

Smart Contract Templates: foundations, design landscape and research directions

Christopher D. Clack
Centre for Blockchain Technologies
Department of Computer Science
University College London

Vikram A. Bakshi
Investment Bank CTO Office
Barclays

Lee Braine
Investment Bank CTO Office
Barclays

August 3, 2016

Abstract

In this position paper, we consider some foundational topics regarding smart contracts (such as terminology, automation, enforceability, and semantics) and define a smart contract as an agreement whose execution is both automatable and enforceable. We explore a simple semantic framework for smart contracts, covering both operational and non-operational aspects. We describe templates and agreements for legally-enforceable smart contracts, based on legal documents. Building upon the Ricardian Contract triple, we identify operational parameters in the legal documents and use these to connect legal agreements to standardised code. We also explore the design landscape, including increasing sophistication of parameters, increasing use of common standardised code, and long-term academic research. We conclude by identifying further work and sketching an initial set of requirements for a common language to support Smart Contract Templates.

1 Introduction

The aim of Smart Contract Templates [1] is to support the management of the complete lifecycle of “smart” legal contracts. This includes the creation of legal document templates by standards bodies and the subsequent use of those templates in the negotiation and agreement of contracts by counterparties. They also facilitate automated execution of the contract and, in the event of dispute, provide a direct link to the relevant legal documentation.

The templates and agreements may (or may not) be agnostic to the method by which a contract is executed – that is a design choice for the template issuer, counterparties, network, etc. The intention is to interface with a wide range of execution platforms. Smart legal contracts could potentially be executed as software agents operating on distributed ledgers (such as Corda [2], Ethereum [5], Hyperledger [12], etc.).

Here we aim to make a practical contribution of relevance to financial institutions. We consider how contracts are written, how they are executed, how they are enforced, and how to ensure that the execution of a contract is faithful to the meaning of the legal documentation. We discuss these issues using reasonably straightforward language, so that it is accessible not only to financial institutions but also to lawyers, regulators, standards bodies, and policy makers. We hope that the issues and views raised in this paper will stimulate debate and we look forward to receiving feedback.

Acknowledgements: We would like to thank Clive Ansell (ISDA), Ian Grigg (R3) and Darren Jones (Barclays) for their helpful feedback.

2 Foundations

In order to lay the foundation for subsequent discussion, there are several topics that require elaboration. Here we consider the key topics of terminology, automation, enforceability, and semantics.

2.1 Terminology — “smart contracts”

In [15], Stark gives an overview of the two different ways that the term “smart contract” is commonly used:

1. The first is entirely operational in nature, involving the execution of software agents, typically but not necessarily on a shared ledger. The word “contract” in this sense indicates that these software agents are executing certain obligations and may take control of certain assets within a shared ledger. There is no clear consensus on the definition of this use of the term “smart contract” — each definition is different in subtle ways [18, 17, 16, 3]. Stark renames these agents as *smart contract code*.
2. The second focuses on how legal contracts can be expressed and executed in software. This therefore encompasses operational aspects, issues relating to how legal contracts are written and how the legal prose should be interpreted. There are several ideas and projects which focus on these aspects such as the Ricardian Contract [7], CommonAccord [4] and Legalese [13]. Stark renames these as *smart legal contracts*.

Given that there is no clear consensus on the terminology being used, it is important that we should be clear in this paper. Here we prefer that the term “smart contract” should cover both versions, so we adopt a higher-level definition based on the two topics of automation and enforceability, that are explored in depth in sections 2.2 and 2.3:

A smart contract is an agreement whose execution is both automatable and enforceable. Automatable by computer, although some parts may require human input and control. Enforceable by either legal enforcement of rights and obligations or tamper-proof execution.

This is sufficiently abstract to cover both “smart legal contracts” (where the agreement is a legal agreement, which is then capable of automatic execution in software) and “smart contract code” (which may not necessarily be linked to a formal legal agreement, yet must be executed automatically). It simply states a requirement that the contract must be enforceable without specifying *what* is the aspect being enforced; for smart legal contracts these might be complex rights and obligations, whereas for smart contract code what is being enforced may simply be the execution of the code.

We focus on smart legal contracts, with the expectation that it will be possible for execution to occur using smart contract code. In addition to our definition of smart contract given above, throughout the rest of this paper we also for clarity adopt Stark’s terms *smart contract code* and *smart legal contract*.

2.2 Automation

We have chosen to say that a smart contract “is automatable” rather than that it “is automatically executed” because in practice there are parts of a legal agreement whose execution might not be automatic and will require human input and control. However, to be a “smart contract” we require that some part of the execution is capable of being automated (otherwise it is not “smart”).

