategies as given. Gans and Halaburda (2023) generalize and extend the analysis of the majority attack. They find that the cost of an attack may be lower when honest miners’ endogenously response by adjusting their computer power. Saleh (2021) studies Proof-of-Stake consensus protocols and finds that attacking the blockchain is not profitable if the market capitalization of the blockchain-native coin is sufficiently large and the settlement lag sufficiently long. John et al. (2025)

Our contribution beyond our framework is to demonstrate that incorporating a simple history-dependence in the form of checkpoints into the longest chain rule can effectively prevent double-spending attacks. The profitability of such attacks is therefore a feature of a particular consensus protocol (or mining strategy) and not a concern about the technology of blockchain itself. Our general model enables a deeper analysis of the incentive structures underlying different consensus mechanisms and allows us to identify additional reasons why consensus may fail. We highlight a fundamental trade-off between conditioning consensus on all available information and achieving consensus.

Garratt and van Oordt (2023) argue that a double-spending attack may not be profitable if the hardware required is specialized and cannot be repurposed, raising the fixed cost of the attack. Moroz et al. (2020) show that when the victim of a double-spending attack can counterattack in the same way as the attacker, then this results in a variant of the ‘War of Attrition’ game. The threat of a counterattack induces a subgame perfect equilibrium of this game in which no attack occurs in the first place. Chiu and Koepl (2022) argue that more intensive miner competition (i.e., more widely distributed mining power) and long settlement lags, which increase in transaction size, can help render double-spending attacks unprofitable as the cost to create a competing, longer chain becomes excessive. We highlight that neither large fixed costs, the ability to counterattack, nor extensive settlement lags are necessary to prevent double-spending attacks.⁵

Our paper shares its objectives with Halaburda et al. (2022). They also develop a game-theoretic model of blockchain, although with reduced form payoffs, emphasizing Knightian uncertainty to capture the spirit of Byzantine Fault Tolerance in the computer science literature. More recently, Leshno et al. (2024) highlight the ‘community response’ to override nodes with a corrupted ledger and present a protocol which formalizes this feature. One as-

study security properties of Proof-of-Work and Proof-of-Stake consensus protocols when blockchain capacity constraints are alleviated.

⁵Pagnotta (2022) studies equilibrium multiplicities that arise if blockchain security depends on the real value of blockchain-native coins. Makarov and Schoar (2021) study the Bitcoin blockchain ecosystem. One of their findings is that mining power is highly concentrated among a small number of mining pools. Cong et al. (2021) show that the rise of centralized mining pools does not necessarily undermine decentralization, which, as frequently argued, needs to be sufficiently high for blockchains to be secure. Auer et al. (2021), Amoussou-Guenou et al. (2024), and Benhaim et al. (2023) study consensus on committee-based or permissioned blockchains. Bakos and Halaburda (2021) compare the security properties of permissioned and permissionless blockchains. Similar to Budish (2025), attacks on permissionless blockchains are profitable if the value of an attack is large relative to block rewards. See, among others, Li (2023) for a study on the security of blockchain scaling solutions. Kang (2023) studies a reputation-based mechanism to address double-spending attacks. Merchants delay the delivery of consumption goods if the payment is done using a wallet which has been found to have double-spent in the past.

pect of the protocol is that nodes finalize transactions once they are certain that the ledger has not been corrupted, a feature which bears similarities with the checkpoint equilibrium described in this paper.⁶ Our contribution relative to these papers lies in our framework, which allows us to study how the specific data recorded on the blockchain, such as transactions, impact incentives of rational nodes across different consensus protocols (or strategies). We then present strategies which prevent double-spending attacks, and characterize a fundamental trade-off between consensus and information efficiency.

Outline. Section 2 introduces the framework. Sections 3 and 4 analyze the longest chain rule and checkpoint strategies, respectively. We study latency in Section 5 and highlight the trade-off between consensus and information efficiency in Section 6. Section 7 concludes. All proofs are in the appendix.

2 A Model of Blockchain

In this section, we develop a model to analyze blockchain consensus. In this model, in each period, miners add a block of data to an existing graph of blockchain data. A block includes units of account as well as, in principle, other data. The baseline model features no latency in the sense that each individual perfectly observes each addition to the blockchain.

Preliminaries. There are $N \in \mathbb{N}$ miners, each infinitely lived and with a rate of time preference $\delta \in (0, 1)$. Time is discrete. In each period t , each miner i proposes a location to add a *block*, $b_{i,t}$, of data. A block consists of three components: hash data, mining rewards, and data entries, e.g. transaction data (as in the Bitcoin blockchain) or computations to be conducted (as in the Ethereum blockchain).⁷ The hash data is determined technologically and is not relevant for our model beyond the fact that it implies a chained data structure. We let $R_{j,b_{i,t}}$ denote the mining rewards in block $b_{i,t}$ for miner j . We assume that mining rewards have the property that $R_{j,b_{i,t}} = \bar{R} \in (0, \infty)$ if $i = j$ and $R_{j,b_{i,t}} = 0$ for $j \neq i$. This implies that only miner i earns a reward if block $b_{i,t}$ is added to the blockchain. In addition, we represent the data for each miner j in any block $b_{i,t}$ proposed by miner i by $Y_{j,b_{i,t}} \in \mathbb{R}$. We assume that the data in a given time period t are exogenous and identical across all miners'

⁶Other papers follow the approach typically taken in the computer science literature to study protocols involving checkpoints, e.g. Buterin and Griffith (2017) and Neu et al. (2021). Karakostas and Kiayias (2021) consider a consensus protocol with external checkpoints and discuss how to decentralize the checkpointing process. Sankagiri et al. (2021) develop a protocol which incorporates checkpoints into the longest chain rule.

⁷Technically, mining rewards are simply a transaction, but it is useful for us to separate them.

blocks, and hence we write $Y_{j,b_{i,t}} = Y_{j,b_t}$. Miners' blocks therefore only differ in terms of block rewards.

A blockchain, in the language of graph theory, is an arborescence. It is a directed graph in which from the genesis block b_0 to any other block b there is exactly one directed path from b_0 to b . Let $\mathcal{B}(G_t)$ denote the set of all blocks in the graph G_t . Let (b', b) denote the edge from block b to block b' , leading away from the genesis block. Denote by $\mathcal{E}(G_t)$ the set of all edges that link the blocks in graph G_t . Let \mathcal{G}_t represent the set of all possible graphs with t blocks and $\mathcal{G} = \bigcup_{t=0}^{\infty} \mathcal{G}_t$. Let $H_t \in \mathcal{H}_t = \bigcup_{\tau=0}^t \mathcal{G}_\tau$ denote the history of the graph at time t , and \mathcal{H}_t the set of all possible histories at time t .

Each miner's action in period t is to choose a location to attempt to add block $b_{i,t}$. A location choice of miner i in period t is a mapping $a_{i,t} : G_t \rightarrow \mathcal{B}(G_t)$. We provide the following definition for *consensus* based on these location choices: Consensus is achieved if $a_{i,t} = a_{j,t}$ for every $(i, j) \in \{1, \dots, N\}^2$ and after every history H_t .

Miners' location choices stochastically determine the state of the graph in the subsequent period. Specifically, we assume that each miner's block is added (in the location of choice chosen by miner i) probabilistically with at most one miner adding a block in a given period. Let $p_i \in (0, 1)$ denote the probability that miner i successfully adds a block to the existing graph with $\sum_{i=1}^N p_i = 1$. This probability represents the mining power of miner i and we treat it as exogenous.

Given a graph G_t and the location choices of miners $(a_{i,t})_{i=1}^N$, the graph in the subsequent period is $G_{t+1} = G_t \cup (b_{i,t}, (b_{i,t}, a_{i,t}))$ with probability p_i . In words, the graph G_{t+1} is the same as the graph G_t but includes a new node $b_{i,t}$ and a new edge from $b_{i,t}$ to $a_{i,t}$.

Chains. Before turning to the structure of preferences and payoffs, it is useful to create notation to describe the various databases represented in a graph, G_t . We interpret each path through the graph, from the origin node to any other node, as a *chain*. Note that each chain may represent a different database than any other chain. Furthermore, recall that the blockchain protocol imposes that every block has a unique parent block (although it may have more child blocks). Hence, the path from any block to the genesis block is unique.

For any graph G_t and block $b \in \mathcal{B}(G_t)$, define the chain $C(b, G_t)$ as the unique path from block b back to the genesis block b_0 . Let $\mathcal{C}(b, G_t) \subseteq \mathcal{B}(G_t)$ denote the set which contains the blocks on the chain from b to b_0 .⁸ We say that block b_n is on the chain $C(b, G_t)$ if $b_n \in \mathcal{C}(b, G_t)$. Furthermore, define $\#C(b, G_t)$ as the number of blocks in the chain. We refer to this number as the *length* of the chain. Finally, we refer to blocks with only one edge as

⁸Formally, we write $\mathcal{C}(b, G_t) = \{\{b, b_n, \dots, b_1, b_0\} \in \mathcal{B}(G_t) \mid (b, b_n), (b_n, b_{n-1}), \dots, (b_1, b_0) \in \mathcal{E}(G_t)\}$.

terminal blocks, and define $\mathcal{T}(G_t)$ as the set of terminal blocks in graph G_t .

Preferences. We now propose a specific functional form for the period payoff that has two components: first, a flow payoff derived from the data contained on the blockchain; and second, a flow payoff derived from consumption of real goods.

We assume that miners derive a linear flow utility from the weighted sum of their data entries on the blockchain, given by $(1 - \delta) \sum_{b \in \mathcal{B}(G_t)} q_{i,b,t} (Y_{i,b} + R_{i,b})$, where $q_{i,b,t}$ denotes the weight of block b for miner i at time t . We link these weights to miners' actions and computing power: if a miner chooses a location in $\mathcal{B}(G_t)$, we say that the miner works on the chain from the origin to that existing block. If more miners work on the same chain, then the data on that chain have a larger weight. In particular, we assume that blocks are weighted according to the other miners' computational mining power allocated to those blocks: $q_{i,b,t} = \sum_{\{j \neq i: b \in \mathcal{C}(a_{j,t}, G_t)\}} p_j / (1 - p_i)$. This expression captures the notion that data are more valuable if they are written in blocks on which more (other) miners agree. Miner i receives value for any data written in blocks that are on the blockchain associated with some other miner's location choice $a_{j,t}$. To the extent there is disagreement, miners obtain value from their data as long as some other miners apply their mining power to these blocks. When there is full consensus and all miners choose the same location, then all data entries in blocks on that chain receive their full value of 1.

It is natural to think of the data Y and R as representing units of account held on the graph G_t . These transaction data could be positive or negative with the interpretation that positive data represent payments received while negative data represent payments sent. Going forward, we refer to these units of account as *coin balances*. Then $Y_{i,b} + R_{i,b}$ represents miner i 's coin balances in block b . The miner's flow utility from her coin balances in this block increases as more miners recognize the block as valid.

Given this interpretation, the second component of miners' period payoffs explicitly links consumption to negative data entries, or *spend transactions*. We assume consumption goods are fairly priced and that settlement occurs with a delay, reflecting existing blockchain norms. For instance, Bitcoin recommends finalizing spend transactions after waiting for six additional confirmed blocks (i.e., six blocks appended on a single chain which achieves full consensus). For simplicity, we assume a one-block delay (on a single chain). Miners derive linear flow utility from consumption. Fair pricing then requires that each unit of coin balances purchases $1/\delta$ units of consumption. Let the absolute value of spend transactions be denoted by $Y_{i,b}^- = |Y_{i,b}| \cdot \mathbb{1}\{Y_{i,b} < 0\}$. Consumption as well as the flow utility derived from consumption at time t are then given by $\sum_{b \in \mathcal{B}(G_t)} \frac{Y_{i,b}^-}{\delta} \cdot \lambda_t(b, H_t)$, where $\lambda_t(b, H_t)$ is an indicator that takes

the value of one if settlement takes place in time period t . We provide the precise definitions of $\lambda_t(b, H_t)$ in Sections 3 and 4 below.

Miner i 's flow utility at time t can then be represented by

$$u^i(\mathbf{q}_{i,t}; H_t) = \sum_{b \in \mathcal{B}(G_t)} \left[(1 - \delta)q_{i,b,t} (Y_{i,b} + R_{i,b}) + \frac{Y_{i,b}^-}{\delta} \cdot \lambda_t(b, H_t) \right], \quad (1)$$

where $\mathbf{q}_{i,t} = (q_{i,b,t})_{b \in \mathcal{B}(G_t)}$. Note that in this formulation, preferences are a function of the value of coin balances (given the current actions of miners) and the history of the graph.⁹

As an example, consider miner m called ‘Satoshi.’ Suppose Satoshi has 1 unit of account on the genesis block $R_{m,b_0} = 1$. Satoshi also has a spend transaction $Y_{m,b_1} = -1$ in the second block b_1 , which is added to the blockchain at the end of period 1. Suppose further that there is a single chain in the graph in every period. In period 2, Satoshi’s flow utility over the graph is 0 because her balance aggregated over the two blocks is zero. Unless future blocks contain more transactions for Satoshi, her utility over the graph (excluding consumption) will continue to be 0 in all future periods. However, when a third block is appended to block b_1 at the end of period 2, Satoshi’s spend transaction vests. She then earns the consumption flow utility equal to $1/\delta$ in period 3. Aggregating and discounting these payoffs from the perspective of period 1 or of period 2, her lifetime utility is 1. Of course, once Satoshi has derived the consumption flow utility in period 3, her discounted lifetime utility is 0. She would then benefit from the construction of an alternative path through the blockchain where her total balance on the consensus chain is 1.

Strategies and Equilibrium. Since miners take transactions as given, they only choose the location of their new block. We focus on *public strategies*, which only depend on the publicly observed sequence of graphs. Formally, a public strategy for miner i , σ_i , is a sequence of mappings from the set of all possible public histories into a set of pure actions, $\sigma_i = (\sigma_{i,t})_{t=0}^\infty$, where $\sigma_{i,t} : \mathcal{H}_t \rightarrow \mathcal{B}(G_t)$. Our equilibrium concept is *perfect public equilibrium*, that is, subgame perfect equilibrium in public strategies. We therefore insist that a strategy profile $\sigma = (\sigma_i)_{i=1}^N$ is an equilibrium if and only if each miner’s strategy is a best response to

⁹Our specification of preferences can be microfounded within a monetarist framework (Lagos and Wright, 2005; Rocheteau and Wright, 2005). In this setting, money trades at a premium since money (or coin) balances represent past “work,” creating an option value tied to exchanging goods for money in frictional goods markets. In our framework, this premium is reflected in the flow utility derived from coin balances, and it is larger the more miners consider these balances as valid. Meanwhile, the consumption flow utility reflects the utility derived from realized consumption purchases in both frictional and frictionless goods markets using coin balances.

other miners’ strategies for each history $H_t \in \mathcal{H}_t$ at each date $t \geq 0$.

Due to their distributed nature, blockchain databases occasionally generate conflicting chains accidentally. (For example, it is well understood that with Bitcoin, two miners may occasionally find a valid block at roughly the same time generating an accidental fork from the perspective of other miners.) For this reason, we view the robustness of strategies that are subgame perfect as an important feature of equilibrium analysis of blockchains.¹⁰

An advantage of studying public perfect equilibria with discounting in our environment is that we may apply and use the one-shot deviation principle. The literature on Nakamoto consensus has routinely studied complex, multi-period deviations and the incentives individuals miners may have to pursue these (see Carlsten et al. (2016) and Eyal and Sirer (2018) for leading examples). The one-shot deviation principle allows us to study these complex strategies as one-shot deviations from particular subgames. But more powerfully, we need not worry about more complex deviation strategies once we construct a strategy that is immune to one-shot deviations from all histories.

3 Longest Chain Rule

In this section, we analyze Bitcoin’s proposed equilibrium strategy, the *longest chain rule* (Nakamoto, 2008). After defining transaction settlement and miners’ strategies under the longest chain rule, we present the general conditions under which it can be sustained as an equilibrium. To highlight the limitations of this strategy, we then turn to three special cases that reveal when and why the longest chain rule breaks down as an equilibrium.

The gist of the longest chain rule is that miners choose the block that defines the longest chain as the predecessor for their potential block. This is a simple coordination mechanism in that it depends only on the current graph G_t . Let $\mathcal{B}^{LC}(G_t) = \operatorname{argmax}_{b \in \mathcal{B}(G_t)} \#C(b, G_t)$ denote the set of (terminal) blocks in the graph G_t such that the chain to these blocks has the largest number of blocks.

Settlement. We assume that merchants deliver consumption goods that correspond to spend transactions contained in block b at time t (and hence $\lambda_t(b, H_t) = 1$) if two conditions are met. First, at least one block has been appended to b . Second, b is contained in the

¹⁰Our focus on public equilibria is natural given the assumptions we have made that mining locations are public information. In practice, at least for short periods of time, miners may be able to hide their mining activity. In such a case, one would want to also permit private actions and study equilibria with private monitoring.

unique longest chain for the first time.¹¹ Let $\Lambda_t(b, H_t) = \prod_{s=0}^t (1 - \lambda_s(b, H_s))$ denote an indicator function that takes the value of 1 if a spend transaction has not vested yet, and 0 otherwise.

Strategies. We now formalize the notion that the longest chain rule calls for miners to extend the longest chain. If $\mathcal{B}^{LC}(G_t)$ is a singleton, miners are called choose the only block in this set as predecessor. If the graph features multiple longest chains so that $\mathcal{B}^{LC}(G_t)$ is not a singleton, the necessary tie-breaking rule for the longest chain rule to be a candidate equilibrium strategy is intuitive and prescribes miners to choose the terminal block on their most preferred longest chain:

Lemma 1. *If the longest chain rule is a (perfect public) equilibrium, then for any graph G_t such that $\mathcal{B}^{LC}(G_t)$ is not a singleton, the longest chain rule must satisfy*

$$\sigma_{i,t}^{LC}(H_t) = b_{i,t}^* \equiv \operatorname{argmax}_{b \in \mathcal{B}^{LC}(G_t)} \sum_{b' \in \mathcal{C}(b, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b, H_t) \right). \quad (2)$$

Intuitively, starting from a graph with multiple longest chains and only positive transactions (and block rewards), if the miner has strictly higher coin balances on one of the longest chains, than by mining in that location she strictly increases the likelihood that this chain becomes the single longest chain and thus the consensus chain in next period. She then earns the flow utility associated with her balances in perpetuity. A similar logic applies to unvested spend transactions. Once a block is added to one of the longest chains and consensus is achieved, the spend transactions on that chain vest immediately and generate a utility of $1/\delta > 1$ per unit of spending. Miners therefore earn a net benefit when spend transactions vest. Hence, miners most preferred longest chain contains the largest sum of positive transactions, block rewards, and the net benefit from unvested spend transactions.

