From "Blockchain hype" to a real business case for Financial Markets

Massimo Morini[1]

massimomorini.research@gmail.com

Bocconi University and Banca IMI

21/03/2016

Introduction: Blockchain Hype vs Blockchain Seclusion?

There has been a lot of noise in the press about the great potential uses for financial markets of Bitcoin-related technology, that could be extracted from the Bitcoin world and applied to existing markets to increase efficiency dramatically. Later, there has been a lot of noise about the fact that there is no actual use but all boils down to a generic enthusiasm called Blockchain Hype, and Bitcoin is the only reality where such technology can be fruitfully used.

This paper shows that there are real business cases for improving financial markets based on the lesson learnt from cryptocurrencies, but, differently from what the hype-enthusiasts say, they are *not* application of a technology to the existing business model of financial markets. They are reforms of the business model itself. What needs to be exported from the world of cryptocurrencies are aspects of the market organization, *inspiration* for a different accounting and legal system, and some aspects of the technology. These can be a huge contribution towards more robust, efficient and stable markets, but the process cannot be immediate and effortless, and can only be achieved within a market-wide strategic view.

One crucial misunderstanding here is the idea that Blockchain Technology can be exported to financial markets *as they are* to make them *more efficient*. This is meaningless; Blockchain technology was created to change some trust-based business processes to make them less reliant on trust; without structural changes in this direction the best of Blockchain technology is lost and just the inefficiencies are left. This misunderstanding is the perfect partner of the idea that Blockchain technology cannot be used outside the Bitcoin world. This is equally meaningless; Bitcoin was created to attempt a level of independence from trust sufficient to allow players to be anonymous and do without any legal protection; other business solutions based on a level of trust intermediate between Bitcoin and traditional finance can use similar technology and yet be very different from Bitcoins. But we must ready to use the concept of trust in a totally different way, as a way to analyze the different parts of a business process and the reason for its current inefficiencies and risks.

In the next we develop these concepts first in a parallel analysis of cryptocurrencies and financial markets. Then we focus on a specific business case regarding the collateralization of financial derivatives, that we describe bottom-up including quantifiable benefits in reducing costs, capital and risk. It is an example where the use of cryptocurrency technology is not more important than the business ideas developed in the analysis of cryptocurrencies; yet it was unconceivable before examples of distributed ledgers, smart contracts and oracles were visible in marketplaces. In fact, it was first presented in Morini and Sams (2015), in an introduction of the Blockchain innovation for the derivatives world.

The misunderstanding about "trust"

Notice that the term "trust" is often used in the Bitcoin debate in a radical way, opposing a totally trustless anarchist model to a corporative model based on absolute trust. None of them really exists. Even Bitcoin has created peculiar elements of trust in new players like a stable group of core developers or miners; and financial markets have never been based on absolute trust in counterparties or central bodies. The

[1] Fruitful conversation with Robert Sams, Giacomo Zucco and Alex Lipton is gratefully acknowledged. The second part of this paper is just an extension of "Smart Derivatives can cure XVA headaches" by Massimo Morini and Robert Sams, Risk Magazine, August 2015. I also thank all those – too many to be mentioned by name - that asked me the questions that form the backbone of this paper. This work expresses the views of its author and does not represent the opinion of his employers, which are not responsible for any use which may be made of its contents.

radicalism of the debate has hidden the fact that different business models are associated to different levels of trust; trust can be hidden in many passages in the working of a market, and can be eliminated or reduced in some without disappearing from others. More than a generic term for ideological debate, *trust* can be used as a precise concept to understand the features of a business model, and how that model can be positively reformed; without forgetting that any removal of trust creates some form of disintermediation, of some institutions or of some functions within institutions, and in this way it requires changes to the business model, and often to the legal, regulatory and accounting framework.

An example of unnecessary element of trust is the reliance on the agreement between two counterparties about the exact representation of a deal without any automation enforcing this agreement, not even in critical cases. Many markets are still crippled by this feature. This can be addressed with elements of distributed automation similar to those seen in cryptocurrencies.

What are the problems of financial markets that we want to solve?

Financial market transactions are still based on the logic of “consensus-by-reconciliation”. Every player gives its own representation of a transaction in its own accounting systems (ledger) and its own IT systems. The only proof that this representation is correct is coincidence with the representation given by the counterparties. The confidence in the legal validity of the contract in all its aspects is crucially dependent on trust in this coincidence, which in practice needs to be verified more than once. This verification requires a number of steps, such as confirmation, affirmation, communication to central bodies, and other reconciliation passages along clearing and settlement.