Automation is generally taken to mean being executed by one or more computers. The phrase “by electronic means” is a synonym. Our definition of smart contracts does not require that this automatic execution occurs on a shared ledger, though that is certainly a possible and even probable method of execution.

As an example of how automation might be achieved using smart legal contracts, Grigg [9] presents the Ricardian Contract triple of “prose, parameters and code”.¹ The legal prose is linked via parameters (name-value pairs) to the smart contract code that provides execution. We might for example envisage that an executable software agent has been developed that will be instantiated on a shared ledger and, once execution has started, will proceed to undertake various transfers of value in accordance with the legal prose. The parameters are a succinct way to inform the code of the final operational details.

The code in this case would be suitable for execution on a specific platform but we can imagine in the future that multiple platforms could be targetted from a single contract.²

2.3 Enforceability

Given a smart contract must be “enforceable”, what are the elements that must be enforced? And how? First we consider *what* must be enforced:

2.3.1 What to enforce

What needs to be enforced is different for smart contract code and smart legal contracts:

- For *smart contract code*, the key requirement is that the code should execute successfully and accurately to completion, within a reasonable time. If the execution platform is in complete control of all of the actions that the smart contract code wishes to perform, then these actions should be executed faithfully and with reasonable performance. Things that can go wrong (and therefore require “enforcement”) would either be technical issues within the platform, or issues that take place outside of the execution platform — an obvious example would be the physical delivery of goods.
- For *smart legal contracts*, things can be considerably more complex. Typically a legal contract would have a large number of rights and obligations that accrue to the different parties to the agreement and are legally enforceable. These are often expressed in complex, context-sensitive, legal prose and may cover not just individual actions but also time-dependent and sequence-dependent sets of actions. There may also be overriding obligations on one or more of the parties such that a *lack* of action could be deemed to be a wrong-performance or non-performance.

¹https://en.wikipedia.org/wiki/Ricardian_Contract

²This could for example be achieved by using the list of parameters to connect the legal prose to a *set* of smart software agents, e.g. one agent per execution platform.

2.3.2 How to enforce

Enforcement might be achieved via traditional or non-traditional methods:

- **Traditional** means of enforcement include a variety of dispute resolution methods such as binding (or non-binding) arbitration, or recourse to the courts of law. There is an established body of law, and the methods by which parties can resolve disputes are well known. The traditional methods are backed by the power of government as embodied in the law, law-enforcement agencies and the courts. For illegal acts, courts are for example empowered (to different extents, according to jurisdiction) to impose fines, sequester assets, or deprive the wrong-doer of liberty. For disputes relating to contracts, the courts have extensive experience of adjudicating on issues of wrong-performance and non-performance, of awarding damages or other reliefs as appropriate, and in some cases assisting in the enforcement of payment of damages.
- **Non-traditional** methods of enforcement may also be imagined. For example, there is currently debate and experimentation on the possibility of enforcing the execution of smart contract code at a network level without the need for dispute resolution either via arbitration or via the courts. This is a fundamentally different notion of enforcement that is often expressed in terms of “tamper-proof” technology, with the assumption that in a perfect implementation of the system wrong-performance or non-performance become impossible.

“Tamper-proof” execution is typically described in terms of distributed networks of computers that are unstoppable and in a technological sense cannot fail regardless of malicious acts, power cuts, network disruption, natural catastrophies or any other conceivable event.³ With such a system, it is assumed that a software agent, once started, could not be stopped. For truly “unstoppable” software agents, code must be embodied to take the appropriate action in response to various dynamic states that might occur (such as another party not having sufficient funds to execute a required payment). In a normal system, the software agent might abort and the wrong-performance or non-performance by a party would be enforced by traditional means; but in a truly unstoppable “tamper-proof” version of the system, all such possibilities would have to be anticipated and appropriate actions determined in advance, so they are no longer deemed wrong-performance or non-performance but are instead anticipated states of the system with known resolution.

Although some groups are actively pursuing tamper-proof smart contract code, our preference is for smart legal contracts that are enforceable by traditional legal methods for reasons including:

- In a system with enforcement by tamper-proof network consensus, there would be no “execute override” provisions. Agreements, once launched as smart contract code, could not be varied. But it is quite common for provisions of an agreement to be varied dynamically — for example, to permit a favoured client to defer paying interest by a few days, or to permit a payment holiday, or to permit the rolling-up of interest over a period. Unless every possible variation is coded in advance, none of this would be possible in a tamper-proof system.

³A tamper-proof network might be used to run a “permissionless” shared system [10] — i.e. where anyone can access the system and trusting a single party is not required. Swanson [17] gives a good overview of many of the complex issues that arise with permissioned and permissionless distributed consensus systems.