Equilibrium. Define a set of blocks, $\mathcal{B}^{-1}(G_t)$, for which one of two things is true. First, a block is a terminal block on a fork which is one block shorter than the longest chain. Or second, a block is the parent block of the terminal block of a longest chain.¹²

Under the tie-breaking rule implied by (2), the only relevant one-shot deviations are those to some block $b \in \mathcal{B}^{-1}(G_t)$. For any other deviation, if the miner successfully adds her block and then reverts to the candidate equilibrium strategy, she immediately abandons her

¹¹Formally, $\lambda_t(b, H_t) = 1$ if $t = \inf \{ \tau \geq 0 : \exists b' \neq b \text{ s.t. } b \in \mathcal{C}(b', G_t), b' \in \mathcal{B}^{LC}(G_t), |\mathcal{B}^{LC}(G_t)| = 1 \}$, and $\lambda_t(b, H_t) = 0$ otherwise.

¹²Formally, the set is defined as $\mathcal{B}^{-1}(G_t) = \{ b' \in \mathcal{B}(G_t) : \#C(b', G_t) = \max_{b \in \mathcal{B}(G_t)} \#C(b, G_t) - 1 \}$.

block which has not become part of a longest chain. Since no other miner is working on the fork either, she thus forgoes the opportunity to have earned the rewards and transactions associated with mining that block to a longest chain.

To compare incentives to append a new block to the end of the longest chain or to some block in $\mathcal{B}^{-1}(G_t)$, consider a thought experiment where miner i adds block $b_{i,t}$ to the graph for sure in period t .

If she adds her block to her preferred longest chain, $b_{i,t}^*$, then that chain becomes the single longest chain and thus the consensus chain in the subsequent period. Consensus is achieved from time $t + 1$ onwards. The miner earns the balances in block $b_{i,t}$ as well as all her balances on preferred longest chain in perpetuity, starting in period $t + 1$. Furthermore, she derives the consumption flow utility $Y_{i,b}^-/\delta$ due to unvested spend transactions contained in all blocks b which lie on the chain running to $b_{i,t}^*$ at time $t + 1$. Miner i also derives the consumption flow utility $Y_{i,b_t}^-/\delta$ once any spend transaction in block $b_{i,t}$ vests at time $t + 2$. Thus, miner i enjoys the following *continuation utility* from time $t + 1$ onwards based on the data present on the blockchain at the end of time t (after miner i has added block $b_{i,t}$):

$$U_{t+1}^i(b_{i,t}, b_{i,t}^*; H_t) = Y_{i,b_t} + \bar{R} + \delta \cdot \frac{Y_{i,b_t}^-}{\delta} + \sum_{b' \in \mathcal{C}(b_{i,t}^*, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \Lambda_t(b', H_t) \right). \quad (3)$$

Suppose instead she adds her block to a chain which is one block shorter. This deviation extends a lack of consensus into period $t + 1$ as the number of longest chains in that period increases by one. Of course, since all miners including miner i will then follow the longest chain rule from time $t + 1$ onwards and only one miner will successfully append a block at time $t + 1$, consensus will be achieved from time $t + 2$ onwards.

At time $t + 1$, miner i derives flow utility from the value of balances on each longest chain on which at least one miner $j \neq i$ is working. This value is given by

$$F_{t+1}^i(b_{i,t}, \hat{b}; H_t) = \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} \frac{p_j}{1 - p_i} \cdot \left(Y_{i,b_t} + \bar{R} + \sum_{b' \in \mathcal{C}(\hat{b}, G_t)} (Y_{i,b'} + R_{i,b'}) \right) \quad (4)$$

$$+ \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t+1}^* = b\}} \frac{p_j}{1 - p_i} \cdot \sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}).$$

Recall that the value of miner i 's balances is proportional to the computing power of (other) miners working on each chain. When other miners work to extend miner i 's preferred

longest chain, the block she mined in period t as well as all other transactions and block rewards along that chain have value given by the first line in (4). When other miners work on other chains, transactions in those chains—which necessarily exclude the block miner i added in period t —have value given by the second line in (4).

As discussed above, the miners achieve consensus in period $t + 2$. Miner i 's expectation over the continuation utility derived from time $t + 2$ onwards based on data present in the blockchain at the end of period t is given by

$$\begin{aligned} \mathbb{E}_t \left[U_{t+2}^i \left(b_{i,t}, \hat{b}; H_t \right) \right] &= \left\{ p_i + \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} p_j \right\} \sum_{b' \in \mathcal{C}(\hat{b}, G_t) \cup b_t} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \\ &+ \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t+1}^* = b\}} p_j \cdot \sum_{b' \in \mathcal{C}(b, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \end{aligned} \quad (5)$$

With probability $\left\{ p_i + \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} p_j \right\}$, miner i 's preferred chain is the consensus chain. The miner then enjoys the value of rewards and transactions in the blocks on this chain as well as the value of newly vested spend transactions, reflected in the first line of (5). Instead, each other chain to $b \in \mathcal{B}^{LC}(G_t)$ becomes the new consensus chain with probability $\sum_{\{j \neq i: b_{j,t+1}^* = b\}} p_j$ yielding similar payoffs, captured by the second line of (5).

We are now ready to state our first main result.

Proposition 1. *The longest chain rule is a perfect public equilibrium if and only if for every history H_t , for every block $\hat{b} \in \mathcal{B}^{-1}(G_t)$, and for every miner i :*

$$U_{t+1}^i \left(b_{i,t}, b_{i,t}^*; H_t \right) \geq (1 - \delta) \cdot F_{t+1}^i \left(b_{i,t}, \hat{b}; H_t \right) + \delta \cdot \mathbb{E}_t \left[U_{t+2}^i \left(b_{i,t}, \hat{b}; H_t \right) \right] \quad (6)$$

Miner i only finds it profitable to follow the longest chain rule if the benefit from doing so—turning the preferred longest chain into the consensus chain immediately and vesting the new transactions and block rewards for sure—outweighs the benefit from deviating. That is, it outweighs the benefit associated with creating a new preferred longest chain which yields some flow utility and becomes the consensus chain in the subsequent period with some probability. Note that, since there is consensus both after following the longest chain rule and after a one-shot deviation once a block has been appended at the end of time t , all data in blocks added from time $t + 1$ onwards do not affect incentives at time t .

It is useful to study a simpler case to help understand the condition in (6). Figure 1

depicts the blockchain as graph G_t with one longest chain with terminal block b_{l_2} and one fork, which is one block shorter than the longest chain, with terminal block b_f .




Figure 1: An illustration of the condition in (6).

Suppose miner i is mining some block $b_{i,t}$ such that should miner i append it to b_f , all other miners find it optimal to work on b_{l_2} and thus on the previously longest chain. For simplicity, we maintain this assumption for the remainder of this section. Since we only need to consider one-shot deviations to either b_f or b_{l_1} , all data in the chain leading to b_r are contained in any candidate consensus chain and do not affect incentives. Then, miner i does not face a profitable one-shot-deviation from the longest chain rule to b_f if

$$\begin{aligned}
 & Y_{i,b_t} + \bar{R} + \delta \cdot \frac{Y_{i,b_t}^-}{\delta} \\
 & \geq p_i \delta \cdot \left(\sum_{b' \in \{b_f, b_t\}} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \right) - \left(\sum_{b' \in \{b_{l_1}, b_{l_2}\}} (Y_{i,b'} + R_{i,b'}) + \frac{Y_{i,b_{l_2}}^-}{\delta} \right) \right)
 \end{aligned} \tag{7}$$

The LHS of (7) is the opportunity cost of a deviation. If miner i appends block $b_{i,t}$ to the longest chain, then this block becomes part of consensus with certainty. The block's balances are then written on the ledger in perpetuity, and any spend transaction vests and the corresponding consumption goods are received in the following time period.

The RHS of (7) is the expected benefit of deviating. Recall that no other miner works on block $b_{i,t}$ after it was appended to b_f . Miner i then successfully turns the fork into the unique longest chain and thus the consensus chain in the following period with probability p_i . The balances on the fork including $b_{i,t}$ —rather than the balances on the previously longest chain

leading to b_{l_2} —are then part of consensus in perpetuity. Similarly, spend transactions vest on the fork but they never vest on the previously longest chain. The benefit from deviating conditional on succeeding is given by the difference. With the complementary probability $1 - p_i$, miner i is unsuccessful and the previously longest chain becomes the consensus chain. The benefit from deviating conditional on failing is then zero.

We proceed by explaining the three distinct manners in which the condition in (6) fails, continuing the example of Figure 1 with particular transaction data and block rewards.

Negative transaction data and double-spending attacks. We now illustrate that the longest chain rule is not robust to negative transaction data and is therefore vulnerable to double-spending attacks. Our argument proceeds in two steps. First, we show that the condition in (6) fails if a miner has a large, vested spend transaction on the longest chain and a competing fork exists that is only one block shorter. Since (6) provides a sufficient condition for the longest chain rule to be a perfect public equilibrium, its failure is a necessary condition for a profitable double-spending attack. Otherwise, the attacker never finds it profitable to extend the fork they created during the attack beyond the longest chain. We then construct a double-spending attack using a two-period strategy. In the first period, the attacker tries to create a fork and, if successful, tries to extend it in the second period. We then identify the conditions under which this more complex strategy is a strictly profitable deviation, i.e., under which a double-spending attack is profitable.

Reconsider Figure 1. Suppose that all blocks b_{l_1} , b_{l_2} and b_f have been appended by miner i and thus contain block rewards \bar{R} for her. Further suppose block b_{l_1} contains a spend transaction for miner i : $Y_{i,b_{l_1}} < 0$. Assume that block b_f was appended after block b_{l_2} . Hence, the spend transaction in b_{l_1} has vested and miner i has derived the associated consumption flow utility. Equation (7) then simplifies to $\bar{R} \geq p_i \delta \cdot |Y_{i,b_{l_1}}|$.

In this example, miner i prefers to deviate from the longest chain rule and extend the fork if the spend transaction is sufficiently large relative to the level of block rewards, given her mining power p_i . This is especially true if the longest chain also contains vested spend transactions or the fork contains balances for other miners, who may then also work on the fork in the subsequent period and help miner i establish it as consensus chain. In other words, forking may be profitable even if the inequality above holds. Importantly, the longest chain rule fails in this case because miners may seek to *remove* transaction data from the consensus chain by working on a shorter fork.

We now describe the two-period strategy, $(\sigma'_{i,t-1}, \sigma'_{i,t})$, that amounts to a double-spending attack. Recall that block b_{l_2} was appended before the fork b_f was created in the example

above. Now consider the time period $t - 1$, i.e., the period before b_f was added and the blockchain contains one single chain $(b_r - b_{l_2})$. Suppose miner i is mining block b_f . Consider $\sigma'_{i,t-1}(H_{t-1}) = b_r$, and $\sigma'_{i,t}(H_t) = b_f$ if $b_f \in \mathcal{B}(G_t)$ and $\sigma'_{i,t}(H_t) = b_{i,t}^*$ otherwise.

Under this strategy, miner i works on block b_r at time $t - 1$ and thus attempts to create a fork. If successful, she continues working on the fork at time t . Miner i reverts to following the longest chain rule at time t if she is unsuccessful at time $t - 1$. Miner i also follows the longest chain rule from time $t + 1$ onwards. Note that miner i 's strategy induces the blockchain of Figure 1 if she is successful in appending her block b_f at time $t - 1$.

We now provide an example in which this strategy constitutes a strictly profitable deviation from the longest chain rule. Suppose that all transactions for all miners in all time periods are zero other than miner i 's spend transaction in block b_{l_1} . Suppose further that miner i has solved blocks b_{l_1} and b_{l_2} and that all other miners continue working on b_{l_2} at time $t + 1$ if miner i successfully creates and extends the fork at times $t - 1$ and t .

Proposition 2 (Profitable double-spending attack). *Suppose $Y_{j,b} = 0$ for all miners $j \in \{1, 2, \dots, N\}$ and all blocks b except for $Y_{i,b_{l_1}} < 0$. The strategy profile $(\sigma'_{i,t-1}, \sigma'_{i,t})$ constitutes a strictly profitable deviation from the longest chain rule if*

$$(1 + p_i \delta) \cdot \bar{R} < p_i^2 \delta^2 \cdot |Y_{i,b_{l_1}}|. \quad (8)$$

To form an intuition, consider the thought experiment where miner i adds a block at time $t - 1$ for sure. She now trades off the opportunity cost of an attack against the expected benefits from a successful attack. The benefit is the discounted absolute value of the spend transaction in block b_{l_1} which the attack would undo at time $t + 1$. Her probability of success is given by p_i^2 , i.e., the probability of both extending the fork and turning into the unique longest chain. The expected cost are the block rewards in block b_f —which she would earn with probability 1 if she extended the longest chain at time $t - 1$ —and in the subsequent block, which she appends with probability p_i at time t and would earn with probability 1 on the longest chain. Note that (8) implies that it is profitable to extend the fork which was created as part of the attack.

Importantly, double-spending attacks may be profitable in practice even if (8) is not satisfied. In our example above, deviating becomes more profitable if other miners also work on the fork and if miner i does not have block rewards on the longest chain. Furthermore, if creating a fork is profitable in the first place, it may well be optimal to continue working on the fork even if miner i is unsuccessful in extending the fork at time t . This is especially

true since she holds strictly positive balances in the form of block rewards on the fork.

Our findings align with Budish (2025), who shows that blockchains become susceptible to double-spending attacks under the longest chain rule unless transactions are sufficiently small relative to block rewards. In the language of Budish (2025), the block rewards on the LHS of (8) are the *flow* benefit of following the longest chain rule. They describe the opportunity cost of a double spending attack. The RHS of (8) captures miner i 's expected *stock* benefit from attacking the blockchain by attempting to omit a negative transaction. Proposition 2 suggests that double-spending attacks become profitable if the absolute value of spend transactions is large. Since the size of transactions in the real world is endogenous to blockchain security, this result calls the usability of blockchain technology into question. More precisely, under the longest chain rule, either transactions need to be small, or the transfer to miners in the form of block rewards—and thus the cost of using a blockchain—need to scale with transactions. Otherwise, the blockchain becomes susceptible to double-spending attacks.

Positive transaction data and consensus redirection attacks. We now explain a second, distinct reason why miners may find it optimal to deviate from the longest chain rule. To visualize the constraints that arise from (7), consider again Figure 1. Suppose now that miner 1 has earned the mining rewards in block b_f but miners 2 and 3 have the mining rewards on blocks b_{l_1} and b_{l_2} , respectively. Otherwise, set $Y_{i,b} = 0$ for every block. Since $\mathcal{B}^{LC}(G_t) = \{b_{l_2}\}$, the longest chain strategy calls for all miners to choose location b_{l_2} .

Consider the net benefit to miner 1 of deviating from the longest chain to block b_f . Given the mining rewards described above, the weight of miners who would like to see fork b_f extended is simply p_1 . Hence, the condition in (7) requires $\bar{R} \geq \delta p_1 2\bar{R}$. Should forks appear and miner 1 have too much weight (say if $\delta \rightarrow 1$ and $p_1 > 0.5$), then she can likely direct consensus to her most preferred chain. And since her most preferred chain does not coincide with the longest chain, she has incentives to deviate.

More generally, we argue that Proposition 1 likely imposes stringent limits on the distribution of mining power and these limits are likely to be violated (or provide miners with incentives to acquire mining power such that they are violated). Indeed, the condition in (7) imposes an upper bound on p_i for miner i should she be the only miner with larger balances on a fork to some $\hat{b} \in \mathcal{B}^{-1}(G_t)$ than on the longest chain. Of course, this upper bound may not suffice should other miners have positive transaction data on the fork, suggesting that the condition in (7) is likely to fail in general. Thus, Proposition 1 reveals that even when a blockchain only features mining rewards, the longest chain rule may not be robust as an

equilibrium (in the perfect public sense) to general distributions of mining power. This is particularly true if we also consider positive transaction data. Importantly, mining power in the Bitcoin network is highly concentrated in practice (Makarov and Schoar, 2021).

Interestingly, and in contrast to the case of negative transaction data where miners may find it optimal to create forks to remove transaction data, the longest chain rule fails because miners want to *add* transaction data to the consensus chain. In Figure 1, miner 1’s balances are not contained in the longest chain and this is precisely the reason why she might find it profitable to deviate from the longest chain rule and *extend* their preferred fork instead.

Lemma 2. *Suppose $\#C(b, G_t) = \#C(b', G_t)$ for every $b, b' \in \mathcal{T}(G_t)$ and $Y_{i,b_t} \geq 0$. Then the condition in (6) is satisfied for miner i .*

However, it is not profitable to create a fork if transaction data are positive. Lemma 2 shows that if only longest chains exist (and thus no fork which is one block short) and the transaction data for miner i are weakly positive, then she has no strictly profitable deviation from the longest chain rule.

Negative transaction data and saving on consensus. We now explain the third distinct reason why miners may find it optimal to deviate from the longest chain rule. Suppose that miner i is mining block b_t which contains a spend transaction for her, $Y_{i,b_t} < 0$. Miner i may now face a profitable deviation to *create* a fork by working on the parent block of the terminal block of one of the longest chains.

The reason is the delay with which transaction goods are delivered. The spend transaction is associated with a flow disutility once it is included in the blockchain—and thus before the miner receives the corresponding goods. This disutility is larger the more other miners work on the chain that includes the spend transaction. If miner i creates a fork and few other miners work on it, then the flow disutility is reduced. If miner i also holds a large amount of mining power, she has a good chance of turning this newly created fork into the consensus chain in the following time period. She then earns the full transaction benefit $Y_{i,b_t}^-/\delta$ but has reduced the associated cost when the transaction is first included in the blockchain. In sum, miners may want to create forks if they have a lot of mining power and initially other miners will not work on the newly created fork.