This is an objective bottle-neck towards more efficient and reliable markets. Current reconciliation steps slow the process down even if the technology enables very fast communication. They also drive costs up. Furthermore, the need for this kind of reconciliation leaves open the risk of disagreement and litigation, making the process uncertain and increasing the capital requirements for members. It is a system intrinsically inefficient that has never been seriously reformed in decades, for lack of incentives and no visibility of a technological and organizational stack suitable for a change. Even if many bits of the fundamental technology to solve it were already available in the past decades, just think of the internet giving a really shared information platform, this had never been applied to changing the foundations of some transactions. Now there is visibility of a different business model in the cryptocurrency example, together with a full technology package enabling it.

What exactly are the solutions that may come out of the Bitcoin experience?

Bitcoin and the other experiments of crypto-currencies or crypto-transactions are based on a single accounting and reporting system, a *Distributed Ledger*. With a Distributed Ledger, the reconciliation bottleneck is avoided since there is a consensus algorithm that verifies transactions and gives to them a unique representation on the ledger, collapsing all reconciliation steps into a a single initial passage. Further reconciliation steps are unnecessary when there is a single authoritative deal representation for all the parties. It is this business model that makes transactions so fast for bitcoins, more generally than any specific piece of technology. This insight is useful for ledgers in financial markets too, even if financial players may need distributed ledgers different from the Blockchain, which is a peculiar implementation where all transactions are reported together, visible to all, and their time-order is defined through a sequence of blocks.

For advanced financial markets, distributed consensus can be applied also to a deal made up of many payments, like a derivative or a bond, through the concept of a Smart Contract, which is a piece of program code, in a given computer language, executing the transaction agreed at inception between the parties. This guarantees the enforcement of consensus, namely that the deal will follow the agreement taken at inception between the parties. Bitcoin have only basic smart contracts, but other cryptocurrencies like Ethereum have smart contracts written in a Turing-complete language, which means it can do everything that a normal computer does.

Notice that this is a further step towards a different and more advanced model of the market. Not only the accounting/reporting of the transaction moves from individual representation to an authoritative distributed representation, but also the contract stops being two pieces of papers to be implemented and represented in separate ways but becomes a unique manager of the transaction signed (cryptographically) by the interested parties. Financial contracts are already translated by parties into software running on IT systems; what was missing were working examples of a technology where the piece of code could become the contract itself and not one of the many representations of it given by the parties. When the unique smart contract signed by the parties manages directly the flow of the transaction, there is a further reduction of delays and risks of disagreements and misalignments.

But all these goals can be obtained just via a central database and computing grid on one server.

For many of the above goals, the answer is: of course. But a computer/database shared among all the players of a market is a centralized solution, with all the well-known limits of centralization. These limits are a central topic in the *state machine replication approach*: centralized systems are usually more efficient from a technological point of view, but they are not fault-tolerant. In abstract terms, this means that failure of the central server is failure of the entire system. In economic terms, this unpleasant fact has additional consequences. In case of centralization, there will be an administrator of the database/hardware, and this institution would bear a great operational risk, the risk of the entire network, thus demanding an equally great power on controlling and changing unilaterally the rules. Centralized solutions create monopolies that drive the business costs up because the monopolist has not the right incentives to contain them; additionally, in finance centralized solutions also generate a concentration of financial risk that drives up - correctly – both the regulatory burden and the amount of risk-management provisions such as collateral.

A centralized database also raises the likelihood of legal disputes; it would be easy to accuse the administrator of tampering with the ledger. Since the ledger must report the situation of everyone and yet belong to no-one, a distributed ledger appears a more natural solution. It avoids the need for a central body and also reduces the legal uncertainties. The ledger downloaded by one party *is* the official ledger as much as the version downloaded by someone else. They are all *replications* of the ledger, there is not one central database and many *duplications*, which is a situation providing ground for uncertainty, reconciliation delays, and legal disputes.

And what if the database was "fully replicated and distributed"?

The technology of Distributed Services (and the state machine replication approach) that developed in the last decades are certainly a crucial part of the solution. There exists database technologies that try to keep away from the risks of centralization and date back to much earlier than Distributed Ledger technology. One can find works on fully replicated distributed databases that date back to as early as the early 90's, like [1]. The evolution of the technology has brought to popular distributed solutions like DVCS (Distributed Version Control Systems), of which Linus Torvalds' Git is the best-known.

