http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–19
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

On the Financing Benefits of Supply Chain Transparency and Blockchain Adoption

Jiri Chod,[a] Nikolaos Trichakis,[b,c] Gerry Tsoukalas,[d] Henry Aspegren,[e] Mark Weber[f,g]

[a] Carroll School of Management, Boston College, Chestnut Hill, Massachusetts 02467; [b] MIT Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142; [c] Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; [d] The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104; [e] Google Inc., Mountain View, California 94043; [f] IBM Research, Cambridge, Massachusetts 02142; [g] MIT-IBM Watson Artificial Intelligence Laboratory, Cambridge, Massachusetts 02142

Contact: chodj@bc.edu, http://orcid.org/0000-0001-5283-710X (JC); ntrichakis@mit.edu, http://orcid.org/0000-0002-8324-9148 (NT); gtsouk@wharton.upenn.edu, http://orcid.org/0000-0003-2011-3646 (GT); henryaspegren@google.com (HA); mark.weber@ibm.com (MW)

Received: July 16, 2018
Revised: January 30, 2019; May 20, 2019
Accepted: May 28, 2019
Published Online in Articles in Advance:
May 14, 2020

https://doi.org/10.1287/mnsc.2019.3434

Abstract. We develop a theory that shows signaling a firm's fundamental quality (e.g., its operational capabilities) to lenders through inventory transactions to be more efficient—it leads to less costly operational distortions—than signaling through loan requests, and we characterize how the efficiency gains depend on firm operational characteristics, such as operating costs, market size, and inventory salvage value. Signaling through inventory being only tenable when inventory transactions are verifiable at low enough cost, we then turn our attention to how this verifiability can be achieved in practice and argue that blockchain technology could enable it more efficiently than traditional monitoring mechanisms. To demonstrate, we develop b_verify, an open-source blockchain protocol that leverages Bitcoin to provide supply chain transparency at scale and in a cost-effective way. The paper identifies an important benefit of blockchain adoption—by opening a window of transparency into a firm's supply chain, blockchain technology furnishes the ability to secure favorable financing terms at lower signaling costs. Furthermore, the analysis of the preferred signaling mode sheds light on what types of firms or supply chains would stand to benefit the most from this use of blockchain technology.

History: Accepted by Victor Martínez-de-Albéniz, operations management.

Funding: Financial support from the Fishman-Davidson Center for Service and Operations Management, the Inter-American Development Bank, the Mack Institute for Innovation Management, the MIT Digital Currency Initiative, the MIT Legatum Center for Entrepreneurship and Development, the MIT Sloan Junior Faculty Research Assistance Program, and the Wharton Dean's Research Fund is gratefully acknowledged.

Supplemental Material: The electronic companion is available at https://doi.org/10.1287/mnsc.2019.3434.

Keywords: blockchain • supply chain financing • supply chain transparency • information asymmetry • signaling

1. Introduction

Firms seeking the capital needed to efficiently run their operations are often impeded by the vexing problem of information asymmetry. Unable to readily ascertain their fundamental operational capabilities and gauge their risk, prospective lenders frequently command prohibitively high financing rates, which lead to operational distortions. Information asymmetry can be especially problematic for small- and medium-sized enterprises (SMEs) and startups, which are likely to be engaged in innovative operations and lack track record and reputation, particularly in developing economies plagued by trust and fraud issues. To overcome this problem, extant literature argues that firms can credibly signal private information to their lenders by distorting *loan requests*. For instance, a cash-strapped entrepreneur with an innovative prototype can try to signal good demand prospects by seeking restrictive covenants or shorter maturities, or requesting a larger loan. Alternatively, firms can also signal through their *inventory transactions*—for example, by sourcing high-quality materials from established suppliers, or distorting inventory trading quantities.

In this paper, we first develop a theory that shows that signaling a firm's fundamental quality (e.g., its operational capabilities) to lenders through inventory transactions to be more efficient—it leads to less costly operational distortions—than signaling through loan requests. Of course, signaling through inventory is only tenable when a firm's supply chain is transparent to its lenders—that is, its inventory transactions are verifiable by lenders at low enough monitoring costs. The paper's second goal is to argue that blockchain technology has the potential to enable supply chain transparency more efficiently than traditional monitoring mechanisms. To this end, we introduce b_verify, an open-source software protocol we developed to demonstrate how this technology can be implemented in a way that is accessible to SMEs in developing economies. In particular, b_verify, which uses the Bitcoin network, can secure supply chain transactions at large scale (thousands per second) and at very low cost per transaction (fractions of a cent). Taken together, our paper identifies an important benefit of blockchain adoption—by opening a window of transparency into a firm's supply chain, blockchain technology furnishes the ability to secure favorable financing terms at lower signaling costs.

1.1. Signaling Operational Capabilities: Cash vs. Inventory

In the presence of information asymmetry between firms and their lenders regarding firms' creditworthiness, high-quality firms have an incentive to signal their operational capabilities to obtain more favorable credit terms. Among the potential signaling mechanisms that have been studied, the majority involve distorting loan terms (Ross 1977, Besanko and Thakor 1987, Milde and Riley 1988). Indeed, signaling through loan requests, which we refer to here as *cash signaling*, is the de facto mechanism as lenders observe loan requests directly without incurring monitoring costs. Firms can also signal by distorting an actual physical investment—for example, an inventory transaction. What we refer to as *inventory signaling* has been studied primarily in the context of signaling to equity investors, who observe delayed and audited financial reports (Bebchuk and Stole 1993, Lai et al. 2012, Lai and Xiao 2018), and suppliers (Cachon and Lariviere 2001, Özer and Wei 2006, Chod et al. 2019a). Because lenders such as banks generally do not costlessly (and perfectly) observe inventory transactions, this setting has received far less attention.

Although both cash and inventory signaling are well established in their own separate literature streams, we are not aware of any work that would pit them against each other. Our work bridges this gap by identifying and resolving the tradeoffs that firms face when given a choice between these two mechanisms and deriving conditions on firm operational parameters under which inventory signaling is more efficient. In particular, in the absence of monitoring costs, we show that inventory signaling is generally preferable and characterize the conditions under which this preference is strict. Then, in the presence of monitoring costs, we characterize the conditions under which cash signaling strictly dominates inventory signaling and vice versa: Our theory predicts that (i) inventory perishability, (ii) inventory illiquidity, (iii) firm's creditworthiness, (iv) smaller firm market's size, (v) lower monitoring costs, and (vi) higher operating costs all favor inventory signaling and, therefore, blockchain adoption.

1.2. Supply Chain Transparency Through Blockchain: The b_verify Protocol

The proposed benefit of supply chain transparency is only tenable when firms' individual inventory transactions are verifiable by lenders at low enough costs, which then raises the question of how this can be achieved in practice. We argue that blockchain technology could provide such transparency in a more efficient way relative to traditional monitoring mechanisms.

Banks do not generally observe inventory transactions, and, if they do (e.g., through letter of credit issuance or asset-based lending), such monitoring is costly and imperfect (Stiglitz and Weiss 1981, Diamond 1991, Burkart and Ellingsen 2004, Fabbri and Menichini 2016). For example, financing through a letter of credit, whereby the bank pays directly the supplier upon the shipment of goods, involves intermediaries and a significant paper trail and can be costly, time consuming, and subject to fraud.[1] Similarly, an inventory-based loan requires intermediaries to issue and audit warehouse receipts, which is again costly and prone to fraud (Trichakis et al. 2015).[2]

Trade credit financing is one setting where inventory transparency arises naturally, thanks to the dual role of suppliers as creditors. The use of supplier financing, however, has limits, as suppliers tend to face a considerably higher cost of capital than banks (Chod 2017). Furthermore, suppliers are often capital constrained themselves, and, to extend credit, they usually borrow from banks—that is, they act as intermediaries.

Blockchain technology has the potential to mitigate many of these issues. Blockchains are cryptographically secure, distributed ledgers that can enable decentralized verifiability of *digital-goods transactions*.[3] The advantages of the technology are relatively well known—at least as they relate to storing records of digital transactions—including (1) strong security; (2) disintermediation—for example, the ability to provide trust in the absence of a trusted party;[4] (3) record integrity, by providing a chain of audit that reduces fraud opportunities; and (4) automation, so that tasks such as making loan payments can be automated (Babich and Hilary 2018). In brief, these features can be leveraged to address many of the aforementioned shortcomings, including the paper-trail inefficiencies, the need for costly intermediaries, and issues of fraud.

There are, however, some potential obstacles that may inhibit the use of blockchain technology in supply chains, or may make it too expensive to deploy. First, it is not obvious to what extent and how exactly blockchains can be successfully ported to provide verifiability of *physical-goods transactions*, such as procuring inventory.

In particular, physical transactions involve a certain amount of human intervention and, therefore, are more susceptible to mistakes or deliberate misrepresentation. The ensuing "garbage-in–garbage-out" problem could completely negate the main purpose of blockchain adoption in providing transparency.

Second, even if transaction verifiability can be successfully ported, it is not clear whether blockchain can be deployed in a way that keeps implementation and operating costs low enough to make it relevant to SMEs in developing economies. To this point, there are substantial differences to consider between *private* and *public* blockchain implementations, each having their own advantages and disadvantages in the context of supply chains. On the one hand, private blockchains have some desirable properties in terms of privacy, but (1) are usually not fully decentralized and do not fully eliminate intermediation; (2) have difficulties scaling to achieve adequate security guarantees; and (3) have (relatively speaking) large infrastructure costs. On the other hand, public blockchains, such as the one backing the Bitcoin network, do not suffer from these issues, but they do lack some of the desired properties that are important to supply chains, such as identification of verified parties, privacy of data, and transaction costs that are controlled in-network.

To demonstrate how these technological and cost issues can be overcome in practice, we developed an open-source software protocol termed b_verify. The protocol uses the Bitcoin network to ensure that recorded data cannot be retroactively modified or altered. At a high level, the protocol is designed to leverage the infrastructure benefits of public blockchains, while taking advantage of several innovations that mitigate some of the aforementioned privacy, identification, and transaction-cost issues. Unlike previous systems, b_verify was designed so that it can provide data integrity at scale and at low cost; it is capable of processing thousands of transactions per second at a fraction of a cent each. More details about the protocol and its key innovations are included in Section 3.

In this paper, we consider a specific use case in agricultural supply chains, demonstrating how the protocol facilitates transaction verifiability in the context of warehouse inventory. Warehouses often play a central role in supply chains, being frequented by suppliers who deposit inventory, buyers who procure inputs, and banks who utilize warehousing receipts and transactions to process loans (Trichakis et al. 2015). An enterprise-grade implementation of b_verify at this nexus of stakeholder interaction can securely provide, depending on the need, relevant supply chain transactional information to stakeholders, along with a cryptographic proof that the records are authentic and that no record has been omitted. According to our theory, this, in turn, should enable high-quality firms to use inventory as a credible signaling device and thereby unlock access to favorable financing terms with smaller operational distortions.

1.3. Related Literature

Signaling Models. The literature on signaling goes back to the seminal paper by Spence (1973). Whereas Spence (1973) considers only one signaling mechanism (education), herein, we allow the informed players to choose between two alternative signaling mechanisms and study which one prevails in a least-cost separating equilibrium. As we shall see, inventory signaling incurs higher unit-signaling costs than cash signaling. Although casual intuition could suggest that signaling through the costlier mechanism is automatically more efficient, a deeper analysis within the classical framework laid out by Spence (1973) reveals that this may or may not be the case. Using a generic game, we formally show that if two signaling mechanisms are available that differ only in the unit-signaling costs, the high type can prefer the costlier or the cheaper mechanism, depending on the circumstances. In particular, our analysis brings to light the following tradeoff. Because the costlier mechanism allows the high type to separate with a smaller distortion, the choice of signaling mechanism trades off the lower per-unit signaling cost of one mechanism against a smaller distortion required to separate with the other. In the context of Spence's education signaling, if, say, liberal arts education is costlier than engineering education, a first degree in the former while a postgraduate degree in the latter could be required for the high type to separate. We demonstrate that, in general, the equilibrium choice of the signaling mechanism by the high type depends on the relation between the cost premia of the costlier mechanism for the two types. Within our specific context, we find that the inventory/cash-signaling cost premia are such that inventory signaling dominates.