To illustrate, consider again Figure 1. Miner i is mining block $b_{i,t}$, which contains a spend transaction for miner i . As before, all other miners continue working on b_{l_2} chain at time $t + 1$ if miner i has successfully extended fork b_f . Suppose further that miner i has no

transaction data or block rewards in any other block. Equation (7) then simplifies to

$$\bar{R} \geq \delta p_i \cdot \left(Y_{i,b_t} + \bar{R} + \frac{|Y_{i,b_t}|}{\delta} \right). \quad (9)$$

Miner i trades off earning the block reward for sure (as well as deriving zero net utility from the spend transaction) against earning the block rewards and a strictly positive net utility from the spend transaction with probability p_i in the following period. Rearranging, we find that this deviation is not profitable if (9) is satisfied. Reversely, if the inequality in (9) fails, the deviation does become profitable in this example.

In practice, the deviation may not be profitable even if the inequality in (9) fails. If all other miners work on miner i 's fork, then the fork becomes the consensus chain with probability 1 in the following time period. Since miner i cannot avoid the flow disutility when including the transaction on the blockchain, she is better off working on her preferred longest chain in the first place. More generally, the deviation becomes less profitable the more other miners work on the fork, as the immediate cost of the spend transaction increases. Furthermore, if miner i has balances in the terminal block of her preferred longest chain, this also reduces the profitability of this deviation.¹³ The following lemma formalizes this discussion:

Lemma 3. *Suppose $\#C(b, G_t) = \#C(b', G_t)$ for every $b, b' \in \mathcal{T}(G_t)$ and $Y_{i,b_t} < 0$. Then the condition in (6) is satisfied for miner i if (9) is satisfied.*

Intuitively, the equality in (9) is a sufficient condition such that this deviation is not profitable. As an example, if $p_i = 0.5$ and $\delta = 0.99$, then (9) becomes $\bar{R} \geq |Y_{i,b_t}|/101$. Given the block rewards of 3.125 Bitcoins (as of January 2025, ignoring transaction fees), creating a fork is not profitable for all spend transactions below 300 Bitcoins in this numerical example. If the benefit of delaying consensus is low ($\delta \rightarrow 1$), this deviation becomes unprofitable. Indeed, if the discount factor is sufficiently large, then creating this short-lived disagreement is no longer profitable as miner i values the present relatively less.

Interestingly, relative to the previous motives to deviate from the longest chain rule, this third type of profitable deviation arises because miners disagree on how to *add new* data. In particular, miners disagree on how to include spend transactions before they have vested.

¹³If the consumption goods are unfairly priced in the sense that miner i derives a higher marginal flow utility than $1/\delta$ for each unit spent, then risking that the spend transaction never vests also becomes less appealing to miner i , decreasing the profitability of such a deviation.

4 Checkpoint Strategies

Biais et al. (2019), Budish (2025), and Proposition 2 above show that under the longest chain rule, blockchains are vulnerable to double-spending attacks, implying that such attacks should have been observed. Yet, Bitcoin’s consensus has largely remained robust to these types of attacks. We reconcile this apparent discrepancy by studying equilibrium strategies that coincide with the longest chain rule on-path but differ off-path in how deviations are treated. Our proposed resolution to miners’ incentives to double-spend is to consider history-dependent strategies. The basic idea is that for every graph, agents determine a reference block—a *checkpoint*—and restrict attention to all chains containing this block. All other chains are ignored regardless of their length.

A trivial solution that rules out double-spend behavior is to simply impose that the block most recently added to the previous checkpoint becomes the new checkpoint. In essence, this proposal rules out all possible forks on the blockchain. If no forks are permitted, then it is impossible for any one agent to omit data from the blockchain. We find this resolution to the double spend problem implausible for real-world implementations. In reality, some forks are non-malicious and occur due to latency—within the unit of time agents observe updates to the blockchain, it is possible to observe multiple blocks being added in the same period. We also show that such a strategy is not robust to latency shocks in Section 5.

We therefore proceed by assuming such strategies are infeasible and looking for checkpoint rules that admit the possibility of forks. More formally, for any history H_t , let $b^{CP}(H_t)$ denote the checkpoint which selects a specific block on the current graph, G_t .

Assumption 1. $b^{CP}(H_t) \notin \mathcal{T}(G_t)$ for all histories H_t .

Assumption 1 states that checkpoint rules may not select a terminal block in the graph for any history. Such a restriction ensures that forks of at least length one are always feasible.¹⁴

Checkpoint Settlement. We now explain settlement with checkpoints. We continue assuming that consumption goods are fairly priced. Under the checkpoint strategies we suggest below, blocks become checkpoints and thus part of consensus one period after they have been appended to the blockchain. The price of consumption goods therefore remains at $1/\delta$. Settlement of spend transactions contained in block b takes place at time t if this is the first time period that block b lies on a chain to the checkpoint. This includes the possibility

¹⁴We discuss latency periods and thus forks of greater length in Section 5.

that the block b itself is the checkpoint.¹⁵

Checkpoint Strategies. To specify a candidate equilibrium strategy with checkpoints, it is helpful to introduce two pieces of notation. First, let $J(b', G_t) \subseteq G_t$ denote the subgraph associated with some root block $b' \in \mathcal{B}(G_t)$. This subgraph contains the block b' , all its child blocks, the child blocks' child blocks, and so on. It also contains all edges connecting these blocks. Let $\mathcal{J}(b', G_t)$ denote the set of blocks on the subgraph. To illustrate, reconsider Figure 1. The subgraph $J(b_r, G_t)$ contains the blocks b_r, b_{l_1}, b_{l_2} and b_f as well as the edges $(b_r, b_{l_1}), (b_r, b_f)$ and (b_{l_1}, b_{l_2}) . Subgraph $J(b_{l_1}, G_t)$ contains the blocks b_{l_1} and b_{l_2} as well as the edge (b_{l_1}, b_{l_2}) . Second, let $M(b', G_t)$ denote the parent block of block b' . For example, in Figure 1 the parent block of b_{l_2} is given by $M(b_{l_2}, G_t) = b_{l_1}$.

We now define checkpoint blocks as follows. In any period, we consider the subgraph with the checkpoint block from the previous period as the common root, $J(b^{CP}(H_{t-1}), G_t)$. Next, we find the set of terminal blocks of the longest chains on this subgraph,

$$\mathcal{B}^{CP}(b^{CP}(H_{t-1}), G_t) = \underset{b \in \mathcal{J}(b^{CP}(H_{t-1}), G_t)}{\operatorname{argmax}} \#C(b, G_t). \quad (10)$$

If this set is a singleton, we choose the new checkpoint to be the parent block of the terminal block on the longest chain: $b^{CP}(H_t) = M(\mathcal{B}^{CP}(b^{CP}(H_{t-1}), G_t), G_t)$.

If there are multiple longest chains ahead of the checkpoint, then the checkpoint randomly updates to the parent of a terminal block of one of these longest chains.¹⁶ Given this new checkpoint, the set of terminal blocks of the longest chains on the new subgraph is then given by $\mathcal{B}^{CP}(b^{CP}(H_t), G_t)$. Genesis block b_0 is the initial checkpoint.

To illustrate the checkpoint selection, reconsider again Figure 1. Suppose that at the beginning of time t the checkpoint is given by block b_r . Since the chain leading to b_{l_2} is the unique longest chain containing the checkpoint, the checkpoint updates to b_{l_1} . If there was a block b_{f_2} chained to block b_f , then there would be two longest chains containing the checkpoint b_r . The checkpoint would then randomly update to either b_{l_1} or b_f .

¹⁵Formally, redefine $\lambda_t(b, H_t) = 1$ if $t = \inf \{ \tau \geq 0 : b \in \mathcal{C}(b^{CP}(H_\tau), G_\tau) \}$, and $\lambda_t(b, H_t) = 0$ otherwise.

¹⁶More formally, if the set $\mathcal{B}^{CP}(b^{CP}(H_{t-1}), G_t)$ is not a singleton, we say that $b^{CP}(H_t) = M(b, G_t)$ with probability $\pi_b \in [0, 1]$ for every $b \in \mathcal{B}^{CP}(b^{CP}(H_{t-1}), G_t)$ and insist that $\sum_{b \in \mathcal{B}^{CP}(b^{CP}(H_{t-1}), G_t)} \pi_b = 1$. Note that it is not difficult to randomly select blocks in blockchain environments. Recall that each block contains hash data, which come in the form of fixed-size values. One method of randomization is to choose the checkpoint candidate block with the lowest hash. Further note that our formulation includes the possibility that the checkpoint does not update, $b^{CP}(H_t) = b^{CP}(H_{t-1})$, if all longest chains ahead of the checkpoint consist of only one block.

Given checkpoint selection, the checkpoint rule satisfies

$$\sigma_{i,t}^{CP}(H_t) = b_{i,t}^* \equiv \operatorname{argmax}_{b \in \mathcal{B}^{CP}(b^{CP}(H_t), G_t)} \left(Y_{i,b} + R_{i,b} + \frac{Y_{i,b}^-}{\delta} \right), \quad (11)$$

where we have abused notation by redefining $b_{i,t}^*$. The strategy calls miners to work on their preferred longest chain ahead of the checkpoint, which has the highest sum of positive transactions, block rewards, and unvested spend transactions. The checkpoint rule corresponds closely to the longest chain rule but is limited to the subgraph following the new checkpoint $b^{CP}(H_t)$. Hence, miners only consider blocks in $\mathcal{B}^{CP}(b^{CP}(H_t), G_t)$.

Checkpoint Equilibrium. The checkpoint strategy has one key feature. The following proposition implies that it is never profitable to append new blocks to any block *behind* the checkpoint. Hence, miners cannot remove transactions in order to double-spend.

Proposition 3. *For any history H_t and for any miner i , there exists no (weakly) profitable one-shot deviations from the checkpoint rule to any block $b \notin \mathcal{J}(b^{CP}(H_t), G_t)$.*

Intuitively, miners ignore all blocks which lie on chains that do not include the checkpoint block. Thus, a newly mined block is immediately abandoned if it is appended to the parent of the checkpoint block. The same is true for a block appended to any chain which branches off behind the checkpoint. In Figure 1, any blocks appended to b_r or b_f are ignored by all miners once block b_{l_1} has become the checkpoint. As a consequence, spend transactions in block b_{l_1} cannot be removed after this block has become the checkpoint block and the corresponding consumption goods have been delivered. Since the first spend only occurs once blocks with spend transactions are behind the checkpoint, double-spending cannot occur under checkpoints strategies.

Furthermore, it is costly to append a block behind the checkpoint. Any positive transaction data and block rewards contained in the new block are then lost. Spend transactions in the new block would never vest. Deviations to blocks behind the checkpoint are therefore strictly unprofitable.

With this result in hand, we can characterize conditions such that the checkpoint rule is indeed a perfect public equilibrium:

Proposition 4. *The checkpoint rule is a perfect public equilibrium if, for every miner i and block b , either $Y_{i,b} \geq 0$ or (9) is satisfied.*

Proof. Given checkpoint selection, any chain on the subgraph $J(b^{CP}(H_t), G_t)$ has length two. Proposition 3 then implies that we only need to consider deviations to the checkpoint block. The claim then immediately follows from Lemmas 2 and 3. \square

In sum, introducing a simple form of history-dependence—as implemented, for example, on the Ethereum blockchain—eliminates incentives for double-spending, which has been identified as a central limitation of blockchain technology (Biais et al., 2019; Budish, 2025). Specifically, incorporating checkpoints into the longest chain rule achieves consensus for all blocks behind the checkpoint. An important consideration for checkpoint strategies is possible network latency, which we discuss next.

5 Latency

We now provide a microfoundation for Assumption 1 by studying a simple latency shock and its interaction with checkpoints. The latency shock is modeled as a one-off unexpected partition of the miners for one time period, which results in an accidental fork. We show that if checkpoints update to terminal blocks, then this creates disagreement among miners on the checkpoint. This disagreement incentivizes deviations and thus stands in the way of consensus. We show that disagreement is prevented if checkpoints lag by one block and discuss our results.

Let $G_{i,t}$ and $H_{i,t}$ denote the blockchain configuration observed and history held by miner i at time t , respectively. Consider some time $t \geq 1$ and suppose that miners hold the same history up to this point: $H_{i,s} = H_{j,s} = H_s$ and hence $G_{i,s} = G_{j,s} = G_s$ for all $(i, j) \in \{1, \dots, N\}^2$ and $s \in \{1, \dots, t\}$.

Consider then the following notion of latency. At the end of time t , communication between miners is interrupted and the set of miners is partitioned into two sets \mathcal{N} and \mathcal{N}' , with $0 < |\mathcal{N}| < N$. Given the breakdown of communication, two miners add a block to the blockchain at the end of time t : one miner belonging to set \mathcal{N} , and a second miner belonging to set \mathcal{N}' . The block added at time t by the miner in \mathcal{N} is not observed at time $t+1$ by any miner belonging to the set \mathcal{N}' , and vice versa. As a consequence, miners $i \in \mathcal{N}$ and $j \in \mathcal{N}'$ observe different blockchain configurations, $G_{i,t+1} \neq G_{j,t+1}$, and thus hold different histories $H_{i,t+1} \neq H_{j,t+1}$.

Latency is resolved at the end of time $t+1$. All miners then observe the same blockchain configuration, including the two blocks added at the end of time t , in all time periods going

forward: $G_{i,s} = G_{j,s} = G_s$ for all $(i, j) \in \{1, \dots, N\}^2$ and $s \geq t + 2$. We also assume that miners learn the blockchain configuration observed by other miners' at time $t + 1$.¹⁷

The latency event is unanticipated by all miners, does not occur again, and is expected by all miners not to occur again.

We now contrast two different checkpoint strategies. The first strategy corresponds to the checkpoint strategy in Section 4, allowing for different histories across miners. Starting point is the set of terminal blocks of what miner i perceives to be the longest chains containing miner i 's previous checkpoint, $\mathcal{B}^{CP}(b^{CP}(H_{i,t-1}), G_{i,t})$. If this set is a singleton, the new checkpoint is selected to be this block's parent:

$$b^{CP}(H_{i,t}) = M\left(\mathcal{B}^{CP}(b^{CP}(H_{i,t-1}), G_{i,t}), G_{i,t}\right). \quad (12)$$

If there are multiple longest chains ahead of the checkpoint, then the checkpoint randomly updates to one block in the set $\mathcal{B}^{CP}(b^{CP}(H_{i,t-1}), G_{i,t})$, as in the main analysis. As before, miners are called to work on their preferred longest chain containing the checkpoint.

The second strategy is exactly as the first strategy, except that (12) is replaced by

$$b^{CP}(H_{i,t}) = \mathcal{B}^{CP}(b^{CP}(H_{i,t-1}), G_{i,t}), \quad (13)$$

if the set is a singleton; otherwise, the checkpoint randomly updates to one of the blocks in the set. Hence, the checkpoint updates to the *terminal block* of the longest chain containing the previous checkpoint, as opposed to its parent. Assumption 1 is thus violated.

The following lemma states that if checkpoints update to terminal blocks, then miners disagree on the checkpoint under our specification of latency from time $t + 1$ onwards. In contrast, miners agree on the checkpoint with parents of terminal blocks as checkpoints:

Lemma 4. *Consider the checkpoint strategy. With checkpoint selection as in (12), we have $b^{CP}(H_{i,s}) = b^{CP}(H_{j,s})$ for all $(i, j) \in \{1, \dots, N\}^2$ and all $s \geq 0$. With checkpoint selection as in (13), we have $b^{CP}(H_{i,s}) = b^{CP}(H_{j,s})$ for all $(i, j) \in \{1, \dots, N\}^2$ and $s \in \{1, \dots, t\}$, but $b^{CP}(H_{i,s}) \neq b^{CP}(H_{j,s})$ for any $i \in \mathcal{N}$, $j \in \mathcal{N}'$ for all $s \geq t + 1$.*

Figure 2 illustrates. Suppose checkpoints update to terminal blocks. Suppose further that at the beginning of time t , all miners agree on block b_{l_3} as the checkpoint. At time t , miner $i \in \mathcal{N}$ appends block b_{l_4} . Miner $j \in \mathcal{N}'$ appends block b_f simultaneously. Miners

¹⁷As a consequence, all miners' histories $\{H_{i,t+1}\}_{i=1}^N$ become common knowledge. This assumption is needed to study the equilibria of the continuation game from time $t + 2$ onwards.

belonging to the set \mathcal{N} only observe block b_{l_4} but not block b_f . Hence, they update their checkpoint to block b_{l_4} . Miners belonging to the complimentary set \mathcal{N}' only observe block b_f but not block b_{l_4} , and thus update their checkpoint to block b_f . Hence, miners in the two groups disagree on the checkpoint block.




Figure 2: An illustration of disagreement on checkpoints.

This divergence in checkpoints does not occur if terminal blocks cannot be checkpoints. To illustrate, consider the checkpoint selection of (12)—that is, under which the checkpoint updates to the *parent of the terminal block* of the longest chain that contains the previous checkpoint. Suppose that at the beginning of time t , all miners agree on block b_{l_2} as the checkpoint. After the block b_{l_4} has been solved, the block b_{l_3} becomes the new checkpoint for all miners in the set \mathcal{N} that observe this block. Similarly, after the block b_f has been solved, the block b_{l_3} becomes the new checkpoint for all miners in the complimentary set \mathcal{N}' that observe this block. Hence, miners in the two groups agree on the checkpoint block, despite communication having been interrupted between the two groups temporarily.

Importantly, consensus is not robust to latency shocks with disagreement on checkpoints:

Proposition 5. *Consider the checkpoint strategy with checkpoint selection as in (12) and the continuation game for histories $\{H_{i,t+2}\}_{i=1}^N$ induced by this strategy. For any such histories, the strategy remains a perfect equilibrium of this game under the conditions described in Proposition 4. Consider then the checkpoint strategy with checkpoint selection as in (13) and the continuation game for histories $\{H_{i,t+2}\}_{i=1}^N$ induced by this strategy. The strategy is not a perfect equilibrium.*

Intuitively, under the checkpoint strategy satisfying Assumption 1, all miners agree on the checkpoint and hence also on the consensus chain. The incentives to deviate from this strategy are then exactly as described in Proposition 4. In contrast, miners disagree on the checkpoint if terminal blocks become checkpoints. Miners are then called to work on different

chains, with computing power split between the two chains. The value of block rewards is then larger on the chain with more computing power. As a consequence, miners called to work on the chain with lower block reward values have incentives to deviate.