- Enforcement by network consensus can only apply to the execution of obligations, or the exercising of rights, that are under the control of the network. However, objects and actions in the physical world are unlikely to be under full (if any) control of the network.
- Mainelli and Milne [14] observe that smart contract code “that involved payments would require posting collateral to be completely automated. This locking-up of collateral would lead to a serious reduction in leverage and pull liquidity out of markets. Markets might become more stable, but the significant reduction in leverage and consequent market decline would be strongly resisted by market participants.”

2.4 The semantics of contracts

Part of our remit is to consider the semantic construction of a contract — i.e. what is the “meaning” of a contract? Does it have more than one meaning? How should a contract be interpreted? We start with a simple semantic framework and view a legal contract as having two interpretations:⁴

1. The *operational semantics*: this is the operational interpretation of the contract, which derives from consideration of precise actions to be taken by the parties. Thus, this is concerned with *executing* the contract.⁵
2. The *denotational semantics*: this is the *non-operational* legal interpretation (or “meaning”) of the entire contract, including all of its obvious constituent parts and including any other legal documents that it references. This is the meaning that would be given to a contract when a lawyer *reads* the contract.

These two semantics do not consider different parts of the contract — they are both interpretations of the whole contract, but with different aims.⁶

A contract may comprise several documents, and the process by which these documents are agreed may be complex. The denotational semantics of even quite straightforward contracts can be very large and complex, yet by contrast the operational semantics might be simple and easily encoded for automatic execution.

The operational semantics dictate the successful execution of the contract to completion. If a dispute arises, then the denotational semantics of the contract typically dictate what happens next⁷ — i.e. in the context of the rights and obligations of the parties, the specification of what remedies shall be applied in the case of partial-performance or non-performance by one party.

The greater part of a legal contract may often be devoted to defining the obligations and liabilities of the parties in the event of a problem with execution. Sometimes, the actions to be taken in case of a material breach of contract are expressed precisely; however, this

⁴The consideration of denotational and operational semantics is well established in the academic discipline of the theory of programming languages. Here we use the terms “interpretation”, “meaning” and “semantics” as synonyms.

⁵Although a contract is executed by a computer program, there may be elements of physicality which are outside the control of the computer code e.g. physical delivery and manual confirmation. Thus, as an additional objective there may be a further level of semantics to define the interaction between the automated and the physical worlds.

⁶Other semantics may also be applied, e.g. to derive different forms of risk associated with a contract.

⁷Although some remedial actions could be specified in the operational semantics.

is not always the case and dispute resolution may require a protracted process of negotiated settlement, arbitration or court proceedings.

Furthermore, it is important to realise the important role of law. It is not possible to take literally the doctrine that all one needs to know about a contract is contained within “the four corners of the document” [6]. A lawyer would read and understand the contract *in the context of* the governing law — i.e. each legal document must be interpreted according to the relevant law (corporate law, consumer law, etc) of its stated or inferred jurisdiction, and therefore the semantics of that law must also be understood. It should be noted that the issue of law relates not only to the denotational semantics but also to the operational semantics — for example, trading with certain countries may be illegal due to government-imposed sanctions.

Given this semantic framework for the legal contracts that underpin financial instruments, we can derive a different perspective of smart contracts:

- smart contract code focuses exclusively on execution and therefore concerns itself only with the execution of those *operational semantics* that are expressed in the code, whereas
- smart legal contracts consider *both the denotational and operational semantics* of a legal contract, whose operational semantics must then be executed (possibly by smart contract code).

This idea was previously expressed in a slightly different way by Grigg [8], displayed as a chart where the y-axis was denoted the “Ricardian axis” of increasing semantic richness (i.e. increasingly capturing the denotational and operational semantics of the legal prose) and the x-axis was the “Smart axis” of increasing performance richness (primarily concerned with the *execution* of the operational semantics): see Figure 1. Both are important, yet they are orthogonal issues and with appropriate interfacing developments can potentially proceed along both axes simultaneously.

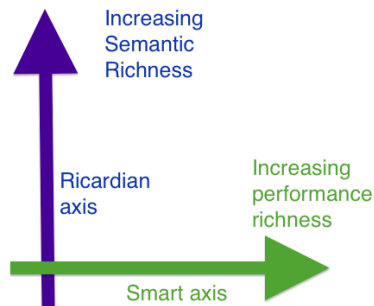


Figure 1: From Grigg [8], the y-axis represents an increasing ability to capture the semantics of a smart legal contract, whereas the x-axis represents the increasing performance of smart contract code.