The Bitcoin Blockchain evolved in the same stream of technological advances, partially based on the same cryptographic solutions. It is a relevant example of radical economic application of this form of technology, and in this way it showed how this technology applied logically to a market brings about a fundamental change in market organization.

Bitcoin found a decentralized solution for chronological tracking and time-stamping that was suitable for its peculiar context of building a market from scratch based on pseudonymous players. Even if this solution cannot be exported rigidly to different contexts like current and foreseeable financial markets, Blockchain is the natural turning point of distributed technology to take inspiration from when building Distributed Ledgers for financial markets, without ideological distinctions between distributed ledgers with blocks and proof-of-work, and distributed ledgers that may be different in these respects. An additional reason to keep more than an eye on Blockchain in evolving existing financial markets is to keep a standard compatible with other Distributed Ledger solutions that have different privacy and validation requirements, cryptocurrencies included.

Would the current technology for Distributed Ledgers be ready to provide this?

No. First of all, there is a scalability issue. The logic of distributed consensus across the entire network limits the amount of transactions that can be managed in a block. Solutions can exist for financial markets, but they are not tested yet.

Furthermore, the most tested market, Bitcoin, has got only basic smart contracts. Large-scale application of smart contracts is exactly the test that distributed ledgers for financial markets need to perform.

Finally, neither Bitcoin nor other solutions like Ethereum have a focus on privacy and identity as needed for financial markets. Identity is an unavoidable issue for any legal recognition; privacy is a concept which is evolving in financial markets, with regulators demanding more and more transparency, and may find various solutions: complex data-encryption, interlinked bilateral ledgers or regulated exploiting of pseudonymity. But in any case these are all elements showing that the process will take time.

So far only the ledger, the Blockchain, is used from the Bitcoin stack. What else is useful?

In Bitcoin there is also a fundamental use of cryptographic techniques such as asymmetric cryptography and hashing, both for ledger management and inside the incentive/selection method called Proof-of-work.

Asymmetric public-private key cryptography is important also for extension to financial markets; as it is already in many fields. This form of cryptography can be used to eliminate a level of intermediation; for example Bitcoin use it to disintermediate the role of banks as providers of cash deposits. In financial markets the main players, including banks, have a different role as structurers, traders, issuers of deals and securities, lenders and managers of credit and market risks. There is less fear of cryptographic disintermediation here, since the layers that can be eliminated or disintermediated in financial markets with no loss of security and a gain of efficiency and transparency are mostly not banks nor their business counterparties.

Furthermore, cryptography may enable at the same time identity and privacy. Other applications of cryptography are emerging now. For example the use of cryptography made here [3] is interesting: cryptography is applied in order to provide a cryptographic guarantee that an operation has been executed. It is a way to enforce a contract with a computer or website being guaranteed the contract has been executed exactly; that’s another bit of technology, related to the concept of Oracles, that developed around Blockchain even if it is not part of it. This can be used for example to secure the process of importing data from outside the Distributed Legder for internal use, something very common in financial markets; excessive reliance on trust here would create a single point of failure outside the control of those that have a stake on the ledger. A similar logic is behind also [4], with the additional feature that the logic is embedded in the hardware itself and not only in the software.

Is Proof-of-work to be exported to financial markets?

Proof-of-work as we see it in Bitcoins may not be applicable to financial market because it is designed to solve a specific problem: finding a way to make players update the blockchain in a honest way even if not forced to it neither by a reputation incentive (because they can be anonymous) neither by any legal framework. This is a very extreme concept of disintermediation and lack of trust that does not apply in a context where players are not anonymous and where fraud is legally prosecuted. This is a first reason to expect proof-of-work not to be used in financial markets for a very long time: the core motivation for its use is missing.

But there is something more, that requires us to get into some more details about proof-of-work. The clever idea behind the mechanism is to require the players that want to get the high remuneration associated to updating the Blockchain (miners) to solve a complex computational problem. This forces them to make an off-ledger investment in energy and computational power that makes it antieconomical to fraud the system. In fact, double spending is the only fraud that miners could implement easily in Bitcoins, since asymmetric cryptography and the public ledger protects, in its own peculiar way, past transactions and possessions. The loss of credibility of the network coming from a fraud would be, for those who have made the off-ledger investment in energy and computational power, a loss higher than the easy gain from double-spending. It is important here to understand a practical point not enough stated in theoretical analyses: that the investment in computational power is dominant over the investment in energy, and that the former is more relevant also because it is a *long-lasting* one: mining technology is very expensive and difficult to reuse for other purposes. This is crucial to understand why it works in making frauds antieconomical, and also why alternatives like proof-of-stake did not work: they did not guarantee an *off-chain*, *long-lasting*, *capital* investment.

