A Proposal For Risk Distribution Over Intelligent Credit Networks

Gaurav Singhal*

Laurence E. Day†

Abstract

Peer to peer lending is a rapidly growing market and is expected to reach US$800 billion by 2025. However, peer-to-peer lending is currently fraught with risk: risk assessment itself is tricky, loans often have no collateral, and the mechanisms for enforcement are often weak.

In a nutshell, we are creating an upgraded peer to peer system that significantly derisks loans using artificial intelligence and decentralized networks, where the amount of a loan held by any one lender is repeatedly novated amongst colleagues who have established explicit lines of trust. Repeated across a network, any one loan becomes 'fragmented', reducing the exposure that any one person bears, with local AI agents performing optimal allocations for each person based on their risk appetite and dollar value in play to maximise their risk adjusted return.

In this brief overview, we will explain the underlying theory and concepts behind this model, and illustrate the computational difficulties which require further attention in order to be viable for real-life application.

* *MSc Information Science & Economics @ Michigan*

† *PhD Computer Science @ Nottingham; MSc Financial Engineering @ WorldQuant*

Contents

1 Network Behaviour

Our goal is to create a variant of the peer-to-peer lending model which we refer to as *network lending*. In this model, agents who are connected to each other via a network represented as a weighted directed acyclic graph (Thyfronitis Litos & Zindros, 2017) distribute funds in an approximately Pareto efficient manner (Mock, 2011) to agents demanding said funds (borrowers) through other agents on the network (lenders).

1.1 Network Characteristics

The basic characteristics of these agents and the network are:

- Agents (representing the vertices V of the network graph) first deposit funds (Q_*) that are held in a global escrow account.
- Any two agents $A, B \in V$ can be connected to each other by a directed weighted edge $e_{AB} \in E$ where the weight $w_{e_{AB}}$ denotes the credit capacity offered by A to B.
- The credit capacity of an edge must be less than or equal to the funds deposited by the agent at the tail of the edge: i.e. $w_{e_{A\cdot}} \leq Q_A \; \forall \; e_{A\cdot}$.
- We define 'trust' to be a synonym for the (extension of) credit capacity between two distinct agents.
- The existence of an edge does not imply the existence of a transpose: i.e. $e_{AB} \nRightarrow \exists \; e_{BA}$.
- The credit capacities of an edge and its transpose need not be equal: i.e. $w_{e_{AB}} \neq w_{e_{BA}}$.
- Since trust (as defined above) is not reciprocal, an agent A has two sets of neighbours:
  1. Trustors ($\{ \; B \parallel \text{B} \in V, \; e_{BA} \in E\}$): the set of vertices that explicitly trust agent A, providing incoming lines of credit.
  2. Trustees ($\{ \; B \parallel \text{B} \in V, \; e_{AB} \in E\}$): the set of vertices that agent A explicitly trusts, and to whom flow outgoing lines of credit.
  3. The trustor/trustee relationship is an inverse one - if agent B is a trustor of agent A, then agent A must be a *trustee* of agent B, and vice versa.

1.2 Agent Characteristics

Any given agent A has a known *risk of default* (p_A) on any funds which they have borrowed, represented as a beta distribution parameterised by scale terms α_A and β_A. Only members of an agents trustor set are able to directly view these parameters; other agents can only infer them by asking other agents until the request reaches a trustor. The resulting message is then propagated back to the requestor, not unlike in a game of telephone.

Hence, any agent B that is not a trustor of A has a *subjective view* of the default risk posed by A, described by parameters α_{BA} and β_{BA} (Sun, Yu, Han, & Liu, 2006a).

Further characteristics that are local to agents are:

- Each agent A has a known risk tolerance represented by two parameters:
  1. z_A - the maximum allowable loss on deposited funds Q_A, e.g. 1%.
  2. k_A - the maximum probability that projected losses exceed z_A, e.g. 5% probability that losses exceed 1%.
- A third parameter λ_A representing the expectation-variance trade-off can be inferred from the above two parameters.
- Borrowers must pay some interest rate r on their loan (parameterised by the loan amount D and the repayment period T), and they can offer a percentage of the loan amount c as collateral.
- Each agent wishes to maximise the expected value of their profits from lending their funds to other agents on the network whilst satisfying the above risk parameters.
- Each agent is limited to interacting solely with their trustor/trustee sets. No agent has a global view of the network, and hence all funds, messages and responses must be routed through neighbours.