Signaling to Financiers/Suppliers. Signaling *t*hrough loan *requests* to financiers has been well studied in the finance literature, going back to Ross (1977), who shows that high-quality firms, concerned with short-term valuation, can signal to investors by requesting larger loans. In Besanko and Thakor (1987), lenders screen borrowers using a credit policy consisting of interest rate, loan amount, collateral, and the credit-granting probability, and high-quality firms signal by borrowing more in equilibrium than they would under full information. Milde and Riley (1988) model a game in which banks screen borrowers by offering higher loans at higher interest rate, and, depending on project characteristics, high-quality borrowers may signal by choosing larger or smaller loans in equilibrium. Duan and Yoon (1993) show that when borrowers choose between spot market borrowing and a loan commitment, high-quality borrowers signal by using larger loan commitments.

Signaling through inventory has been studied primarily in the operations management literature in the context of signaling to suppliers and signaling to equity investors. Inventory signaling to suppliers is the subject of Cachon and Lariviere (2001) and Özer and Wei (2006), both of whom examine how a privately informed manufacturer can credibly share demand forecast with a supplier, which then uses this forecast to build capacity. Taking the perspective of the manufacturer and that of the supplier, respectively, Cachon and Lariviere (2001) and Özer and Wei (2006) show that the manufacturer can signal by overordering. Chod et al. (2019a) study supplier diversification in a model that features inventory signaling to suppliers who are also trade creditors.

The literature on inventory signaling to equity investors draws upon Bebchuk and Stole (1993), who show that when firms are concerned with the short-term valuation and investors observe the investment level, firms can signal high productivity by over-investing. Building on the same premise, but in the supply chain context, Lai et al. (2012) show that inventory overinvestment due to signaling can be prevented by using a menu of buyback contracts; Lai and Xiao (2018) find that the first-best inventory decisions can be also achieved in equilibrium when the manager's short-termism is endogenous; and Schmidt et al. (2015) focus on characterizing pooling equilibria.

Our paper contributes to the above literatures by contrasting signaling with loan requests and signaling with inventory and by establishing the conditions under which the latter dominates. By studying the effect of inventory-transaction observability in the context of lending, our paper is intimately related and contributes to the literature on supplier financing.

Supplier Financing. The literature on supplier financing is vast, and we only review papers here that are closest to ours. Burkart and Ellingsen (2004) show that observability of the input transaction by the supplier reduces the borrower's diversion opportunities. Considering multiple inputs, Fabbri and Menichini (2016) and Chod (2017) show that transaction observability also reduces the asset-substitution problem. Although these papers focus on moral hazard, our work focuses on information asymmetry, where it affords new insights. In particular, the insights in these papers are less relevant in settings in which opportunistic behavior can be alleviated through other means, such as debt covenants, strong legal institutions, and so forth (Iancu et al. 2017). Our theory holds irrespective of buyer opportunism.

Closer to our work, within information asymmetry models, Burkart and Ellingsen (2004, p. 571) state that "there is no obvious distinction between lending cash and lending inputs." Our work demonstrates that such a distinction does exist: It characterizes the conditions under which an input transaction (inventory signaling) transmits private information about the borrower quality more efficiently than a cash transaction (cash signaling).

The explanation offered by our model is distinct from, and possibly more robust than, existing explanations that rationalize supplier financing based on the assumption that suppliers have an a priori informational advantage over banks (Emery 1984, Biais and Gollier 1997, Jain 2001). Albeit certainly the case in some instances, it is unlikely that financial institutions specialized in developing lending relationships and assessing creditworthiness are systematically disadvantaged relative to suppliers. Our proposed explanation is immune to this criticism because our model assumes a level playing field between banks and suppliers and identifies a monitoring advantage of supplier financing that emerges endogenously from the very nature of the transaction and the supply chain transparency it provides. Our theory, although robust to the foregoing criticisms, admits its own limitation—namely, that it is relevant only in the presence of information asymmetry between buyers (borrowers) and lenders regarding the former's creditworthiness.

Blockchain Literature. Most existing literature on blockchains is in computer science, starting with the original white paper by Nakamoto (2008). Given that blockchain technology is very new and still being actively developed, the management literature on the topic is scarce. Babich and Hilary (2018) provide a qualitative discussion of the technology's potential to improve production and distribution networks and implications for operations management researchers. Several papers examine the economics of mining and/or optimal design of blockchain systems—for example, Biais et al. (2019), Huberman et al. (2017), Budish (2018), Cong et al. (2019), Hinzen et al. (2019), Saleh (2019), and references therein. To the best of our knowledge, ours is the first research paper to explore both practical and theoretical implications of blockchains for supply chain finance and operations management.

A few recent papers, however, have started exploring the impact of blockchain in other areas. For instance, Yermack (2017) examines implications of blockchains for corporate governance, arguing that the transparency of ownership offered by blockchains may upend the balance of power in traditional governance structures. Catalini and Gans (2017) study how blockchain technology could shape innovation by reducing transaction verifiability costs and bypassing intermediaries. Chod and Lyandres (2018), Gan et al. (2019), and Chod et al. (2019b) study financing of entrepreneurial ventures by issuing crypto-tokens (initial coin offerings) on existing blockchain platforms. Falk and Tsoukalas (2020) study crowdsourcing—and, more specifically, token-weighted voting—for blockchain-based systems, such as token-curated registries. Other relevant studies include Halaburda (2018), Cong et al. (2018), and Cong and He (2019).

We contribute to this literature by examining how the existing blockchain technology—and, in particular, the transaction verifiability it provides—can be used by firms to transmit information about their inherent quality. Although the literature has recognized the benefits of blockchain-enabled verifiability of asset ownership (Biais et al. 2019), we are not aware of any papers that would connect transactional verifiability afforded by blockchain to transparency regarding the firm fundamental quality. As we show in this paper, the leap from one to the other is possible, but can be very subtle.

Other Relevant Literatures. Our paper is related to the operations literature on information sharing and signaling, which spans different areas, including signaling operational capabilities in supply chain finance (Tang et al. 2018), signaling quality in experimentation and innovation settings (Acemoglu et al. 2017, Bimpikis et al. 2018), signaling content accuracy in social networks (Candogan and Drakopoulos 2019), information sharing in the context of crowdfunding (Babich et al. 2020, Belavina et al. 2020, Chakraborty and Swinney 2020), screening firm production capabilities (Chick et al. 2016), and many others. Because the key innovation of the blockchain technology that we focus on is that it enables information sharing when transacting parties do not trust each other, our work also contributes to the literature on trust in supply chains (Özer et al. 2014).

2. Signaling Operational Capabilities: Cash vs. Inventory

As alluded to in the Introduction, cash signaling can take many forms. For example, a firm could signal by seeking tight covenants or short maturities. Similarly, inventory signaling could involve the choice of high-quality suppliers or materials. To facilitate comparison, we study a parsimonious model in which cash signaling involves requesting a larger loan and inventory signaling sourcing more inventory.

Consider a firm that can be one of two types: low-quality or high-quality, denoted by subscripts L and H, with probability $1-h$ and h, respectively. The two types differ in their operational capabilities, which manifest in different demand curves the firms face in the output market. Firm of type i, or simply firm i, sells its output at price $\tilde{\alpha}_i - x$, where x is the quantity sold and $\tilde{\alpha}_i$ is a demand shock or price intercept that follows a two-point distribution

$$\tilde{\alpha}_i := \begin{cases} \alpha_i \text{ with probability } 1-b_i, \\ 0 \text{ with probability } b_i, \end{cases} \qquad i \in \{L, H\},$$

where $\alpha_H > \alpha_L$ and $b_H < b_L$. When $\tilde{\alpha}_i = \alpha_i$ ($\tilde{\alpha}_i = 0$), we say that firm i's product is a success (failure). If the product is a success, the firm generates sales revenue $(\alpha_i - x)x$; if the product is a failure, it generates no revenue, and we assume that the unsold output has no residual value. Thus, a high-quality firm has a higher probability of success and faces a larger market size conditional on success.[5] A firm's type constitutes its private information.

Before the firm finds out whether its product is a success or a failure, it purchases Q units of an input, referred to as "inventory" at a unit cost c. Subsequently, but still before success/failure is revealed, the firm transforms x units of this input into output and brings this output to the market. Any linear internally funded cost of production can be subsumed by the demand-curve intercept α_i, which can thus be also interpreted as an (inverse) measure of the firm's operating costs. Inventory units not processed spoil and have zero salvage value.

Inventory is financed entirely by credit, which is priced competitively—that is, the lender charges fair interest, at which it expects to break even.[6] For example, when the lender provides credit D and expects no repayment in the bankruptcy state, whose probability it believes to be b, and full repayment otherwise, it charges interest r that satisfies

$$(1-b)(D+r) = D. \tag{1}$$

One can easily see that firm i goes bankrupt if its product is a failure.

We assume that the production lead time is longer than the grace period granted to the firm by the input supplier. In other words, the firm needs to secure input financing before it completes production and, as a result, cannot use output (whether it is ultimately observable by the lender or not) as a signaling device vis-à-vis the lender. We also assume that firms are not allowed to pay dividends unless they repay lenders, they cannot repay early, and they deposit all excess cash at the risk-free rate, which is normalized to zero. The equity value of a firm of type i that purchases Q units of input, transforms x units of this input into output, and faces interest r, equal to

$$V_i(Q, x, r) := (1-b_i)((\alpha_i - x)x - cQ - r), \quad i \in \{L, H\}.$$

Firms make decisions so as to maximize their equity value.

We make a technical assumption that the difference between the failure probabilities of the two types is large enough so that

$$\frac{b_L}{1-b_L}(\alpha_L - c) > \frac{b_H}{1-b_H}(\alpha_H - c). \tag{2}$$

This condition ensures that if high and low types request the same loan amount under full information, the low type is charged a higher interest. This condition is necessary to rule out the unlikely, but theoretically conceivable, scenario in which the lender charges the low type a more favorable interest, knowing that, given its smaller market size, the low type will put at risk (by investing into inventory) a smaller fraction of the amount borrowed. Most important, this condition provides the high (low) type with the incentive to signal (imitate).

Our analysis proceeds as follows. We first consider a benchmark "full-information" case, wherein the lender knows the firm's true type. We then consider the asymmetric information case under two alternative scenarios that differ with respect to the firm's supply chain transparency vis-à-vis the lender. Section 2.2 deals with the cash-signaling game that ensues when the lender does not observe the firm's inventory order before setting the credit terms. Under the assumption of zero monitoring costs, Section 2.3 deals with the inventory-signaling game that arises when the lender does observe the firm's inventory order before setting the credit terms. Then, in Section 2.4, we pit the two signaling games against each other and characterize the equilibrium choice of signaling mechanism. Section 2.5 provides a sensitivity analysis. The assumption of zero monitoring costs is relaxed in the electronic companion (EC), alongside other robustness checks. In our analysis of signaling games, we focus on pure-strategy perfect Bayesian Nash equilibria.

Note that, in this section, we use blockchain simply as a running example of a monitoring mechanism that can provide supply chain transparency to lenders; our analysis remains applicable to all other alternative mechanisms we discussed in the Introduction. Accordingly, we denote the equilibrium outcomes of the cash and inventory signaling games using subscripts Ø and Ƀ, respectively, where Ƀ is a mnemonic for blockchain-enabled supply chain transparency.

2.1. Full Information

Suppose for now that the lender knows the firm's true type. Under full information, the lender can always infer the firm's inventory order from the loan amount. Furthermore, absent any signaling incentives, the firm has no reason to borrow more than it will spend on inventory, or to order more inventory than it will process into output—that is, $D = cQ$ and $x = Q$.

When financing Q units of inventory, firm of type i faces interest $cQb_i/(1-b_i)$ according to (1). Consequently, its optimal inventory decision is

$$Q_i^{\text{fb}} := \underset{Q \geq 0}{\arg\max}\, V_i\left(Q, Q, cQ\frac{b_i}{1-b_i}\right), \quad i \in \{L, H\}. \tag{3}$$

The corresponding loan amount is $D_i^{\text{fb}} = cQ_i^{\text{fb}}$, and let V_i^{fb} denote the resulting equity value.