The disagreement on checkpoints is precisely why different miners are called to work on different chains, and hence why consensus is not robust to latency shocks with terminal blocks as checkpoints. Disagreement is averted under Assumption 1.

The possibility for disagreement re-emerges if latency shocks last longer or occur repeatedly under the rule that checkpoints update to parents of terminal blocks. Longer latency shocks can be addressed by introducing longer checkpoint lags, as long as the time until latency is resolved remains bounded. With longer checkpoint lags, the consensus redirection and saving on consensus attacks from Section 3 re-emerge (ahead of the checkpoint).

Repeated latency shocks require a different modification. Leshno et al. (2024) study such a setting in which, using the language of our framework, checkpoints update to blocks for which it has become common certainty that they lie on the longest chain containing the previous checkpoint. A similar modification of the checkpoint rule would eliminate the threat of permanent disagreement under repeated latency in our framework as well.

6 The trade-off between consensus and informational efficiency

We now argue that with longer checkpoint lags to accommodate latency and the resulting attack vectors, a fundamental trade-off between informational efficiency—conditioning consensus on all available information—and consensus itself arises.

Consider the following class of candidate equilibrium strategies to which we refer as *full-information checkpoint rules*. Assume a general checkpoint lag of $k \geq 1$ blocks. Fair pricing (on a single chain) then requires that each unit of account buys $(1/\delta)^k$ units of consumption goods. Define i 's data on the chain leading to terminal block $b \in \mathcal{T}(J(b^{CP}(H_t), G_t))$ on the subgraph starting at the checkpoint (excluding the checkpoint itself) as

$$D_i(b, H_t) = \sum_{b' \in \mathcal{C}(b, G_t) \setminus \mathcal{C}(b^{CP}(H_t), G_t)} \left(Y_{i,b'} + R_{i,b'} + \delta^{n(b', b, H_t)} \cdot \frac{Y_{i,b'}^-}{\delta^k} \right), \quad (14)$$

where $n(b', b, H_t)$ denotes the number of blocks that need to be appended to the terminal

block b and its future child blocks before spend transactions in b' vest.¹⁸ The miners' data are aggregated according to $S(b, H_t) = \sum_{i=1}^N f(D_i(b, H_t)) \cdot g_i(p_i)$ for each terminal block b , where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous and strictly increasing and $g_i : (0, 1) \rightarrow (0, 1)$ is continuous and weakly increasing for all i . Examples for this aggregation include computing averages where $f(D_i(b, H_t)) = D_i(b, H_t)$, e.g. using equal weights $g_i(p_i) = \frac{1}{N}$ for all i or probability weights $g_i(p_i) = p_i$ for all i . Note that $S(b, H_t)$ has a natural interpretation as a welfare criterion, where the Pareto weight for miner i is given by $g_i(p_i)$.

Define $b_t^F \equiv \operatorname{argmax}_{b \in \mathcal{T}(J(b^{CP}(H_t), G_t))} S(b, H_t)$. Under full-information checkpoint rules, the checkpoint is chosen to be the k -ancestor of b_t^F (rather than the parent as in Section 4).¹⁹ Set $\sigma_{i,t}^F(H_t) = b_t^F$. As in Section 4, settlement is linked to the checkpoint.

Proposition 6. *No full-information checkpoint rule achieves consensus.*

Proof. We prove the result by constructing a graph for which consensus fails. Consider Figure 3 and suppose that the current checkpoint is given by b_r . Suppose miner 1 has solved blocks b_{l_1} to b_{l_k} and miner 2 has solved blocks b_{f_1} to b_{f_k} . Suppose $g_1(p_1) \geq g_2(p_2)$. Block b_{f_1} contains one transaction $Y_{2,b_{f_1}} = Y \in \mathbb{R}_+$ that satisfies $f(k\bar{R}) \cdot g_1(p_1) = f(k\bar{R} + Y) \cdot g_2(p_2)$. No other block on the graph contains any transaction. Suppose further that the next block to be appended, b_t , contains the only one transaction $Y_{3,b_t} = \varepsilon$, where $\varepsilon > 0$.

Given the strategies, the checkpoint updates to b_{l_1} if the current block is appended to b_{l_k} , and to b_{f_1} if it is appended to b_{f_k} . Since all miners return to playing equilibrium strategies from the next time period onwards, miner 1's payoff from working on b_{l_k} exceeds the payoff from working on b_{f_k} by $k\bar{R}$; similarly, miner 2's payoff from working on b_{f_k} exceeds the payoff from working on b_{l_k} by $k\bar{R} + Y$. \square

Proposition 6 shows that when checkpoint strategies fully incorporate all available information, consensus cannot be sustained in equilibrium. Even small changes in recent blocks can generate large shifts in expected payoffs, creating strong incentives for miners to deviate. While such full-information rules are welfare maximizing in principle, they fail to align individual incentives to sustain consensus.

We now prove the existence of equilibrium strategies which do achieve consensus, but at the cost of informational efficiency. These strategies called *partial-information checkpoint rule* only consider data in blocks at fork points, such as blocks b_{l_1} and b_{f_1} in Figure 3.

¹⁸Note that $n(b', b, H_t) = k - (\#C(b, H_t) - \#C(b', H_t))$. To illustrate, suppose block b' is chained to the checkpoint, which is k blocks behind the corresponding terminal block b . Then b is $k - 1$ blocks ahead of b' , and hence $n(b', b, H_t) = 1$. If b' is the terminal block, then $n(b', b, H_t) = k$.

¹⁹Formally, b' is a k -ancestor to b if $b' \in \mathcal{C}(b, G_t)$ and $\#C(b', G_t) = \#C(b, G_t) - k$, with $k \geq 1$.




Figure 3: An illustration of incentives under the full-information checkpoint rule.

To build these strategies, define $\mathcal{P}(b, G_t)$ as the set containing block b 's child blocks.²⁰ Consider the operator $T : \mathcal{B}(G_t) \rightarrow \mathcal{B}(G_t)$ which selects b 's child block with the highest balances across miners if b has at least one child block, and otherwise selects block b itself:

$$T(b) = \mathbb{1}_{\{\mathcal{P}(b, G_t) \neq \emptyset\}} \cdot \operatorname{argmax}_{b' \in \mathcal{P}(b, G_t)} \sum_{i=1}^N \left(Y_{i, b'} + R_{i, b'} + \frac{Y_{i, b'}^-}{\delta^k} \right) + \mathbb{1}_{\{\mathcal{P}(b, G_t) = \emptyset\}} \cdot b. \quad (15)$$

Define $b_t^P \equiv \lim_{s \rightarrow \infty} T^s(b^{CP}(H_t))$. As above, the checkpoint is chosen to be the k -ancestor of b_t^P , settlement is linked to the checkpoint, and $\sigma_{i,t}^P(H_t) = b_t^P$.²¹

Proposition 7. *The partial-information checkpoint rule achieves consensus. However, there exist histories for which $S(b_t^P, H_t) < S(b_t^F, H_t)$.*

We illustrate the result using Figure 3. Suppose again that miner 1 has solved blocks b_{l_1} — b_{l_k} and miner 2 has solved blocks b_{f_1} — b_{f_k} . Suppose the graph only contains two transactions, both for miner 3: $Y_{3, b_{l_1}} = Y \in (0, \infty)$ and $Y_{3, b_{f_k}} = 2Y$. Since the balances in b_{l_1} exceed the balances in b_{f_1} , the strategy calls miners to work on the upper fork even though it has lower balances than the lower fork. However, working on the lower fork cannot change the balances in its first block, and hence neither miner 2 nor miner 3 can induce other miners to switch to their preferred chain. As a consequence, miners extend the chain containing lower balances across miners. And if one considers the sum of balances as measure of welfare, then miners achieve consensus but at the cost of efficiency.

The existence of equilibrium strategies that attain consensus by conditioning only on partial information highlights the tradeoff between informational efficiency and consensus

²⁰Formally, $\mathcal{P}(b, G_t) = \{b' \in \mathcal{B}(G_t) : (b', b) \in \mathcal{E}(G_t)\}$.

²¹Whenever there are blocks with equal balances across miners, one can use the hash data to break ties.

itself. This trade-off, which we view as a blockchain analog of the Grossman–Stiglitz paradox, underscores the fundamental economic limits to fully decentralized consensus.

7 Conclusion

In this paper, we develop a novel, game-theoretic framework to study blockchain consensus. We demonstrate that incorporating a simple history-dependence in the form of checkpoints into strategies can effectively prevent double-spending attacks. This is because miners collectively agree to disregard forks that attempt to reverse previously confirmed transactions. As a result, once a transaction is part of consensus, the corresponding balances cannot be spent a second time. In practice, various protocols already feature some form of history-dependence. When checkpoint strategies must accommodate the possibility of network latency, then achieving consensus ‘ahead of the checkpoint’ presents a greater challenge. We formalize this notion by showing that strategies, which respond to all information on the blockchain and hence (seek to) maximize welfare, fail to achieve consensus.

References

- AMOUSSOU-GUENOU, Y., B. BIAIS, M. POTOP-BUTUCARU, AND S. TUCCI-PIERGIOVANNI (2024): “Committee-based blockchains as games between opportunistic players and adversaries,” *The Review of Financial Studies*, 37, 409–443.
- AUER, R., C. MONNET, AND H. S. SHIN (2021): “Permissioned distributed ledgers and the governance of money,” *Available at SSRN 3770075*.
- BAKOS, Y. AND H. HALABURDA (2021): “Permissioned vs Permissionless Blockchain Platforms: Tradeoffs in Trust and Performance,” *NYU Stern School of Business working paper*.
- BENHAIM, A., B. H. FALK, AND G. TSOUKALAS (2023): “Scaling blockchains: Can committee-based consensus help?” *Management Science*, 69, 6525–6539.
- BIAIS, B., C. BISIERE, M. BOUVARD, AND C. CASAMATTA (2019): “The Blockchain Folk Theorem,” *The Review of Financial Studies*, 32, 1662–1715.
- BUDISH, E. (2025): “Trust at Scale: The Economic Limits of Cryptocurrencies and Blockchains,” *The Quarterly Journal of Economics (forthcoming)*.

- BUTERIN, V. AND V. GRIFFITH (2017): “Casper the friendly finality gadget,” *arXiv preprint arXiv:1710.09437*.
- BUTERIN, V., D. HERNANDEZ, T. KAMPHEFNER, K. PHAM, Z. QIAO, D. RYAN, J. SIN, Y. WANG, AND Y. X. ZHANG (2020): “Combining ghost and casper,” *arXiv preprint arXiv:2003.03052*.
- CARLSTEN, M., H. KALODNER, S. M. WEINBERG, AND A. NARAYANAN (2016): “On the instability of bitcoin without the block reward,” in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 154–167.
- CHIU, J. AND T. V. KOEPPL (2022): “The economics of cryptocurrency: Bitcoin and beyond,” *Canadian Journal of Economics/Revue canadienne d’économique*, 55, 1762–1798.
- CONG, L. W., Z. HE, AND J. LI (2021): “Decentralized mining in centralized pools,” *The Review of Financial Studies*, 34, 1191–1235.
- EYAL, I. AND E. G. SIRER (2018): “Majority is not enough: Bitcoin mining is vulnerable,” *Communications of the ACM*, 61, 95–102.
- FISCHER, M. J., N. A. LYNCH, AND M. S. PATERSON (1985): “Impossibility of distributed consensus with one faulty process,” *Journal of the ACM (JACM)*, 32, 374–382.
- GANS, J. S. AND H. HALABURDA (2023): ““Zero Cost” Majority Attacks on Permissionless Blockchains,” *NBER Working Paper No. 31473*.
- GARRATT, R. J. AND M. R. VAN OORDT (2023): “Why fixed costs matter for proof-of-work-based cryptocurrencies,” *Management Science*, 69, 6482–6507.
- GROSSMAN, S. J. AND J. E. STIGLITZ (1980): “On the impossibility of informationally efficient markets,” *The American economic review*, 70, 393–408.
- HALABURDA, H., Z. HE, AND J. LI (2022): “An economic model of consensus on distributed ledgers,” *NBER Working Paper No. 29515*.
- JOHN, K., T. J. RIVERA, AND F. SALEH (2025): “Proof-of-work versus proof-of-stake: A comparative economic analysis,” *The Review of Financial Studies*, 38, 1955–2004.
- KANG, K.-Y. (2023): “Cryptocurrency and double spending history: Transactions with zero confirmation,” *Economic Theory*, 75, 453–491.

- KARAKOSTAS, D. AND A. KIAYIAS (2021): “Securing proof-of-work ledgers via checkpointing,” in *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, IEEE, 1–5.
- LAGOS, R. AND R. WRIGHT (2005): “A Unified Framework for Monetary Theory and Policy Analysis,” *Journal of Political Economy*, 113, 463–484.
- LESHNO, J. D., E. SHI, AND R. PASS (2024): “On the viability of open-source financial rails: Economic security of permissionless consensus,” *arXiv preprint arXiv:2409.08951*.
- LI, J. (2023): “On the security of optimistic blockchain mechanisms,” *Available at SSRN 4499357*.
- LYNCH, N. A. (1996): *Distributed algorithms*, Elsevier.
- MAKAROV, I. AND A. SCHOAR (2021): “Blockchain Analysis of the Bitcoin Market,” *Available at SSRN 3942181*.
- MOROZ, D. J., D. J. ARONOFF, N. NARULA, AND D. C. PARKES (2020): “Double-spend counterattacks: Threat of retaliation in proof-of-work systems,” *arXiv preprint arXiv:2002.10736*.
- NAKAMOTO, S. (2008): “Bitcoin: A peer-to-peer electronic cash system,” *Bitcoin White paper*.
- NEU, J., E. N. TAS, AND D. TSE (2021): “Ebb-and-flow protocols: A resolution of the availability-finality dilemma,” in *2021 IEEE Symposium on Security and Privacy (SP)*, IEEE, 446–465.
- PAGNOTTA, E. S. (2022): “Decentralizing money: Bitcoin prices and blockchain security,” *The Review of Financial Studies*, 35, 866–907.
- ROCHETEAU, G. AND R. WRIGHT (2005): “Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium,” *Econometrica*, 73, 175–202.
- SALEH, F. (2021): “Blockchain without waste: Proof-of-stake,” *The Review of Financial studies*, 34, 1156–1190.
- SANKAGIRI, S., X. WANG, S. KANNAN, AND P. VISWANATH (2021): “Blockchain cap theorem allows user-dependent adaptivity and finality,” in *Financial Cryptography and*

Data Security: 25th International Conference, FC 2021, Virtual Event, March 1–5, 2021, Revised Selected Papers, Part II 25, Springer, 84–103.

STEWART, A. AND E. KOKORIS-KOGIA (2020): “Grandpa: a byzantine finality gadget,” *arXiv preprint arXiv:2007.01560*.

XU, J., C. WANG, AND X. JIA (2023): “A survey of blockchain consensus protocols,” *ACM Computing Surveys*, 55, 1–35.

8 Proofs

Lemma 1. Suppose that the strategy profile σ is such that all miners follow the longest chain rule as in equation (2) from time t onwards. That is, $\sigma_{i,\tau}(H_\tau) = b_{i,\tau}^*$ for all i , after every history H_τ , and for all $\tau \geq t$. Suppose $\mathcal{B}^{LC}(G_t)$ is not a singleton. Since all miners follow the longest chain rule and only one miner appends a block at time t , there will be a single longest chain at time $t + 1$. All miners then append all their new blocks to this chain going forward, and all blocks on this chains have a value of 1 for every miner.

Miner i 's expected payoff from following the longest chain rule is given by:

$$\begin{aligned}
V_{i,t}(\sigma; H_t) &= \sum_{b \in \mathcal{B}(G_t)} \left((1 - \delta)q_{i,b,t}(Y_{i,b} + R_{i,b}) + \frac{Y_{i,b}^-}{\delta} \cdot \lambda_t(b, H_t) \right) \\
&+ \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t}^* = b\}} p_j \cdot \delta \cdot \left(Y_{i,b_t} + \delta \cdot \frac{Y_{i,b_t}^-}{\delta} + \sum_{b' \in \mathcal{C}(b, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \right) \\
&+ p_i \cdot \delta \cdot \left(Y_{i,b_t} + \bar{R} + \delta \cdot \frac{Y_{i,b_t}^-}{\delta} + \sum_{b' \in \mathcal{C}(b_{i,t}^*, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \right) \\
&+ \delta^2 \mathbb{E}_t \left(Y_{i,b_{t+1}} + R_{i,b_{t+1}} + \delta \frac{Y_{i,b_{t+1}}^-}{\delta} \right) + \mathbb{E}_t \sum_{\tau=0}^{\infty} \delta^{3+\tau} \left(Y_{i,b_{t+\tau+2}} + R_{i,b_{t+\tau+2}} + \delta \frac{Y_{i,b_{t+\tau+2}}^-}{\delta} \right)
\end{aligned} \tag{16}$$

For detailed explanations of all payoff expressions in the proofs, see the Online Appendix.

Now consider an arbitrary deviation to $b \in \mathcal{B}^{LC}(G_t)$, $b \neq b_{i,t}^*$, by miner i , for this time period only. Denote this new profile by σ' and the payoff by $V_{i,t}(\sigma'; H_t)$. As one chain becomes the single longest chain under this profile σ' , it remains true that consensus is

achieved after the next block has been appended. Hence, the only difference in the expression of $V_{i,t}(\sigma'; H_t)$ is, as miner i selects a different longest chain to become the consensus chain if she is successful in appending the next block, in the third term on the RHS of (16), where $b_{i,t}^*$ is replaced by the block b , to which i deviates.

By definition of $b_{i,t}^*$, we have $V_{i,t}(\sigma; H_t) \geq V_{i,t}(\sigma'; H_t)$. Therefore, there is no strictly profitable one-shot deviation. Conversely, should a candidate equilibrium strategy specify a different tie-breaking rule, then a one-shot deviation applying the rule specified in Equation (2) immediately yields a strictly profitable deviation. \square

Proposition 1. Consider a one-shot deviation by miner i to some $\hat{b} \in \mathcal{B}^{-1}(G_t)$ before reverting to the longest chain rule. Denote this new profile by σ' . If miner i is successful in appending the next block, then the number of longest chains at time $t + 1$ increases by one. Consensus is then stalled until period $t + 2$, at which point one of the longest chains including the new one becomes the consensus chain. If miner i is not successful in appending the next block, then consensus is achieved at time $t + 1$ as under the longest chain rule.