Now we go back to current financial markets. While proof-of-work is not a waste of resources in Bitcoin since it is the only off-chain long-lasting investment of the crucial players, in financial markets it would be a real waste since the existence of off-chain economic commitment for crucial players is already proved; they have already a strong incentive to maintain the credibility of the whole financial system. This state of things may not last for ever; but it is the reality we start from.

So, should transaction visibility and validation be left to the counterparties only?

In principle, a basic extension of the current reality is a consensus algorithm where just the two parties involved sign the smart contract and validate the transaction, potentially on a private distributed ledger. This is already an improvement in terms of efficiency and finality of financial markets, removing much of the need for reconciliation and the risk of litigation. This is sufficient, for example, for the practical business case described in the second part.

Yet, it would be shortsighted to depart from the cryptocurrency experience so much to use a solution bound to be bilateral. There can be many services benefiting from *multilateral* reliable and efficient distributed transaction validation and recording. For example, in some extensions of the business case described next collateral may be provided or guaranteed by a third party; in this case consensus, speed and transparency like those allowed by a multiplayer ledger are particularly useful. Other examples are the use of the techniques for compression of exposures that are possible on a network, see for example [2], or the possibility to give regulators a broader and deeper vision of the market.

This can lead to a range of possible consensus algorithms, not excluding something more similar to what we have seen for cryptocurrencies. In fact, there is something more to say. The validation algorithm used in the Bitcoin world is mainly required to avoid double spending. This possibility seems of marginal importance in financial markets with trusted, or at least known, members. This is not true; just the same economic problem takes a different appearance. In fact double spending, or spending of non-existing resources, happens in financial markets too and in fact it is considered the main risk by players and regulators: just, it is named *default*.

Default risk is where we see clearly that trust is not unlimited in financial markets. Financial markets are made up of trustable parties, but no party is completely trustable because even the largest parties can default, as history shows. And default is by definition a form of double spending: a player has promised payments for an amount higher than the funds he had actually available, as the history of defaults like Enron and Parmalat, and partially Lehman, shows clearly. Thus thinking of methods for assessing fund availability within the network beyond pure unilateral confirmation is relevant, and this may introduce the case for more advanced validation, involving regulators or custodians or some other players not directly related to the transaction.

But a similar business case is still out-of-sight. Still difficult, but nearer, are solutions to some credit-related issues that can be solved using some aspects of the Distributed Ledger model that we discussed above, and that are robust to different choices about consensus and ledger visibility.

Ok. Can we move from general problems to a specific one?

It is about time. Saying in general that a business reform eliminates reconciliation or makes settlement faster is not enough. There can be cases in which reconciliation steps are a real business need or faster settlement is prevented by regulators for good reasons. One needs also to show business cases where this sort of worries are overweighted by the risk and cost savings coming from less reconciliation and faster settlement.

As seen in [5] and later in [9], a relevant case regards collateral and default management in the derivatives market, a market as big as around 7-8 times the world GDP in terms of notional; in terms of value it is as large as the US GDP or the global bond market.

Credit risk is an issue particularly for *Over-The-Counter* (not *listed*) derivatives, that are the dominant part of the market. The issue reached dramatic levels after the financial crisis started in 2008. The Lehman's default marked a crucial change in the derivatives markets. From an aggressive market with high leverage, little attention to risk, and a disordered multiplication of complex payoffs, we moved to a market of strong standardization, heavy regulations, and overly attention to risks. This has made the financial world a safer place under many points of view; but some negative side-effects are now clear. First, derivatives users like funds and corporates are increasingly unhappy with a market in which prices do not express the intrinsic market risk of a financial product (interest rate risk, commodity risk etc.), but are skewed by charges that are all more or less related to default risk.

This includes Credit Valuation Adjustment, or CVA, a valuation adjustment made by financial dealers for the risk of default of banks' counterparties, an adjustment called FVA (Funding Value Adjustment) for the cost of funding of banks, that increased when the banks' default risk increased, and KVA (Capital Value Adjustment), an adjustment for the amount of extra-capital that banks have to hold to contrast this increased default risk. Additional costs to users of derivatives as financial services come from the recent increase of the margin requirements for market players (these are part of funding costs and generate another value adjustment, the MVA or Margin Value Adjustment), which also is a response to increased default risk.

