

NBER WORKING PAPER SERIES

THE FINE PRINT IN SMART CONTRACTS

Joshua S. Gans

Working Paper 25443

<http://www.nber.org/papers/w25443>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

January 2019

Thanks to Kevin Bryan, Gillian Hadfield, Richard Holden, Arun Majumdar and Kamil Shafiq for useful comments on an earlier draft of this paper. Responsibility for all views are my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w25443.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Joshua S. Gans. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Fine Print in Smart Contracts
Joshua S. Gans
NBER Working Paper No. 25443
January 2019
JEL No. D86,K12

ABSTRACT

One of the purported benefits of blockchain technologies is the ability to house what have been termed ‘smart’ contracts. Such contracts are potentially self-executing depending on the state of information recorded on a blockchain ledger. This paper examines the capabilities of smart contracts from an economic perspective. It is demonstrated that by improving observability and reducing the costs of verification of contract obligation performance, the space of feasible contracts can be enlarged. Moreover, by providing commitments to various monetary payments, a blockchain can potentially create a foundation to house certain mechanisms that have been shown to overcome difficulties of contractual incompleteness. This is demonstrated using a simple international trade environment. Thus, even though smart contracts must respect the incentives of decision-makers in their obligations, they have the potential to use easily verifiable elements to create incentives to reduce hold-up and other contractual difficulties.

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
CANADA
and NBER
joshua.gans@gmail.com

I Introduction

Beyond its initial application in creating digital tokens (aka Bitcoin), one of the more intriguing possibilities from Blockchain technologies is that they could support what have been termed ‘smart contracts.’ These contracts are based on an idea by Nick Szabo (1994, 1996). “A smart contract is a computerized transaction protocol that executes the terms of a contract.” (Szabo, 1994) Szabo is quick to point out that automated contract execution is not new. He argues that a vending machine does just that. You indicate what you want, pay the price and, following that, the machine dispenses the product. What is new is the idea that the contract involved could be both purely digital (the vending machine is not) and complicated in that it may involve complex calculations, multipart deals, transfer of rights and various forms of encryption.

Contracting has always been about the fine print. Contracts are a set of obligations, many of which arise in particular circumstances. To be sure, at a high level a contract is representing a transaction, trade or exchange of payment for goods and services. But things often do not go to plan and contracts anticipate these. If there is a disruption, whose responsibility is it to bear the costs? If there is a dispute, how will it be resolved? And when is it the case that it is just better for everyone to walk away from their obligations?

Smart contracts promise to automate all of this. To be sure, code can capture the fine print. But what code cannot do is determine whether a particular event or state that gives rise to obligations has arisen. That relies on interfacing with the ‘real’ world. Moreover, it often relies on a human element to make sure the right information is injected at the right time. As with every other situation involving humans, that is a weak link in full automation. We want to have automated contracts that we can trust and rely upon but to make them work properly there is often human choice in the fine print. Importantly, to understand fully the prospects for smart contracting today we have to pay attention to that fine print. We can’t leave it unread.

So how do Blockchain technologies factor into this? Blockchain technologies involve a ledger of information which represents the consensus of a decentralised mechanism. In other words, the information in the ledger is not vouched for by one agent (as would be the case in a centralised system). Instead, it is vouched for by many, distributed agents who agree that the information on the Blockchain today was the same as the information placed on the Blockchain

at an earlier time. Moreover, the information on the Blockchain is *valid* in that it conforms to a set of constraints (perhaps based on what was there earlier). There is nothing that implies that the information on the Blockchain is ‘The Truth’ but instead that, whatever it was, it was placed there at a specific time and has not been altered.²

It is at this point where the usefulness of the Blockchain in contracting becomes apparent. A contract is an agreement between two or more parties. When such an agreement is reached, the fact of that agreement as well as the terms of it can be put on the Blockchain. That way, if there is a dispute as to what was agreed to and when, it is very easy to verify that information.³ To be sure, in the paper economy, we have relied on paper and copies to provide a similar function (usually distributing it amongst different parties so that later a consensus can be identified). But in the digital world, there is concern that such information could be easily altered. Hence, it is critical to have an immutable record.

Compared to paper, however, the Blockchain offers the possibility of more. Why simply record a contract when you can code the contract? In Bitcoin, two parties agree to the transfer of bitcoins from one to another at a specific time. However, the protocol then does an extra step, it transfers the ownership of the bitcoins as specified in the agreement. To do this actually takes a number of steps. First, the payor notifies the network they wish to send bitcoin to the payee and provides their private key. Second, the network then checks that the payor does, in fact, own those bitcoin (by verifying they possess the correct private key) and that they have not already told the network of their intention to send them to someone else. Third, only then, is the intention converted into a reassignment of ownership in the next block on the blockchain. All of these steps are performed when we give cash to one another (it is more complicated for checks and wire transfers). Having the cash in hand signifies the first two steps and then giving it to the payee completes the process. In other words, the Bitcoin protocol is able to automate the execution of the agreement to transfer bitcoin ownership once the payor validly initiates it.

² This statement is not quite true. What the information and time on the Blockchain tell us is that the current consensus, which is that the information was on the Blockchain at that time. In principle, the Blockchain may have forked at a subsequent time with a potential re-write of history. That said, forks are public as are the changes they bring about. Thus, the path from the initial vouching for the information to the present can be traced.

³ Catalini and Gans (2016) highlight that it is this potential reduction in the cost of such verification that represents the principle economic effect arising from Blockchain technologies.

If the Bitcoin protocol can automate a digital property rights transfer, then it is not hard to imagine that an automated protocol could execute more complex contracts. One such complexity may be to include a contingency. For Bitcoin, execution takes place at the discretion of the payor. But one can also imagine that the transfer might be contingent on some event. For instance, perhaps the payment should only be made if it rains in a particular location next Thursday. That could be coded along with a data source for the event of rain (which has the property that it either does or does not occur). Then having agreed to the contract and provided the relevant signatures (i.e., private keys), the process would be automated from that point forward.