3 Smart Contract Templates

Smart Contract Templates provide a framework to support complex legal agreements for financial instruments, based on standardised templates. Following Grigg’s Ricardian Contract triple [9], they use parameters to connect legal prose to the corresponding computer code, with the aim of providing a legally-enforceable foundation for smart legal contracts.

Complex sets of legal documentation can be augmented with the identification of operational parameters that are key to directing the executable behaviour of the smart contract

code (in this paper we call these “execution parameters”) — the smart contract code is assumed to be *standardised* code whose behaviour can be controlled by the input of such parameters.

Here we explore the design landscape for the implementation of Smart Contract Templates. We observe that the landscape is broad and that there are many potentially viable sets of design decisions. We therefore propose that a new domain-specific language should be developed to support the design and implementation of Smart Contract Templates, and we conclude this section by sketching some requirements for that language. The design of that language has already begun, and we have given it the name “CLACK” — a Common Language for Augmented Contract Knowledge.

3.1 Templates and Parameters

A template is an electronic representation of a legal document as issued by a standards body — for example, by the International Swaps and Derivatives Association (ISDA). As illustrated in Figure 2, a template contains both legal prose and parameters, where each parameter has an identity (a unique name), a type, and may (but need not) have a value.

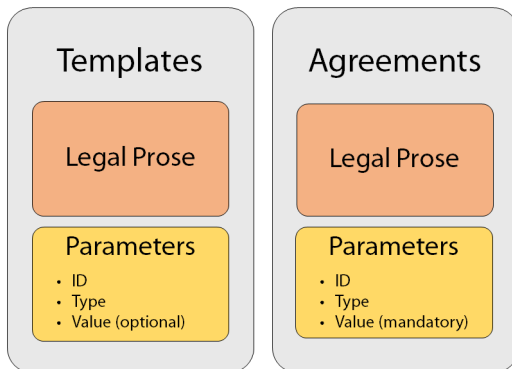


Figure 2: A template may contain both legal prose and parameters. Each parameter has an identifier (a name), a type, and an optional value. Agreements are derived from templates, and both the legal prose and parameters may be customised during negotiation. Values are mandatory for all parameters in a signed agreement.

An agreement is a fully-instantiated template (including any customised legal prose and parameters). The customisation of legal prose and parameters at this stage is commonplace and results from negotiation between the counterparties. We also observe that it is common for agreements to comprise multiple documents such as Framework Agreements (e.g. a Master Agreement) with various Annexes (e.g. a Schedule) and Credit Support Documentation (e.g. a Credit Support Annex). Thus, the legal prose of an agreement will be derived from that of the template, but need not be identical, and similarly the parameters of the agreement will be derived from the template but need not be identical.

Deriving the set of execution parameters may be complicated by three factors:

1. It is common for execution parameters to be embedded in the legal prose — identification of such parameters would initially be undertaken by visual inspection and be aided by a graphical user interface.

2. Some of the values identified as “parameters” in the agreement (and in the template) may not have an operational impact and therefore should not be included in the set of execution parameters.
3. It is possible for a parameter to be defined (given a name) in one document, given a value in a second document, and used (e.g. in business logic) in a third document.

Although parameters need not have values in a template, they must have values in a signed agreement. All of an agreement’s parameter values are a critical part of the contract as they directly reflect the business relationship between parties and those that are execution parameters influence the operation of the contract.

3.2 The design landscape for Smart Contract Templates

In this section, we consider the possible areas for future development of Smart Contract Templates. We do this by considering three distinct areas of development relating to: (i) the sophistication of parameters, (ii) the standardisation of code, and (iii) long-term research. We conclude by sketching a roadmap for future development.

3.2.1 Increasing the sophistication of parameters

Most parameters in existing legal document templates have simple types, such as date, number, etc. These are “base” or “primitive” types⁸ and, as an example, Figure 3 illustrates the identification of a date in a master agreement; once highlighted and annotated, the name (“Agreement Date”), type (“Date”) and value (“16-Mar-2016”) of this parameter will be passed to the executable code.

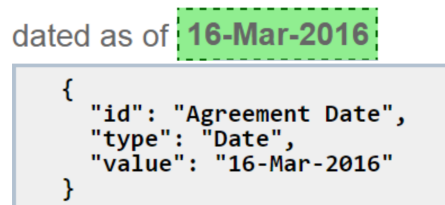


Figure 3: From a Barclays demonstration of Smart Contract Templates: an editor permits a date in the legal prose to be highlighted, and then annotated to denote a simple parameter. The parameter has a name “Agreement Date”, type “Date” and value “16-Mar-2016”.

It is not necessary for parameters to be restricted to base types. It is very likely that values of more complex types, such as lists, will also need to be transferred to the executable code.