Miner i 's expected payoff from the one-shot deviation is given by

$$\begin{aligned}
V_{i,t}(\sigma'; H_t) &= \sum_{b \in \mathcal{B}(G_t)} \left((1 - \delta) q_{i,b,t} (Y_{i,b} + R_{i,b}) + \frac{Y_{i,b}^-}{\delta} \cdot \lambda_t(b, H_t) \right) \\
&+ \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t}^* = b\}} p_j \cdot \delta \cdot \left(Y_{i,b_t} + \delta \cdot \frac{Y_{i,b_t}^-}{\delta} + \sum_{b' \in \mathcal{C}(b, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \right) \\
&+ \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t}^* = b\}} p_j \cdot \delta^2 \cdot \mathbb{E}_t \left(Y_{i,b_{t+1}} + R_{i,b_{t+1}} + \delta \cdot \frac{Y_{i,b_{t+1}}^-}{\delta} \right) \\
&+ p_i \cdot \delta \cdot (1 - \delta) \cdot \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} \frac{p_j}{1 - p_i} \cdot \left(Y_{i,b_t} + \bar{R} + \sum_{b' \in \mathcal{C}(\hat{b}, G_t)} (Y_{i,b'} + R_{i,b'}) \right) \\
&+ p_i \cdot \delta \cdot (1 - \delta) \cdot \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t+1}^* = b\}} \frac{p_j}{1 - p_i} \sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}) \\
&+ p_i^2 \cdot \delta^2 \cdot \left(\mathbb{E}_t \left[Y_{i,b_{t+1}} + \bar{R} + \delta \cdot \frac{Y_{i,b_{t+1}}^-}{\delta} \right] + Y_{i,b_t} + \bar{R} + \frac{Y_{i,b_t}^-}{\delta} \right. \\
&\quad \left. + \sum_{b' \in \mathcal{C}(\hat{b}, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \right) \\
&+ p_i \cdot \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} p_j \cdot \delta^2 \cdot \left(\mathbb{E}_t \left[Y_{i,b_{t+1}} + \delta \cdot \frac{Y_{i,b_{t+1}}^-}{\delta} \right] + Y_{i,b_t} + \bar{R} + \frac{Y_{i,b_t}^-}{\delta} \right. \\
&\quad \left. + \sum_{b' \in \mathcal{C}(\hat{b}, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b_t}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \right) \\
&+ p_i \cdot \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t+1}^* = b\}} p_j \cdot \delta^2 \cdot \left(\mathbb{E}_t \left[Y_{i,b_{t+1}} + \delta \cdot \frac{Y_{i,b_{t+1}}^-}{\delta} \right] \right. \\
&\quad \left. + \sum_{b' \in \mathcal{C}(b, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \right) \\
&+ \mathbb{E}_t \sum_{\tau=0}^{\infty} \delta^{3+\tau} \left(Y_{i,b_{t+\tau+2}} + R_{i,b_{t+\tau+2}} + \delta \cdot \frac{Y_{i,b_{t+\tau+2}}^-}{\delta} \right)
\end{aligned} \tag{17}$$

Taking differences between (16) and (17) reveals that $V_{i,t}(\sigma; H_t) \geq V_{i,t}(\sigma'; H_t)$ if and only if the expression in (6) is satisfied. \square

Proposition 2. Since $Y_{j,b} = 0$ for all miners j and all blocks b (with the notable exception of $Y_{i,b_{l_1}} < 0$) and $R_{i,b_{l_1}} = R_{i,b_{l_2}} = \bar{R}$, the expression in (16) simplifies to

$$\begin{aligned} V_{i,t-1}(\sigma; H_{t-1}) &= (1 - \delta)(Y_{i,b_{l_1}} + 2\bar{R}) + (1 - p_i) \cdot \delta \cdot (Y_{i,b_{l_1}} + 2\bar{R}) \\ &\quad + p_i \cdot \delta \cdot (Y_{i,b_{l_1}} + 2\bar{R} + \bar{R}) + \delta^2 \cdot \mathbb{E}_t(R_{i,b_{t+1}}) + \delta^3 \cdot \mathbb{E}_t \sum_{\tau=0}^{\infty} \delta^\tau (R_{i,b_{t+\tau+2}}) \end{aligned} \quad (18)$$

Note that $\mathbb{E}_t(R_{i,b_{t+1+s}}) = p_i \bar{R}$ for all $s \geq 0$. To ease notation, define $D \equiv Y_{i,b_{l_1}} + 2\bar{R}$. Plugging in and simplifying, we have $V_{i,t-1}(\sigma; H_{t-1}) = D + \delta \cdot \frac{p_i \bar{R}}{1-\delta} \equiv V$. The payoff from following the longest chain rule is therefore given by miner i 's current balances D plus the expected present value of future block rewards. The deviation payoff is given by

$$\begin{aligned} V_{i,t-1}(\sigma'; H_{t-1}) &= (1 - \delta)D + (1 - p_i) \cdot \delta \cdot V + p_i \cdot \delta \cdot (1 - \delta)D \\ &\quad + p_i \cdot (1 - p_i) \cdot \delta^2 \cdot V + p_i^2 \cdot \delta^2 \cdot (1 - \delta)D \\ &\quad + p_i^2 \cdot (1 - p_i) \cdot \delta^3 \cdot V + p_i^3 \cdot \delta^3 \cdot [3\bar{R} + (V - D)] \end{aligned} \quad (19)$$

Differencing we have $V_{i,t-1}(\sigma; H_{t-1}) - V_{i,t-1}(\sigma'; H_{t-1}) = p_i \delta \left[p_i^2 \delta^2 Y_{i,b_{l_1}} + (1 + p_i \delta) \bar{R} \right]$. The expression is negative if the inequality in (8) is satisfied. \square

Lemma 2. If $Y_{i,b_t} \geq 0$, the condition in (6) becomes

$$\begin{aligned}
& Y_{i,b_t} + \bar{R} + \sum_{b' \in \mathcal{C}(b_{i,t}^*, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \tag{20} \\
& \geq (1 - \delta) \cdot \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} \frac{p_j}{1 - p_i} \cdot \left(Y_{i,b_t} + \bar{R} + \sum_{b' \in \mathcal{C}(\hat{b}, G_t)} (Y_{i,b'} + R_{i,b'}) \right) \\
& \quad + (1 - \delta) \cdot \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t+1}^* = b\}} \frac{p_j}{1 - p_i} \cdot \sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}) \\
& \quad + \delta \cdot \left\{ p_i + \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} p_j \right\} \cdot \left(Y_{i,b_t} + \bar{R} + \sum_{b' \in \mathcal{C}(\hat{b}, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \right) \\
& \quad + \delta \cdot \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t+1}^* = b\}} p_j \cdot \sum_{b' \in \mathcal{C}(b, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right)
\end{aligned}$$

Since there are only longest chains as $\#C(b, G_t) = \#C(b', G_t)$ for every $b, b' \in \mathcal{T}(G_t)$, and by the definition of $b_{i,t}^*$, we have for every $\hat{b} \in \mathcal{B}^{-1}(G_t)$

$$\sum_{b' \in \mathcal{C}(\hat{b}, G_t)} (Y_{i,b'} + R_{i,b'}) \leq Y_{i,b_t} + \bar{R} + \sum_{b' \in \mathcal{C}(b_{i,t}^*, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \tag{21}$$

Plugging in, the LHS of (20) is an upper bound to the RHS of (20). \square

Lemma 3. If $Y_{i,b_t} < 0$, the condition in (6) becomes

$$\begin{aligned}
& \bar{R} + \sum_{b' \in \mathcal{C}(b_{i,t}^*, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \tag{22} \\
& \geq (1 - \delta) \cdot \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} \frac{p_j}{1 - p_i} \cdot \left(Y_{i,b_t} + \bar{R} + \sum_{b' \in \mathcal{C}(\hat{b}, G_t)} (Y_{i,b'} + R_{i,b'}) \right) \\
& + (1 - \delta) \cdot \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t+1}^* = b\}} \frac{p_j}{1 - p_i} \cdot \sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}) \\
& + \delta \cdot \left\{ p_i + \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} p_j \right\} \cdot \left(Y_{i,b_t} + \bar{R} + \frac{Y_{i,b_t}^-}{\delta} + \sum_{b' \in \mathcal{C}(\hat{b}, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right) \right) \\
& + \delta \cdot \sum_{b \in \mathcal{B}^{LC}(G_t)} \sum_{\{j \neq i: b_{j,t+1}^* = b\}} p_j \cdot \sum_{b' \in \mathcal{C}(b, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta} \cdot \Lambda_t(b', H_t) \right)
\end{aligned}$$

We now use the bounds derived in the proof of Lemma 2. Since there are only longest chains as $\#C(b, G_t) = \#C(b', G_t)$ for every $b, b' \in \mathcal{T}(G_t)$, and by the definition of $b_{i,t}^*$, we obtain the following sufficient condition for (22) to hold:

$$\bar{R} \geq (1 - \delta) \cdot \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} \frac{p_j}{1 - p_i} (Y_{i,b_t} + \bar{R}) + \delta \left\{ p_i + \sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} p_j \right\} \left(Y_{i,b_t} + \bar{R} + \frac{Y_{i,b_t}^-}{\delta} \right) \tag{23}$$

The condition is satisfied if $\sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} p_j = 1 - p_i$. Since the RHS of (23) is linear in all terms, it is either maximized or minimized at $\sum_{\{j \neq i: b_{j,t+1}^* = b_{i,t}\}} p_j = 0$. If it is minimized, then the condition is always satisfied. If it is maximized, then (9) as stated in the lemma is a sufficient condition for (22). \square

Proposition 3. The payoff of following the candidate equilibrium strategy is given by (16), replacing the set $\mathcal{B}^{LC}(G_t)$ with $\mathcal{B}^{CP}(b^{CP}(H_t), G_t)$. Consider a deviation to some block $b \notin \mathcal{J}(b^{CP}(H_t), G_t)$. Since all other miners play checkpoint strategies in all time periods, and since miner i herself reverts to playing checkpoint strategies from the next time period, no miner is working on the chain containing block b_t at time $t + 1$ and any other future time period. Hence, we have $q_{i,b_t,\tau} = 0$ and $\lambda_\tau(b_t, G_\tau) = 0$ for all $\tau \geq t + 1$. Denoting the strategy profile for this deviation by σ' , we have

$$\begin{aligned}
V_{i,t}(\sigma'; H_t) &= \sum_{b \in \mathcal{B}(G_t)} \left((1 - \delta) q_{i,b,t} (Y_{i,b} + R_{i,b}) + \frac{Y_{i,b}^-}{\delta} \cdot \lambda_t(b, H_t) \right) \tag{24} \\
&+ \sum_{b \in \mathcal{B}^{CP}(b^{CP}(H_t), G_t)} \sum_{\{j \neq i: b_{j,t}^* = b\}} p_j \cdot \delta \cdot \left(Y_{i,b,t} + \delta \cdot \frac{Y_{i,b,t}^-}{\delta} + \sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}) + \frac{Y_{i,b}^-}{\delta} \right) \\
&+ (1 - p_i) \cdot \delta^2 \cdot \mathbb{E}_t \left(Y_{i,b_{t+1}} + R_{i,b_{t+1}} + \delta \cdot \frac{Y_{i,b_{t+1}}^-}{\delta} \right) \\
&+ p_i \cdot \delta \cdot (1 - \delta) \cdot \sum_{b \in \mathcal{B}^{CP}(H_t, G_t)} \sum_{\{j \neq i: b_{j,t}^* = b\}} \frac{p_j}{1 - p_i} \sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}) \\
&+ p_i \cdot \delta \cdot \sum_{b \in \mathcal{B}^{CP}(b^{CP}(H_t), G_t)} \sum_{\{j \neq i: b_{j,t}^* = b\}} p_j \delta \left(\mathbb{E}_t \left(Y_{i,b_{t+1}} + \delta \cdot \frac{Y_{i,b_{t+1}}^-}{\delta} \right) + \sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}) + \frac{Y_{i,b}^-}{\delta} \right) \\
&+ p_i^2 \cdot \delta^2 \cdot \left(\mathbb{E}_t \left[Y_{i,b_{t+1}} + \bar{R} + \delta \cdot \frac{Y_{i,b_{t+1}}^-}{\delta} \right] + \sum_{b' \in \mathcal{C}(b_{i,t}^*, G_t)} (Y_{i,b'} + R_{i,b'}) + \frac{Y_{i,b_{i,t}^*}^-}{\delta} \right) \\
&+ \delta^3 \cdot \mathbb{E}_t \sum_{\tau=0}^{\infty} \delta^\tau \left(Y_{i,b_{t+\tau+2}} + R_{i,b_{t+\tau+2}} + \delta \cdot \frac{Y_{i,b_{t+\tau+2}}^-}{\delta} \right)
\end{aligned}$$

Note that for every $b \in \mathcal{B}^{CP}(b^{CP}(H_t), G_t)$ we have

$$\sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}) \leq \sum_{b' \in \mathcal{C}(b, G_t)} (Y_{i,b'} + R_{i,b'}) + \frac{Y_{i,b}^-}{\delta} \leq \sum_{b' \in \mathcal{C}(b_{i,t}^*, G_t)} (Y_{i,b'} + R_{i,b'}) + \frac{Y_{i,b_{i,t}^*}^-}{\delta}.$$

The first inequality follows since we only add (weakly) positive objects, the second by definition. Hence $V_{i,t}(\sigma; H_t) \geq V_{i,t}(\sigma'; H_t) + p_i \delta \left(Y_{i,b_t} + \bar{R} + \delta \cdot \frac{Y_{i,b_t}^-}{\delta} \right) > V_{i,t}(\sigma'; H_t)$. \square

Lemma 4. Consider the checkpoint strategy with checkpoint selection as in (12). For all $1 \leq s \leq t$, we have $G_{i,s} = G_s$ and $H_{i,s} = H_s$ for all $i \in \{1, \dots, N\}$, as well as $b_{s-1} = M(b_s, G_s)$. That is, all miners observe the same blockchain configuration, hold the same history, and given the strategy to extend the longest chain containing the checkpoint there is a single chain in which block b_s is chained to block b_{s-1} . Hence $b^{CP}(H_{i,s}) = b_{s-2}$ for all $2 \leq s \leq t$ and $b^{CP}(H_{i,1}) = b_0$ for all $i \in \{1, \dots, N\}$.

Let b_t and b'_t denote the blocks to block b_{t-1} . added at time t by $i \in \mathcal{N}$ and $j \in \mathcal{N}'$, respectively. At time $t + 1$, we thus have $G_{i,t+1} = G_t \cup (b_t, (b_t, b_{t-1}))$ and $H_{i,t+1} = H_t \cup G_{i,t}$ for any $i \in \mathcal{N}$, and $G_{j,t+1} = G_t \cup (b'_t, (b'_t, b_{t-1}))$ and $H_{j,t+1} = H_t \cup G_{j,t}$ for any $j \in \mathcal{N}'$. Since $b^{CP}(H_t) = M(b_{t-1}, G_t) = b_{t-2}$ and $b_{t-1} = M(b_t, G_{i,t}) = M(b'_t, G_{j,t})$ for all $i \in \mathcal{N}$ and $j \in \mathcal{N}'$, we have $b^{CP}(H_{i,t+1}) = b^{CP}(H_{j,t+1}) = b_{t-1}$ for all $i \in \mathcal{N}$ and $j \in \mathcal{N}'$.

Recall that $G_{i,s} = G_{j,s} = G_s$ for all $(i, j) \in \{1, \dots, N\}^2$ and $s \geq t + 2$. Under the checkpoint strategy, block b_{t+1} is chained to either b_t or b'_t . If $b_t = M(b_{t+1}, G_{t+2})$, and since $b^{CP}(H_{i,t+1}) = b^{CP}(H_{j,t+1}) = b_{t-1} = M(b_t, G_{t+2})$, we have $b^{CP}(H_{i,t+2}) = b^{CP}(H_{j,t+2}) = b_t$ for all $(i, j) \in \{1, \dots, N\}^2$. Similarly, if $b'_t = M(b_{t+1}, G_{t+2})$, then $b^{CP}(H_{i,t+2}) = b^{CP}(H_{j,t+2}) = b'_t$ for all $(i, j) \in \{1, \dots, N\}^2$. Then under the checkpoint strategy we have $b_{t+1} = M(b_{t+2}, G_{t+3})$ and hence $b^{CP}(H_{i,t+3}) = b^{CP}(H_{j,t+3}) = b_{t+1}$. This updating process continues in perpetuity, and thus $b^{CP}(H_{i,t+s+1}) = b^{CP}(H_{j,t+s+1}) = M(b_{t+s}, G_{t+s+1}) = b_{t+s-1}$ for all $s \geq 2$.

Consider then the checkpoint strategy with checkpoint selection as in (13). As above, we have $G_{i,s} = G_s$ and $H_{i,s} = H_s$ for all $i \in \{1, \dots, N\}$ for all $1 \leq s \leq t$, as well as $b_{s-1} = M(b_s, G_s)$. Hence $b^{CP}(H_{i,s}) = b_{s-1}$ for all $1 \leq s \leq t$. At time $t + 1$, we have $G_{i,t+1} \neq G_{j,t+1}$ and thus $b^{CP}(H_{i,t+1}) = \mathcal{B}^{CP}(b_{t-1}, G_{i,t+1}) = b_t \neq b^{CP}(H_{j,t+1}) = \mathcal{B}^{CP}(b_{t-1}, G_{j,t+1}) = b'_t$ for any $i \in \mathcal{N}$ and $j \in \mathcal{N}'$. Given the checkpoint selection as in (13) and by the definition of $\mathcal{B}^{CP}(b^{CP}(H_{i,s-1}), G_{i,s})$, any future checkpoint for miner $i \in \mathcal{N}$ must lie on the subgraph induced by the current checkpoint b_t , $J(b_t, G_s)$. Similarly, any future checkpoint for miner $j \in \mathcal{N}'$ must lie on the subgraph induced by the current checkpoint b'_t , $J(b'_t, G_s)$. Since both b_t and b'_t are chained to b_{t-1} , none of these subgraphs have a common block: $\mathcal{J}(b_t, G_s) \cap \mathcal{J}(b'_t, G_s) = \emptyset$ for all $s \geq t + 1$. Hence for all $s \geq t + 1$ and for any $i \in \mathcal{N}$ and $j \in \mathcal{N}'$, we have $b^{CP}(H_{i,s+1}) \neq b^{CP}(H_{j,s+1})$. \square

Proposition 5. Consider the checkpoint strategy with checkpoint selection as in (12) and the continuation game from time $t + 2$ onwards induced by this strategy. By the proof of Lemma 4, we have $b^{CP}(H_{i,t+2}) = b^{CP}(H_{j,t+2})$ for any $i \in \mathcal{N}$ and $j \in \mathcal{N}'$. Note that histories $\{H_{i,t+2}\}_{i=1}^N$ are common knowledge among miners; hence it is common knowledge that all miners have the same checkpoint and are called to work on the same chain. The first part of the claim then follows directly from Proposition 4.