The validation of an agreement and its contents is critical to being able to automate an agreement's execution. Without that validation, the parties could not be sure what was being executed was what they intended. For this reason, the Blockchain provides the means to validate the code of an agreement and also to verify its performance. However, this does not necessarily imply that a Blockchain can fully automate a contract perfectly. That, as we will see, depends on the nature of the agreement.

The purpose of this paper is to consider when precisely smart contracts will be useful if hosted on a Blockchain. In particular, it will highlight what I will call "the fundamental smart contract challenge" that results because smart contracts are inherently digital but must, in order to be useful, interact with non-digital activity. This is something that has been neglected in discussions of the Blockchain even if, as I will argue, economists have concentrated almost exclusively on such issues over the past two decades of contract theory. It is the fine print we cannot neglect.

In the end, a smart contract is an automated contract. The key issue will be to identify what aspect of contracting is automated. Varian (2010) noted that an increasing number of transactions were becoming computer mediated as a result of the Internet. He noted that it was possible to use data to make contracts more nuanced. For instance, a car rental company could monitor the speed of driving and adjust a driver's rate accordingly. Car finance companies could put in connected sensors into cars and prevent operation if payments were missed. Online advertisers, like Google, could run auctions in real time based on the submitted bids of advertisers. In each case, data was collected after the contract was signed which then determined

the payments and sometimes obligations of parties. Computer-mediation allowed these to be automatically executed.

More ambitiously, Casey and Niblett (2017) see the possibility of “self-driving contracts.” These contracts do not even go the way of computer-mediated transactions in specifying the full details and contingencies of the contract. Instead, they see a future where parties could specify intentions or goals and then an artificial intelligence will fill the gaps in the contract in line with these goals as more is learned about the environment the parties are operating in. The way they see this as working is that artificial agents can fill in contractual gaps in the way a court or a mediator might do today. However, while those gaps may be filled in requires explicit judicial enforcement. They see this as requiring various changes in how contract law operates today.

Such contracts do not necessarily require a Blockchain and, as such, are not really part of the scope of smart contracting as it is postulated. That said, they do involve the notion of automation which is a key part of smart contracting.

II Obligations and Trust

Contracts are made up of obligations. For instance, a buyer and seller may strike an agreement whereby the buyer is obliged to pay the seller some money while the seller is obliged to supply a good or service to the buyer. An agreement will facilitate trade or other valuable interactions when an agent can be reasonably certain that an obligation on another party will be met.

Trust is a means of creating reasonable expectations that obligations will be met when the performance of obligations is separated in time. To see this, trust-less performance of obligations can occur when exchange is simultaneous. For instance, a buyer and seller meet and, at the same moment, fulfill payment and product supply obligations. This is why hostages might be exchanged on a bridge which allows the obligations to be fulfilled simultaneously. Or, alternatively, why the first experiments in kidney exchanges worked with donors having operations simultaneously. If trust is present, then simultaneity is not required and obligations can be performed at different times.

There are, of course, social mechanisms that allow trust to arise so that obligations can be time separated. Some people might be considered trustworthy and there are social cues that allow people to assess that. In other situations, parties have an on-going relationship and the prospect of repeated interactions causes them to honour and expect to have honoured obligations.

Beyond such social forces, there have been institutional mechanisms that have developed to allow obligations to be performed sequentially. The most ubiquitous of these is contract law. Contract law bundles obligations together and then sets forth conditions under which the performance of those obligations will be required. To ensure performance, an independent third party — usually a Court — will assess obligation performance and impose penalties or other requirements if obligations are not met. The expectation of such verification and sanction serves to cause agents to meet their obligations. In the process, the obligations become trusted.

Another critical institution that allows trust to develop is reputation. Reputation allows a relationship to extend beyond the parties that can directly observe the performance of past obligations. By having a means of communicating that performance in a public manner, an agent can have an incentive to perform obligations in order to preserve their reputation and enhance opportunities for future dealings. As an institution, reputations require a means of accounting for past performance in a manner that can be verified, be secure and be communicated widely.

In order for a Blockchain to have value it must work to enhance trust in situations where neither current social nor institutional mechanisms are present. Blockchains are sometimes held to be trust-less technologies but this refers to their potential ability to work independently of social mechanisms for trust. In order to fulfill an economic function, they need to substitute for technologies — both legal and other — that provide institutional bases for trust. In this regard, smart contracts that are supported by the Blockchain will still be measured by their ability to enhance trust in terms of the set of obligations that comprise those contracts.

III The Contract Workflow

In order to understand the potential value of the Blockchain in facilitating smart contracts it is useful to explore, in somewhat laborious detail, the workflow associated with contract performance, execution and verification. Much legal and economic work on contracts focusses

on the contract formation process — its terms, whether an agreement is possible, etc. Here, instead, I focus on what happens after a contract and its terms have been agreed upon.

Consider two parties to a contract — the erstwhile Alice and Bob. Each has obligations — x_A and x_B — that they have agreed to perform. If this occurs, Alice receives $v_A(x_B) - c_A(x_A)$ while Bob receives $v_B(x_A) - c_B(x_B)$. The values and costs may include some form of payment between them. In each case, if an obligation is not performed, it costs the performer nothing and reduces the value received by the other to 0.