The (subjective) risk for a given lender increases as their distance from a borrower increases. Further, no funds can be routed at a loss. Ergo, incentive compatibility requires a loan to be either derisked or yield higher interest as it is propagated from primary lenders (direct trustors of the borrower) to secondary and higher lenders (trustors of primary lenders et al). This lends itself to our first lemma:

Lemma 1.1. *All payment paths from lenders to a given borrower must satisfy risk monotonicity.*

Loan propagation occurs through *novation*, which we discuss in the next section.

2 Agent Behaviour

2.1 Novation

Novation is the central behaviour of agents on a network such that we have defined above; it is the process by which funds are propagated from borrowers to many lenders on said network. In a financial context, novation is the act of purchasing debt, amending the contract and then reselling some or all of the debt to additional parties, as illustrated below:

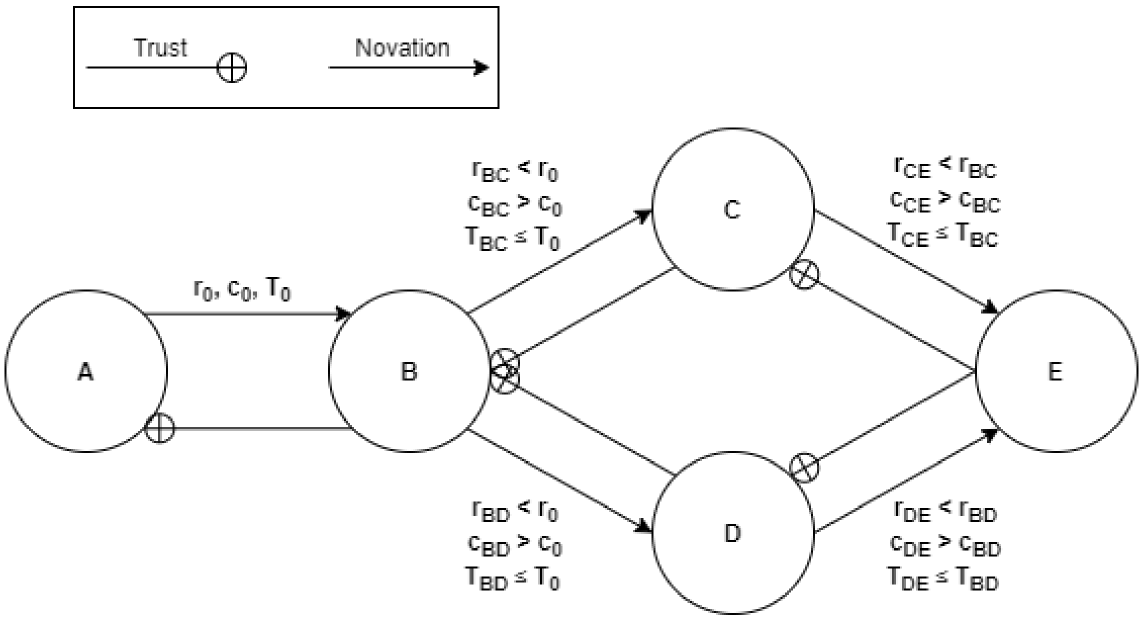

Figure 1: The Novation Process

In Figure 1 above, agent B purchases debt from a borrower agent A in their capacity as a trustor of A. As B is a trustee for agents C and D, they sell part of this debt to both of these trustor agents (provided the terms are within their risk profile) at a decreased interest rate and collateral, for a time period less than or equal to than the initial loan. In turn, agents C and D are both trustees of agent E, and so sell parts of their own newly-purchased debt to E on their own terms (a further decreased interest rate etc.). At settlement of the initial loan from B to A, provided that no default occurs, all secondary agents involved in the chain profit from the difference in terms.

With respect to the aforementioned credit network, lenders always perceive borrowers as riskier - due to the dispersion of their perceived risk profile - the further they are from the borrower on the network, per Lemma 1.1. As rational novation cannot occur at a loss, an additional lemma emerges from incentive compatibility:

Lemma 2.1. *For a loan originator to make a profit, $r_n \leq r_{n+1}$. As interest cannot increase, Lemma* 1.1 *necessitates that the collateral on the novated debt monotonically increases and/or the tenor decreases as we move further away from the borrower, i.e.* $c_n > c_{n+1}$ *and* $T_n \geq T_{n+1}$.