The inventory level that maximizes the value of equity in (3) also maximizes total value of equity and debt—that is, $Q_i^{\text{fb}} = \arg\max_{Q\geq 0}[-cQ + (1-b_i)(\alpha_i - Q)Q]$, $i \in \{L, H\}$. Therefore, we refer to Q_i^{fb} as the first-best inventory order and to D_i^{fb} as the first-best loan amount of firm i. Because the first-best order quantity increases with demand-curve intercept α_i, and decreases with bankruptcy probability b_i, the high type chooses a larger order quantity than the low type—that is, $Q_H^{\text{fb}} \geq Q_L^{\text{fb}}$.

2.2. Cash Signaling: Borrowing in the Absence of Blockchain

Recall that in the absence of blockchain, the lender cannot observe the firm's inventory order before pricing the loan. It can only observe the loan amount requested D, based on which it forms its belief about the firm's type, $\beta_\varnothing(D) \in \{L, H\}$. We follow Spence (1973) in assuming that this belief follows a threshold structure:

$$\beta_\varnothing(D) := \begin{cases} H & \text{if } D \geq d, \\ L & \text{o/w.} \end{cases} \tag{4}$$

The lender believes that a firm is of the high type if, and only if, it requests a loan amount $D \geq d$ for some (endogenously determined) threshold $d > 0$. Associating a larger loan request with the high type is reasonable because the first-best loan amount of the high type is above that of the low type.[7] The interest the lender sets is then a function of the loan amount—that is, $r_\varnothing = r_\varnothing(D)$.

The firm faces a three-stage decision problem. In the first stage, it chooses the loan amount D. The lender issues the loan and sets the interest based on its belief $\beta_\varnothing(D)$ regarding the firm type. In the second stage, the firm chooses the amount of inventory Q to purchase, subject to the loan obtained covering the purchasing cost, $D \geq cQ$. In the third stage, the firm decides how much of the purchased inventory to process, x. We solve the firm's problem by backward induction.

Because the lender does not observe the order quantity, the firm cannot use inventory to signal and, therefore, has no incentive to buy more inventory than it will eventually process. Therefore, the third-stage production decision simplifies into $x = Q$. However, the lender observes the loan amount, which thus becomes a signaling device. As a result, the firm may overborrow—that is, borrow more than its first best—or even more than it will eventually invest in inventory—to signal its type. Having obtained a loan D at interest $r_{\varnothing}(D)$, the firm of type i purchases inventory

$$Q_i(D) := \underset{0 \le Q \le D/c}{\arg\max}\, V_i(Q, Q, r_{\varnothing}(D)),\ \ i \in \{L, H\}.$$

It can be shown that the optimal inventory order takes the form $Q_i(D) = \frac{1}{c}\min\{D, \overline{D}_i\}$, where $\overline{D}_i$ is an "investment cap," above which it is not economical to invest cash into inventory—that is, any borrowing above the investment cap sits idle.

Because the firm has no other investment or diversion opportunities, the lender rationally anticipates that any amount borrowed above the investment cap will not be put at risk. Therefore, it charges interest based only on the amount that it expects to be invested in inventory, $c\,Q_{\beta_{\varnothing}(D)}(D)$, in accordance with its belief $\beta_{\varnothing}(D)$. The fair interest is then given by

$$r_{\varnothing}(D) = \begin{cases} c\,Q_H(D)\dfrac{b_H}{1-b_H} & \text{if } D \ge d, \\ c\,Q_L(D)\dfrac{b_L}{1-b_L} & \text{o/w.} \end{cases} \tag{5}$$

In the first stage, the firm requests a loan amount D so as to maximize the value of equity

$$V_{i\varnothing}(D) := V_i(Q_i(D), Q_i(D), r_{\varnothing}(D)), \quad i \in \{L, H\}.$$

A separating equilibrium (SE) is characterized by the optimal loan amounts $\{D_L^{se}; D_H^{se}\}$ and a consistent belief structure given by (4) that satisfies the following necessary and sufficient conditions:

$$\max_{D<d} V_{H\varnothing}(D) \le \max_{D\ge d} V_{H\varnothing}(D), \quad \text{and} \tag{6}$$

$$\max_{D<d} V_{L\varnothing}(D) \ge \max_{D\ge d} V_{L\varnothing}(D). \tag{7}$$

Condition (6) ensures that the high-quality firm borrows an amount at or above the threshold d, whereas condition (7) ensures that the low-quality firm borrows below d. Because these conditions may lead to multiple SEs, we adopt the Cho and Kreps (1987) intuitive criterion refinement, which eliminates any Pareto-dominated equilibria. We refer to any equilibria that survive as least-cost separating equilibria (LCSEs).

Intuitively, at the LCSE, the high type borrows either the minimum amount of money necessary to separate from the low type, or his first best, whatever is larger. The minimum loan amount that the high type needs to borrow to separate from the low type is the maximum loan amount that the low type is willing to borrow to imitate the high type. This amount determines the least-cost equilibrium belief threshold d, and it is such that the low type is indifferent between ordering d, while being perceived as high type, and ordering his first best while being perceived as low type. For the presentation of the equilibrium results, we make use of two useful thresholds, $b^{\varnothing}$ and b^{se}, which satisfy $b^{\varnothing} > b^{se}$. (These and other useful quantities hereafter are defined in the EC.)

Proposition 1. *Absent blockchain-enabled supply chain transparency, if $b_H > b^{se}$, there exists a unique LCSE, which is given by loan amounts $D_L^{se} = D_L^{fb}$ and $D_H^{se} = \max(D_H^{fb}, d)$, where d takes a different form depending on whether $b_H \ge b^{\varnothing}$ or not. If $b_H \le b^{se}$, no separating equilibrium exists.*

All proofs are included in the EC. When a SE exists, the low type follows its first best, whereas for the high type, there are two scenarios. When $d \le D_H^{fb}$, the low type is not willing to borrow up to the high type's first best, and so the high type can borrow its first-best D_H^{fb} without being imitated. The second and more salient scenario takes place when $d > D_H^{fb}$—that is, the low type is willing to mimic the high type's first best. In this case, the high type needs to overborrow up to d to separate. This ultimately distorts its inventory order as well in the sense that $Q_H(d) > Q_H^{fb}$.[8]

The existence and form of the separating equilibrium depend critically on the bankruptcy probability of the high type, b_H, which determines the strength of low type's incentives to imitate.

(1) If $b_H \ge b^{\varnothing}$, the high type's bankruptcy probability is relatively high, which limits the low type's willingness to imitate. Specifically, the low type is not willing to borrow beyond its investment cap in order to imitate—that is, $d \le \overline{D}_L$.

(2) If $b^{se} \le b_H < b^{\varnothing}$, the advantage of being perceived as the high type is significant enough for the low type to be willing to borrow beyond its investment cap—that is, $d > \overline{D}_L$.

(3) If $b_H \le b^{se}$, the high type's bankruptcy probability is so low, that in order to be perceived as the high type, the low type is willing to borrow up to the *high type's* investment cap $\overline{D}_H$. Thus, the high type cannot separate by borrowing less than $\overline{D}_H$. Recall that any borrowing beyond $\overline{D}_H$ would not be invested in inventory by either type. Without further distorting a firm's operations and thereby increasing its signaling cost, borrowing beyond this point would be a cheap talk. Therefore, borrowing beyond $\overline{D}_H$ does not allow the high type to separate either. As a result, no separating equilibrium can exist in this case.

2.3. Inventory Signaling: Borrowing in the Presence of Blockchain

We assume for now that the use of blockchain technology comes at no monitoring cost, and then relax this assumption in the EC. Because blockchain allows lenders to observe and verify firms' inventory orders, firms can use inventory as a signaling device. Let $\beta_{\text{Ƀ}}(Q) \in \{L, H\}$ be the belief regarding the firm type that the lender forms upon observing an order quantity Q. We borrow again from Spence (1973) in assuming that the lender's belief has a threshold structure—that is,

$$\beta_{\text{Ƀ}}(Q) := \begin{cases} H & \text{if } Q \geq q, \\ L & \text{o/w.} \end{cases} \tag{8}$$

Thus, if a firm orders inventory Q above some (endogenous) threshold q, the lender believes the firm to be of the high type. Otherwise, it believes that the firm is of the low type. This is reasonable given that the high type has a higher first-best order quantity—that is, $Q_H^{\text{fb}} > Q_L^{\text{fb}}$.

Because the lender forms its belief based on the actual order quantity rather than the loan request, the firm has no incentive to borrow more than what it invests in inventory. The firm's decision problem thus simplifies into a two-stage optimization. In the first stage, it decides the order quantity Q, or, equivalently, the loan amount $D = cQ$. The lender provides the loan and charges interest

$$r_{\text{Ƀ}}(Q) := \begin{cases} cQ\,\dfrac{b_H}{1 - b_H} & \text{if } Q \geq q, \\ cQ\,\dfrac{b_L}{1 - b_L} & \text{o/w,} \end{cases} \tag{9}$$

according to (1) and (8). In the second stage, the firm decides how much of the purchased inventory Q to transform into output x. We solve the two-stage problem via backward induction.

Because the firm may have incentive to overorder in the first stage to signal that it is of the high type, it may end up with more inventory than it will be willing to process into output in the second stage. Therefore, we can no longer take for granted that $x = Q$. Instead, we need to allow the firm to choose its output level optimally. Given an inventory amount Q and facing interest $r_{\text{Ƀ}}(Q)$, the firm of type i chooses output $x_i(Q) := \arg\max_{x \leq Q} V_i(Q, x, r_{\text{Ƀ}}(Q))$.

Because the sales revenue starts decreasing in output at $\overline{Q}_i := \frac{1}{2}\alpha_i$, the firm will never produce beyond this "production cap," even if it means not using the entire inventory. In other words, the firm's optimal output takes the form $x_i(Q) = \min\{Q, \overline{Q}_i\}$. In the first stage, the firm selects the order quantity Q to maximize the value of equity $V_{i\text{Ƀ}}(Q) := V_i(Q, x_i(Q), r_{\text{Ƀ}}(Q))$.

A separating equilibrium is characterized by the low type's and the high type's optimal order quantities $\{Q_L^{\text{se}}; Q_H^{\text{se}}\}$ and a consistent belief structure given by (8) that satisfies the following necessary and sufficient conditions:

$$\max_{Q<q} V_{H\text{Ƀ}}(Q) \leq \max_{Q \geq q} V_{H\text{Ƀ}}(Q), \quad \text{and} \tag{10}$$

$$\max_{Q<q} V_{L\text{Ƀ}}(Q) \geq \max_{Q \geq q} V_{L\text{Ƀ}}(Q). \tag{11}$$

Condition (10) ensures that a high-quality firm orders a quantity at or above the threshold q. Condition (11) ensures that a low-quality firm orders a quantity below this threshold.

Similar to d, the LCSE belief threshold q is given by the low type's indifference point. In this case, it is the maximum amount of inventory that the low type is willing to order so as to imitate the high type. As before, we use a threshold bankruptcy probability, denoted by $b^{\text{Ƀ}}$.

Proposition 2. *In the presence of blockchain-enabled supply chain transparency, there always exists a unique LCSE, which is given by order quantities $Q_L^{\text{se}} = Q_L^{\text{fb}}$ and $Q_H^{\text{se}} = \max(Q_H^{\text{fb}}, q)$, where q takes a different form depending on whether $b_H \geq b^{\text{Ƀ}}$ or not.*

In the LCSE, the low type always orders its first best Q_L^{fb}, being unable to imitate the high type. For the high type, there are two possible scenarios depending on the low type's willingness to overorder. If the low type is not willing to imitate the high type's first best—that is, $q \leq Q_H^{\text{fb}}$—the high type can order its first best. The second scenario is more interesting. If the low type is willing to imitate the high type's first best—that is, $q > Q_H^{\text{fb}}$—the high type has to inflate its order up to q units to avoid imitation.[9]

The functional form of the belief threshold q depends on the high type's bankruptcy probability, which affects the low type's incentive to imitate.