Consider next the checkpoint strategy with checkpoint selection as in (13) and the continuation game from time $t + 2$ onwards induced by this strategy. By the proof of Lemma 4, miners disagree on the checkpoint at time $t + 2$: $b^{CP}(H_{i,t+2}) \neq b^{CP}(H_{j,t+2})$ for any $i \in \mathcal{N}$ and $j \in \mathcal{N}'$. Note that histories $\{H_{i,t+2}\}_{i=1}^N$ are common knowledge among miners; hence it is common knowledge that miners belonging to different sets disagree on the checkpoint.

We now construct a strictly profitable deviation from the candidate strategy. Suppose $\sum_{i \in \mathcal{N}} p_i \leq \sum_{j \in \mathcal{N}'} p_j$ and consider some miner $i \in \mathcal{N}$ at some time $T \geq t + 2$.²² Suppose $Y_{i,b_i,s} = 0$ for all $s \geq 0$. Define $q_{i,\mathcal{N}} = \frac{\sum_{k \in \mathcal{N}, k \neq i} p_k}{1 - p_i}$ and $q_{i,\mathcal{N}'} = \frac{\sum_{j \in \mathcal{N}'} p_j}{1 - p_i}$, and note that $q_{i,\mathcal{N}} < q_{i,\mathcal{N}'}$. Suppose all other miners follow the candidate strategy in every period $s \geq t + 2$, and miner i also follows the strategy in all these periods other than T . Hence, each subgraph branching of the respective checkpoints b_t and b'_t is a single chain at time T .

Consider the deviation to work on the terminal block of the chain branching off b'_t , $\mathcal{T}(J(b'_t, G_T))$. Hence, each subgraph branching of the respective checkpoints b_t and b'_t is a single chain without forks also in all future times $s > T$. Thus, in every period $s > T$, we have $q_{i,b,s} = q_{i,\mathcal{N}}$ for every block on the subgraph starting at b_t ; and $q_{i,b,s} = q_{i,\mathcal{N}'}$ for every block on the subgraph starting at b'_t .

Following the candidate strategy at time T implies an expected discounted life-time payoff from the block reward for the time- T block given by $p_i \delta q_{i,\mathcal{N}} \bar{R}$: miner i adds the block with probability p_i and earns block rewards \bar{R} , and if so starts enjoying the corresponding flow utility at rate $(1 - \delta)q_{i,\mathcal{N}}$ in every period from the next time period onwards. Deviating from this strategy by working on $\mathcal{T}(J(b'_t, G_T))$ yields an expected discounted life-time payoff from the block reward for the time- T block given by $p_i \delta q_{i,\mathcal{N}'} \bar{R} > p_i \delta q_{i,\mathcal{N}} \bar{R}$. This is a strictly profitable deviation. \square

Proposition 7. The payoff from the candidate equilibrium strategy profile σ reads

$$\begin{aligned}
V_{i,t}(\sigma; H_t) &= \sum_{b \in \mathcal{B}(G_t)} \left((1 - \delta) q_{i,b,t} (Y_{i,b} + R_{i,b}) + \frac{Y_{i,b}^-}{\delta^k} \cdot \lambda_t(b, H_t) \right) \tag{25} \\
&+ p_i \delta \left[\sum_{b' \in \mathcal{C}(b'_t, G_t)} \left(Y_{i,b'} + R_{i,b'} + \delta^{n(b', b'_t, H_t) - 1} \cdot \frac{Y_{i,b'}^-}{\delta^k} \cdot \Lambda(b', H_t) \right) + Y_{i,b_t} + \bar{R} + \delta^k \cdot \frac{Y_{i,b_t}^-}{\delta^k} \right] \\
&+ (1 - p_i) \delta \left[\sum_{b' \in \mathcal{C}(b'_t, G_t)} \left(Y_{i,b'} + R_{i,b'} + \delta^{n(b', b'_t, H_t) - 1} \cdot \frac{Y_{i,b'}^-}{\delta^k} \cdot \Lambda(b', H_t) \right) + Y_{i,b_t} + \delta^k \cdot \frac{Y_{i,b_t}^-}{\delta^k} \right] \\
&+ \delta^2 \cdot \mathbb{E}_t \sum_{\tau=0}^{\infty} \delta^\tau \left(Y_{i,b_t+\tau+1} + R_{i,b_t+\tau+1} + \delta^k \cdot \frac{Y_{i,b_t+\tau+1}^-}{\delta^k} \right).
\end{aligned}$$

Consider a deviation to some block $b \neq b'_t$ such that no miner including miner i works

²²Since the labels \mathcal{N} and \mathcal{N}' are arbitrary, there always exists a set of miners that satisfies this inequality.

on $b_{i,t}$ at time $t + 1$: $\sigma_{j,t+1}^T(H_t) = b_t^P$ for all $j \in \{1, \dots, N\}$. Denoting this strategy profile by σ' , the payoff $V_{i,t}(\sigma'; H_t)$ is given by (25), only that the second term on the RHS reads

$$p_i \delta \left[\sum_{b' \in \mathcal{C}(b_t^P, G_t)} \left(Y_{i,b'} + R_{i,b'} + \delta^{n(b', b_t^P, H_t)} \cdot \frac{Y_{i,b'}^-}{\delta^k} \cdot \Lambda(b', H_t) \right) \right]. \quad (26)$$

It immediately follows that $V_{i,t}(\sigma; H_t) > V_{i,t}(\sigma'; H_t)$. Consider then a deviation to some block \hat{b} such that $\sigma_{j,t+1}^T(H_t) = b_{i,t}$ for all $j \in \{1, \dots, N\}$. Denoting this strategy profile by σ'' , the payoff $V_{i,t}(\sigma''; H_t)$ is given by (25), only that the second term on the RHS reads

$$p_i \delta \left[\sum_{b' \in \mathcal{C}(\hat{b}, G_t)} \left(Y_{i,b'} + R_{i,b'} + \frac{Y_{i,b'}^-}{\delta^{k-(n(b', \hat{b}, H_t)-1)}} \cdot \Lambda(b', H_t) \right) + Y_{i,b_t} + \bar{R} + \frac{Y_{i,b_t}^-}{\delta^{k-k}} \right]. \quad (27)$$

By the construction of (15), it must be that $\mathcal{C}(\hat{b}, G_t) \subset \mathcal{C}(b_t^P, G_t)$ and $n(b', \hat{b}, H_t) > n(b', b_t^P, H_t)$. It then follows that $V_{i,t}(\sigma; H_t) \geq V_{i,t}(\sigma''; H_t)$, with the equality strict unless $Y_{i,b} = 0$ for all $b \in \mathcal{C}(b_t^P, G_t) \setminus \mathcal{C}(b^{CP}(H_t), G_t)$ and $R_{i,b} = 0$ for all $b \in \mathcal{C}(b_t^P, G_t) \setminus \mathcal{C}(\hat{b}, G_t)$.

The second claim follows from the example in the main text for $g_i(p_i) = \frac{1}{N}$ for all i . For other weights, one can always include additional transactions for arbitrary miners in block b_{f_k} to ensure that the inequality is satisfied. \square

Chapter II & III

Are cybersecurity provider platforms too concentrated? And do we need to worry?

Abstract

We study the equilibrium distribution of firms cybersecurity choices when organizations can either develop in-house security capabilities or outsource these functions to specialized provider platforms. We analyze how the presence of cybersecurity platforms affects aggregate security outcomes and welfare. Firms choose between three options: maintaining in-house security units with endogenous investment levels, purchasing standardized security services from a profit-maximizing platform, or forgoing security investment entirely. We characterize equilibria under two scenarios: random cyber incidents with exogenous probability, and strategic attackers who optimally choose their technological sophistication, thus endogenizing the probability of a successful attack.

Our analysis reveals that platform concentration creates negative externalities that harm both platform clients and others. When a significant share of firms outsource security to platforms, strategic attackers respond by increasing their technological sophistication to exploit the concentrated attack surface. This raises the optimal security investment level for firms maintaining in-house capabilities, imposing costs on larger firms that would otherwise prefer self-protection. Moreover, we demonstrate that under certain parameter values, smaller firms using platform services experience higher rates of successful attacks than larger firms with in-house security.

In the third chapter we construct a novel dataset by exploiting a recent Securities and Exchange Commission rule requiring publicly traded companies to disclose material cybersecurity incidents within four business days after they determine that such incident has taken place. The incidents are reported as part of the form 8-K (current events). Our empirical analysis reveals that smaller firms experience cybersecurity incidents significantly more frequently than larger firms. This finding is consistent with theoretical predictions that smaller firms face greater vulnerability. To assess market perceptions of cyber risk, we conduct event studies examining stock price reactions to cybersecurity disclosures for firms in our sample. Using both market-adjusted returns and the CAPM, we calculate abnormal returns during a two-week window surrounding disclosure dates. Our analysis reveals negative abnormal returns with significant t-test values observed in the post

disclosure period, while cumulative abnormal returns do not follow a conventional event study structure. The magnitude of responses is larger for smaller firms, suggesting that capital markets perceive cyber incidents as more damaging to smaller companies. This differential market reaction likely reflects smaller firms' more limited resources for incident response, recovery, and reputation management. The combination of higher incident frequency and larger market reactions suggests that smaller firms face a double burden: they are both more frequently targeted and suffer worse consequences when incidents occur.

Contents

| | |
|--|-----------|
| 1 Introduction | 3 |
| 2 Literature Review | 4 |
| 3 Model | 7 |
| 3.1 Optimal investment levels | 8 |
| 3.2 Cybersecurity provider platform optimization problem | 10 |
| 4 Strategic attacker's optimization problem | 11 |
| 4.1 Optimal investment levels | 11 |
| 4.2 Platform's Optimization Problem | 12 |
| 4.3 Attacker's Problem | 13 |
| 4.4 Comparative Statics at Equilibrium | 14 |
| 5 Data | 18 |
| 5.1 Event Study | 19 |
| 5.1.1 Event Study without CAPM | 20 |
| 5.1.2 Event Study with CAPM | 20 |
| 6 Conclusion | 22 |

1 Introduction

The rapid digitalization of economic activity has fundamentally transformed how organizations manage cybersecurity risks. As cyber threats become increasingly sophisticated and persistent, firms face a critical strategic decision: should they develop and maintain in-house cybersecurity capabilities, or should they outsource these functions to specialized platform providers? This question has gained particular urgency in recent years as the cybersecurity market has witnessed substantial consolidation, with a relatively small number of large platforms commanding significant market share and serving thousands of client organizations. This concentration in the cybersecurity platform market raises important economic and policy questions. When numerous firms rely on the same provider platform for their security infrastructure, does this create systemic vulnerabilities? Are smaller firms, which may lack the resources to develop sophisticated in-house capabilities, disproportionately exposed to risk through their dependence on shared platforms? How do strategic attackers respond to this market structure, and does platform concentration inadvertently create attractive targets for sophisticated cyber adversaries? This paper addresses these questions by developing a model of firm security choices in the presence of platform providers. We analyze how firms of different sizes optimally allocate resources between in-house security investments and platform outsourcing. First, we consider a baseline scenario where cyber incidents occur randomly with exogenous probability. This allows us to characterize the fundamental economic trade-offs that determine firms security choices and platform pricing strategies. Second, we extend the model to incorporate a strategic attacker who endogenously chooses its technological sophistication in anticipation of the distribution of firm security choices. We demonstrate that the existence of cybersecurity platforms creates a negative externality on firms that choose to maintain in-house security capabilities. When platforms attract a significant share of the market, strategic attackers respond by increasing their technological sophistication, which raises the optimal level of security investment required for self-protection. Consequently, larger firms that would prefer in-house security face higher costs due to platform-induced changes in attacker behavior. Furthermore, we show that under certain conditions, the expected measure of successful attacks on smaller firms using platform services exceeds that of larger firms with in-house capabilities. This finding is consistent with what we observe in the data: that smaller firms face greater vulnerability.

Lastly, we conduct event studies examining stock price reactions to cybersecurity disclosures for firms in our sample. We observe negative abnormal returns with significant t-test values in the post disclosure period with larger magnitude of responses for smaller firms, suggesting that capital markets perceive cyber incidents as more damaging to smaller companies. The combination of higher incident frequency and larger market reactions suggests that smaller firms face a double burden: they are both more frequently targeted and suffer worse consequences when incidents occur.

These results contribute to ongoing policy debates about cybersecurity regulation. Our analysis suggests that while platforms may offer economies of scale and specialized expertise, their prevalence

can inadvertently reshape the threat landscape in ways that increase aggregate risk exposure. These implications are particularly stark for small enterprises, which may face a constrained choice set between expensive in-house security and potentially vulnerable platform dependence.

2 Literature Review

Our paper contributes to several interrelated strands of literature in cybersecurity economics, industrial organization, and platform markets. This review organizes the relevant literature into five main themes: the economics of cybersecurity investment, platform markets and network effects, strategic attacker models, concentration and systemic risk, and the outsourcing decision in security contexts.

The economic analysis of cybersecurity investment decisions has developed substantially since the foundational work of [Gordon and Loeb \(2002\)](#), who demonstrated that optimal security investment levels are not monotonically increasing in the value of assets at risk. Their model showed that firms should invest most heavily in protecting assets with intermediate vulnerability levels, as very low vulnerability assets require little protection and very high vulnerability assets may not justify the investment required to secure them adequately. This counterintuitive result has important implications for understanding heterogeneous security choices across firms and asset classes.

Subsequent research has extended this framework to incorporate various forms of interdependence and externalities. [Kunreuther and Heal \(2003\)](#) analyzed cybersecurity as a weakest-link public good problem, where the overall security of an interconnected system depends on the least-protected participant. This creates free-rider problems and potential coordination failures, as individual firms may underinvest in security relative to the social optimum. [Varian \(2004\)](#) further explored system reliability when individual components have different security levels, showing that liability rules and mandatory security standards can mitigate underinvestment problems.

More recent contributions have incorporated behavioral factors and information asymmetries into security investment models. [Cavusoglu et al. \(2008\)](#) demonstrated that firms systematically undervalue security breaches that occur with low probability but high impact, leading to persistent underinvestment. [Anderson and Moore \(2006\)](#) argued that many security failures result not from insufficient investment but from misaligned incentives between different stakeholders in complex value chains. Our paper builds on this tradition by examining how platform intermediation affects the incentive structure for security investment.

The economics of platform markets has received extensive attention in recent industrial organization literature. [Rochet and Tirole \(2003\)](#) provided the foundational analysis of two-sided markets, where platforms facilitate interactions between distinct user groups and must carefully balance pricing across sides to maximize participation and welfare. Their work demonstrated that platforms may subsidize one side of the market while extracting surplus from the other, and that platform competition

can lead to inefficient fragmentation or excessive concentration depending on the strength of network effects.

[Rochet and Tirole \(2006\)](#) further developed the theory of two-sided markets, examining how pricing structures affect platform adoption and welfare. They showed that the allocation of fees between platform sides depends critically on the relative price elasticities of demand and the strength of cross-side network effects. In markets with strong network effects, platforms may compete aggressively for market share, potentially leading to winner-take-all outcomes.

[Armstrong \(2006\)](#) extended this analysis to incorporate multi-homing behavior, where users may simultaneously participate on multiple competing platforms. The possibility of multi-homing fundamentally changes competitive dynamics and can reduce platforms' ability to exercise market power. However, in the cybersecurity context, multi-homing may be prohibitively costly due to integration challenges and operational complexity, potentially strengthening lock-in effects and platform market power.

[Evans and Schmalensee \(2016\)](#) synthesized the platform markets literature and emphasized the importance of understanding platform pricing strategies, entry dynamics, and the role of data in creating competitive moats. In cybersecurity platforms specifically, data aggregation across clients may generate positive spillovers through improved threat detection and response capabilities. However, as our model demonstrates, concentration may also create negative spillovers by attracting more sophisticated attacks.

Recent work has examined the competitive effects of dominant platforms in technology markets. [Cabral \(2020\)](#) analyzed the welfare implications of platform dominance and market tipping, showing that winner-take-all dynamics can emerge even when multiple platforms could coexist in principle. This literature informs our analysis of concentration in cybersecurity platform markets and the welfare consequences for firms at different points in the size distribution.

The incorporation of strategic, profit-maximizing attackers represents an important development in cybersecurity economics. Traditional models often treat attack probability as exogenous or as a simple decreasing function of defense expenditure. However, realistic attacker behavior requires optimizing over both target selection and attack sophistication in response to the observed distribution of defenses.

[Grossklags et al. \(2008\)](#) developed game-theoretic models of security investment with strategic attackers and defenders, demonstrating that Nash equilibria often feature inefficiently low defense and inefficiently high attack expenditure relative to the social optimum. The presence of strategic attackers fundamentally alters the nature of security investments from a purely technical problem to a strategic interaction with multiple equilibria and potential coordination failures.

[Fultz and Grossklags \(2009\)](#) examined attacker behavior under uncertainty about defender capabilities, showing that information asymmetries can lead to both over- and under-investment in attacks depending on the attacker's beliefs about the distribution of security levels. This insight motivates

our analysis of how platform concentration affects attacker information and strategic choices.

More recently, [Liu et al. \(2015\)](#) analyzed dynamic security investment games where attackers learn about defense effectiveness through repeated interactions. They demonstrated that reputational concerns and learning dynamics can significantly affect equilibrium security levels. Our model contributes to this literature by examining how market structure—specifically the presence of large platform providers—shapes attacker incentives and optimal attack sophistication.