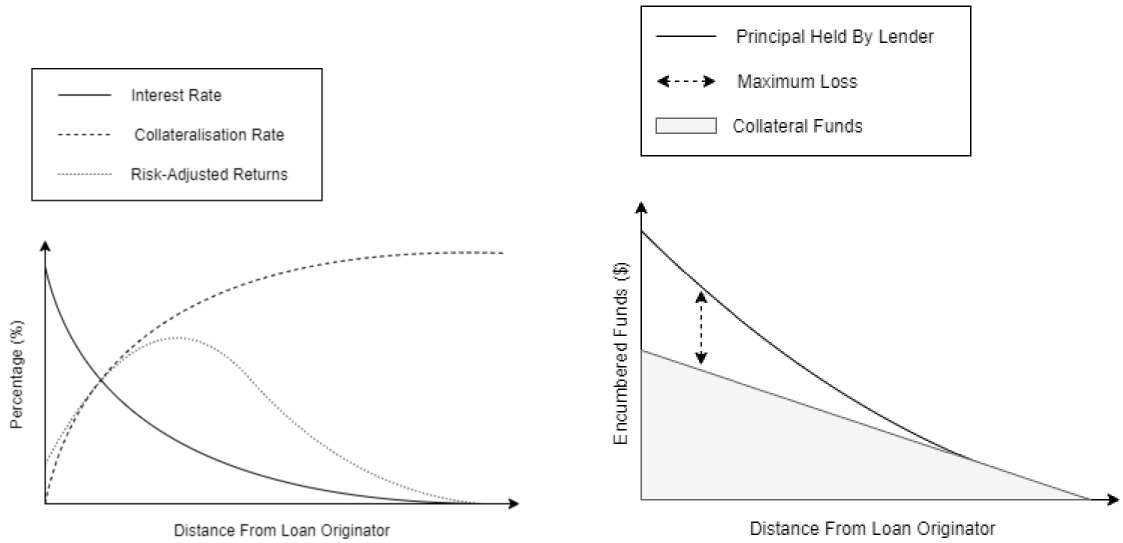

Figure 2: Consequences Of Lemma 2.1 As Functions Of Distance From Loan Originator

In Figure 1 we saw a relatively simple example of a loan being propagated by novation. Consider agent E, the agent at the end of a novation chain. This agent has no trustors that they can novate debt further to, and so from their perspective, the issue of participating in the original loan becomes a tractable optimisation problem, which we describe in Section 5.

3 An Extended Example

The underlying concepts explained thus far are best illustrated by seeing a loan - and its novation - in action. Consider the network below:

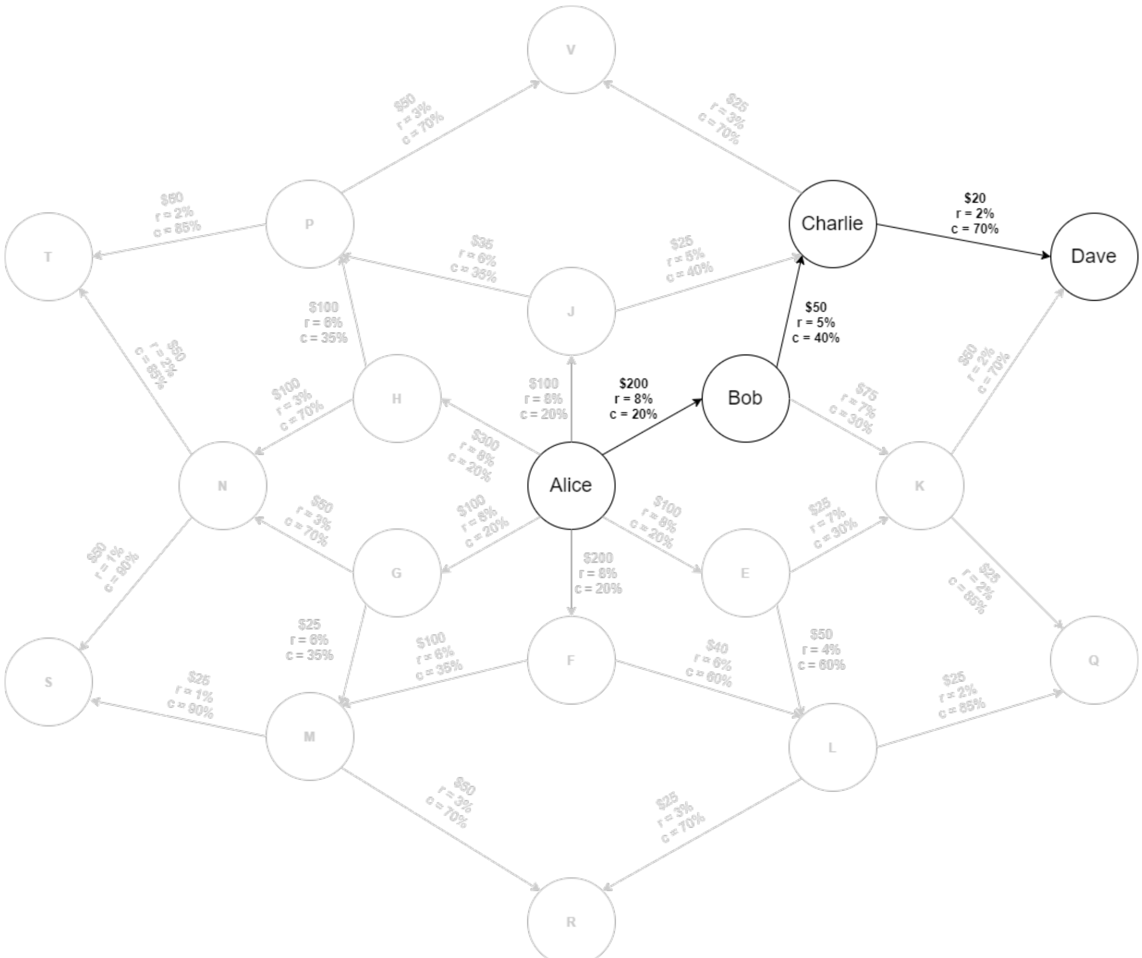

Figure 3: Explicit Novation Chain From Borrower Alice To Lender Dave

In this network, a borrower Alice has requested (and received) a loan for $1,000. Her loan was initially purchased by six of her trustors, one of whom is Bob. Bob is in turn trusted by Charlie, who is trusted by Dave. Whilst Figure 3 above shows example paths for all of Alice's participating trustors, we will be explicitly focusing on the chain novation from Alice through to Dave, exploring how each participant profits from their involvement.

All agents on the above network have already performed the prerequisite steps for novation; namely, they have:

1. Determined a *subjective default risk profile* for Alice (see Section 7).

2. Coalesced their trustors' demand functions (if any).

3. Broadcast their own demand functions (for Alice's loan) to their trustees.

3.1 Following Alice Down The Rabbit Hole

Based on the results that Alice received from her trustors, she requested that her loan of \$1,000 is taken with an interest rate of 8% and a 20% collateral deposit, and of this principal, Bob contributed \$200. Now, we describe the process by which we propagate this segment of the loan through Bob's own network of trustors.

Bob - who has knowledge of the demand for Alice's loan from Charlie, sells him \$50 of this newly acquired \$200 debt. This \$50 is sold at a lower interest rate than the initial debt (5% rather than 8%), but Bob has doubled the collateralisation rate, bringing it up to 40%. Here is where Bob and Charlie now stand overall:

- Bob holds \$150 of Alice's debt, at 8% interest with 20% collateralisation.
- Charlie holds \$50 of Alice's debt, at 5% interest with 40% collateralisation.
- Bob now also holds a *swap note* on Alice's debt worth \$50 at 3% interest (8% - 5%) provided Alice repays her loan in full, but may cost him up to \$10 (\$50 × (40% - 20%)) if Alice defaults.

Let us consider what this change in circumstance has done to Bob's potential profits:

- Bob initially placed \$160 at risk (\$200 × (100% - 20%)) in order to earn \$16 in interest, yielding an *effective* interest rate (EIR) of \$16/\$160 = 10%.
- After novating \$50 of this \$200 debt to Charlie, Bob is now risking \$130 (the result of \$150 × (100% - 20%) + \$50 × (40% - 20%)) in order to earn \$13.50 in interest (\$150 × 8% + \$50 × (8% - 5%)), yielding an EIR of \$13.50/\$130 = 10.38%.

It seems counterintuitive, but this novation has *increased* Bob's overall risk-adjusted return, whilst simultaneously exposing Charlie to an asset at terms he is comfortable with.

We now repeat this process with Charlie, who – knowing his trustor Dave's demand for Alice's principal – novates \$20 of the \$50 debt he has purchased from Bob at a further reduced interest rate of 2%, adding an additional 30% collateral.

As was the case with Bob, we compare Charlie's pre-and-post-novation returns, noting that Charlie now holds a swap note worth \$30 at 3% interest with a maximum potential loss of \$9 (\$30 × (70% - 40%)):

- Pre-novation: Charlie risks \$30 (\$50 × (100% - 40%)) to earn \$2.50 in interest, for an effective interest rate of 8.33%.
- Post-novation: Charlie now risks \$24 (\$30 × (100% - 40%) + \$20 × (70% - 40%)) in order to earn \$2.10 in interest (\$30 × 5% + \$20 × (5% - 2%)), yielding an EIR of \$2.10/\$24 = 8.75%.