1. If $b_H \geq b^{\text{Ƀ}}$, the high type's bankruptcy probability is considerable, so the advantage of being perceived as a high type is not sufficient to justify for the low type to order above its production cap in an attempt to imitate—that is, $q \leq \overline{Q}_L$.
2. If $b_H < b^{\text{Ƀ}}$, the reward from imitating the high type is so significant that the low type is willing to order beyond its own production cap in order to imitate—that is, $q > \overline{Q}_L$.

Finally, note that a SE always exists with blockchain. Because increasing inventory of illiquid and perishable goods beyond a firm's first best is always costly, overordering inventory, unlike overborrowing cash, is never cheap talk. Furthermore, because such overordering always is costlier for the low type, which derives less value from each unit of inventory, separation is always possible.

In the next section, we examine the firm's choice whether to adopt blockchain technology.

2.4. Comparison of Signaling Modes: Equilibrium Adoption of Blockchain

Suppose that, in addition to making the operations and financing decisions that we have examined thus far, firms need to decide whether to adopt blockchain. In this case, the set of potential separating equilibria comprises four possible classes: $\text{Ƀ} - \emptyset$, $\text{Ƀ} - \text{Ƀ}$, $\emptyset - \emptyset$, and $\emptyset - \text{Ƀ}$, where the left (right) entry represents the low (high) type's technology choice. For example, $\text{Ƀ} - \text{Ø}$ means that in equilibrium, the low type uses blockchain, whereas the high type does not.

We continue to assume that each lender holds a threshold belief structure, where q and d denote the equilibrium order quantity and loan amount thresholds, respectively. The necessary and sufficient conditions characterizing a separating equilibrium in this case are

$$\max\left\{\max_{D<d} V_{H\text{Ø}}(D), \max_{Q<q} V_{H\text{Ƀ}}(Q)\right\}$$
$$\leq \max\left\{\max_{D\geq d} V_{H\text{Ø}}(D), \max_{Q\geq q} V_{H\text{Ƀ}}(Q)\right\},$$
$$\max\left\{\max_{D<d} V_{L\text{Ø}}(D), \max_{Q<q} V_{L\text{Ƀ}}(Q)\right\}$$
$$\geq \max\left\{\max_{D\geq d} V_{L\text{Ø}}(D), \max_{Q\geq q} V_{L\text{Ƀ}}(Q)\right\}.$$

These conditions in some sense combine conditions (10)–(11) and (6)–(7) from our previous analysis. The first inequality ensures that in any SE, the high type chooses to signal high—with blockchain or without—by ordering or borrowing above the corresponding belief threshold. The second inequality ensures that the low type chooses to order or borrow below the corresponding belief threshold. Both firms' actions are thus consistent with their lenders' beliefs in equilibrium.

Note that in any SE, the low type chooses its first-best financing/inventory. Because its operations/financing decisions are not distorted by signaling, the low type is indifferent between using and not using blockchain—that is, for any SE of class $\text{Ƀ} - \text{Ƀ}$ ($\text{Ƀ} - \text{Ø}$), there is an equivalent equilibrium of class $\text{Ø} - \text{Ƀ}$ ($\text{Ø} - \text{Ø}$), in which both types make the same inventory/financing decisions. Because in practice, the use of blockchain comes at a cost, in what follows, we restrict our attention to SE, in which the low type does not use blockchain.

Proposition 3.

1. *A separating equilibrium always exists.*

 a. *If $b_H \leq b^{\text{se}}$, all separating equilibria belong to class $\text{Ø} - \text{Ƀ}$. Among these, there exists a unique LCSE, which is given by $D_L^{\text{se}} = D_L^{\text{fb}}$ and $Q_H^{\text{se}} = \max(Q_H^{\text{fb}}, q)$, with $d = \infty$ and q as in Proposition 2.*

 b. *Otherwise, there are additional separating equilibria, which belong to class $\text{Ø} - \text{Ø}$. Among these, there exists a unique LCSE, which is given by $D_L^{\text{se}} = D_L^{\text{fb}}$ and $D_H^{\text{se}} = \max(D_H^{\text{fb}}, d)$, with $q = \infty$ and d as in Proposition 1.*

2. *No pooling equilibria survive the intuitive criterion.*

Depending on the high type's failure probability b_H, one of two scenarios arises:

(a) If $b_H \leq b^{\text{se}}$, the high type cannot separate itself without using blockchain because borrowing large amounts of cash without being able to show how it was spent is cheap talk. If this is the case, the high type has no choice but to use blockchain in order to separate itself. The low type follows its first best without the use of blockchain. Thus, all separating equilibria fall into the $\text{Ø} - \text{Ƀ}$ class, and the least-cost among them has a structure analogous to the LCSE characterized in Section 2.3.

(b) If $b_H > b^{\text{se}}$, the high type is able to separate even without blockchain, and, therefore, additional separating equilibria emerge in the $\text{Ø} - \text{Ø}$ class. The least-cost among them is identical to the LCSE described in Section 2.2. In other words, there are two candidates for the LCSE of the full game: one of class $\text{Ø} - \text{Ø}$, in which the high type does not adopt blockchain; and one of class $\text{Ø} - \text{Ƀ}$, in which it does. Which of these two equilibria leads to lower signaling costs and thus emerges as the LCSE is what we examine next.

We formally define the signaling costs for any given SE as the difference between the high type's equity value under the first best and under that SE.[10] The high type's equity value under the $\text{Ø} - \text{Ƀ}$ and $\text{Ø} - \text{Ø}$ equilibria identified in Proposition 3 is $V_{H\text{Ƀ}}^{\text{se}} := V_{H\text{Ƀ}}(Q_H^{\text{se}})$ and $V_{H\text{Ø}}^{\text{se}} := V_{H\text{Ø}}(D_H^{\text{se}})$, respectively. Therefore, the signaling costs under equilibrium of class $\text{Ø} - j$ are

$$\mathscr{C}_j := V_H^{\text{fb}} - V_{Hj}^{\text{se}}, \quad j \in \{\text{Ƀ}, \text{Ø}\}.$$

When both technology modes allow the high-quality firm to separate, using blockchain is preferable, and the LCSE is of class $\text{Ø} - \text{Ƀ}$ if, and only if, $\mathscr{C}_{\text{Ƀ}} \leq \mathscr{C}_{\text{Ø}}$.

Whether the last inequality holds true or not depends on the cost of overborrowing cash relative to the cost of overordering inventory, as well as on the total amount by which the high type needs to overborrow or overorder in order to separate itself. Overborrowing by \$1 is generally cheaper than overordering by \$1 worth of inventory because the former affords the option not to convert the cash into goods. Consequently, the low type is willing to overborrow more than it is willing to overorder in order to imitate. As a result, signaling with cash generally requires a *larger distortion* by the high type than signaling with inventory. In particular, when both signaling mechanisms allow separation, we have $d > cq$ if $b^{\text{se}} < b_H < b^{\text{Ø}}$, and $d = cq$ if $b_H \geq b^{\text{Ø}}$. Interestingly, when $b_H \geq b^{\text{Ø}}$, the low type is not willing to overborrow beyond its investment cap, and the aforementioned option is "out of the money." In this case, overborrowing is equally costly to the low type as overordering, and both signaling games result in the same equilibrium distortion.

What the previous discussion also makes apparent is that preference for one signaling mechanism or the other cannot be explained by the mere observation that overordering is costlier than overborrowing. We illustrate this point with an in-depth discussion in the EC, in which we analyze a generic signaling game with two mechanisms.

To sum up, in order to separate itself without the use of blockchain, the high type needs to overborrow by the same or a larger amount than it has to overorder with blockchain. At the same time, overordering a unit of inventory is equally or more costly than overborrowing the equivalent amount of cash. Which of the two effects dominates determines whether blockchain adoption increases or decreases the total signaling cost and, therefore, whether it emerges as the LCSE. To resolve the tradeoff, we use a critical value b^{cr}, which satisfies $b^{\text{se}} < b^{\text{cr}}$. The following result answers the key questions as to when, and why, a high-quality firm prefers to use the blockchain technology.

Theorem 1. *Preference for the use of blockchain depends on the failure probability b_H as follows*:

i. *If $0 \le b_H \le b^{\text{se}}$, the high-quality firm prefers to use blockchain because in its absence, it cannot separate itself from the low type.*

ii. *If $b^{\text{se}} < b_H < b^{\text{cr}}$, the high-quality firm prefers to use blockchain because it allows separation from the low type at a lower cost—that is, $\mathscr{C}_{\text{Ƀ}} < \mathscr{C}_{\varnothing}$.*

iii. *If $b^{\text{cr}} \le b_H$, the high-quality firm is indifferent regarding the use of blockchain—that is, $\mathscr{C}_{\text{Ƀ}} = \mathscr{C}_{\varnothing}$.*

Whether the high type can separate and, if so, at what cost, depends, among others, on how strongly the low type is willing to imitate. The low type's incentive to imitate then depends on the benefit of being perceived as the high type, which in turn depends on b_H.

i. When b_H is small, the low type is so eager to imitate that, in a cash-signaling game, it is willing to overborrow all the way up to the high type's investment cap. This makes it impossible for the high type to separate because borrowing beyond this cap is a cheap talk. In contrast, overordering illiquid inventory is never a cheap talk, and inventory signaling thus always allows separation. Therefore, the benefit of blockchain adoption for a firm of "very high" quality is that it makes separation possible.

ii. When b_H is moderate, both inventory signaling and cash signaling allow the high type to separate. However, because the low type finds it costlier to overorder inventory than to overborrow an equivalent amount of cash, the high type is able to separate with a relatively smaller inventory distortion. Importantly, unlike the low type, the high type finds that it always optimal to invest all cash in inventory and is therefore indifferent between overborrowing cash and overordering the equivalent amount of inventory. As a result, the high type always prefers a smaller inventory distortion to a larger credit distortion. To put it differently, the high type prefers signaling with inventory because it has a comparative advantage in monetizing inventory. The benefit of blockchain adoption for a firm of "moderately high" quality is thus a smaller signaling cost.

iii. When b_H is large, both inventory and cash signaling allow the high type to separate at the same signaling costs. In this regime, the low type's incentive to imitate is so low, that in a cash-signaling game, it is not willing to borrow more than it is willing to invest in inventory. Overborrowing cash and overordering inventory is thus equally costly, even for the low type, and there is no difference between cash signaling and inventory signaling. Thus, a firm of "somewhat high" quality does not benefit from blockchain adoption in the context of signaling to lenders.

The three regimes identified in Theorem 1 are illustrated in Figure 1. In more informal terms, the main message of the theorem can be summed up as follows.

Main Result. *Blockchained-enabled supply chain transparency allows firms to convey private information about operational capabilities to lenders more efficiently.*

Next, we conduct a sensitivity analysis. In the EC, we consider robustness checks, in which inventory processing could be costlier for the high-quality firm, monitoring costs are nonzero, and SE under a general, nonthreshold belief structure.

2.5. Sensitivity Analysis

In this section, we discuss how the characteristics of a high-quality firm affect its preference for the signaling mode and thus for blockchain adoption.

Failure Probability b_H. It follows directly from Theorem 1 that as the high type's failure probability increases, its preference for blockchain adoption fades. Intuitively, as the high type becomes less creditworthy, the low type's incentive to imitate weakens, and it becomes easier for the high type to separate even absent supply chain transparency.

Demand-Curve Intercept α_H. Recall that a larger α_H can capture a larger market size as well as lower operating cost of the high type. To understand the effect of α_H, we can reformulate Theorem 1 so as to link the equilibrium outcome to the value of this parameter. To that end, we utilize two thresholds, α^{se} and α^{cr}, such that $\alpha^{\text{se}} < \alpha^{\text{cr}}$. Noting that threshold $b^{\varnothing}$ is independent of α_H, we can write the following corollary to Theorem 1.