It is useful to note at this stage that most contractual obligations are created by an event — or what I will term a **state**. The most obvious state is a signature on an agreement. However, some states may be that an obligation to ship a good will only be created once payment is received or, alternatively, a payment will only be made once a good is received. In other situations, the state may represent an external event; e.g., an insurance payout only arises in the case of an accident. In other words, it is a state that triggers the creation of an obligation even within an agreed upon contract. Thus, one common form of conflict may arise is a dispute over whether the relevant state underlying an obligation has occurred. If it has, then a party will expect an obligation to be performed. If not, then no such expectation is warranted. Thus, if there is uncertainty or misunderstanding over whether an underlying state has arisen, this can cause a discord in expectations regarding the performance of an obligation.

For this reason, a party will likely have to monitor whether the underlying state triggering an obligation owed to them is actually arisen. In some situations, that state may be the performance of an obligation themselves. This may involve checking on a delivery, examining work quality or ensuring a payment has reached an account. For each state, suppose that it costs the monitoring party, m , to confirm with certainty whether a state has occurred or not. In some situations, this may be a cost to confirm whether an obligation itself has been performed. A party may want to incur this cost in order to know whether they should challenge performance or not.

In addition to monitoring underlying states and performance between the parties, obligation performance will often be verified by a third party (such as a Court). This is distinct from monitoring as that third party requires what is termed ‘hard evidence’ that performance has taken place. This may be some form of validated receipt or the results of a trusted audit. This too is

costly and, critically, involves the cooperation of the party performing the obligation. That party can, at cost z , provide hard evidence if it exists.

It is useful to note here that digitization provides potential costs and benefits with regard to this workflow. Digital technologies often imply a reduction in m — in that it is easier for a party receiving an obligation to check the underlying state or performance to their own satisfaction. By contrast, a purely digital indication of a state or performance may not be sufficient to verify performance to third parties. This is because digital assets can typically be easier to manipulate, tamper with or misrepresent. Thus, when there is a digital workflow, this means that the verification cost, z , may be higher.

To see the roles these costs play in contractual performance, let's focus on Alice and her incentives to undertake her contractual obligations. If Alice chooses to perform her obligations, then Bob faces a choice as to whether to confirm that performance or not. As Bob does not know whether Alice has chosen to perform or not, this is a difficult choice. Importantly, expending those costs allows Bob to challenge Alice's performance and require her to produce hard evidence. If she has performed, that is possible and she incurs z , and the matter is settled. If not, then some dispute resolution process would arise. Note that if there is no dispute resolution process, Alice has no incentive to produce hard evidence, Bob has no incentive to confirm performance and Alice has no incentive to perform her obligations. Thus, the contract agreement would be infeasible.

On the other hand, suppose that if Bob is successful in his dispute, he receives d while Alice pays d . In this case, so long as $z < d$, Alice will want to produce hard evidence if she has it and so long as $c_A(x_A) + z < d$, Alice will want to perform her obligations. In addition, Bob will gain $d - m$ if he confirms non-performance and $-m$ if he confirms performance. This means that Bob will play a mixed strategy at this node in any equilibrium. He will want the probability that he monitors, p , to be such that Alice has an incentive to perform: that is, so that $c_A(x_A) < p(d - z)$. Thus, Bob's expected costs of monitoring would be pm which is at least $mc_A(x_A)/(d-z)$. These costs represent a weight reducing the set of contracts that would otherwise be feasible. Notice that the higher is d , the lower are these expected costs but that if z is too high, then regardless of other factors, the contract will not be feasible.

Given this analysis we can prove the following:

Proposition 1. *The contract will only be feasible if $d > z$ and*

$$v_A(x_B) + v_B(x_A) > (c_A(x_A) + c_B(x_B))(d - z + m)/(d - z)$$

What this analysis shows is that to improve the scope for contracting, it is important to both reduce the costs of observability (m) as well as the costs of verification (z). Moreover, these have to be reduced for each obligation.

IV Impact of Blockchain Technology

With this model in place, we can now turn to consider the impact of blockchain technology. Blockchain technology involves three elements. First, and at the core of the technology, is a **distributed ledger**. While some argue that are, at essence, a database that records the timing and nature of states when a ledger is distributed, it has a significant security advantage over centralized databases. Typically, a centralized database has a single point of attack for access to all of its entries while a distributed ledger's entries must be individually attacked in order to be altered. The trade-off for this security is that centralized databases are more adaptable and easier to scale than distributed ledgers.⁴

What this means is that the distributed ledger — that is, the blockchain itself — is a place where states can be recorded. Their security means that they can then be relied upon to create obligations. In effect, this means that these represent a reduction in the monitoring costs (m) that might otherwise be necessary. Moreover, to the extent that states are used to create obligations outside of any particular set of parties, say, for instance, a regulatory reporting requirement or tax liability, when a state is securely recorded on the distributed ledger, it obviates the need for human auditors to confirm those states. This is especially so when the entries on two or more secure ledgers can be read to confirm one another (Cao et.al., 2018).

The second element of blockchain technologies is a **virtual machine**. All blockchains have a codebase that indicates the actions that can be coded with respect to both ledger entries and token assignment. Bitcoin's is fairly limited, allowing for token transfer and some other simple functions. By contrast, Ethereum has a virtual machine that is Turing complete — meaning that it

⁴ See Abadi and Brunnermeier (2018) and Arrunada and Garicano (2018).

is rich enough to perform the functions of any classical computer. This, of course, allows the creation of obligations to be automated and so has a role in reducing monitoring costs (m), but critically, the fact that contracts are coded allows them to adhere to standards of clarity and compliance that, say, a lawyer would otherwise provide, and also to allow interdependency — e.g., by allowing obligations in different contracts with a different set of parties to also be created. In other words, a common code can reduce the costs of sharing state information more widely.