As was the case with Bob, Charlie has reduced his exposure whilst increasing his risk-adjusted returns. Finally, Dave - the end of the chain - now holds an asset representing a debt of \$20 in order to earn \$0.40 in interest, with a maximum potential loss of \$6.

As a rule, novation always increases risk-adjusted return, allowing loan risk to propagate throughout a network as a consequence of agents making rational, self-interested decisions.

3.2 The Bigger Picture

In the interests of completeness, below we provide a table of the *true* - considering all chains - figures of funds exposed, interest and returns pre-and-post-novation for *all* loan participants, leaving the intermediate calculations to the interested reader:

Participant	Pre-Novation Funds Exposed	Post-Novation Funds Exposed	Interest	Risk-Adj. Return
Bob	\$160.00	\$77.50	\$8.25	10.64%
E	\$80.00	\$42.50	\$4.25	10.00%
F	\$160.00	\$79.00	\$7.60	9.62%
G	\$80.00	\$48.75	\$5.00	10.25%
H	\$240.00	\$145.00	\$15.00	10.34%
J	\$80.00	\$42.25	\$4.65	11.00%
Charlie	-	\$31.50	\$2.60	8.25%
K	-	\$51.25	\$5.50	10.73%
L	-	\$24.75	\$3.15	12.72%
M	-	\$63.75	\$5.75	9.01%
N	-	\$32.50	\$3.00	9.23%
P	-	\$65.25	\$5.60	8.58%
Dave	-	\$21.00	\$1.40	6.67%
Q	-	\$7.50	\$1.00	13.34%
R	-	\$22.50	\$2.25	10.00%
S	-	\$7.50	\$0.75	10.00%
T	-	\$15.00	\$2.00	13.34%
V	-	\$22.50	\$2.25	10.00%
Total	**\$800.00**	**\$800.00**	**\$80.00**	-

As we can see, the funds exposed - loan principal not covered by collateral - and interest due do not change, but are rather dissipated amongst the various network participants. Moreover, the novation of Alice's loan permits twelve network participants to act as lenders despite not knowing Alice directly.

4 The Consequences Of Default

Whether it be through bad luck, poor choices, or malicious intent, not all loans are repaid in full - if they're repaid at all. The word *default* is typically used in this case, and in the world of traditional financial institutions, usually comes associated with a whole host of troubles; bank account seizures, professional license revocation and withholding of tax refunds, to name but three.

In the context of existing peer-to-peer lending platforms with a centralised point of control, to default on a loan is often to render yourself ineligible for further participation, or at best be restricted to borrowing trivial amounts for the foreseeable future. Attempts have been made to model the probability and determinants of default within peer-to-peer networks (Carmichael, 2014; Serrano-Cinca, Gutiérrez-Nieto, & López-Palacios, 2015), but such attempts typically require data far in excess of the means of most people within financially underserved communities (i.e. home improvement, credit card repayment, state unemployment et al). Indeed, in most developing nations, the institutional frameworks which would give rise to this data is often non-existent.

With this said, peer-to-peer loans in their current format offer little recourse for lenders that fund a loan that is later defaulted upon, often due to the lack of an enforcement mechanism. On most peer-to-peer platforms, loans are considered to be between the lender and the borrower, putting the onus of legal action - if feasible - on the former.

Nonetheless, peer-to-peer lending is one of the most common forms of credit access for the unbanked and those seeking smaller amounts that can be justifiably requested from friends and family. In the case of small businesses, this is especially prevalent, as loans to businesses that are just starting up are considerably riskier than to established enterprises from the perspective of traditional financial institutions (Mills & McCarthy, 2014).

So, what differentiates microfinance - the term which most institutions would use to describe peer-to-peer lending - to traditional banking? In the remainder of this section, we briefly delve into some of the literature surrounding the topic, discuss the impact of a borrowers social circle on their loan trustworthiness, and explain how we propose to handle defaults within the context of the model proposed in Section 3.

4.1 Microfinance & The Financially Underserved

It has been long established that wealth distribution amongst individuals globally roughly follows a Pareto distribution, where 80% of the profits accrue to 20% of the population (Stiglitz, 1969). This pattern is also true of banking clients, where the vast majority of revenue comes from a fat left tail. The long right tail which constitutes the remainder is typically not served by financial institutions, due to the low return on investment compared to focusing on their existing - profitable - client base. Their reasoning is sound, albeit unfortunate for those falling within said tail: the costs of administering small loans are high, expending resources that could be better used elsewhere.