Buy-side clients still need financial markets and derivatives for their investment and diversification needs, and to hedge their costs and risks, in terms of cashflows and under an accounting point of view. For these clients, the above transformation meant first of all a sharp cost increase.

In the current market situation, what is used to reduce credit risk? Collateral?

The mainstream approach to reduce the size of these charges is to reduce the losses in case of default, of banks or their counterparties, through collateralisation.

Collateral for derivatives is of two kinds. First, we have Variation Margin. The derivative is revaluated *every day* by party A using its pricing model f^A that takes in input the current value M_t^A of the relevant market variables from the info provider chosen by A, and gives current derivative value

$$V_t^A = f^A(M_t^A)$$

If V_t^A is negative to A, which means that A expects to pay to party B in the transaction more than what A will receive, so that A is a net debtor, A will ensure that an amount of cash (less often an amount of other assets, bonds or equities, with a *haircut* rule) equal to V_t^A is available for the counterparty B in a collateral account. Party B does the same thing but with its model f^B and its data M_t^B. Hopefully

$$V_t^A \approx -V_t^B$$

and the process proceeds smoothly. When there is a remarkable difference between V_t^A and $-V_t^B$, the two counterparties talk to each other for a reconciliation. In some cases it is the net creditor that makes a margin call, but this does not change the general picture.

In many cases, particularly when a party is a non-financial corporate that has difficulties to move cash quickly or to compute quickly the right amount of collateral, this process of collateral update happens less often than daily. It may be that longer period is stated in the agreements, or that the agreements accept explicitly to leave a part of the exposure not collateralized (via defining *thresholds* or a *minimum transfer amount*). These are inferior Variation Margin agreements that contrast with with the top-class agreements between banks, characterized by daily updates, no minimum transfer amounts and zero threshold.

Is this all the needed Margin? Or additional margin, such as the Initial Margin, can be useful?

Even in case of Variation Margin, there is always an expected delay between the last collateral update and the closeout for liquidation of a defaulting counterparty's positions, leaving risk of default still open. This delay is called Margin Period of Risk (MPOR), and comes from summing the collateral frequency with the delay between default time and the computation of a closeout amount. The total delay is estimated to be rather large by regulators since, when a default happens, there is no guarantee that the valuation of the residual derivatives, V_τ^A and $-V_\tau^B$, with τ being the default time, coincide for the two parties. The current process assume disagreement and potential litigation, and a reconciliation procedure driven by the liquidators that involve asking various third parties to give a valuation of the residual deal before arriving at a closeout amount. This pushes MPOR to range from 5 to 40 days.

Thus, on top of Variation Margin, there can be an additional amount of collateral called Initial Margin to cover the risks due to the length of the MPOR. In an Initial Margin agreement, Counterparties use their risk models to make a conservative estimate (worst case scenario or Value at Risk computation) of the difference between the amount of collateral available at the beginning of the MPOR (*last collateral update*) and the actual default closeout amount computed at the end of MPOR (*closeout day*). This computation needs to take into a account the impredictability of market movements along the MPOR and the uncertainties on how the closeout amount will be computed. Under a long MPOR, Initial Margin can be very high.

Is this solution fully satisfactory to market players?

It has some very relevant limits.

1. First of all, collateral management is not, in the current market, so easy for non-financial players. Computing, finding and moving the necessary margin liquidity can be an obstacle even just to agree on a top-class Valuation Margin procedure.
2. Secondly, even a top-class Variation Margin procedure is tampered by uncertainty on the different valuation models *and* market data *and* computations *and* accounting representations from the two parties, an uncertainty that can create misalignments and makes the process never faster than daily.
3. Third, the margin period of risk is very long: summing together collateral frequency and the period for the agreement on closeout, on average one considers 10 days. It is a delay sufficiently high for having still remarkable credit risk and capital cost (KVA) even in presence of VM.
4. Initial Margin on top of Variation Margin can reduce these costs dramatically, but only at the cost of incurring in a fourth problem: setting up a conservative initial margin agreement is a high cost in terms of funding costs. Initial Margin stays in a secluded account and due to its size, that in turn depends on the length of the MPOR, it drains a large amount of liquidity from institutions.

Any additional solution in place towards a reduction of credit risk? What about CCPs?