Corollary 1. *If $b_H \ge b^{\varnothing}$, the high-quality firm is indifferent regarding the use of blockchain—that is, $\mathscr{C}_{\text{Ƀ}} = \mathscr{C}_{\varnothing}$—for any α_H. Otherwise, preference for the use of blockchain depends on α_H as follows*:

i. *If $\alpha_H \le \alpha^{\text{se}}$, the high-quality firm prefers to use blockchain because in its absence it cannot separate itself from the low type.*

ii. *If $\alpha^{\text{se}} < \alpha_H < \alpha^{\text{cr}}$, the high-quality firm prefers to use blockchain because it allows separation from the low type at a lower cost—that is, $\mathscr{C}_{\text{Ƀ}} < \mathscr{C}_{\varnothing}$.*

Figure 1. First-Best Outcomes (Dashed Line); Outcomes Without Blockchain (Dot-Dashed Line); Outcomes with Blockchain (Solid Line) for the High Type Under $a_L = 5.1, a_H = 6, c = 3, b_L = 0.4$

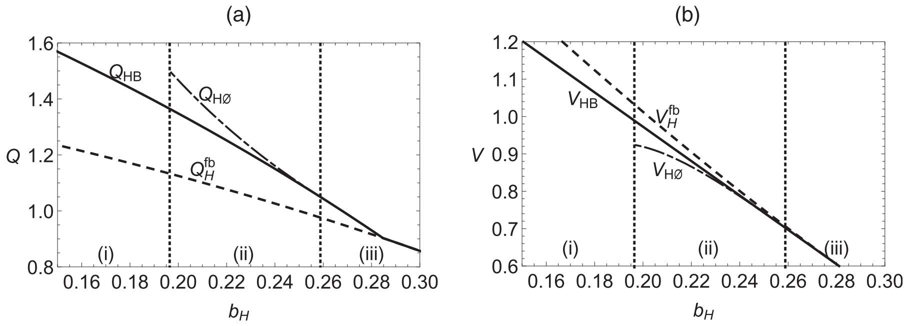

Notes. (a) Order quantities. (b) Firm values.

iii. *If $\alpha_H \geq \alpha^{\text{cr}}$, the high-quality firm is indifferent regarding the use of blockchain—that is, $\mathscr{C}_{\mathbb{B}} = \mathscr{C}_{\emptyset}$.*

According to the corollary, the high type's preference for blockchain is negatively related to α_H. As α_H becomes larger—that is, as the high type's market size increases or its operating costs decrease—its first-best input order as well as loan request increase. Thus, it becomes more difficult for the low type to imitate, which makes it easier for the high type to separate even absent supply chain transparency.

Salvage Value. Finally, recall our assumption that input inventory is perishable or illiquid (e.g., due to its customization) and thus has no salvage value. In the next corollary, we consider the other extreme case, in which the input goods can be resold for the full procurement cost.

Corollary 2. *If the input inventory can be salvaged without a loss, there is no difference between cash signaling and inventory signaling, and, therefore, preference for the use of blockchain vanishes.*

If the firm can resell input inventory and recover its full procurement cost, buying inventory is cheap talk, and supply chain transparency makes no difference from a signaling perspective.

3. Supply Chain Transparency Through Blockchain: The b_verify Protocol

The proposed benefit of supply chain transparency is only tenable when firms' individual inventory transactions are verifiable by lenders at low enough costs. Our discussion in the Introduction highlighted that blockchain represents a promising technology that could provide supply chain transparency securely and more efficiently than existing mechanisms.

To be precise, herein we use the broad term "blockchain technology" to describe the family of distributed ledger consensus protocols derived from Bitcoin. Consensus is defined as agreement regarding the state of data among a group of parties who may not trust each other. Updates, such as peer-to-peer transactions of Bitcoin with notes attached to them, are batched by consensus protocols as new "blocks" appended to previous ones in a manner that creates an interdependent "blockchain." This interdependence, combined with game-theoretic design and cryptography, makes a blockchain secure and extremely difficult to change when it is kept by a large, distributed network such as in Bitcoin. In the hype cycle of the technology, this extreme difficulty became known as "immutability"—the alluring idea of uncompromising permanence in record keeping.

At the same time, as we discussed in the Introduction, there are several potential obstacles, related to operational costs, privacy, and data integrity, that could make the use of the technology in supply chains inefficient. To address these, we developed b_verify, a novel blockchain-based protocol for securing and updating large amounts of inventory transaction data in supply chains.[11]

Next, we first give a high-level overview of the system. Then, we present an example use case in agricultural supply chains. We conclude by discussing in some detail the system's innovative features, illustrating how b_verify has the potential to overcome the aforementioned obstacles related to implementation and spur adoption.

3.1. Overview

We designed b_verify to be a low-cost, open-source, lightweight, and flexible protocol, whose main objective is to enable a network of mutually distrusting parties, such as firms and their lenders, to securely record and verify transactions on an immutable distributed ledger. Implementation and operating costs are kept low via a "hybrid" design, whereby the system leverages some of the existing low-cost advantages of permissionless public blockchains (infrastructure, security, and immutability), while also implementing novel solutions to preserve other desirable properties often associated with more expensive private blockchain implementations, such as permissioning and privacy.

In particular, b_verify does not create its own blockchain; rather, it stores hashed transaction data on an existing public blockchain ledger, such as Bitcoin. Importantly, the transactions are processed by a low-cost disposable and "untrusted" server, which uses a cryptographic data structure to allow for public commitments to the bitcoin ledger.[12] There are two notable innovations here: The first is that server and data-processing protocols leverage the security guarantees provided by the Bitcoin blockchain—specifically, *nonequivocation*[13]—to create a partially ordered, timestamped record of transactions that can be verified by interested parties. Barring the technical details, which are left for the appendix, the end result is a protocol in which a user's device can present a cryptographic proof that a set of records is complete and authentic.

The second notable innovation is that the protocol requires only a short "signature" or "hash-key" of the transaction data to be committed to the public ledger, therefore lowering bandwidth-related costs, while providing privacy. That is, parties need to verify only a subset of the ledger that pertains to a specific user or type of asset. We provide more details on these design choices and other core innovations in Section 3.3, after we present a use case.

3.2. Agriculture Warehouse Implementation

To test and demonstrate the features of an application servicing the b_verify protocol, we consider the use case of warehouse receipts in agricultural supply chains. This use case is motivated and informed by our research collaborations with public and private organizations in Mexico and Ukraine, which have identified the secure digitization of warehouse receipts as an economic priority where blockchain technology shows great promise. As discussed in the Introduction, warehouses often play a central role in supply chains (Trichakis et al. 2015), and, therefore, installing b_verify at this nexus of stakeholder interaction could afford multiple benefits.

Figure 2 provides a schematic overview that exemplifies a simplified warehouse receipt issuance implementation. A farmer deposits or withdraws produce from a warehouse. The produce is weighed by a warehouse employee, optionally, by using a digital Internet-of-Things (IoT) scale. Using applications servicing the b_verify protocol, the farmer, the warehouse employee, and the IoT digital scale use their respective private keys to sign the transaction, which is then validated and committed to the Bitcoin blockchain by an untrusted server. The information from this transaction can then be confirmed by another party, such as a lender interested in verifying the warehouse receipt as collateral (for example, the lender might want to ensure the receipt is authentic, up-to-date, and has no other outstanding liens against it). Should a loan be issued, this update is signed by the involved parties with a new transaction in the same manner as the first. The nonequivocation property achieved by b_verify ensures that the warehouse receipt cannot be pledged multiple times, or sold if a lien is outstanding.

3.3. Implementation Challenges and Solutions

We discuss in some detail below some of the key innovative features of the system that we implemented to overcome obstacles related to verification, operational costs, privacy, and data integrity, which could inhibit the use of the technology for supply chain transparency.

Verification of Physical Transactions. As mentioned, blockchain technology was originally developed in the context of the Bitcoin currency, involving merely digital transactions. How exactly this technology has to be adapted to provide verifiability of physical-goods transactions, such as procuring inventory, is currently an open question that is actively being pursued by practitioners (e.g., see IBM TrustChain 2017).

In providing verifiability, there could be many differentiating factors between digital- and physical-goods transactions. In our view, the paramount factor, which b_verify can help to address, is that operational transactions in supply chains involve a certain amount of human intervention or manual entries and, therefore, are more susceptible to mistakes or deliberate misrepresentation. In an agricultural supply chain, for example, transactions involving commodities usually require warehouse employees to weigh physical goods to verify the quantity traded. In addition, laboratory employees use various tools, including advanced electronic devices, to assess quality variables such as moisture, oil levels, and protein content. In this process, an employee could mishandle (by accident or not) the incoming inventory shipment and misrepresent the quantity received or its quality. These types of issues undermine verifiability and could negate any of the potential advantages of implementing the technology in the first place.

To mitigate this issue, b_verify solicits independent attestations from multiple clients before committing a transaction to the blockchain. Clients, who could be people or Internet-connected devices, cryptographically sign and attest to the details of the transaction. In the warehousing implementation, for instance, the IoT digital scale has a direct feed to the b_verify system, and the transaction is committed only if the

Figure 2. Example of b_verify Use Case for Warehouse Operations

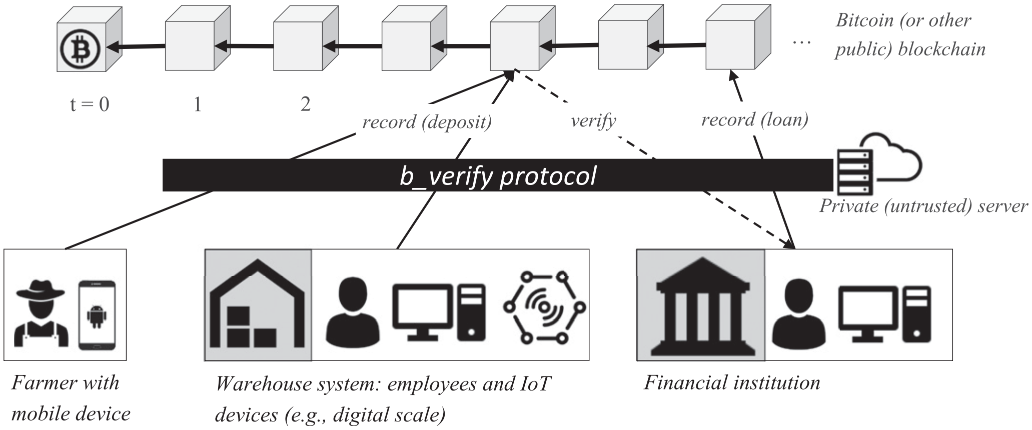
t = 0
1
2
…
Bitcoin (or other public) blockchain
record (deposit)
verify
record (loan)
b_verify protocol
Private (untrusted) server
Farmer with mobile device
Warehouse system: employees and IoT devices (e.g., digital scale)
Financial institution

information sent to the server by all three parties (warehouse employee, depositor, and the digital device) is consistent. Once the system verifies the consistency of all the attestations, it hashes the data and commits it to the public Bitcoin blockchain, where it is immutably stored to provide a permanent audit chain. Arguably, this combination of signed statements from both human users and IoT devices, along with the transparency and security of the Bitcoin-based ledger, would provide much stronger assurances of data integrity for physical goods.

Privacy and Permissioning. Two key issues that arise when public ledgers are used to store private data are (1) to be able to identify the parties involved; and (2) to ensure that the stored data are encrypted in such a way that only permissioned parties are able to recover it. To deal with the first issue, b_verify requires a public key infrastructure to identify participants and verify their signatures. Regarding the second issue, it is important to recognize that blockchains require data to be shared among a set of participants. In the context of a supply chain, this could represent a problem because firms have a strong incentive to guard proprietary data and prevent it from falling into the hands of competitors. In the agriculture sector that b_verify targets, transaction records are often required by law to be public, so this need not always be an issue. More generally, this problem can be addressed at a technological level by using encryption to selectively hide or obfuscate data, or by committing to the ledger only a "hash signature" of the information—something b_verify accomplishes by using Merkle Prefix Trie data structures (see appendix). Privacy is a major concern in practice and is currently an area of active research.