However, there is business to be done in serving this thin tail, as credit access is inelastic - regardless of the interest rates involved, borrowers who are in sufficient need will take it up (Dehejia, Montgomery, & Morduch, 2012). It was this realisation that led to the setting up of formal *microfinance institutions* (MFIs) such as the Grameen Bank (Morduch, 2000; Yunus & Jolis, 2003). MFIs - whilst not banks - are bank-*like* in that they consider risk and revenue margin as two of their key concerns, but focus primarily on alleviating poverty rather than maximising profit.

Without the infrastructure of traditional banks, MFIs typically encounter higher administrative costs, with nearly fifty cents spent on fees for every dollar loaned (Aleem, 1990), with the knock-on effect being that without government or private assistance, the interest rates that are charged are typically well above average (Mendoza, 2011). Moreover, some MFIs struggle to receive funding from donors or their recipients due to central bank restrictions on who can gather deposits (Morduch, 2000), and despite the fact that small financial institutions such as MFIs nowadays form a significant credit source for micro/small/medium enterprises (Weston & Strahan, 1996), these funding stressors occasionally lead MFIs to deviate from their initial mission - delving into predatory practices, usury, or physical threat (Serrano-Cinca & Gutiérrez-Nieto, 2014).

4.2 The Impact Of Social Networks

With the advent of the internet and increasing global adoption, MFIs have in some cases been supplanted by crowdfunding or peer-to-peer lending platforms, wherein the loans which are granted are funded by a cohort of individual lenders operating via a centralised point of access, such as Lending Club and Kiva (Fenwick, Vermeulen, & McCahery, 2017; Bachmann et al., 2011). Within these platforms, potential borrowers submit their requests, and their loans are often graded as a function of their default risk, allowing lenders to do so according to their risk profile (Emekter, Tu, Jirasakuldech, & Lu, 2015).

There is a substantial body of literature concerning the risk factors which best indicate the default risk of a loan, due to the information asymmetry between a potential lender and borrower - the borrower knows far more about their capability to pay back a loan than the lender does, *prima facie*.

Crucial to our central thesis is the observation that proximity to a lender - in the sense of real world connections - is a potential indicator of credit risk (Morse, 2015).

5 Portfolio Optimisation

We now return our focus to the issue of network participants deciding how much debt they are willing to purchase when presented with a new lending opportunity. The setup for this problem is as follows:

- A lending agent A has a maximum amount of funds q_{AB} that they are willing to allocate to the debt of a borrower agent B, where $q_{AB} \leq Q_A$.
- A has a current portfolio Γ consisting of multiple risky assets $\gamma_{1,\ldots,N}$. Each constituent asset γ_i is parameterised by the following:
  - x_i - the amount of the asset held in the portfolio.
  - r_i - the interest rate of the asset.
  - c_i - percentage of collateral covering the asset in the event of default.
  - T_i - tenor of asset (time remaining until funds are returned).
  - α_i, β_i - scaling terms for the beta distribution parameterising risk of default.
- A has known risk aversion parameters z_A, k_A and λ_A.

The lender must decide how much of a borrower agent B's debt they wish to add to their portfolio as a new asset γ_{N+1}, where $\gamma_{N+1} = x_{AB} \leq q_{AB}$, given r_{AB}, c_{AB}, T_{AB}, α_{AB} and β_{AB}.

5.1 Simplifying Assumptions

With the following two assumptions, we can approach the above problem - the calculation of γ_{N+1} - in a similar manner to traditional Markowitz portfolio optimisation (Markowitz, 1952).

Assumption 1: without loss of generality, we can approximate the *beta distribution* via the *Kumaraswamy distribution* (Jones, 2009), which is defined as -

$$K(x) = ab\ x^{a-1}(1 - x^a)^{b-1}$$

- where $a > 0$, $b > 0$ are the scale parameters of the distribution. Approximating the beta distribution in this way circumvents the fact that its analytical presentation has no closed form and is thus undifferentiable.

Assumption 2: we ignore time dynamics and assume that all debt has the same settlement time (i.e. $T_* = 1$), thus eliminating tenor from the equations to follow.

5.2 The Problem Made Concrete

In making the two assumptions from the previous subsection, we arrive at the following optimisation problem that a lender with no trustors in a novation chain (i.e. Dave in Section 3) must solve.