The final element of blockchain technologies is a **token**. Tokens are a digital asset that can only appear in one entry in the ledger at a time. In this respect, they are capable of being made rival even though they are purely digital. Obviously, this provides a means of creating a currency that can be used for payments (as in Bitcoin) including to those that provide resources to the blockchain infrastructure. However, such tokens may also be used to underpin measures of the performance of obligations (such as a reputation score) and also as a means of communicating the performance of obligations themselves.⁵ In this regard, they provide a means of commitment — the possession of a token can generate certain decision and control rights. Such rights can, in turn, assist in creating favourable incentives for obligations to be performed. However, we leave the details of this to a discussion below and turn next to consider why such incentives are a challenge for smart contracts. In particular, you will note that while blockchain technologies can potentially reduce m , there is nothing inherent in these elements that reduces verification costs (z). That is, unless you can find a way to use the blockchain to obviate the need for a human third party that undertakes verification functions.

V The Fundamental Smart Contract Challenge

Recall that a smart contract is not necessarily a legal agreement — the parties may have an underlying contract which is legally binding. Instead, the smart contract is the code for the contract's execution. It is a means by which obligations can be recorded triggering certain other obligations that may operate in an automated way.

⁵ The seminal work in this regard is Kocherlakota (1998).

As already noted, one way in which smart contracts can improve contracting opportunities is by making the performance of obligations more easily confirmed. In the context of our model this would allow a reduction in m — the costs to any party of monitoring contractual states. This is most obviously useful with regard to payments but one can imagine situations where important information is digitised — such as the delivery of a product — and this could be recorded and made known at a low cost. Of course, it is digitisation that generated the opportunities for such transmission of information. Blockchain technologies may facilitate these but because the information required only has to be confirmed by the parties to the contract and not outside of the relationship, it is not clear whether such technologies add additional value.

The challenge for smart contracts comes with regard to performance obligations that need to be verified — that is, to provide hard evidence that an outside party would be able to use to confirm contract performance. In some situations, hard evidence — or its equivalent — can come when the party receiving the obligation confirms that the obligation has been performed. But for the most part the challenge arises because a party performing an obligation has an incentive to claim an obligation has been performed while those receiving it may have an incentive to claim it has not been performed; this is especially the case if their own obligations (e.g., making a payment) is contingent upon the obligation being reported as performed. Importantly, with transparency (low or incurred m) the parties themselves have common knowledge of obligation performance or non-performance.

To emphasize this further, note that hard evidence has the quality that it only exists if an obligation has actually been performed but who possesses it can be either party. Following Bull and Watson (2004), suppose that Alice is able to produce a document in the event that Bob performs his obligation but not otherwise. Then that document is positive evidence of performance while its absence is negative evidence of non-performance. Suppose, in a smart contract, there will be a payment from Alice to Bob if Bob performs his obligation but not otherwise. The question is: will Alice disclose the document to trigger the smart contract payment? If that payment is positive, it is easy to see that she will not. If Bob performs his obligation, Alice will receive the benefit of that and not have to pay if the document is not disclosed. Of course, if Alice were to be paid if Bob performs his obligation, then the incentives

would be aligned. More generally, a party should receive a positive payment if state 1 arises rather than state 2, if that party has positive evidence of state 1 or other parties have positive evidence of state 2. In our case, if the smart contract is going to trigger a payment from Alice, then either Alice needs to have a document that proves non-performance or Bob needs to have a document that proves performance. Without either of these, there is no hard evidence that can support the smart contract envisaged here.

Note that this places a lot of weight on the technical environment that exists. Not only must hard evidence of obligation performance exist, that evidence must be in the hands of the ‘right’ agent. Moreover, that agent is a different agent depending upon whether the obligation is been performed or not. In other words, for a smart contract to work and reduce the costs associated with verification, it must be the case that the costs of verifying the right evidence at the right time by the right person must be reduced. For purely digital documents this may be a challenge as digital documents may be easier to fabricate than other documents — that is, technically, a digital document might be able to be produced by both performance and non-performance in which case it is cheap talk rather than hard evidence.

The ability to generate hard evidence is the fundamental smart contract challenge. To the extent that contract obligations rely on evidence to be provided outside of a purely digital realm where evidence may be hard coded, any smart contract needs to create incentives for disclosure of the truth of contract performance. The remainder of this paper will examine candidates for this that might arise from Blockchain technologies including security that allows sensors to be embedded in smart contracts, machine learning predictions as a substitute for hard evidence, and mechanisms to elicit truth telling.

VI Working through the Fine Print

Given this discussion, I now turn to work through the fine print of a particular contractual situation that might be assisted with the support of smart contracts. It has the quality of being difficult to achieve as well as having human elements that underpin that difficulty. In doing this, I will show how innovations in economics — particularly, mechanism design — might enable the performance of smart contracts.

A Simple International Trade Example

To ground this discussion of smart contracting and what it can do, consider the following, very simple, example of international trade. A buyer (B) in one country wants to purchase a product from a seller (S) in another country. We assume that the buyer has a value (V) on the product while the seller has costs (C) of producing the product and shipping it to the buyer. They agree on a price (P) for the buyer to pay the seller in return for shipping the product. This price is less than the buyer's value and higher than the seller's cost; that is $V > P > C$. Thus, both the buyer and seller are better off from the deal.

If their contract was easily enforceable, all would be well. However, let's assume that this isn't the case. For instance, suppose that if the seller ships the product to the buyer but the buyer does not pay, the seller ends up with $-C$ rather than $P - C$ while the buyer gets V rather than $V - P$. Thus, if the buyer does not have to pay before the product is delivered, they have an incentive to not pay. On the other hand, suppose that payment is required upfront. That mitigates the seller's risk but there may be a new risk for the buyer. For instance, suppose the seller can ship an inferior product (worth $v < V$ to the buyer) for a smaller cost ($c < C$). In this case, if the buyer has already paid, the seller earns $P - c$ rather than $P - C$ by shipping the inferior product. This results in the buyer getting $v - P$ rather than $V - P$. If $v < P$, the buyer would not want to take that risk. They would return rather than keep the inferior product for that price.