The passing of parameters to the executable code is necessary because of the desire to use *standardised* code. It would, for example, be theoretically possible to generate entirely new code for every trade and in this case there would be no need for parameters. The number of parameters, and the complexity of the types of those parameters, will typically increase as the code becomes more generic.

Beyond parameters with base types and more complex types such as lists, parameters can also be expressions containing references to other parameter names. Unless those other

⁸https://en.wikipedia.org/wiki/Data_type

parameter names are defined within the expression, the expression is effectively a function. Where a function is passed as a parameter, this is known as a “higher-order” parameter and the receiving code is known as a “higher-order” function.⁹ An example of a higher-order parameter is illustrated in Figure 4 where some business logic in the legal prose has been highlighted and annotated to be a parameter with name “DailyInterestAmount” of type “Expression”¹⁰ and with a value that is an encoding of the business logic in a format that is easily understandable by a computer.

(x) the amount of cash in such currency on that day; multiplied by
 (y) the relevant Interest Rate in effect for that day; divided by
 (z) 360 (or, in the case of pounds sterling, 365)

```
{
  "id": "DailyInterestAmount",
  "type": "Expression",
  "value": "(CashAmount * InterestRate) / (if Currency == 'GBP' then 365 else 360)"
}
```

Figure 4: From a Barclays demonstration of Smart Contract Templates: an editor permits business logic from the legal prose to be highlighted, and then annotated to denote a higher-order parameter. The parameter has a name “DailyInterestAmount”, type “Expression”, and value corresponding to an arithmetical expression.

The business logic refers to three things whose values are unknown. The first two are simple: “the amount of cash in such currency on that day” (in the expression this is called CashAmount), and “the relevant Interest Rate in effect for that day” (in the expression this is called InterestRate). The third occurs in the phrase “in the case of pounds sterling”, which requires some analysis to determine that it is referring to the prevailing currency (hence the name Currency used in the expression) and that the normal abbreviation for pounds sterling is “GBP”. Since this expression contains three parameter names whose values are unknown, it will be stored in the set of execution parameters as a function taking three arguments (CashAmount, InterestRate, and Currency) and returning a numeric value that is the result of the expression.

It should be noted that when business logic is converted into an expression this may involve the creation of new parameter names (e.g. a new name for the expression itself, and for unknown quantities). Sometimes the business logic may refer to a name that is already defined as a parameter, and sometimes it may refer to a value provided by an “oracle” — i.e. a value, such as an interest rate, that is provided from a trusted source of data and is available to the code while it is executing.

The use of parameters may not only be used to support greater standardisation of code. In the far future, we may see an increasing use of a formally structured style of expression embedded in legal prose; if all business logic in legal prose could be replaced with arithmetical or logical expressions, such as the higher-order parameters discussed in the previous paragraph, this would lead to reduced ambiguity in legal prose and fewer errors in translating legal prose into execution parameters. Such adoption of formal logic into legal prose would require such formal constructs to gain acceptance in the courts and to be admissible as evidence of the intentions of the parties.

⁹https://en.wikipedia.org/wiki/Higher-order_function

¹⁰In a conventional type system this expression would have an associated function type such as: (Decimal, Decimal, Currency) -> Decimal

Figure 5 illustrates our view of how the sophistication of parameters and their role in Smart Contract Templates may evolve in the future.

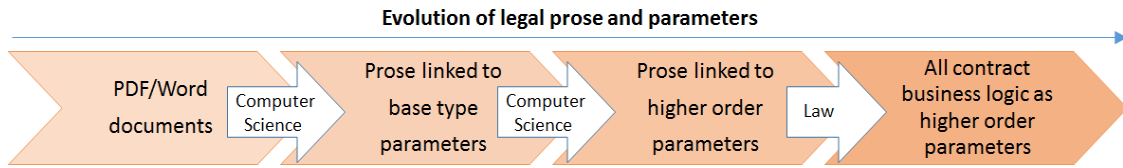


Figure 5: Execution parameters may become more sophisticated in the future, evolving from just simple base type parameters to also include more complex higher-order parameters. In the far future, if the encoding of business logic used in the parameters becomes acceptable to lawyers and admissible in court, then it could potentially replace the corresponding legal prose.

3.2.2 Increasing the use of common standardised code

In the previous subsection, we observed that the passing of parameters to executable code is necessary because of the desire to use *standardised* code. This is important for efficiency reasons as different smart contract code would otherwise have to be built, tested, certified and deployed for *every* different trade. The effort is reduced if such code can be standardised with parameters being passed to each invocation of that code.

This therefore drives a desire for greater genericity of code, which can be enabled by passing more parameters, and/or more sophisticated parameters (with more complex types). Yet despite the gains of standardised and more generic code, there remains the problem that each bank currently manages its own distinct codebases. If smart contract code could be *common* (i.e. *shared*) then it could be built, tested and certified once — and then utilised by every counterparty.