One solution is trading through central counterparties (CCPs), which can reduce credit risk through *trade compression*. Consider a situation where bank A owes 100 to bank B, bank B owes 100 to C, C owes 100 to A. If all the players trade through a central counterparty, the three above payments cancel out with each other, reducing settlement and credit risk.

CCPs do something even more important beyond compression: by pooling risks together, they reduce the size of potential losses through the *netting* effects. When a bank defaults, its obligations towards a counterparty are usually netted with those of the counterparty towards the bank. This reduces the closeout amount to be paid thus reducing potential losses. When there is a unique counterparty like the CCP, this netting effect is stronger.

CCPs, however, are an intrinsically centralized solution. Centralization has the advantages just mentioned, but also symmetric disadvantages like creating a central institution whose default, however unlikely, would spread losses to the entire market at an unprecedented speed and scale. See on this the literature by experienced and popular researchers like Jon Gregory or Darrell Duffie. This also means that the regulatory burden is particularly high on such institutions, increasing also collateral cost and demand, since a CCP is such a single point of failure that it needs to be massively overcollateralized.

We also have to remember that a centralized body lacks some of the competitive pressure to optimize collateral costs to members. Excessive collateral demand does worry regulators since it can strain the market's liquidity conditions.

Finally CCPs, as a natural corollary to this business, decide unilaterally the rules for variation and initial margin. The rules are also changed unilaterally quite often, particularly for Initial Margin.

Can we think of an alternative or complementary solution without the costs of centralization?

Here is where Distributed Ledgers come into play, but they can be useful only if we are eager to take from cryptocurrencies not only some technology but also inspiration on how the process can be designed, making a change that needs to be technological, regulatory, legal, and organizational.

On the technology side, smart contracts suitable for derivatives can be implemented within a distributed ledger system if the consensus algorithm contains what is known as a Turing-complete state-transition function – for example, it must support if-then-else-branching, enabling the conditional features of a derivative to be executed.

A smart contract transaction might, for example, instruct the network to transfer:

max (S_{1Y}– X, 0)

from account A to account B a year from now, where S_{1Y} is the price of a given security one year later, provided a certain sum – the value of the contract – is transferred from account B to account A of the distributed ledger now. This is a sketch of the implementation for a cash-settled call option.

Once knowledge of S_{1Y} is provided in real time to the smart contract through an oracle managing the access to trusted data providers, the contract can take care of the terminal settlement, transferring the right amount of money automatically. The smart contract can be much more detailed than the simple example provided above, incorporating more complex contractual features such as breakups, American exercise, legal requirements and International Swaps and Derivatives Association standards. And the Smart Contract can take care of Collateral regulation.

What are the real savings of smart derivative contracts on a distributed ledger?

The main savings are seen when we consider collateral. The smart contract can include the implementation of a model that computes the amount of collateral to exchange, as a subroutine or from an external source communicating with the network through a precise, cryptographically signed agreement with the contract itself (an Oracle in the extended sense of [3]).

After the above reasoning, there is no much to say about which changes are brought by a distributed ledger logic coupled with a smart contract that uses a single cryptographically secure implementation of a model.

All the uncertainties we have seen before are eliminated. There can be no differences due to the model or the data or the computation or the accounting rules: the agreement was taken not on a generic paper contract, but on a single smart contract managing the quantification of the payments through a single model implementation, and recording the exchanges on a single ledger.

Since on a distributed ledger precise rules for collateral payments have been agreed and validated from the start, and are then managed by a digital contract, the need for reconciliations and the risk of litigation are minimized here. This reduces credit risk in two ways. First of all, collateral can match exposures much more precisely than now. Second, slashing down the time required for reconciliation means that much faster collateral update becomes possible: collateral exchange frequency can be taken down from the 1 day delay of the best today's agreements to a fraction of a hour. This makes, collateral, eventually, a precisely, real-time guarantee. This eliminates problem 2) seen above.

Additionally, in an environment where transactions are naturally automated, and collateral is quantified and managed by the smart contract, also problem 1) can be reduced.

But this means counterparties should agree on a valuation model, moving from f^A, f^B to f

Extending the range of what is contractually agreed and validated at the beginning, reducing the scope for trust and future reconciliation, is a core point of this possible business evolution. It is the price to pay for efficiency, risk-reduction and cost saving. But this specific price may not be seen as so high in these days.