The relationship between market concentration and systemic risk has been extensively studied in financial economics, with important lessons for cybersecurity contexts. The financial crisis literature emphasized how interconnected institutions and common exposures can amplify shocks and create systemic vulnerabilities. [Acemoglu et al. \(2015\)](#) developed network models showing that financial systems with more concentrated connection patterns are more susceptible to cascading failures, even when individual institutions appear well-capitalized.

In the context of cybersecurity, [Böhme and Kataria \(2006\)](#) were among the first to analyze correlated security failures and systemic risk in interconnected networks. They demonstrated that diversity in security implementations can serve as a form of portfolio diversification, reducing the probability of widespread simultaneous failures. This insight suggests that excessive concentration in cybersecurity platforms may increase systemic vulnerability by creating common mode failures.

[Geer et al. \(2003\)](#) made an influential argument that software monoculture—specifically the dominance of particular operating systems and applications—creates systemic vulnerability by allowing exploits to propagate widely. Their analysis emphasized that security through diversity can be a more robust strategy than relying on the security of any single implementation, however well-designed. This argument extends naturally to the concentration of firms on a small number of cybersecurity platform providers.

Recent empirical work has begun to quantify the extent and consequences of cybersecurity platform concentration. [Romanosky et al. \(2019\)](#) analyzed data breach patterns and found that incidents affecting third-party service providers have substantially larger downstream impacts than direct breaches, supporting concerns about concentration risk. Our theoretical model provides a framework for understanding the mechanisms behind these empirical patterns.

The decision to outsource security functions involves classic make-or-buy trade-offs familiar from transaction cost economics and organizational theory. [Coase \(1937\)](#) and [Williamson \(1979\)](#) established that firms choose to internalize activities when transaction costs, hold-up problems, or asset specificity make market contracting inefficient. Security represents a particularly interesting application of these principles, as it involves both specialized expertise that may be efficiently provided by external specialists and firm-specific knowledge that may be difficult to transfer.

[Grossman and Hart \(1986\)](#) and [Hart and Moore \(1990\)](#) developed property rights theories of firm boundaries that emphasize how ownership affects investment incentives in the presence of incomplete

contracts. When security investments are non-contractible or difficult to verify, the allocation of residual control rights determines who has appropriate incentives to invest in prevention and response capabilities. Our model implicitly assumes that platform contracts are incomplete, with platforms choosing security levels that maximize their profits rather than client welfare.

In the specific context of IT and security outsourcing, [Clemons and Row \(1992\)](#) analyzed how information technology affects vertical integration decisions, showing that IT can both enable efficient outsourcing through better monitoring and coordination, and create switching costs that lock firms into particular providers. The cybersecurity context features both effects: platforms may offer genuine efficiency gains through specialization and scale, but also create dependencies that limit firms' ability to switch or multi-home.

[Herath and Herath \(2008\)](#) specifically examined organizational decisions to outsource information security functions, finding that firms weigh the benefits of specialized expertise and scale economies against concerns about loss of control, vendor opportunism, and the difficulty of writing complete contracts for security services. Their survey evidence suggests that these concerns are particularly acute for firms in regulated industries or those handling sensitive data.

Our paper synthesizes insights from these various literatures to analyze how platform concentration emerges endogenously from firm optimization, how it affects aggregate security outcomes, and whether the resulting equilibria are socially efficient. By incorporating both firm heterogeneity and strategic attacker behavior, we provide new insights into the welfare implications of cybersecurity platform markets and the potential justification for regulatory intervention.

3 Model

We consider the following model. At time 0 a cybersecurity platform set its price p . At time $t = 1$ firms choose to have in-house cybersecurity, outsource it to a provider, or not invest in their security at all. We have a continuum of firms $\pi_i \in [0, 1]$. At $t = 3$ a cybersecurity incident takes place with probability δ and payoffs are realized.

- In-house division has fixed cost f . Additionally, $2sm$ is the marginal cost to increase the level of security per unit of firm size.
So, for $s > 0$ level of security, the in-house cost is $c = f + s^2m\pi$. f and m are fixed numbers.
- Outsourcing to a provider has cost $c = p$ per unit of firm size that provides fixed $s = l$ level of security. p and l are fixed numbers.
- No security has $c = 0$ cost and $s = 0$.

Firm has expected profit π . $\pi > 0$ is fixed. This is what characterizes the firm.

There is δ probability of an incident. The success probability of the attack is a decreasing function of s , $\rho(s)$. If the attack succeeds, firm loses $\alpha\pi$. Here $\delta, \alpha >$ are fixed numbers.

Then the expected payoff of the firm is

$$\pi - \rho(s)\delta\alpha\pi - (\text{cyber investment})$$

Without loss of generality we can ignore profit. Therefore, payoffs for each of the options could be characterized as:

- in-house cybersecurity: $-\rho(s)\delta\alpha\pi - (f + s^2m\pi)$
- outsourcing to the platform $-\rho(l)\delta\alpha\pi - p\pi$
- no cybersecurity investment $-\rho(0)\delta\alpha\pi$

Therefore, each firm solves a piecewise maximization problem. We take $\rho(s) = 1 - s$ that is suitable and tractable. The maximization problem becomes:

$$\max_s \begin{cases} -(1-s)\delta\alpha\pi - (f + s^2m\pi) & s \neq l \\ -(1-l)\delta\alpha\pi - p\pi & s = l \\ -\delta\alpha\pi & s = 0 \end{cases} \quad (3.1)$$

We find the optimal investment levels in the next sub section.

3.1 Optimal investment levels

For $s \neq l$, the FOC is

$$\delta\alpha\pi - 2m\pi = 0 \rightarrow s_i^* = \frac{\delta\alpha}{2m} \quad (3.2)$$

Therefore the optimal profit for firm i is $\hat{\pi}_i = -(1 - \frac{\delta\alpha}{2m})\delta\alpha\pi - (f + (\frac{\delta\alpha}{2m})^2m\pi) = -\delta\alpha\pi[1 - \frac{\delta\alpha}{2m}] - f$. That is if firm i chooses to have an in-house cybersecurity unit. Note that this only holds if $\frac{\delta\alpha}{2m} < 1 \rightarrow \delta\alpha < 2m$. Because we want $s \geq 0$. So if $\delta\alpha \geq 2m$, then the optimal level of investment would be $s^* = 0$ and not investing in cybersecurity at all would make the firm better off.

Hence firm i chooses to have an in-house unit if this optimal profit is greater than the profit in the case of outsourcing and in the case of no security:

$$-\delta\alpha\pi[1 - \frac{\delta\alpha}{2m}] - f \geq -(1-l)\delta\alpha\pi_i - p\pi_i \rightarrow \delta\alpha\pi_i[\frac{\delta\alpha}{2m} - l] + p\pi_i \geq f \rightarrow \pi_i[\delta\alpha(\frac{\delta\alpha}{2m} - l) - p] \geq f \quad (3.3)$$

So the threshold firm size is

$$\hat{\pi} = \frac{f}{\delta\alpha(\frac{\delta\alpha}{2m} - l) - p} \quad (3.4)$$

And for the in-house choice to be better than having no security at all, we need

$$-\delta\alpha\pi[1 - \frac{\delta\alpha}{2m}] - f \geq -\delta\alpha\pi \rightarrow \frac{(\delta\alpha\pi)^2}{2m} \geq f \quad (3.5)$$

Note that if the denominator of $\hat{\pi}$ is negative, that means no firm chooses to have an in-house security unit. Likewise firm i chooses to outsource its security if $\pi_i < \hat{\pi}$ and

$$-(1-l)\delta\alpha\pi - p\pi \geq -\delta\alpha\pi \rightarrow p \leq \delta\alpha l \quad (3.6)$$

Therefore, we have three regions:

$$\begin{cases} s = 0 & \frac{(\delta\alpha\pi)^2}{2m} < f \text{ and } p > \delta\alpha l \\ s = l & \pi_i < \hat{\pi} \\ s = s^* & \pi_i \geq \hat{\pi} \end{cases} \quad (3.7)$$

As we can see, when f or m are larger, $\hat{\pi}$ is larger. This makes sense because these two parameters capture the cost of having an in-house cybersecurity unit. So if they are large, only bigger firms could afford them and a bigger section of firms will opt to use the services of the platform instead.

Similarly if l is larger, $\hat{\pi}$ is also larger. This is because if the security level provided by the platform is sufficiently big, more firms opt to use its services.

On the other hand if δ or α are large, $\hat{\pi}$ is smaller. This indicates that if probability of an attack or the losses incurred in case of an attack are large enough, more firms will opt to have their own cybersecurity unit.

Lastly if p is large, $\hat{\pi}$ is small. This clearly means that if the price platform is charging is high, more firms find it optimal to invest in their own cybersecurity unit.

3.2 Cybersecurity provider platform optimization problem

Given the firms choices, the platform solves its own profit maximization problem by choosing its price. The platform's profit is

$$P = \int_{\pi_{min}}^{\hat{\pi}} p\pi_i di = \frac{p}{2}[\hat{\pi}^2 - \pi_{min}^2] \quad (3.8)$$

The provider's FOC is therefore $\hat{\pi}(p) + 2p\hat{\pi}'(p) = 0$. We get $p = -\frac{1}{2} \times \frac{\hat{\pi}(p)}{\hat{\pi}'(p)}$. Where we have

$$\hat{\pi}'(p) = -\frac{f}{(p + \delta\alpha(\frac{\delta\alpha}{4m} - l))^2} \quad (3.9)$$

Substituting this we get

$$p = -\frac{1}{2} \times \frac{\hat{\pi}(p)}{\hat{\pi}'(p)} = \frac{1}{2}(p + \delta\alpha(\frac{\delta\alpha}{4m} - l)) \quad (3.10)$$

From here we derive the provider's optimal price:

$$p^* = \delta\alpha(\frac{\delta\alpha}{4m} - l) \quad (3.11)$$

Therefore the threshold firm size is

$$\pi^* = \frac{1}{2} \times \frac{f}{\delta\alpha(\frac{\delta\alpha}{4m} - l)} \quad (3.12)$$

Note that if $p^* > \delta\alpha l$, instead of choosing p^* the platform simply chooses $\delta\alpha l$ as its price. That way it can still capture part of the firm distribution, while otherwise it will lose all of its clients. As we can see, when m or l are large, p^* is small and when $\delta\alpha$ is large, p^* is also large.

The equilibrium in this environment is defined as such. The platform chooses its price $p \in (0, \infty)$ such that it maximizes its payoff $P(p, \hat{\pi}(p))$. Firm i chooses its security level $s_i \in S_i$ where $S_i = [0, 1]$ such that it maximizes its profit $profit_i(s_i, p)$. In equilibrium firms choices are as follows:

$$\begin{cases} s = l & \pi_i < \hat{\pi} \\ s = s^* & \pi_i \geq \hat{\pi} \end{cases} \quad (3.13)$$

4 Strategic attacker's optimization problem

In this section, we introduce a strategic attacker. So instead of an attack happening at random, we have an attacker that maximizes its own profit function, by choosing a parameter k which represents the technological sophistication. So at time $t = 0$ the attacker invests in its technological sophistication, and at $t = 3$ carries out an attack and the payoffs realize.

Under this assumption, the environment changes and the probability of a successful attack can be defined as $\min\{k(1-s), 1\}$. Therefore, the firm piecewise maximization problem becomes:

$$\max_s \begin{cases} -k(1-s)\delta\alpha\pi - (f + s^2m\pi) & s \neq l \\ -k(1-l)\delta\alpha\pi - p\pi & s = l, \\ -\delta\alpha\pi & s = 0 \end{cases} \quad (4.1)$$

4.1 Optimal investment levels

For $s \neq l$, the FOC is

$$k\delta\alpha\pi - 2sm\pi = 0 \rightarrow s_i^* = \frac{\delta\alpha k}{2m} \quad (4.2)$$

Therefore the optimal profit for firm i is

$$-k(1 - \frac{\delta\alpha k}{2m})\delta\alpha\pi - (f + (\frac{\delta\alpha k}{2m})^2m\pi) = -k\delta\alpha\pi[1 - \frac{\delta\alpha k}{4m}] - f \quad (4.3)$$

That is if firm i chooses to have an in-house cybersecurity unit.

Hence firm i chooses to have an in-house unit if this optimal profit is greater than the profit in the case of outsourcing:

$$-k\delta\alpha\pi[1 - \frac{\delta\alpha k}{4m}] - f \geq -k(1-l)\delta\alpha\pi - p\pi \rightarrow \pi_i[p + k\delta\alpha(\frac{\delta\alpha k}{4m} - l)] \geq f \quad (4.4)$$

So

$$\hat{\pi} = \frac{f}{p + k\delta\alpha(\frac{\delta\alpha k}{4m} - l)} \quad (4.5)$$

Therefore, we have three regions again:

$$\begin{cases} s = 0 & \frac{(\delta\alpha k\pi)^2}{2m} < f \text{ and } p > \delta\alpha l \\ s = l & \pi_i < \hat{\pi} \\ s = s^* & \pi_i \geq \hat{\pi} \end{cases} \quad (4.6)$$

4.2 Platform's Optimization Problem

Now we proceed to solve the platform's problem again, considering that the attacker is strategically choosing its technological sophistication.

Given the firms choices, the platform solves its own profit maximization problem by choosing its price. The platform's profit is

$$P = \int_{\pi_{min}}^{\hat{\pi}} p\pi_i di = \frac{p}{2}\hat{\pi}^2 \quad (4.7)$$

The provider's FOC is therefore $\hat{\pi}(p) + 2p\hat{\pi}'(p) = 0$. We get $p = -\frac{1}{2} \times \frac{\hat{\pi}(p)}{\hat{\pi}'(p)}$. Where we have

$$\hat{\pi}'(p) = -\frac{f}{(p + k\delta\alpha(\frac{\delta\alpha k}{4m} - l))^2} \quad (4.8)$$

Substituting this we get

$$p = -\frac{1}{2} \times \frac{\hat{\pi}(p)}{\hat{\pi}'(p)} = \frac{1}{2}(p + k\delta\alpha(\frac{\delta\alpha k}{4m} - l)) \quad (4.9)$$

From here we derive the provider's optimal price:

$$p^* = k\delta\alpha(\frac{\delta\alpha k}{4m} - l) \quad (4.10)$$

Therefore the threshold firm size is

$$\pi^* = \frac{1}{2} \times \frac{f}{k\delta\alpha(\frac{\delta\alpha k}{4m} - l)} \quad (4.11)$$

4.3 Attacker's Problem

We define the attacker's utility function as

$$k(1-s)\mathcal{P}(i) - C(k) \quad (4.12)$$

where $\mathcal{P}(\pi_i)$ is the payoff of successfully attacking firm i and $C(k)$ is the cost associated with the sophistication level k . We assume the cost function $C(k) = \gamma k^2$ and assume $\mathcal{P}(i) = 1$.

The attacker's optimization problem becomes

$$\begin{aligned} \max_k \int_0^\pi \rho(l, k) di + \int_\pi^1 \rho(s_i, k) di - \gamma k^2 &= \max_k \int_0^\pi k(1-l) di + \int_\pi^1 k(1-s) di - \gamma k^2 = \\ \max_k k(1-l) \int_0^\pi di + k(1-s) \int_\pi^1 di - \gamma k^2 &= \max_k k(1-l) \times \hat{\pi} + k(1-s)(1-\hat{\pi}) - \gamma k^2 \end{aligned} \quad (4.13)$$

First order condition gives

$$(1-l)\hat{\pi} + (1-s)(1-\hat{\pi}) - 2\gamma k = 0 \quad (4.14)$$

Substituting for $\hat{\pi}(k)$ and $s^*(k) = \frac{\delta\alpha k}{2m}$ we get

$$\left\{ \frac{(\delta\alpha)^2}{4m} \left(2\gamma + \frac{\delta\alpha}{2m} \right) \right\} k^2 - \delta\alpha \left\{ \frac{(\delta\alpha)}{4m} + l \left(2\gamma + \frac{\delta\alpha}{2m} \right) \right\} k - \delta\alpha \left(\frac{f}{4m} - l \right) = 0 \quad (4.15)$$

We multiply the above by $\frac{4m}{\alpha\delta}$ to get

$$\left\{ \delta\alpha \left(2\gamma + \frac{\delta\alpha}{2m} \right) \right\} k^2 - \left\{ (\delta\alpha) + 4m \times l \left(2\gamma + \frac{\delta\alpha}{2m} \right) \right\} k - (f - 4m \times l) = 0 \quad (4.16)$$

So we get

$$k^* = \frac{\left\{ (\delta\alpha) + 4ml \left(2\gamma + \frac{\delta\alpha}{2m} \right) \right\} + / - \sqrt{\left\{ (\delta\alpha) + 4ml \left(2\gamma + \frac{\delta\alpha}{2m} \right) \right\}^2 + 4 \left\{ \delta\alpha \left(2\gamma + \frac{\delta\alpha}{2m} \right) \right\} (f - 4ml)}}{2 \left\{ \delta\alpha \left(2\gamma + \frac{\delta\alpha}{2m} \right) \right\}} \quad (4.17)$$

Which in turn leads to

$$k^* = \frac{\left\{ (\delta\alpha) + 4ml\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\} + \sqrt{\left\{ (\delta\alpha) - 4ml\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}^2 + \left\{ 4\delta\alpha f\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}}}{2\left\{ \delta\alpha\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}} \quad (4.18)$$

We can now calculate the equilibrium values:

$$s^* = \delta\alpha \times \frac{\left\{ (\delta\alpha) + 4ml\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\} + \sqrt{\left\{ (\delta\alpha) - 4ml\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}^2 + \left\{ 4\delta\alpha f\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}}}{4m\left\{ \delta\alpha\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}} \quad (4.19)$$

At equilibrium, first the attacker chooses its technological sophistication $k \in [1, \infty]$. Then the platform chooses $p \in (0, \infty)$ to maximize its payoff $P(p, \hat{\pi})$. Firms choose their level of security $s_i \in S_i$ to maximize their profit. And lastly the attack takes place and the payoffs are realized.