What this demonstrates is that regardless of the order in which obligations (payment or shipping as the case may be) are performed, one side finds entering the arrangement too risky. Thus, absent some other mechanism, valuable international trade will not take place.

It is useful to reflect on why this international trade example results in transactional failure whereas many similar transactions take place within economies every day. The key assumption here is that while the performance of the contract is readily observable to both parties (i.e., m is low or 0), because of a high verification cost (z), it is not possible to rely on a court in either location or somewhere in between to enforce the contract. In other words, transactional failure arises because of **costly verification**. To be sure, social trust could substitute for this lack of verification. While the absence of trust is assumed here, it is also a realistic assumption given

that this is a ‘one shot’ transaction where neither party has information regarding the past history of the other.⁶

The Simplest Smart Contract

If there is one thing that the blockchain has proven that it can do is to hold digitised tokens in a ledger and transfer them between different accounts/owners with the permission of the token-holder. In other words, it is a form of purely digitised payments. This is not to say that such payments do not exist elsewhere; it is just that they are the blockchain’s proven application.

One of the things that proponents of the blockchain realized early on was that payment movements could be easily programmable. It is this aspect that caused those proponents to see blockchain technologies as particularly suited to smart contracts. The ‘smart’ pertained to an agreement that would be executed on the recording of other events on the blockchain.

The simplest smart contract is capable of solving one type of risk in our international trade example: the risk of buyer non-payment. Without a smart contract, the only way a seller can mitigate that risk is asking for payment upfront. However, for a buyer that creates the risk that the seller might not ship the item at all. A smart contract can resolve this issue and it does this by way of an escrow arrangement.

Suppose that the contract asks the buyer to place tokens in a holding (or escrow) account to cover P . The seller cannot immediately access those tokens. Instead, the contract specifies an event or state — recorded on the blockchain — that will trigger the transfer of those tokens to the seller. If that state does not occur within a certain period of time, the tokens revert back to the buyer. In this way, the escrow arrangement can be automated.

The key to this happening is for the state to be capable to blockchain verification. One form of an event may be that the buyer simply acknowledges receipt of the goods. The idea here is that the buyer would not do so if they had not, in fact, received the product. Another event could be for the seller to send evidence that the product had been shipped and received to a third party who was trusted to review these and that third party, upon confirmation, would trigger the event.

⁶ Interestingly, this situation has been at the heart of issues to do with trading on platforms such as eBay. eBay’s solution was to improve procedures that potentially lowered z but z remained positive nonetheless. (See Schmitz and Rule, 2017).

The problem here is that not much is being saved beyond a normal escrow service where the third party did the entire job. The idea of a smart contract is that such processes would be automated.

There is a way such automation could happen. Shipping companies today send shippers receipts that record various events. One could imagine a smart contract plugging directly into the information posted by shippers. Once the product has been acknowledged to them as received, then the event would be triggered and the payment made. It would save the buyer having to take the additional step of acknowledging receipt and it would not require a separate step of having a third party verify this information.

To be sure, these services exist and it is not our task here to work out if the blockchain saves on the broad costs associated with escrow services. It just should be noted that the blockchain could enable that service and it could operate entirely on the blockchain given the appropriate connections with other services — notably shipping company systems.

The Sensor Solution

While a simple smart contract might resolve one risk — buyer nonpayment — the other risk — seller sending a poor quality product — remains. At a first pass, the verification of product quality lies outside of the digital realm and so cannot be automated.

Nonetheless, there are situations where hardware in the form of sensors can measure the quality of certain products. If this hardware can be directly and securely linked to the blockchain, they could form part of the event that would trigger a payment in a smart contract.

The challenge here is that such technologies require a careful re-working of the supply chain. At the moment, this is the purview of larger companies that design supply chains for quality assurance. It is far from clear that there exist areas of international trade where sensor verification could open up what is, in effect, new and distinct supply chains.

Thus, while sensors may improve supply chains, it is not clear they have a distinct role in verification. It may be that because the Blockchain is regarded as more secure, sensor information may be more readily accessible for use in contract performance verification. Nonetheless, there are simply limits as to how far into the real world digitization can progress.

While sensors may be incorporated into smart contracts, they are far from a general purpose solution.

The Machine Learning Solution

One possible way a smart contract could open up new possibilities is if it executes contractual performance based on the outcomes of machine learning algorithms. Such algorithms would take data from sensors and other digital evidence and then be able to predict the likelihood that performance took place based on that data. In this way, if the likelihood of performance was sufficiently high, then contract could automatically execute.

In this way, what artificial intelligence methods combined with a Blockchain housing a smart contract could do is mimic a Court's procedure in verifying contract performance. In other words, the Court itself could be automated and the potential dispute process could be so cheap that it could just run on each contractual obligation rather than being triggered by one of the parties.

The Mechanism Design Solution

We have seen that the key contracting challenge is not the verification of payments — the blockchain and other solutions help there. In addition, when all the relevant parts of a contract can be digitized, verification using the blockchain becomes possible. The question we now turn to ask is: can the blockchain help with verification that is not digitized?

Let us return to the international trade example. Recall, if payments are capable of verification, the primary risk is the buyer receives a product that is inferior to what was promised; that is, it has a value to them, v , that is less than the expected v . Moreover, this poses a risk to trade if $P > v$ in which the buyer would not want to take the risk. Finally, the buyer could hold up payment to the seller if the product is of low quality, however, the buyer could also do that if the product is as expected. Thus, not only do buyers face risks of receiving inferior products, sellers who supply the right product face risk of payment not being made.

This is a problem for which there are no easy solutions. For instance, we have already noted that one potentially easy verifiable process that can be integrated into smart contracts is the

shipment and receipt of goods. This would work to ensure the buyer received the product. However, it could also work in terms of ensuring that the seller received a returned product. In this case, suppose that a buyer received a product that was of low quality. The buyer could simply return it and once that return was received by the seller, the funds in escrow would be paid back to the buyer.