First of all, since banks are already accepting, and in some cases they are even seeking, more consensus about models. Before the crisis, private valuation models were regularly used for complex payoffs, and valuation differences could be seen as drivers of value as much as of risk. Today, the stress in regulations, margins and credit risk has changed the picture, making risk the dominant effect, and valuation differences have already been minimized in many contexts. In the post-crisis years regulations regarding CCPs have already led the market to accept external standardized valuations for margin purposes (initial and variation margin), and ISDA/Iosco have led the market to agree on a common model for part of the margin (the initial margin) even for non-cleared products. But this process towards sharing calculation logic is not only regulations-led, and goes beyond the margins issue: services like Markit Totem [6] are used by banks to reach indirectly a general consensus also on the pricing logic of complex, non-cleared products, and have gained importance in the last years.

Secondly, what would happen on a distributed ledger is much less restrictive than the model standardization banks are already accepting, since it is in principle a bilateral agreement between the same parties that have just agreed on one a price (valuation model at time zero) and on a future collateral exchange (valuation model in future times), not the acceptance of a one-size-fits-all model like in case of CCP or regulatory intervention.

Finally, we get the majority of the benefits even if the counterparties just agree on a valuation model for all those cases when *valuation becomes a payment in the contract*, like collateral regulation (and potentially anticipated closeout, a topic addressed below). If the ledger is used just as a transaction report but is not *the only accounting report*, players can be left free to keep private models for valuation in their own accounting, while binding them to smart contract agreement when valuation is used to quantify payments with the original counterparty. Misalignments between the private accounting valuation model and the agreed collateral valuation model are already existing in dealing with CCPS.

Can we do something more to reduce credit risk in collateralized derivatives?

Faster and precise collateral would already reduce risks and associated costs; but the full gain would be extending the agreement on the re-valuation model from collateral to default closeout. A mutually agreed valuation model can change completely the closeout process – reducing the MPOR to a few hours.

With collateral on a ledger, a missed collateral update (a default warning) is detected in real time, and can trigger the smart contract to terminate itself and provide immediately closeout valuation based on the agreed model. Suppose party B misses a collateral payment at t. We can agree on the smart contract that there is a *grace period* δ, say *δ =few hours*, during which the contract waits for a collateral payment from B. At the end, the difference between deal value and collateral from A point of view is computed with the agreed model

$$\Delta = f(M_{t+\delta}) - f(M_{t-\varepsilon})$$

where ε is the interval between two collateral payments, so that $t - \varepsilon$ is the time of the last collateral update. Party A becomes owner of the collateral, already in his availability, with only Δ be paid by A to B if $\Delta < 0$ to A; Party B is fred by its residual obligation, apart Δ to be paid by B to A if $\Delta > 0$ to A.

In this way, the closeout amount is promptly computed in the network using the agreed bilateral model $f(M_{t+\delta})$, and the Margin Period of Risk is slashed down to the short $\varepsilon + \delta$ time, that can realistically be few hours, an order of magnitude smaller than the current MPOR of few days. The discrepancy between the last collateral update and the closeout amount, Δ, will be as small as the change in few hours of a net present value, computed with a single model.

No more will derivative users have to endure litigation and lengthy procedures involving multiple third parties to arrive at a closeout amount, solving problem 3) and reducing risk and associated regulatory capital. The gap Δ between collateral and close-out amounts can be reduced to much smaller levels, and if we want to close even this risk, we can think of Initial Margin here too; it will have to cover Δ, and it will be much smaller than it is now, creating lees strain on the liquidity of financial players; which solves problem 4).

Could fast reaction to missed payments create more defaults for temporary liquidity problems?

Not necessarily, because on a ledger we can reduce the gap between collateral and close-out amounts to levels sufficiently small to allow us to exclude «on-chain» default: a missed collateral payment can be treated as an unwinding that generates a small balance to be settled in a longer term, when temporary problems, if they were really the issue, will surely be over. Let us see how this can be done.

It is reasonable to worry that a market where everything is faster or more automatic creates more technical defaults, due to problems like a temporary lack of digital cash. But the procedure above for the case of a missed collateral payment needs not to be considered a default in the usual legal sense, since we can *design it contractually*. We increase the risk of "technical" defaults only if we ask B to pay Δ immediately after the grace period. But since the payment is now determined by a precise contractual agreement and it has all the likelihood to be small, being based on a MPOR of few hours rather than 10 days, we can postpone this payment to a later time, to allow the counterparty to get the necessary liquidity. Default in the legal sense is thus driven out of the ledger; if this happens, it will be driven by external reasons, and will affect the network only for the pre-computed amount Δ.