We envisage that the potential economic benefits of using common (shared) code will drive greater adoption in the future. One possible evolutionary route could build upon the use of common utility functions — programs that are already very nearly identical in all counterparties. As the potential economic benefits become clearer and the supporting technologies mature, the size and importance of such common code could increase until, eventually, common business logic may be executed as standardised smart contract code. Figure 6 illustrates how the sharing of code may evolve in the future.

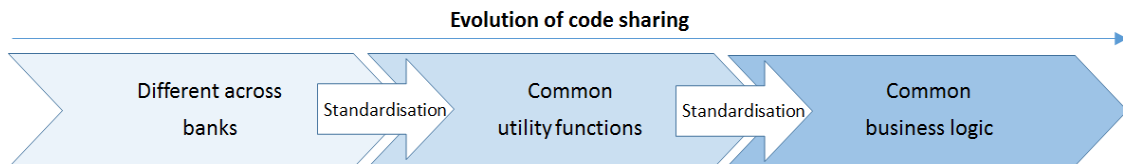


Figure 6: Code may become more standardised in the future through increased sharing, evolving from different codebases across banks to greater adoption of common utility functions to common business logic.

3.3 Long-term research challenges

Several research challenges concerning smart contracts, distributed ledgers, and blockchains are currently being explored in academia. Financial institutions are providing input and inspiration by highlighting relevant business challenges in technology and operations. A good example is the potential for greater straight-through-processing from contract to execution. Currently, lawyers draft legal contracts, which are then negotiated and changed by possibly other teams of lawyers, and then operations staff inspect the contract documents and/or other materials¹¹ to identify the execution parameters that are then passed to code that may have been written some time ago.

This raises several issues:

- Can we be absolutely certain of the meaning of the contract? Are all parties truly agreed on the meaning of the contract, or do they instead each have a subtly different understanding of what the contract means?
- Can we be certain that *all* execution parameters have been identified by the operations staff? Can we be certain that those parameters that have been identified are indeed operationally relevant? And can we be certain that their names, types and values been faithfully transcribed?
- After the parameters have been passed to the code, and the code runs, can we be certain that the code will faithfully execute the operational semantics of the contract? and will it do so under all conditions?

A possible solution would be to develop a formal language in which to write legal documents — i.e. contract documents — such that the semantics would be clear and the execution parameters could automatically be identified and passed to standardised code (alternatively, new code could be generated). This formal language would:

1. derive a number of important qualities from well-designed computer programming languages, such as a lack of ambiguity, and a compositional approach where the meaning of any clause can be clearly deduced without reading the rest of the document; and
2. be simple and natural to use, to such an extent that a lawyer could draft contracts using this formalism instead of using traditional legal language.

The former aspect has already received considerable attention in academia (see survey in [11]) and beyond (e.g. open-source Legalese project [13]). In contrast, the latter aspect is likely to be by far the greater challenge.

Another challenge is whether such a contract, written in a computer-like language, would be admissible in court as a true and faithful representation of the intentions of the parties. Issues of signature and tamper-evident documents are easily solved, yet whether a court would accept the definitions of the meanings of phrases in such a contract is not immediately clear.

As illustrated in Figure 7, this problem could be solved in two ways:

1. As a first step, the language could generate a document version of the contract in a more “natural” legal style, with the expectation that this document would be admissible in court.
2. Eventually, further research in domain-specific languages and law could result in a new formalism itself being admissible in court.

¹¹Other supporting materials may include confirmations, emails, facsimiles, telephone recordings etc.

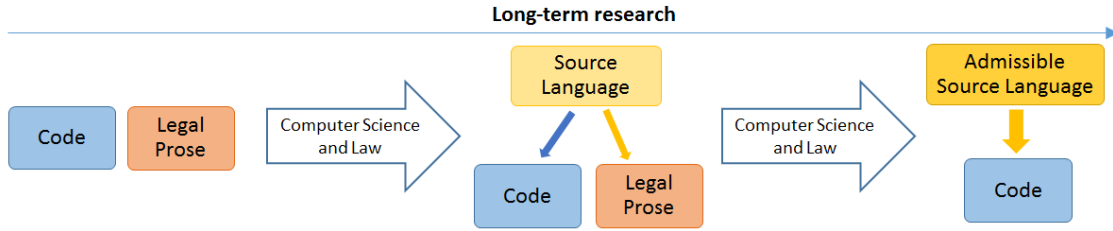


Figure 7: Long-term research may lead from existing separate code and legal prose to source languages which can be automatically translated into both executable code and legal prose, with the prose being admissible in court. Even longer term research could result in formal languages which themselves are admissible in court. Note, this figure omits parameters for clarity.