However, there are risks here as well. First of all, the seller, who knowingly sent a poorer quality product the buyer, may try to get the buyer to accept a lower price rather than a return. If the buyer has their own costs in trying to return a product, they might accept this. But, taken as a whole, the buyer has not got the product they wanted and, instead, has ended up with something that, while they were willing to pay for it rather than return it, is ultimately worse off than expected. This is the sort of thing that causes people to avoid entering into such transactions in the first place.

Second, one way out of this would be for the seller to commit to compensating the buyer should the buyer want to return the product. This is something that smart contracts could help facilitate by requiring the seller to place tokens in escrow; being paid only if it is required. The seller would, thus, have incentives to send the expected product to the buyer. But even here there are risks. There may be unscrupulous buyers who face low costs of trying to return a product, who may take advantage of sellers in order to receive the compensation for returns. In other words, a seller who ships a good quality product may find themselves having to pay for a return regardless.

All this means that a simple — “return it if you do not like it and we will cover the cost” — contract, will not overcome the challenges in ensuring the product of the expected quality is shipped and then accepted by the buyer.

To resolve this, we want to find a mechanism that gives sellers the incentives to send the right product and buyers the incentive to only complain and withhold payment if they receive the wrong product. The good news is that both the buyer and seller actually know the truth — whether the good or bad product has been shipped and received. The bad news is that, by assumption, no one else knows this information.

There is, however, an economic mechanism that can work to resolve this. It is based on the mechanisms studied by Moore and Repullo (1988) and Maskin and Tirole (1999). In particular, it is a “simple sequential mechanism” (Moore, 1992). What these mechanisms involve is a set of rewards/punishments to each party based on the ‘messages’ or claims they make regarding what is going on. Ideally, the rewards are designed so that if one party tells the truth, they earn more than if they lie.

Before introducing the mechanism, it is worth noting that I am not the first to propose this nor to highlight the notion that blockchain technologies might enable it. Holden and Malani (2018). They examine the ‘hold-up’ problem whereby choices that cannot be contracted upon (that is, verified by a third party) could be resolved by a simple sequential mechanism. They are not thinking about the more pedestrian international trade type issues but, instead, the broader issue of incentives for cost reduction and quality in higher stakes business arrangements. They argue that what has limited such mechanisms is the possibility that parties to a contract may renegotiate outcomes. They then explore whether a blockchain could commit the parties not to engage in such renegotiations that can undermine mechanisms by denying them commitment power to shape incentives. It is not clear whether a blockchain could engage in a commitment against renegotiation when the stakes are high (so high that parties would be happy to bear the costs of court adjudication). However, I should note that this is not what is at the heart of my discussion here. In particular, as the costs of verification are otherwise prohibitively high in the international trade example this also implies that renegotiation is infeasible. Instead, here I rely on the feature of smart contracting to make payments triggered by states and obligations.

Here is a mechanism that can achieve a truth telling outcome in the international trade example. It focusses on the buyer’s incentive to report the truth.

1. (Announcement) The buyer announces the quality of the good (V or v)
 - If the quality announced is V , the contract is completed and payment is made.
 - If the quality announced is v , the seller can challenge the claim.
 - If there is no challenge, no payment is made to the seller and the product is returned with the buyer receiving compensation of f .
 - If there is a challenge, we move the ‘Challenge Stage’

2. (Challenge) The buyer is immediately fined, F , to be paid to the seller
 - The seller offers the buyer the choice between:
 1. Keeping the product for a price of p ($< P$)
 2. Returning the product for compensation of f
 - If the buyer chooses option 1, the new price is struck and the contract concludes with relevant payments being made.
 - If the buyer chooses option 2, the product is returned, with the seller paying the buyer f and the seller paying a fine of $2F$ to a third party.

The resulting game tree is as depicted in Figure One. This mechanism has a unique equilibrium outcome whereby the seller always sends the buyer the good quality and the buyer always accepts this.

Proposition 2. *Consider a contract where, if there is no seller challenge, the buyer can return the product and receive compensation of f but otherwise must pay P . If the seller challenges, the outcomes are as specified above. Suppose that (i) $V > p + f > v$; (ii) $P < p + F$ and (iii) $F > f$, then international trade in the good quality product takes place.*

Proof: Suppose the quality of the good is V but B claims v . Suppose S initiates the challenge. B will face a choice between $V - p$ and returning the product for f . So long as $V - p > f$ or $V > p + f$, B will keep the product. In this case, S ultimately receives a payment of $p + F$ while B ultimately pays $p + F$. It is worthwhile for S to initiate the challenge as they receive F .

Suppose the quality of the good is v and B claims this. If S initiates the challenge, B will face a choice between $v - p$ and f . B will choose to return the product if $v - p < f$. In this case, S will be fined and so ultimately receive $f - F$ while B will ultimately pay $F - f$.

First, note that if $V > p + f > v$, B will have an incentive to claim the truth about product quality (paying P for a good quality product as $V - P > f$ by (i) and asking for a return otherwise) prior to any challenge and to match any challenge from S when the product quality is v . Given this, if the product quality is v , S does not have an incentive to launch a challenge.

Second, as noted above, S has an incentive to initiate a challenge if the product quality is V (and is claimed otherwise by B) while B , because of this, does not have an incentive to claim the product quality is not V prior to the challenge.

Finally, S earns $P - C$ by shipping the good quality product and, at most, $-f$, by shipping the bad quality product.

Some remarks are in order. First, there is considerable leeway here in the choice of fines (F), compensation (f) and prices (p). Indeed, f might be zero or even negative while F could be very

small. The key is that both parties want to avoid a challenge being triggered as this results in a loss to each of them.