Scalable Nonequivocation. In a supply chain, as well as many other systems, participants may have a strong incentive to modify or omit information. Ensuring data integrity and consistency in these environments is challenging because of the possibility that a server could equivocate—for example, show two different sequences of events to different users. If a participant in a supply chain system can equivocate by presenting, for example, a different sequence of inventory transaction events to other participants, the system is compromised, rendering our theory impractical. The standard solution is to find a trusted third party to manage data; however, this may be impossible or inconvenient. Recent work has proposed the use of public blockchains such as Bitcoin, relying on the extreme difficulty in compromising (or forking) a large public blockchain like Bitcoin. These protocols, such as a "Bitcoin witnessing scheme" called Catena (Tomescu and Devadas 2017), generally require one Bitcoin transaction per log statement. As a result, these protocols cannot support many users and applications due to limited "block space" for data and high transaction fees imposed by the network, which increase with the popularity and therefore the security of the network. For example, a Catena implementation supporting 10,000 applications (a relatively small figure for a complex supply chain application) would require the entirety of the Bitcoin's network transaction capacity; in other words, it would bring the network down. In contrast, b_verify introduces a new protocol that can support many millions of application logs and many thousands of log statements per second, reducing the cost of each log statement to a fraction of a cent.[14]

Limited Hardware Requirements. The main hardware requirements of the b_verify system are fairly limited, in that the application itself can be launched from a basic smartphone and/or a tablet, and only a generic server is required to do the data processing and encryption, before it is committed to the public ledger. In particular, neither of these devices is required to have state-of-the-art storage or processing capabilities.

Recognizing that SMEs often have limited hardware computing resources, we designed b_verify so that it can accommodate "thin clients" that are not required to download the entire Bitcoin blockchain to participate,[15] providing an easy "plug-and-play" experience for users to facilitate adoption on the ground.

Limited Implementation and Operating Costs. One of the advantages of utilizing a public ledger to store data, such as the one provided by the Bitcoin blockchain, is that the infrastructure needed to ensure the security and visibility of data is "outsourced" to the global network. In other words, b_verify leverages the infrastructure already in place and the security guarantees that it provides. Typically, using this shared resource is costly. For instance, the cost of a Bitcoin transaction has experienced peaks, sometimes reaching $20–$40. However, b_verify amortizes this cost among many users, reducing the cost to the individual user dramatically. This, together with the aforementioned technical innovations, combine to ensure that the system, as a whole, can be run at minimal cost (fractions of a cent in U.S. dollars) to b_verify users.

Taken together, we believe the above features provide a relatively secure, low-cost, and scalable way of implementing transaction verifiability in the context of physical supply chains. For more information on the b_verify system, and for links to the open-source code, please see the appendix.

4. Concluding Remarks and Predictions

Opening a small window of transparency into a firm's supply chain—in particular, its input transactions—could go a long way to alleviate the difficulty of financing operations, which is a systemic problem for SMEs, that is all the more severe in developing economies. As we demonstrate in this paper, blockchain technology could provide an efficient way to accomplish this by furnishing input-transaction verifiability in supply chains in a way that is accessible to SMEs. The potential benefits are vast and global in scale. Although the SME sector represents the backbone of many of the world economies, accounts for over half of jobs, and over a third of GDP worldwide, the World Bank estimates that its global financing shortfall tops $2.6 trillion (Alibhai et al. 2017). Notably, the emergence of trading platforms, such as Amazon and Alibaba, has provided one possible solution to the SMEs' struggles: Being able to observe virtually all business transactions of participants, these platforms have very recently started to provide cheap loans through their financing arms—Amazon Lending and Ant Financial (see, e.g., Dong et al. 2019). Unfortunately, these benefits are confined to members of such closed systems. The use of publicly available blockchain platforms, in contrast, has the potential to redefine global supply chain operations by democratizing supply chain transparency.

By comparing signaling costs in the presence and absence of blockchain, our analysis also provides a way to quantify the benefits that this technology affords to potential adopters. These benefits can then be weighed against the practical costs of adoption, so that firms can make a more informed strategic decision about implementation.

Our analysis also sheds light onto what types of supply chains would profit the most from blockchain adoption and what specific benefits this would provide.

Prediction 1. Propensity to adopt blockchain is positively related to the firm's creditworthiness.

This prediction follows from Proposition 3, according to which only high-quality firms adopt blockchain in equilibrium, and Theorem 1, according to which preference for adoption increases with the firm's success probability. Note that our model assumes that a firm's creditworthiness (success probability) is not observable to outsiders. Therefore, to test this prediction, one would need to measure creditworthiness by the ex-post default rate, while controlling for risk factors observable ex ante, such as firm age, size, profitability, leverage, etc.

Prediction 2. Propensity to adopt blockchain by a high-quality firm is (a) negatively related to the firm's market size, and (b) positively related to the firm's operating costs.

The result follows from Corollary 1, according to which the high type's preference for blockchain adoption is negatively related to α_H, which can relate to the firm's market size as well as its operating costs. Importantly, this negative relation is true only conditional on the firm being of a high type. Thus, we expect the relation to hold empirically only within a sample of relatively creditworthy firms.

Prediction 3. Propensity to adopt blockchain is (a) positively related to perishability of the firm's inputs, and (b) negatively related to liquidity of the firm's inputs.

As stipulated in Corollary 2, the benefit of supply chain transparency disappears if input inventory can be converted to cash without a loss. In general, as the salvage value of inventory increases, overordering becomes a weaker signal, and the benefit of supply chain transparency afforded by blockchain fades. Therefore, the blockchain benefits are expected to be greater when the firm sources perishable or illiquid (e.g., unique or customized) inputs.

Prediction 4. Propensity to adopt blockchain is positively related to the degree of information asymmetry that the firm is subject to.

This prediction follows directly from the fact that the benefit of blockchain identified in our model exists only in the presence of information asymmetry between borrowers and lenders (compare the results in Sections 2.1 and 2.4). Therefore, we expect the technology to be more prevalent in supply chains that are innovation-intensive; dominated by smaller, privately-owned firms; and/or in which output is more differentiated across firms.

Prediction 5. Blockchain adoption decreases the cost of debt financing and operational distortions for high-quality firms.

This prediction follows from Theorem 1, according to which blockchain adoption allows high-quality firms to signal their quality to lenders in some cases, and it reduces signaling cost in others. In other words, blockchain enables high-quality firms to obtain external funds at the cost that reflects their true default risk, which would either be impossible or would require larger operational distortions otherwise. The second part of the prediction follows from the fact that blockchain adoption reduces excess ordering.

Appendix. Additional Specifications for the b_verify System

Securing Receipts via Merkle Prefix Tries

In b_verify, a warehouse receipt is a digital document that can be issued, transferred, or redeemed. To perform these operations, all parties involved (e.g., the warehouse and depositor) must consent by digitally signing ownership changes. Note that the use of digital signatures alone does not guarantee that everyone will see the same owner for each receipt. For example, a nefarious participant could transfer a receipt he owns and then try to redeem or loan the old receipt. To prevent this from happening, the receipts in b_verify have associated cryptographic proofs, which use commitments to the Bitcoin blockchain to guarantee that everyone will agree on who currently owns a receipt. b_verify constructs these proofs efficiently through the use of a central server, shown in Figure A.1. However, unlike in a typical client–server model, b_verify does not require clients to trust the server. Instead, participants use cryptographic proofs to carefully update and verify receipt ownership.

Figure A.2 illustrates this process. Receipt ownership is tracked by using a forest of Merkle trees. Each user's "account" at a given warehouse is mapped to a Merkle tree that contains the receipts issued by the warehouse that are currently owned by the user. This set changes over time as individual receipts are issued, transferred, or redeemed. Any updates to this tree require both the warehouse and the user to sign. The central server is only required to manage the roots (a 32-byte cryptographic digest) of each Merkle tree. It maintains a mapping from the account to the root using a data structure called a Merkle Prefix Trie (MPT). The server commits to this mapping by "witnessing" it to Bitcoin through broadcast of a Bitcoin transaction. Whenever the ownership record changes, some subset of the Merkle trees will change, resulting in new roots. The server will then update its data structures to reflect the changes and witness the new commitment to Bitcoin. Because clients do not trust the server, they require proofs from the server to ensure that the mapping was updated correctly. In b_verify, the server can perform many updates simultaneously in a single transaction on the Bitcoin ledger to increase throughput and amortize transaction costs.

Figure A.1. (Color online) A Basic, Low-Cost "Untrusted" Server Is Used to Pool Transactions and Obfuscate the Data via Hashing Before It Is Sent to the Bitcoin Blockchain

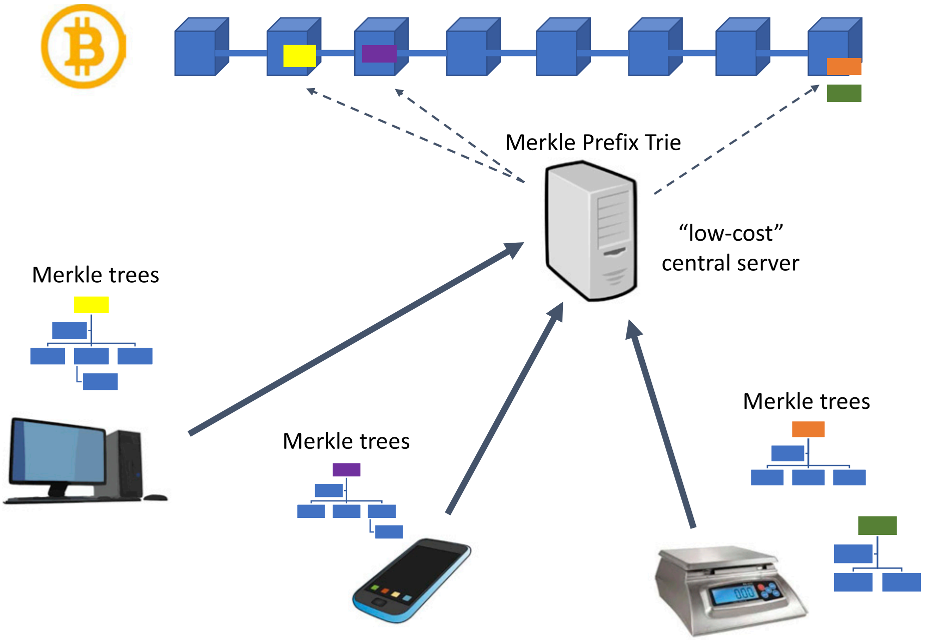

Figure A.2. (Color online) The b_verify Merkle Prefix Trie (MPT) Structure

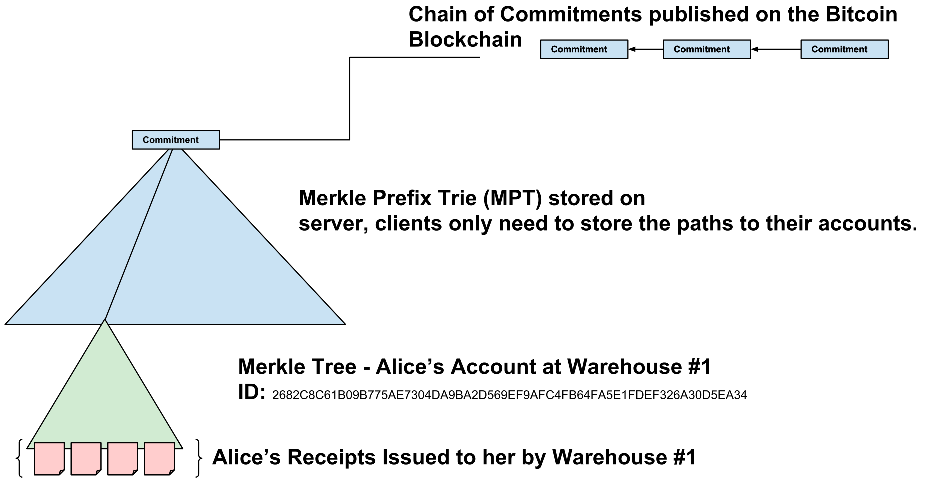
Chain of Commitments published on the Bitcoin Blockchain
Commitment
Commitment
Commitment
Commitment
Merkle Prefix Trie (MPT) stored on server, clients only need to store the paths to their accounts.
Merkle Tree - Alice’s Account at Warehouse #1
ID: 2682C8C61B09B775AE7304DA9BA2D569EF9AFC4FB64FA5E1FDEF326A30D5EA34
Alice’s Receipts Issued to her by Warehouse #1

Users in b_verify participate by downloading a mobile “wallet” application. This application periodically syncs with the Bitcoin network and the central server. Hidden from the user, the application maintains the user’s receipts along with the associated cryptographic proof. As the server broadcasts new commitments, the mobile application updates the proofs as needed. The proof of ownership of a receipt can then be sent to and verified by anyone using the b_verify protocol and does not require the involvement of the central server. Note that these proofs are only intended to be read and evaluated by computers; to a human, the proofs appear to be very long strings of numbers written in hexadecimal—for example, see Figure A.3. In addition to the receipts, the user’s wallet application contains the secret keys necessary to transfer or receive new receipts. When a receipt is issued, transferred, or redeemed, the wallet application ensures that the operation has been performed correctly and then uses these keys to digitally sign updates. The wallet application can then ask for a cryptographic proof from the server.