3.3.1 Future developments and a new supporting common language

The areas of future development described in the preceding sections are brought together in Figure 8, illustrating the potential evolution of important aspects of legally-enforceable smart contracts.

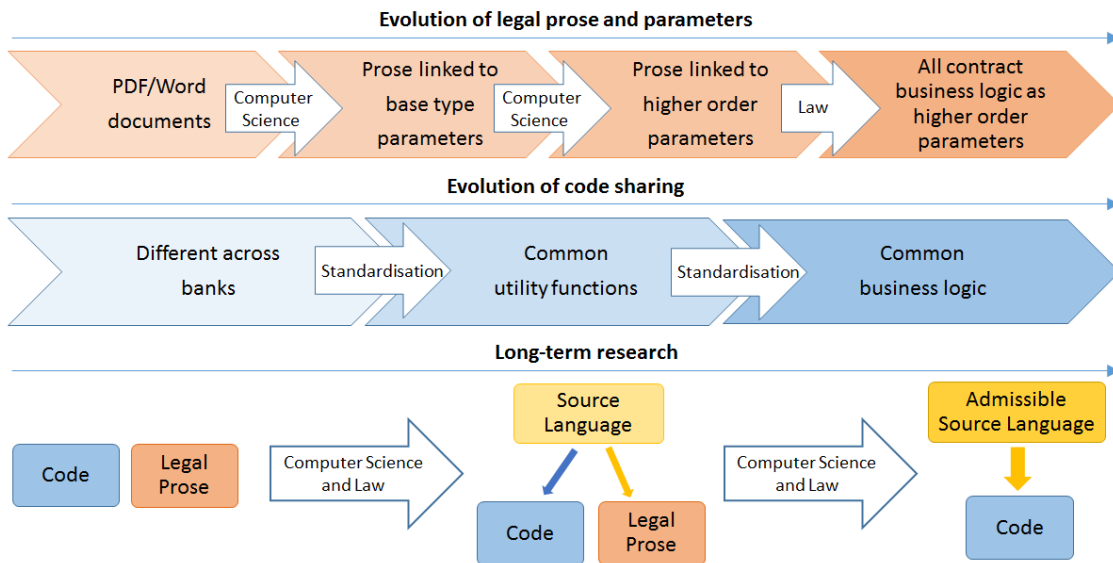


Figure 8: Potential evolution of important aspects of legally-enforceable smart contracts: legal prose and parameters, code sharing, and long-term research.

As a result of this complexity, we were motivated to define a common language to support the specification of different solutions across the design space of Smart Contract Templates.

As a pun on the name of one of the authors of this paper, the new language has been named “CLACK” — the Common Language for Augmented Contract Knowledge. Initially the language will help in the specification of different design choices and in the building of prototypes. In general, the language should be as flexible as possible to support a wide range of requirements. An initial set of requirements has been sketched as follows:

- It should provide support for both legal prose and parameters.
- It should support different internal structured formats, such as XML.
- It should support the output of execution parameters in a variety of formats, such as FpML.
- It should support contracts that comprise multiple documents.
- It should manage multiple agreements being instantiated from a single template (and hierarchies of templates).
- It should permit parameters to be defined in one document, given a value in a second document, and used in a third document.
- It should support a wide range of parameter types, including higher-order parameters.
- It should support increasing standardisation and sharing of common code.
- It should support multiple execution platforms.
- It should support full interaction with, and increasing automation of, legal prose.
- It should support Ricardian Contracts, and therefore for example it should support digital signing of contracts, the construction of a cryptographic hash of the contract, and the use of that hash as an identifier for reference and recovery of the smart contract.

It is a significant challenge to ask a single specification language to do all of the above. We aim to report on progress in a subsequent academic paper, including an abstract syntax and concrete examples.

4 Summary and Further Work

4.1 Summary

This paper began by considering four foundational topics regarding smart contracts: terminology, automation, enforceability, and semantics. We defined a smart contract as an agreement whose execution is both automatable and enforceable. We viewed legal contracts within a simple semantic framework as having two interpretations: the operational semantics concerning execution of the contract and the denotational semantics concerning the non-operational legal interpretation of the contract.

We then described templates for legally-enforceable smart contracts as electronic representations of legal documents containing prose and parameters (with each parameter comprising an identifier, a type and an optional value). Agreements are then fully-instantiated templates (with all parameters having values), including any customised legal prose and parameters. By also selecting the appropriate smart contract code, this approach results in the creation of Ricardian Contract triples.