Second, there may be concern that the challenge mechanism would not work as, once triggered, there is an incentive to limit the damage by agreeing to an alternative arrangement. Indeed, this may even happen by the seller threatening to challenge. While this might be a concern generically, in this example, we have already assumed a situation where, absent a mechanism and set of escrow payments, it is hard for the seller and buyer to contract. How then would they be able to enforce a renegotiation between them? It would seem that the incentives to adhere to the contract would be stronger than seeing if you can fix things up later on.

Third, there is a practical concern that both parties may not want to enter into an arrangement where they could end up worse off if the other party does something stupid or makes a mistake. This is the argument of Aghion et.al. (2012). Indeed, Aghion et.al. (2017) test this mechanism in the lab. While they demonstrate that people may fear losses due to mistakes, the variant of their experiment that is closest to the situation modelled here where the buyer who receives a high quality product does not face a challenge, did perform well in their experiments. Beyond the lab, only real world adoption would constitute a practical test.

What is interesting about this mechanism is that because it requires escrow funds of $P + F$ from the buyer and F from the seller, a blockchain with a smart contract is perhaps uniquely able to provide the confidence that payments will flow as intended. The contract makes all factors transparent and automates all actions. In other words, it provides a substitute for a court verification process that might otherwise be too costly.

Some practicalities

It is reasonable to wonder how such a mechanism might look in practice. For instance, what is the magnitude of the escrow required of buyers and sellers as part of a transaction? To explore this, let's suppose the following: $V = \$30$, $v = \$0$, $P = \$15$, $C = \$10$ and $c = \$5$. In this case, for Proposition 2 to apply $p + f$ must be less than \$30 while $p + F$ must be greater than \$15. So let's suppose that $p = \$10$ while $F = \$6$ and $f = \$6$. Under this deal, if the expected product is shipped, then the buyer (in a challenge round) would keep it as this would get them \$14 in

surplus while returning it would get them \$0. By contrast, if the poor product is shipped, then the buyer will choose to return it than keep it as it is worth nothing to them and not worth paying the discounted \$10. Notice that any positive discounted price would lead to this outcome here. Given this, the seller would accept the challenge if the low quality product was shipped but otherwise offer the buyer a discounted price. Overall, then the buyer who receives their expected product does not challenge (as this would raise their effective price by \$1) and one who does receive a low quality product challenges as they can get a return refunding the potential fine they would otherwise pay.

Seen in this light, the challenge mechanism here is a just a variant on a ‘no questions asked’ return policy. The buyer can always return the item but has to (a) put up \$6 in escrow if they do this and (b) agree to consider a counter-offer from the seller. Therefore, unless the buyer wants to return the item they do not need to put more than the original price of the product in escrow. For the seller, however, to assure the buyer, they would need to keep the original price in escrow plus the difference between that price and the effective challenge price (of $P + f = \$21$). With \$21 in escrow, if there is a challenge the seller can refund P and pay the buyer f for their trouble. Thus, the buyer need not pay any money beyond the agreed price upfront and the seller need only keep that price and put in an additional amount to cover the refund amount to assure the buyer.⁷

VII Future Directions

This paper has looked at the details — or fine print — of smart contracts. The goal was to understand where blockchain technologies — made up of a distributed ledger with a virtual machine and token layer — could enhance the possibilities for efficient contracting. That meant being explicit on where contracts work and do not work today without the assistance of such technologies. This led to a focus on security (as provided by the distributed ledger), clarity (as provided by virtual machines) and commitment (as provided by tokens).

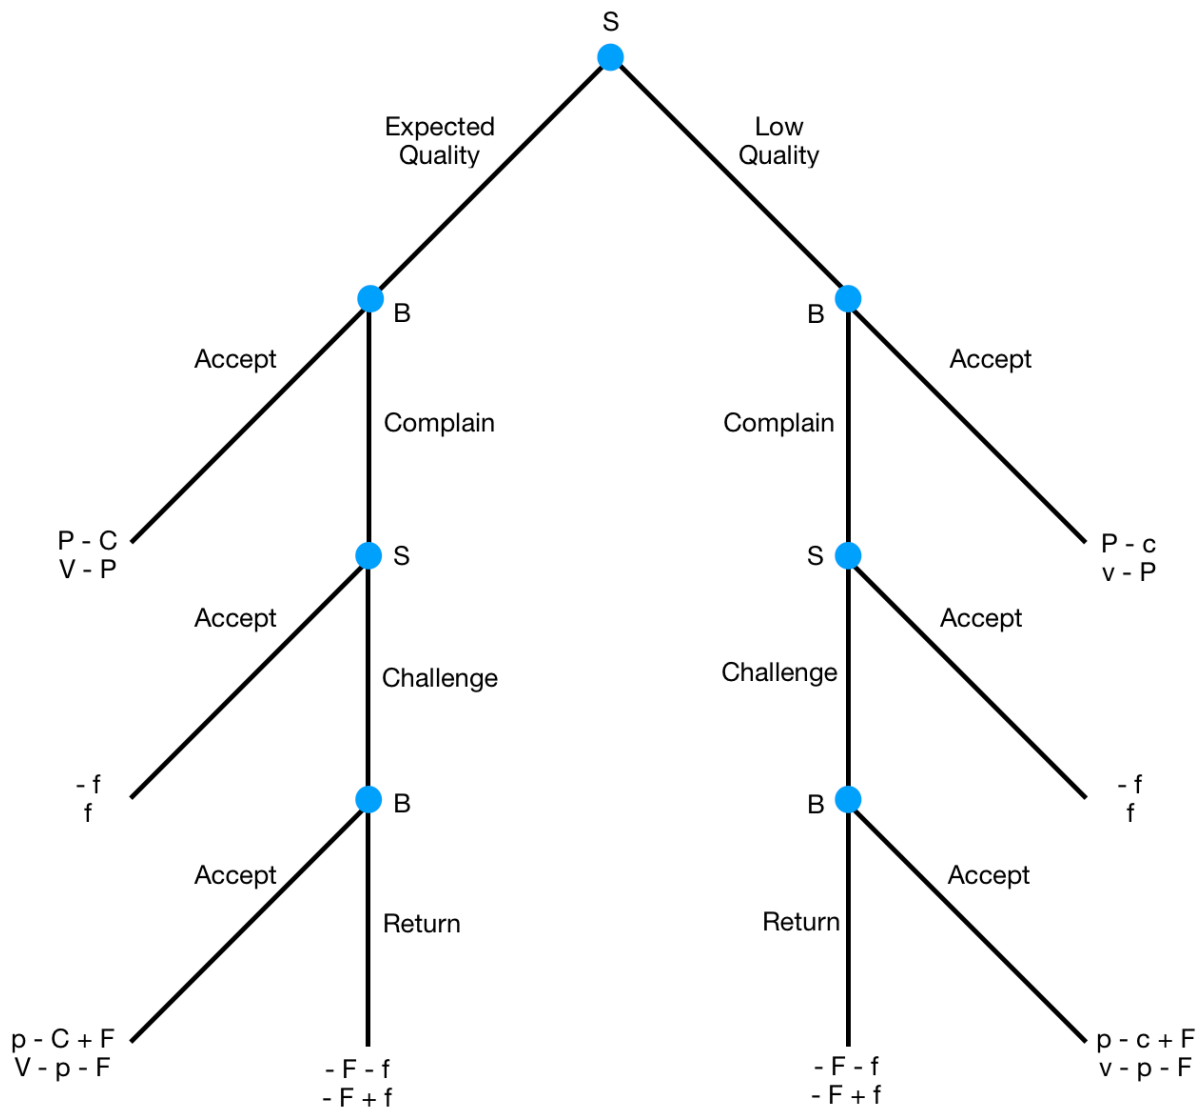
The good news is that blockchains can make the monitoring of states (including verified obligation performance) cheaper precisely because those states can be relied upon due to security