4.4 Comparative Statics at Equilibrium

In this section we study how equilibrium values move with l . Let us start with k . Derivative of k wrt l is

$$\frac{\partial k}{\partial l} = \frac{\left\{ 4m\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\} - \frac{\left\{ (\delta\alpha) - 4ml\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\} \left(4m\left(2\gamma + \frac{\delta\alpha}{2m}\right)\right)}{\sqrt{\left\{ (\delta\alpha) - 4ml\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}^2 + \left\{ 4\delta\alpha f\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}}}}{2\left\{ \delta\alpha\left(2\gamma + \frac{\delta\alpha}{2m}\right) \right\}} \quad (4.20)$$

I show in the appendix that $\frac{\partial k}{\partial l} > 0$. For illustration, take the case where $l = 0$, i.e. there is no platform in the market. In this case the attacker's payoff function becomes

$$\int_0^1 k(1-s)di = k(1-s) = k\left(1 - \frac{\delta\alpha k}{2m}\right) \quad (4.21)$$

Optimal k is derived by solving the FOC: $1 - \frac{\delta\alpha}{m}k = 0 \rightarrow k^* = \frac{m}{\delta\alpha}$. This is smaller than the k^* in the case where we have a platform. Therefore, there is a spillover from existence of platform onto the firms that are not its clients: existence of platforms incentivizes the attacker to invest more in its technological sophistication, and as a result, the firms that are not the platform's clients also suffer from this.

Next, let us take a look at the derivative wrt l of expected measure of successful attacks for the firms that choose to have an in-house unit of cyber security. This expected value is

$$\int_{\hat{\pi}}^1 \rho(k^*, s^*) di = k^* \times (1 - s^*)(1 - \hat{\pi}) \quad (4.22)$$

$$\begin{aligned} \frac{\partial \int_{\hat{\pi}}^1 \rho(k^*(l), s^*(l)) di}{\partial l} &= \frac{\partial \int_{\hat{\pi}}^1 \rho(k^*, s^*) di}{\partial k} \times \frac{\partial k}{\partial l} + \frac{\partial \int_{\hat{\pi}}^1 \rho(k^*, s^*) di}{\partial l} = \\ & \left[\left(1 - \frac{\delta \alpha k}{m}\right)(1 - \hat{\pi}) + k \left(1 - \frac{\delta \alpha k}{2m}\right)(1 - \hat{\pi}') \right] \times \frac{\partial k}{\partial l} + k \left(1 - \frac{\delta \alpha k}{2m}\right) \left(1 - \frac{\partial \hat{\pi}}{\partial l}\right) = \\ & \left[\left(1 - \frac{\delta \alpha k}{m}\right)(1 - \hat{\pi}) + k \left(1 - \frac{\delta \alpha k}{2m}\right) \left(1 + \frac{f(\frac{\delta \alpha k}{2m} - l)}{2\delta \alpha k^2 (\frac{\delta \alpha k}{4m} - l)^2}\right) \right] \times \frac{\partial k}{\partial l} + \\ & k \left(1 - \frac{\delta \alpha k}{2m}\right) \left(1 - \frac{f}{2\delta \alpha k (\frac{\delta \alpha k}{4m} - l)^2}\right) \end{aligned} \quad (4.23)$$

Again we show in the appendix that this expression is positive. Therefore, we can conclude that the measure of successful attacks for the firms with in-house security unit increases with l .

Next, let us take a look at the derivative wrt l of expected measure of successful attacks for the firms that choose to have an in-house unit of cyber security. This expected value is

$$\int_0^{\hat{\pi}} \rho(k^*, l) di = k^* \times (1-l)\hat{\pi} \quad (4.24)$$

$$\frac{\partial \int_0^{\hat{\pi}} \rho(k^*(l), l) di}{\partial l} = \frac{\partial \int_0^{\hat{\pi}} \rho(k^*, l) di}{\partial k} \times \frac{\partial k}{\partial l} + \frac{\partial \int_0^{\hat{\pi}} \rho(k^*, l) di}{\partial l} = \quad (4.25)$$

$$[(1-l)\hat{\pi} + k(1-l)\hat{\pi}'] \times \frac{\partial k}{\partial l} - k\hat{\pi} + k(1-l) \frac{\partial \hat{\pi}}{\partial l} =$$

$$(1-l)(\hat{\pi} + k \frac{f(\frac{\delta\alpha k}{2m} - l)}{2\delta\alpha k^2(\frac{\delta\alpha k}{4m} - l)^2}) \times \frac{\partial k}{\partial l} + k[(1-l) \frac{f}{2\delta\alpha k(\frac{\delta\alpha k}{4m} - l)^2} - \hat{\pi}] =$$

$$(1-l)(\hat{\pi} + k \frac{f(\frac{\delta\alpha k}{2m} - l)}{2\delta\alpha k^2(\frac{\delta\alpha k}{4m} - l)^2}) \times \frac{\partial k}{\partial l} + \quad (4.26)$$

$$k \frac{f}{2\delta\alpha k(\frac{\delta\alpha k}{4m} - l)} \left[\frac{1-l}{\frac{\delta\alpha k}{4m} - l} - 1 \right] \quad (4.27)$$

Note that [4.26](#) at equilibrium is positive. We also know that $k \frac{f}{2\delta\alpha k(\frac{\delta\alpha k}{4m} - l)}$ in [4.27](#) is positive. Therefore, we only need to find the sign of $[\frac{1-l}{\frac{\delta\alpha k}{4m} - l} - 1]$. We show that below:

$$\frac{\delta\alpha k}{4m} - l < \frac{\delta\alpha k}{2m} - l < 1-l \rightarrow \frac{1-l}{\frac{\delta\alpha k}{4m} - l} > 1 \rightarrow \frac{1-l}{\frac{\delta\alpha k}{4m} - l} - 1 > 0 \quad (4.28)$$

Again we show derivative of the measure of successful attacks for firms that outsource their security is positive in wrt l .

Lastly, to compare s^* and l , let's take a look at $\frac{s^*}{l}$:

$$\frac{s^*}{l} = \delta\alpha \times \frac{\left\{ (\delta\alpha) + 4ml(2\gamma + \frac{\delta\alpha}{2m}) \right\} + \sqrt{\left\{ (\delta\alpha) - 4ml(2\gamma + \frac{\delta\alpha}{2m}) \right\}^2 + \left\{ 4\delta\alpha f(2\gamma + \frac{\delta\alpha}{2m}) \right\}}}{4ml \left\{ \delta\alpha(2\gamma + \frac{\delta\alpha}{2m}) \right\}} =$$

$$1 + \delta\alpha \times \frac{\delta\alpha + \sqrt{\left\{ (\delta\alpha) - 4ml(2\gamma + \frac{\delta\alpha}{2m}) \right\}^2 + \left\{ 4\delta\alpha f(2\gamma + \frac{\delta\alpha}{2m}) \right\}}}{4ml \left\{ \delta\alpha(2\gamma + \frac{\delta\alpha}{2m}) \right\}} > 1$$

We see that $\frac{s^*}{l} > l$, and therefore, $s^* > l \rightarrow 1 - s^* < 1 - l$. This shows that in equilibrium, small firms who chose to outsource their security are more susceptible to cyber attacks than their larger counterparts. This result is what we also see in data.

5 Data

On July 26 2023 the SEC adopted the long anticipated final rules on cybersecurity risk management, strategy, governance, and incident disclosure for issuers. The rules became effective on December 18 2023. The new rules are part of the SEC's larger efforts focused on cybersecurity regulation with a growing universe of rules aimed at different types of SEC registrants. The rules introduce three new types of disclosure requirements relating to

- material cybersecurity incidents
- cybersecurity risk management processes
- cybersecurity management and governance

The disclosures are reported as part of the form 8-K (current events). We scraped this data for the past year to build a dataset of firms and their reported cyber incidents. We study if there is a relationship between firms size and the frequency of being targeted.

Form 8-K is a report that publicly traded U.S. companies must file with the Securities and Exchange Commission (SEC) to announce "material events" that may affect their financial condition or operations. Since December 2023 any material cybersecurity risk must also be reported using form 8-K within four days of the incident taking place. Scraping these reports I constructed a dataset that details firms and any reported cyber incident that took place in the past year. I also document if they reported any cyber resilience and defense strategies in their 8-k. The dataset also contains the firms size (revenue in the past 12 months).

We seek to find out, using this dataset, if there is any relationship between a firm's size and how frequently a firm experiences cyber incidents. i.e. whether smaller or larger firms are more likely to experience cyber threats. The model suggests that the probability of successful attack is larger for smaller firms and what we find in the data [\[1\]](#) corroborates this result.

We also conduct an event study using this dataset to find if the stock market responds to these announcements. We find out that the abnormal returns are negative with significant t-test values and that the effect is larger for smaller firms. So small firms are both more frequently targeted, and once attacked, they experience larger losses.

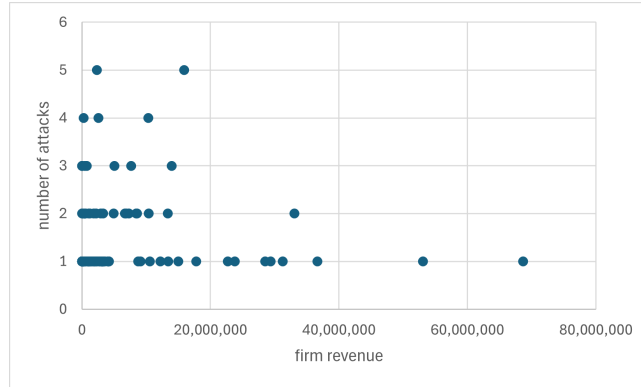


Figure 1: scatter plot of frequency of attacks vs firms revenues

5.1 Event Study

In this section I conduct an event study to evaluate whether filing of report 8-k had any significant influence on the stock prices of the firms that filed this form in the past year. Finance theory suggests that capital markets reflect all available information about firms in the firms' stock prices. Following this premise, one can study how a particular event changes a firm's prospects by quantifying the impact of the event on the firm's stock. In its most common form, it focuses on stock returns, and in less used forms, it focuses on trading volumes and volatilities.

Return event studies quantify an event's economic impact in so-called abnormal returns. Abnormal returns are calculated by deducting the returns that would have been realized if the analyzed event had not taken place (normal returns) from the actual returns of the stocks. While the actual returns can be empirically observed, the normal returns need to be estimated. For this, the event study methodology makes use of expected return models, which are also common to other areas of Finance research. Event studies are concerned with the question of whether abnormal returns on an event date or, more generally, during a window around an event date (called the event window) are unusually large in absolute value. To answer this question we carry out a hypothesis test where the null hypothesis specifies that the expected value of the abnormal return on a certain date (or the cumulative abnormal return during the event window) is zero; if the null hypothesis is rejected, we conclude that the event had an impact.

We compute the abnormal return ($AR_{i,t}$) for firm i for the days in the event window around each instance where a cyber incident was reported. Next, we take the average across all firms and all incidents $AAR_t = \frac{1}{N} \sum_i AR_{i,t}$. This is our random variable in the null hypothesis. Here by an instance we mean a single filing of the form 8-k.

5.1.1 Event Study without CAPM

We take the event window $[t_1, t_2]$ to be a two week period around the date when the 8-k was filed. Therefore, we have $t_1 = -7$ and $t_2 = 7$ and the date of filing is $t = 0$.

For the computation of the abnormal returns of firm i on day t , denoted by $AR_{i,t}$ we use

$$AR_{i,t} = R_{i,t} - R_{m,t} \quad (5.1)$$

We average this across all firms and incidents to get the AAR_t . We add up individual abnormal returns to create a cumulative abnormal return over the event window $[t_1, t_2]$. So at every $t' \in [t_1, t_2]$, we have:

$$CAAR_{t'} = \sum_{t_1}^{t'} AAR_t \quad (5.2)$$

After calculating the abnormal returns and the cumulative abnormal return, we use t-test to see whether the abnormal returns are significant after the event was reported. For $H_0 = E(CAAR_t) = 0$, t-test is

$$t = \frac{CAAR_{t_1, t_2}}{S_{CAAR}} \quad (5.3)$$

If we are interested in the significance of abnormal returns on a specific date, i.e. $H_0 = E(AAR_t) = 0$, the t-test is then

$$t = \frac{AAR_{i,t}}{S_{AAR}} \quad (5.4)$$

What we find out is that usually the returns deviate from their expected values a few days after filing of the form 8-k and the cumulative abnormal returns are significant. This effect is larger for smaller firms.

5.1.2 Event Study with CAPM

We take the event window to be a two week period around the day when the 8-k was filed. We take the estimation period to be the three month period prior to the event window. This study can be used to test both market efficiency and information content of the public announcement of filing 8-k. For the computation of the abnormal returns of firm i on day t , denoted by $AR_{i,t}$ we use

$$AR_{i,t} = R_{i,t} - (\alpha_i + \beta R_{m,t}) \quad (5.5)$$

We add up individual abnormal returns to create a cumulative abnormal return over the event window $[t_1, t_2]$.

$$CAR_{t_1, t_2} = \sum_{t_1}^{t_2} AR_{i,t} \quad (5.6)$$

After calculating the abnormal returns and the cumulative abnormal return, I use t-test to see whether the abnormal returns are significant after the event was reported. For $H_0 = E(AR_{i,t}) = 0$, t-test is

$$t = \frac{AR_{i,t}}{S_{AR_i}} \quad (5.7)$$

Where S_{AR_i} is the standard deviation of the abnormal returns calculated based on the estimation window values.

For $H_0 = E(CAR_{t_1,t_2}) = 0$, t-test is

$$t = \frac{CAR_{t_1,t_2}}{S_{CAR_i}} \quad (5.8)$$

where $S_{CAR_i}^2 = (t_2 - t_1 + 1) \times S_{AR_i}^2$.

In the event study with CAPM, we see again that the returns deviate from their expected values a few days after filing of the form 8-k. But the cumulative abnormal returns are not significantly different from zero.

Next, we plot the cumulative average abnormal returns during the event window [2](#). As we can see, this plot does not follow the regular event study plot that we expected.

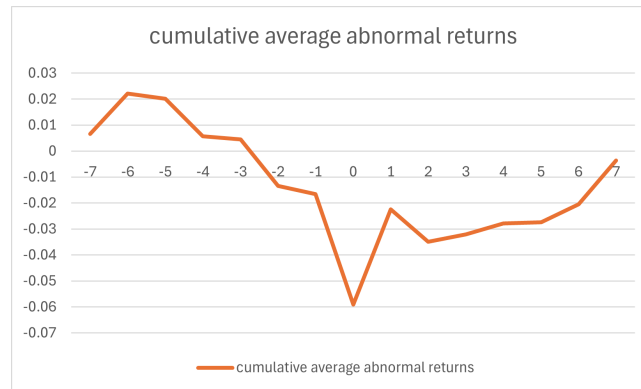


Figure 2: cumulative average abnormal returns during the event window

6 Conclusion

In the theoretical section of the paper, comparative statics reveal that increases in platform security provision lead to greater platform adoption and higher attacker sophistication, which in turn has spillovers to large firms who chose to build their in-house cybersecurity unit. Platform security is forcing the attacker to be more sophisticated. Then large firms who do not want to pay the increased price face a more sophisticated attacker and suffer more losses, despite adjusting their security levels. In addition, we observe that in equilibrium, firms who outsource enjoy a lower level of security compared to their larger counterparts. Therefore, we expect to see higher frequency of realized attacks in smaller firms. Our empirical analysis confirms this.

These findings have important policy implications. The market equilibrium may feature excessive concentration on platforms relative to the social optimum, suggesting potential justification for regulatory intervention. Policy options could include taxing p or subsidizing in-house security capabilities for small and medium enterprises, mandating minimum security standards, requiring platforms to internalize the systemic risks they create, or promoting diversity in security implementations through regulatory incentives. Our results contribute to ongoing debates about cybersecurity regulation, critical infrastructure protection, and the systemic risks associated with technology platform concentration in markets with strategic adversaries.

References

- Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608.
- Anderson, R. and Moore, T. (2006). Security economics and the internal market. *IEEE Security & Privacy*, 4(1):68–72.
- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691.
- Böhme, R. and Kataria, G. (2006). Models and measures for correlation in cyber-insurance. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*.
- Cabral, L. (2020). Merger policy in digital industries. *Information Economics and Policy*, 54:100866.
- Cavusoglu, H., Mishra, B., and Raghunathan, S. (2008). Decision-theoretic and game-theoretic approaches to it security investment. *Journal of Management Information Systems*, 25(2):281–304.
- Clemons, E. K. and Row, M. C. (1992). Information technology and industrial cooperation: the changing economics of coordination and ownership. *Journal of Management Information Systems*, 9(2):9–28.

- Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16):386–405.
- Evans, D. S. and Schmalensee, R. (2016). *Matchmakers: The new economics of multisided platforms*. Harvard Business Review Press.
- Fultz, N. and Grossklags, J. (2009). Blues in the face of attackers: a framework for evaluating information security investments. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*.
- Geer, D., Bace, R., Gutmann, P., Metzger, P., Pfleeger, C. P., Quarterman, J. S., and Schneier, B. (2003). Cyberinsecurity: The cost of monopoly—how the dominance of microsoft’s products poses a risk to security.
- Gordon, L. A. and Loeb, M. P. (2002). The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4):438–457.
- Grossklags, J., Christin, N., and Chuang, J. (2008). Security and insurance management in networks with heterogeneous agents. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 160–169. ACM.
- Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4):691–719.
- Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6):1119–1158.
- Herath, H. S. and Herath, T. C. (2008). Outsourcing information security: A literature review. In *Proceedings of the 41st Annual Hawaii International Conference on System Sciences*, pages 346–346. IEEE.
- Kunreuther, H. and Heal, G. (2003). Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–249.
- Liu, P., Zang, W., and Yu, M. (2015). A game-theoretic analysis of attack and defense in cyber-physical systems. *Decision Support Systems*, 75:1–10.
- Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029.
- Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND Journal of Economics*, 37(3):645–667.
- Romanosky, S., Ablon, L., Kuehn, A., and Jones, T. (2019). Content analysis of cyber insurance policies: how do carriers price cyber risk? *Journal of Cybersecurity*, 5(1):tyz002.

Varian, H. R. (2004). System reliability and free riding. *Economics of Information Security*, pages 1–15.

Williamson, O. E. (1979). Transaction-cost economics: the governance of contractual relations. *The Journal of Law and Economics*, 22(2):233–261.