The players will still be unhappy when a counterparty defaults and, for example, a hedge is lost. But at least now players have as soon as possible as much cash as possible to find a new counterparty for the same deal; the long waiting times and discrepancies are cut out by design.

Margin Period of Risk

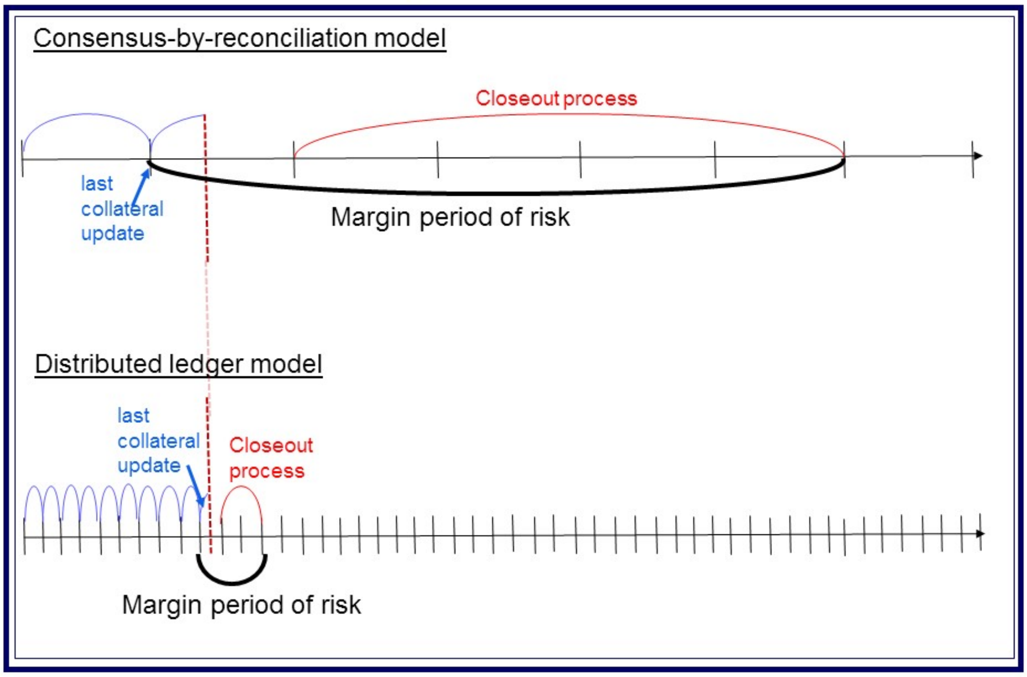

It seems many things have to change to allow this: new (smart) contracts, new (distributed) accounting…

A lot of things have been to be made consistent with such a framework, including the regulatory framework. We also need money with digital representation, may this be a digital currency fully convertible in central bank accounts, an independent crypto-currency, or just one or more currencies issued by financial institution redeemable with fiat currencies or other assets. The first choice is preferable, the other two choices have their own limitations, let's mention just two of them: too much market risk, or volatility risk, for current cryptocurrencies, see for example [7], too much credit risk, or default risk, for banks' money. The network needs to receive a number of inputs from outside: calendar changes, fixings, data for valuation, and potentially valuation from an external engine. The technology for communication between ledger and the external world is the technology of Oracles in the sense of [3]. Standard contract specifications, including ISDA standards, will have to be expressed as template code.

Distributed ledger technology is the natural way to get the cost and risk savings seen above for derivatives, not only because they requires faster clearing and settlement, but more importantly because they require first to move the market logic towards putting the on-ledger smart contract at the center of the transaction, as opposed to the current approach based on two different implementations and two different reports of a paper contract.

Legal and regulatory status could come earlier than expected if regulators see advantages in an architecture which is more transparent and creates less risk than most of the current solutions. This is why I think it useful that we continue with the analysis of advanced business cases: to show the possible advantages for financial markets, and to clarify the hard necessary passages, going once for all beyond the false dichotomy between "blockchain hype" and "blockchain seclusion".

References

[1] http://www.dtic.mil/dtic/tr/fulltext/u2/a272895.pdf

[2] http://www.trioptima.com/

[3] http://www.oraclize.it/

[4] https://software.intel.com/en-us/sgx

[5] http://www.risk.net/risk-magazine/opinion/2422606/-smart-derivatives-can-cure-xva-headaches

[6] http://www.markit.com/Product/Totem

[7] http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2458890

[9] http://www.assiomforex.it/riviste/newsletter