Figure A.3. (Color online) Hashed Client Data Stored in the Server’s Merkle Prefix Trie

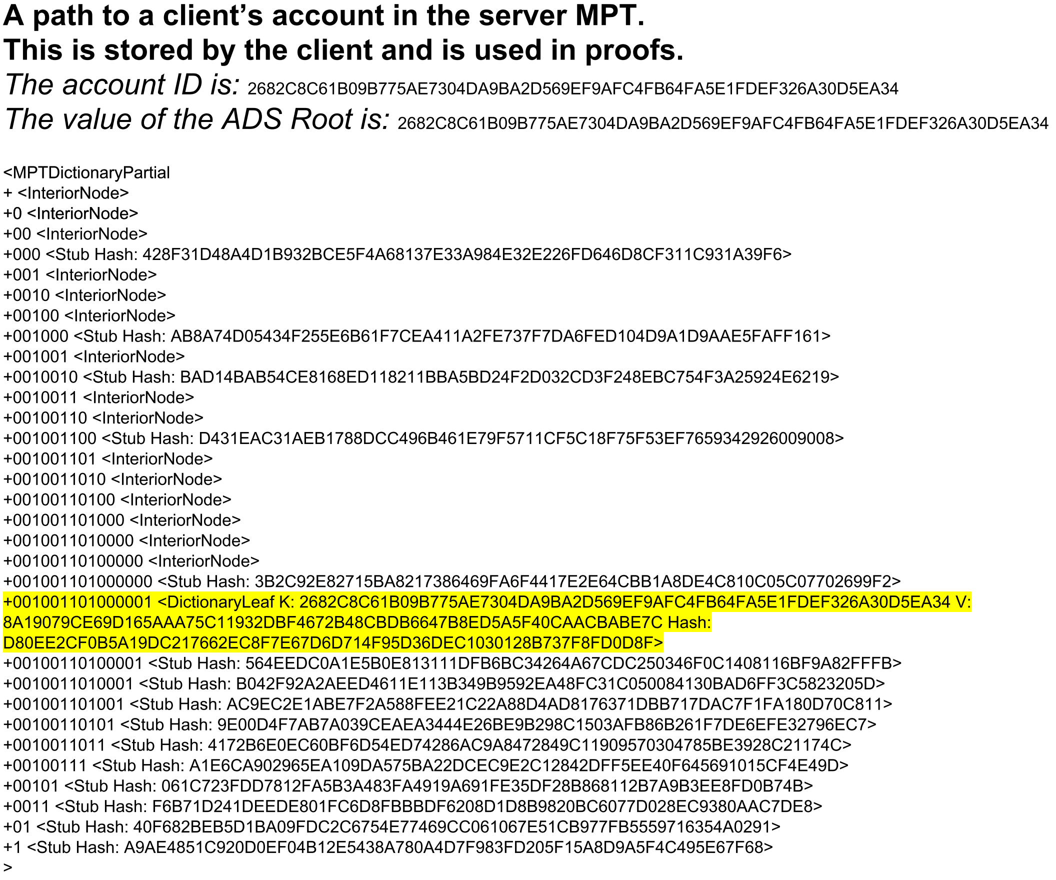

A path to a client’s account in the server MPT.
This is stored by the client and is used in proofs.

The account ID is: 2682C8C61B09B775AE7304DA9BA2D569EF9AFC4FB64FA5E1FDEF326A30D5EA34
The value of the ADS Root is: 2682C8C61B09B775AE7304DA9BA2D569EF9AFC4FB64FA5E1FDEF326A30D5EA34

```
<MPTDictionaryPartial
+ <InteriorNode>
+0 <InteriorNode>
+00 <InteriorNode>
+000 <Stub Hash: 428F31D48A4D1B932BCE5F4A68137E33A984E32E226FD646D8CF311C931A39F6>
+001 <InteriorNode>
+0010 <InteriorNode>
+00100 <InteriorNode>
+001000 <Stub Hash: AB8A74D05434F255E6B61F7CEA411A2FE737F7DA6FED104D9A1D9AAE5FAFF161>
+001001 <InteriorNode>
+0010010 <Stub Hash: BAD14BAB54CE8168ED118211BBA5BD24F2D032CD3F248EBC754F3A25924E6219>
+0010011 <InteriorNode>
+00100110 <InteriorNode>
+001001100 <Stub Hash: D431EAC31AEB1788DCC496B461E79F5711CF5C18F75F53EF7659342926009008>
+001001101 <InteriorNode>
+0010011010 <InteriorNode>
+00100110100 <InteriorNode>
+001001101000 <InteriorNode>
+0010011010000 <InteriorNode>
+00100110100000 <InteriorNode>
+001001101000000 <Stub Hash: 3B2C92E82715BA8217386469FA6F4417E2E64CBB1A8DE4C810C05C07702699F2>
+001001101000001 <DictionaryLeaf K: 2682C8C61B09B775AE7304DA9BA2D569EF9AFC4FB64FA5E1FDEF326A30D5EA34 V:
8A19079CE69D165AAA75C11932DBF4672B48CBDB6647B8ED5A5F40CAACBABE7C Hash:
D80EE2CF0B5A19DC217662EC8F7E67D6D714F95D36DEC1030128B737F8FD0D8F>
+00100110100001 <Stub Hash: 564EEDC0A1E5B0E813111DFB6BC34264A67CDC250346F0C1408116BF9A82FFFB>
+0010011010001 <Stub Hash: B042F92A2AEED4611E113B349B9592EA48FC31C050084130BAD6FF3C5823205D>
+001001101001 <Stub Hash: AC9EC2E1ABE7F2A588FEE21C22A88D4AD8176371DBB717DAC7F1FA180D70C811>
+00100110101 <Stub Hash: 9E00D4F7AB7A039CEAEA3444E26BE9B298C1503AFB86B261F7DE6EFE32796EC7>
+0010011011 <Stub Hash: 4172B6E0EC60BF6D54ED74286AC9A8472849C11909570304785BE3928C21174C>
+00100111 <Stub Hash: A1E6CA902965EA109DA575BA22DCEC9E2C12842DFF5EE40F645691015CF4E49D>
+00101 <Stub Hash: 061C723FDD7812FA5B3A483FA4919A691FE35DF28B868112B7A9B3EE8FD0B74B>
+0011 <Stub Hash: F6B71D241DEEDE801FC6D8FBBBDF6208D1D8B9820BC6077D028EC9380AAC7DE8>
+01 <Stub Hash: 40F682BEB5D1BA09FDC2C6754E77469CC061067E51CB977FB5559716354A0291>
+1 <Stub Hash: A9AE4851C920D0EF04B12E5438A780A4D7F983FD205F15A8D9A5F4C495E67F68>
>
```

Figure A.4. (Color online) Secure Issuance and Commitment to the Bitcoin Blockchain of Warehouse Receipts

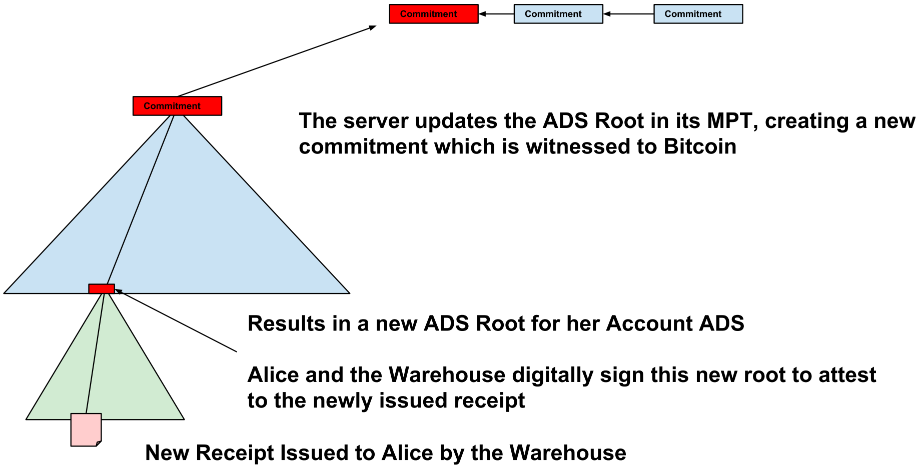
Commitment
Commitment
Commitment
Commitment
The server updates the ADS Root in its MPT, creating a new commitment which is witnessed to Bitcoin
Results in a new ADS Root for her Account ADS
Alice and the Warehouse digitally sign this new root to attest to the newly issued receipt
New Receipt Issued to Alice by the Warehouse

Figure A.5. (Color online) Secure Transfer and Commitment to the Bitcoin Blockchain of a Client-to-Client Transaction

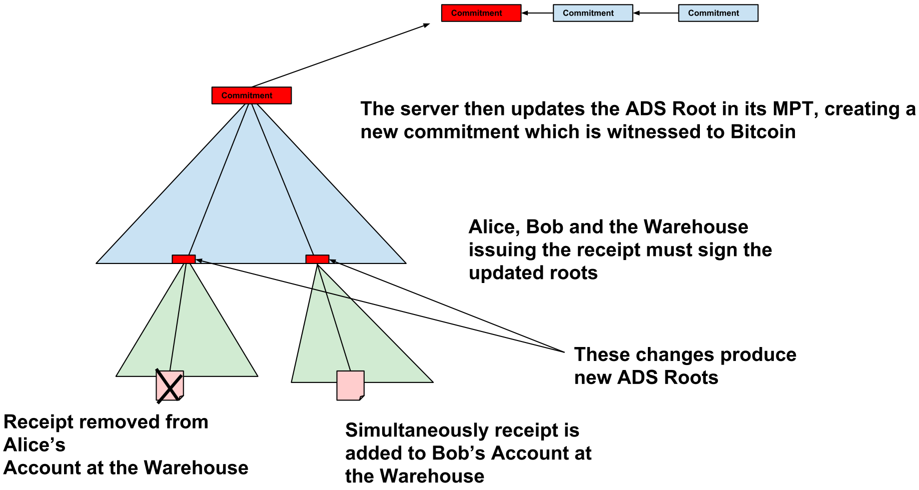
Commitment
Commitment
Commitment
Commitment
The server then updates the ADS Root in its MPT, creating a new commitment which is witnessed to Bitcoin
Alice, Bob and the Warehouse issuing the receipt must sign the updated roots
These changes produce new ADS Roots
Receipt removed from Alice's Account at the Warehouse
Simultaneously receipt is added to Bob's Account at the Warehouse

Issuing a Receipt

If a warehouse wishes to issue a receipt to Alice (a fictitious user), the warehouse adds the receipt to Alice's "account" at the warehouse, by inserting the cryptographic hash of the receipt into the Merkle tree representing her account. Alice and the warehouse must then sign the new root of this data structure, reflecting the addition of the receipt. Finally, Alice and the warehouse send the signed updated root to the server, who updates the root in the MPT stored on the server and commits it to Bitcoin. These actions are illustrated in Figure A.4.