The design landscape was then explored including increasing the sophistication of parameters from base types to complex higher-order types to business logic that could be admissible in court and potentially replace the corresponding legal prose. We also explored increasing the use of common standardised code through greater sharing, evolving from

different codebases across banks to broader adoption of common utility functions to common business logic. Additionally, long-term academic research was outlined which could lead to source languages which can be automatically translated into both executable code and legal prose; even longer term research could result in formal languages which themselves are admissible in court.

To help manage this complexity, we sketched an initial set of requirements for a common language for specifying different solutions across the design space of Smart Contract Templates. These new technologies could benefit not only financial services but also other industries that could use automated legal agreement. This has a potential to provide opportunities for new types of jobs, including highly-skilled roles.

4.2 Further Work

The CLACK language is being specified and prototyped to support Smart Contract Templates. Next steps include fully specifying: (i) intra-document and inter-document referencing including ambiguity and conflict resolution strategies, (ii) syntax and semantics of expressions within higher-order parameters, etc. Our aim is to report on further development of the CLACK language in a subsequent academic paper, including an abstract syntax and concrete examples.

A benefit of looking to the future is that it helps to identify a potential roadmap for applying academic research within industry. In this case, Smart Contract Templates have already demonstrated a way to link *standardised agreements* to *standardised code* and so, in the near term, it may be possible to utilise them within existing infrastructure. In the longer term, they could be utilised on shared ledgers.

There are many open questions for the future. We have explored some of these questions in this paper, but we will end with one more: is it possible to provide straight-through-processing of financial contracts, with full confidence in the fidelity of the automated execution to the operational semantics of the contract? This, of course, will require substantial work from academia working with lawyers, standards bodies and the financial services industry.

References

- [1] L. Braine. Barclays' Smart Contract Templates, 2016. Barclays London Accelerator, <https://vimeo.com/168844103> and <http://www.ibtimes.co.uk/barclays-smart-contract-templates-heralds-first-ever-public-demo-r3s-corda-platform-1555329>.
- [2] R. G. Brown. Introducing R3 Corda: A distributed ledger designed for financial services, 2016. <http://r3cev.com/blog/2016/4/4/introducing-r3-corda-a-distributed-ledger-designed-for-financial-services>.
- [3] R.G. Brown. A simple model for smart contracts, 2015. <https://gandal.me/2015/02/10/a-simple-model-for-smart-contracts/>.
- [4] CommonAccord, 2016. <http://www.commonaccord.org/>.
- [5] Ethereum, 2016. <https://www.ethereum.org/>.
- [6] The Gale Group Inc. *West's Encyclopedia of American Law*. 2nd. edition, 2008.

- [7] I. Grigg. The Ricardian Contract. In *Proceedings of the First IEEE International Workshop on Electronic Contracting*, pages 25–31. IEEE, 2004. http://iang.org/papers/ricardian_contract.html.
- [8] I. Grigg. On the intersection of ricardian and smart contracts, 2015. http://iang.org/papers/intersection_ricardian_smart.html.
- [9] I. Grigg. The sum of all chains — let’s converge!, 2015. Presentation for Coinscrum and Proof of Work. <http://financialcryptography.com/mt/archives/001556.html>.
- [10] M. Hancock and E. Vaizey. Distributed ledger technology: beyond block chain, 2016. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/492972/gs-16-1-distributed-ledger-technology.pdf.
- [11] T. Hvitved. *Contract Formalisation and Modular Implementation of Domain-Specific Languages*. PhD thesis, Department of Computer Science, University of Copenhagen, 2012.
- [12] Hyperledger, 2016. <https://www.hyperledger.org/>.
- [13] Legalese, 2016. <http://www.legalese.com/>.
- [14] M. Mainelli and A. Milne. The impact and potential of blockchain on the securities transaction lifecycle, 2016. SWIFT Institute Working Paper No. 2015-007.
- [15] J. Stark. Making sense of blockchain smart contracts, 2016. <http://www.coindesk.com/making-sense-smart-contracts/>.
- [16] T. Swanson. Great chain of numbers: A guide to smart contracts, smart property and trustless asset management, 2014. <https://s3-us-west-2.amazonaws.com/chainbook/Great+Chain+of+Numbers+A+Guide+to+Smart+Contracts%2C+Smart+Property+and+Trustless+Asset+Management+-+Tim+Swanson.pdf>.
- [17] T. Swanson. Consensus-as-a-service: a brief report on the emergence of permissioned, distributed ledger systems, 2015. <http://www.ofnumbers.com/wp-content/uploads/2015/04/Permissioned-distributed-ledgers.pdf>.
- [18] N. Szabo. Formalizing and securing relationships on public networks. *First mind*, 2(9), 1997.