⁷ If the seller did not want even this risk, f could be set to be \$0. In this case, the buyer may be concerned about being out of pocket if a challenge needed to be mounted because this, by definition, means they pay a higher price in that eventuality. Rational economics alone cannot tell us whether these considerations would matter.

features. The bad news is that most contractual situations that struggle today, do so because of issues to do with the incentives for humans to perform obligations that cannot be easily measured — either digitally or by a third party such as a court.

Nonetheless, because so much can be included in smart contracts running off the blockchain, there is potential to innovate and deploy economic mechanisms that could create incentives for humans to perform and report on obligations in a truthful manner. In other words, the commitment engendered by the blockchain could substitute for a lack of trust in the real world. Whether these mechanisms exist, are practical, can be understood by real people and yield a return are still open questions. But it should be clear that the missing fine print in smart contracting is economic innovation in mechanisms and markets.

Figure One: Extensive Form Game



References

- Abadi, Joseph and Markus Brunnermeier (2018), “Blockchain Economics,” *mimeo.*, Princeton.
- Aghion, Philippe, Drew Fudenberg, Richard Holden, Takashi Kunimoto, and Olivier Tercieux (2012). “Subgame-Perfect Implementation Under Value Perturbations.” *Quarterly Journal of Economics*, 127, 1843–1881.
- Aghion, Philippe, Ernst Fehr, Richard Holden, and Tom Wilkening. “The Role of Bounded Rationality and Imperfect Information in Subgame Perfect Implementation—An Empirical Investigation.” *Journal of the European Economic Association* 16, no. 1 (2017): 232-274.
- Arrunada, Benito and Luis Garicano (2018), “Blockchain: The Birth of Decentralized Governance,” *mimeo.*, IESE.
- Bull, Jesse, and Joel Watson (2007), “Hard Evidence and Mechanism Design,” *Games and Economic Behavior* 58 (1): 75–93.
- Bull, Jesse, Joel Watson (2004), “Evidence Disclosure and Verifiability,” *Journal of Economic Theory*,
- Cao, Sean, Lin William Cong and Baozhong Yang (2018), “Auditing and Blockchains: Pricing, Misstatements and Regulation,” *mimeo.*, Chicago.
- Casey, Anthony J. and Anthony Niblett (2017), “Self-Driving Contracts,” *Journal of Corporation Law*, 42 (1), pp.101-131.
- Catalini, Christian and Joshua S. Gans (2016), “Some Simple Economics of the Blockchain,” *mimeo.*, MIT.
- Holden, Richard and Anup Malani (2018), “Can Blockchains Solve the Holdup Problem in Contracts,” *Working Paper*, No.2018-12, Becker-Friedman Institute, Chicago.
- Kocherlakota, Narayana R. (1998), “Money is Memory,” *Journal of Economic Theory*, 81 (2), pp.232-251.
- Maskin, Eric, and Jean Tirole (1999), “Unforeseen contingencies and incomplete contracts,” *Review of Economic Studies*, 66 (1): 83-114.
- Moore, John (1992), “Implementation, contracts, and renegotiation in environments with complete information,” *Advances in Economic Theory* 1: 182-281.
- Moore, John, and Rafael Repullo (1988), “Subgame perfect implementation,” *Econometrica*, 1191-1220.
- Schmitz, Amy J. and Colin Rule (2017), *The New Handshake: Online Dispute Resolution the Future of Consumer Protection*, American Bar Association: Chicago.
- Varian, Hal (2010), “Computer-Mediated Transactions,” *American Economic Review*, 100 (2), pp.1-10.
- Szabo, Nick (1994), “Smart Contracts,” (Retrieved from <http://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart.contracts.html>)