Once the server has witnessed the commitment to Bitcoin, Alice and the warehouse can ask the server for the path to the issued receipt. Using this information Alice and the warehouse can present and share the proof of the issued receipt with others.

Transferring a Receipt

If Alice seeks to transfer a receipt issued to her by the warehouse to Bob (another fictitious user), she removes the receipt from her account Merkle tree and adds it to Bob's account Merkle tree. The addition and removal of the receipt are reflected in the resulting new Merkle roots for the respective data structures. At this point, Alice, Bob, and the warehouse must sign the new roots for the transfer to become valid. The server then updates its MPT to reflect this and witnesses the update to Bitcoin. These actions are illustrated in Figure A.5.

Once the server has published the commitment to Bitcoin, Alice and Bob can asks for paths to their respective account data structures. Using these paths, both of them can construct proofs for the receipts they own, as well as a proof of provenance for the receipt. Crucially, by broadcasting the new commitment, Alice's previous proof of ownership of the receipt is invalidated. This is critical—she can no longer present a correct proof that she owns the receipt after she has transferred it. Other systems fail to achieve this property without the use of a trusted intermediary. For more information, we refer the reader to Aspegren (2018).

Endnotes

[1] According to the Federal Bureau of Investigation, "Letters of credit frauds are often attempted against banks by providing false

documentation to show that goods were shipped when, in fact, no goods or inferior goods were shipped"

[2] Perhaps the most infamous example of falsifying warehouse receipts in asset-based lending is the De Angelis salad oil swindle, which nearly crippled the New York Stock Exchange (Taylor 2013).

[3] Blockchains were originally designed to solve the infamous double-spending problem for digital currency—that is, to ensure that a digital asset transmitted from one party to another has not already been spent elsewhere.

[4] In the Bitcoin blockchain, for instance, the task of recording transactions is assigned to individual miners who compete through a proof-of-work mechanism at every round, to append blocks to the existing chain, in exchange for compensation (composed of transaction fees and new currency issuance).

[5] The reason we assume that a higher probability of success is associated with a larger market size is that both are likely to result from superior management or operations capabilities.

[6] Fairly priced credit is a standard assumption in the finance literature (see, e.g., Biais and Gollier 1997 and Burkart and Ellingsen 2004). We normalize the lender's cost of capital to zero without a loss of generality.

[7] As we show in Section B.4 in the EC, assuming a threshold-type belief structure is without any loss of generality because starting with an arbitrary belief structure leads to the same *least-cost* separating equilibrium.

[8] Whereas the LCSE loan amounts are always unique, the belief threshold is unique only in this second scenario. When both firms choose first best, any belief threshold in $[d, D_H^{\text{fb}}]$ is consistent with their equilibrium actions.

[9] Whereas the LCSE quantities $\{Q_L^{\text{se}}, Q_H^{\text{se}},\}$ are always unique, the equilibrium belief threshold is unique only in this second scenario. When both firms order first best, any belief threshold in $[q, Q_H^{\text{fb}}]$ is an equilibrium belief structure.

[10] Recall that in any SE, the low type follows its first best, and it is the high type that bears all signaling costs.

[11] We developed b_verify in collaboration with the Digital Currency Initiative of the MIT Media Laboratory, with monetary support from the Inter-American Development Bank and the MIT Legatum Center for Development and Entrepreneurship, as well as expertise and user insights provided by the Government of Mexico, the World Bank Group, and a number of private companies in Mexico and Ukraine.

[12] To motivate the server to witness new log statements and provide proof updates, b_verify introduces a novel incentive design with a penalty smart contract adjudicated on the Ethereum blockchain.

[13] Equivocation is a concept from Computer Science that refers to a situation in which an entity is able to make inconsistent statements to different parties.

[14] b_verify can be used with any public blockchain; we select Bitcoin because it is the oldest and arguably the most secure against forking.

[15] This is important because it means that clients can fully participate without having to download the entire Bitcoin blockchain that is currently (as of January 2019) close to 200 GB.

References

Acemoglu D, Drakopoulos K, Ozdaglar A (2017) Information obfuscation in a game of strategic experimentation. Working paper, Massachusetts Institute of Technology, Cambridge.

Alibhai S, Bell S, Conner G (2017) *What's Happening in the Missing Middle? Lessons from Financing SMEs* (World Bank, Washington, DC).

Aspegren H (2018) b_verify: Scalable non-equivocation for verifiable management of data. Thesis, Massachusetts Institute of Technology, Cambridge.

Babich V, Hilary G (2018) What operations management researchers should know about blockchain technology. Working paper, Georgetown University, Washington, DC.

Babich V, Tsoukalas G, Marinesi S (2020) Does crowdfunding benefit entrepreneurs and venture capital investors? *Manufacturing Service Oper. Management*. Forthcoming.

Bebchuk LA, Stole LA (1993) Do short term objectives lead to under or overinvestment in long term projects? *J. Finance* 48(2):719–730.

Belavina E, Marinesi S, Tsoukalas G (2020) Designing crowdfunding platform rules to deter misconduct. *Management Sci*. Forthcoming.

Besanko D, Thakor AV (1987) Competitive equilibrium in the credit market under asymmetric information. *J. Econom. Theory* 42(1): 167–182.

Biais B, Gollier C (1997) Trade credit and credit rationing. *Rev. Financial Stud.* 10(4):903–937.

Biais B, Bisiere C, Bouvard M, Casamatta C (2019) The blockchain folk theorem. *Rev. Financial Stud*. 32(5):1662–1715.

Bimpikis K, Drakopoulos K, Ehsani S (2018) Disclosing information in strategic experimentation. Stanford Graduate School of Business Working Paper 3635, Stanford University, Stanford, CA.

Budish E (2018) The economic limits of Bitcoin and the blockchain. NBER Working Paper No. 24717, National Bureau of Economic Research, Cambridge, MA.

Burkart M, Ellingsen T (2004) In-kind finance: A theory of trade credit. *Amer. Econom. Rev.* 94(3):569–590.

Cachon G, Lariviere M (2001) Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Sci.* 47(5):629–646.

Candogan O, Drakopoulos K (2019) Optimal signaling of content accuracy: Engagement vs. misinformation. *Oper. Res*. Forthcoming.

Catalini C, Gans JS (2017) Some simple economics of the blockchain. Rotman School of Management Working Paper No. 2874598, Toronto; MIT Sloan Research Paper No. 5191-16, Massachusetts Institute of Technology, Cambridge, MA.

Chakraborty S, Swinney R (2020) Signaling to the crowd: Private quality information and rewards-based crowdfunding. *Manufacturing Service Oper. Management*, ePub ahead of print April 3, https://doi.org/10.1287/msom.2019.0833.

Chick SE, Hasija S, Nasiry J (2016) Information elicitation and influenza vaccine production. *Oper. Res.* 65(1):75–96.

Chod J (2017) Inventory, risk shifting, and trade credit. *Management Sci.* 63(10):3207–3225.

Chod J, Lyandres E (2018) A theory of ICOs: Diversification, agency, and information asymmetry. Working paper, Boston College, Boston.

Chod J, Trichakis N, Tsoukalas G (2019a) Supplier diversification under buyer risk. *Management Sci.* 65(7):3150–3173.

Chod J, Trichakis N, Yang SA (2019b) Platform tokenization: Financing, governance, and moral hazard. Working paper, Boston College, Boston.

Cho IK, Kreps M (1987) Signaling games and stable equilibria. *Quart. J. Econom.* 102(2):179–221.

Cong LW, He Z (2019) Blockchain disruption and smart contracts. *Rev. Financial Stud.* 32(5):1754–1797.

Cong LW, He Z, Li J (2019) Decentralized mining in centralized pools. NBER Working Paper No. 25592, National Bureau of Economic Research, Cambridge, MA.

Cong LW, Li Y, Wang N (2018) Tokenomics: Dynamic adoption and valuation. Columbia Business School Research Paper 18-46, Columbia University, New York.

Diamond DW (1991) Monitoring and reputation: The choice between bank loans and directly placed debt. *J. Polit. Econom.* 99(4):689–721.

Dong L, Ren L, Zhang D (2019) Financing small and medium-size enterprises via retail platforms. Working paper, Washington University in St. Louis, St. Louis, MO.

Duan J-C, Yoon SH (1993) Loan commitments, investment decisions and the signalling equilibrium. *J. Banking Finance* 17(4):645–661.

Emery GW (1984) A pure financial explanation for trade credit. *Quart. J. Econom.* 19(3):271–285.

Fabbri D, Menichini AMC (2016) The commitment problem of secured lending. *J. Financial Econom.* 120(3):561–584.

Falk B, Tsoukalas G (2020) Token weighted crowdsourcing. *Management Sci*. Forthcoming.

Gan R, Netessine S, Tsoukalas G (2019) Inventory, speculators and initial coin offerings. Wharton School Research Paper, University of Pennsylvania, Philadelphia.

Halaburda H (2018) Blockchain revolution without the blockchain. Working paper, New York University, New York.

Hinzen FJ, John K, Saleh F (2019) Proof-of-work's limited adoption problem. Working paper, NYU Stern School of Business, New York University, New York.

Huberman G, Leshno JD, Moallemi CC (2017) Bitcoin's fatal flaw: the limited adoption problem. Working paper, Columbia University, New York.

Iancu DA, Trichakis N, Tsoukalas G (2017) Is operating flexibility harmful under debt? *Management Sci.* 63(6):1730–1761.

IBM TrustChain (2017) Consortium of jewelry industry leaders announce TrustChain, first global blockchain initiative to bring full transparency to consumers. Press release, IBM, Armonk, NY. Accessed September 1, 2019, http://newsroom.ibm.com/announcements?item=122899.

Jain N (2001) Monitoring costs and trade credit. *Quart. Rev. Econom. Finance* 41(1):89–110.

Lai G, Xiao W (2018) Inventory decisions and signals of demand uncertainty to investors. *Manufacturing Service Oper. Management* 20(1):113–129.

Lai G, Xiao W, Yang J (2012) Supply chain performance under market valuation: An operational approach to restore efficiency. *Management Sci.* 57(2):332–346.

Milde H, Riley JG (1988) Signaling in credit markets. *Quart. J. Econom.* 103(1):101–129.

Nakamoto S (2008) Bitcoin: A peer-to-peer electronic cash system. Working paper, Bitcoin, Arlington, VA.

Özer Ö, Wei W (2006) Strategic commitments for an optimal capacity decision under asymmetric forecast information. *Management Sci.* 52(8):1238–1257.

Özer Ö, Zheng Y, Ren Y (2014) Trust, trustworthiness, and information sharing in supply chains bridging China and the United States. *Management Sci.* 60(10):2435–2460.

Ross SA (1977) The determination of financial structure: The incentive-signalling approach. *Bell J. Econom.* 8(1):23–40.

Saleh F (2019) Blockchain without waste: Proof-of-stake. Working paper, McGill University, Montreal.

Schmidt W, Gaur V, Lai R, Raman A (2015) Signaling to partially informed investors in the newsvendor model. *Production Oper. Management* 24(3):383–401.

Spence M (1973) Job market signaling. *Quart. J. Econom.* 87(3):355–374.

Stiglitz JE, Weiss A (1981) Credit rationing in market with imperfect information. *Amer. Econom. Rev.* 71(3):393–410.

Tang CS, Yang SA, Wu J (2018) Sourcing from suppliers with financial constraints and performance risk. *Manufacturing Service Oper. Management* 20(1):70–84.

Taylor B (2013) How the salad oil swindle of 1963 nearly crippled The NYSE. *Bus. Insider* (November 23), https://www.businessinsider.com/the-great-salad-oil-scandal-of-1963-2013-11.

Tomescu A, Devadas S (2017) *Catena: Efficient non-equivocation via Bitcoin. Proc. 2017 IEEE Sympos. Security Privacy* (IEEE Computer Society, Los Alamitos, CA), 393–409.

Trichakis N, Tsoukalas G, Moloney E (2015) Credem: Banking on cheese. Harvard Business School Case 615-046, Harvard Business School, Boston.

Yermack D (2017) Corporate governance and blockchains. *Rev. Finance* 21(1):7–31.