An agent A seeking to optimise their portfolio aims to maximise -

$\Gamma = \mathtt{max}(\theta^\intercal(PR + (1-P)C))$, such that:

(1) $b \times B(1+\frac{1}{a}, b) = \frac{\theta^\intercal(PR+(1-P)C)}{\theta^\intercal R}$

(2) $b \times (B(1+\frac{2}{a}, b) - B(1+\frac{1}{a}, b)) = (\frac{\Sigma(R-C)^\intercal\theta}{\theta^\intercal 1_M})^\intercal \times \Omega \times \frac{\Sigma(R-C)^\intercal\theta}{\theta^\intercal 1_M}$

(3) $(1-(1-k_A)^{\frac{1}{b}})^{\frac{1}{a}} \geq z_A \times \frac{\theta^\intercal}{\theta^\intercal R}$

- where

- θ is the vector $[x_1, \ldots, x_M]$ representing the total amount of each risky asset γ_i in Γ.
- P is the vector [1 - p_i, ..., 1 - p_M] representing the payback probability (the default complement) of each risky asset γ_i in Γ.
- R is the vector [1 + r_1, ..., 1 + r_M] representing the total percentage return realised if a given risky asset γ_i is repaid in full.
- C is the vector $[c_1, \ldots, c_M]$ representing the percentage of collateral in place for each risky asset γ_i in the event of default.
- a and b are the parameters of the Kumaraswamy distribution.
- $B(\cdot,\cdot)$ is the Euler integration *beta function* - not to be confused with the beta distribution mentioned in Assumption 1 of the previous subsection - defined as:

$$B(x, y) = \int_0^1 t^{x-1}(1-t)^{y-1}\ dt$$

- Σ is the vector $[\sigma_1, \ldots, \sigma_M]$ of standard deviations of each risky asset γ_i, where:

$$\sigma_i = \sqrt{\alpha_i\beta_i} \times \frac{1}{\alpha_i+\beta_i} \times \sqrt{1+\alpha_i+\beta_i}$$

- 1_M is the vector containing M iterations of the value 1.
- Ω is the $M \times M$ correlation matrix of risky assets in Γ.
- z_A and k_A are the risk aversion parameters of agent A.

The objective is simply to maximise the expected value of the portfolio Γ. In the event that no defaults occur and all risky assets are repaid in full, the returns on Γ are equal to the dot product $\theta \cdot R$. In contrast, if *every* constituent asset defaults, the 'returns' are simply the aggregate of all collateral, $\theta \cdot C$.

Equations (1) and (2) seek to establish equality between the beta distributions and their corresponding Kumaraswamy distributions. The left hand sides of these equations correlate to the first and second moments of the Kumaraswamy distribution, and the right hand sides to the empirical normalised first and second moments of the portfolio Γ itself. Equation (3) imposes limits on the allowable tail risk of the resulting Kumaraswamy distribution describing the portfolio Γ.

Assuming that a portfolio Γ is optimised for an existing M risky assets, the problem of deciding how much debt agent A should purchase from a potential borrowing agent B becomes one of solving for the unknown θ vector element $x_{M+1} = x_{AB}$ as well as the associated vector entries - i.e. c_{M+1}, r_{M+1} - which ensure that Γ *remains* optimal.

To do this, we seek to frame the above in terms of a *demand function* D_{AB} such that -

$$x_{AB} = D_{AB}(r_{AB}, c_{AB})$$

- where $r_{AB} > 0$ and $0 < c_{AB} < 1$.

6 The Demand Function & Message Passing

The ability of agents to pass these demand functions as messages to their neighbours is at the heart of the propagation-by-novation process. If a potential novating agent B receives a demand function D_{CA} from a trustor C in the form of a message, the problem becomes one of calculating terms $r_{BA}^{Out} = r_{CA}^{In}$ and c_{CA}^{In} at which agent B would be able to novate some of the *initial borrowing agent* A's incoming debt x_0 to C, where $\{r,c\}_{BA}^{Out} = \{r,c\}_{CA}^{In}$. We can illustrate this as follows:

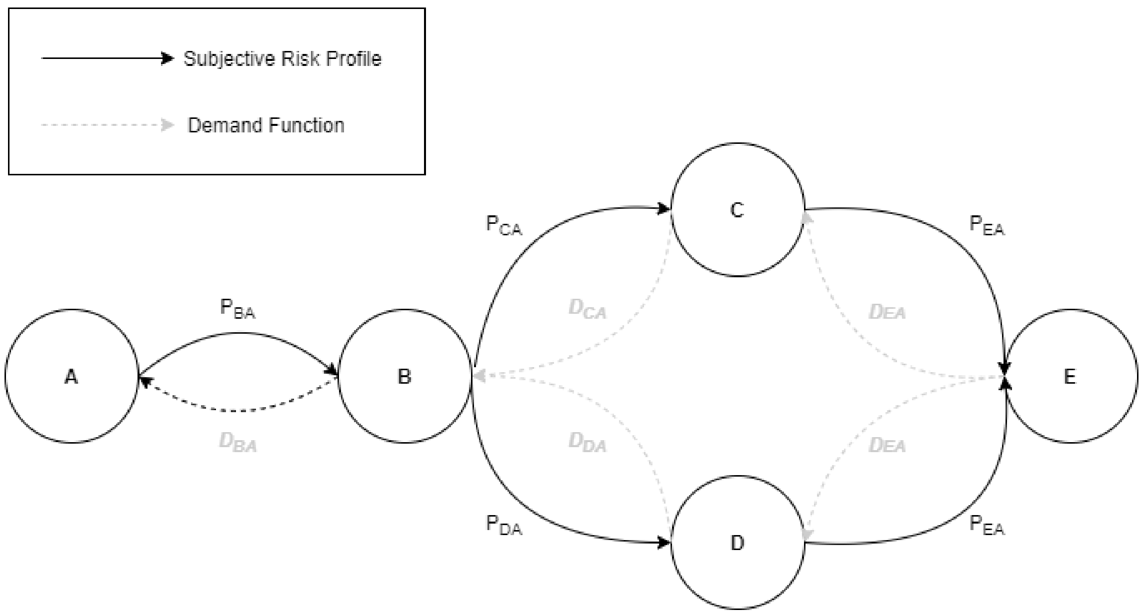

Figure 4: Demand Function Back-Propagation

In the above, the data P_{XY} flowing from trustee to trustor represents information sufficient for the receiving agent Y to calculate their subjective risk profile of the initial borrower X, which we discuss in detail in Section 7. In response to receiving a demand function from a trustor (i.e. D_{CA}), agent B solves for three unknowns x_{BA}, r_{BA}^{Out} and c_{BA}^{Out} and extends their portfolio vectors as follows:

- P = [. . . , 1 - p_{BA}, 1 - p_{BA}] - agent B is effectively creating two new risky assets in their portfolio; the first the debt which B purchases from A, and the second the debt which B sells at novated terms to C. p_{BA} is the probability of default that B assesses A to possess.

- $R = [\ldots, 1 + r_{BA}^{In}, 1 + r_{BA}^{In} - r_{BA}^{Out}]$ - since agent B is selling debt to a willing agent C, the first entry corresponds to the interest rate offered by A, and the second indicates that the difference between the buy and sell rates is the expected rate of return they will earn on the novation.

- C = [. . . , c_{BA}^{In}, -c_{BA}^{Out}] - the first entry corresponds to the collateral of the purchased debt, whereas the second is negative as its encumbered by the additional collateral they have added during the novation process.

- $\theta = [\ldots, x_{BA}, D_{CA}(r^{In}_{CA}, c^{In}_{BA} + c^{In}_{CA})]$ - where D_{CA} is the demand function of C for the debt of borrower A and $r^{In}_{CA} = r^{Out}_{BA}$ and $c^{In}_{CA} = c^{Out}_{BA}$).
- $\Sigma = [\ldots, \sigma_{BA}, \sigma_{BA}]$ - the two assets have the same underlying normalized standard deviation.
- $\Omega = \begin{bmatrix} \ddots & \cdots & \cdots & \cdots \\ \vdots & \ddots & \rho_{AB} & -\rho_{AB} \\ \vdots & \rho_{BA} & \ddots & \vdots \\ \vdots & -\rho_{BA} & \cdots & 1 \end{bmatrix}$ - where ρ_{BA} is the portfolio correlation vector.

6.1 Propagation-By-Novation

In light of everything we have encountered so far, we can now outline the full propagation-by-novation process as follows:

Wave 1: a borrowing agent A broadcasts a message requesting a risk assessment for a particular amount of debt. The message propagates away from A through various levels of trustors, where each intermediate agent combines their information about the trustor which delivered the message with agent A's initial message to create a subjective risk profile of A. As the message moves further away from A, uncertainty around the risk increases.

Wave 2: after a finite number of edge traversals, risk assessment stops (when the message cannot propagate to any additional trustors). At this point, agents begin to propagate their demand functions regarding the initial loan amount back towards agent A. It is important that a strict post-order depth-first search ordering is followed to ensure that an agent has all of their neighbours demand functions prior to computing their own.

Wave 3: finally, agent A coalesces the demand functions of their trustors into a function that proscribes the interest rate at which they can borrow the requested amount as a function of collateral. Once A settles on their terms, their trustors - armed with the knowledge of their own trustors demand functions - begin the novation process, propagating the loan and settling accounts. Funds are encumbered when held for collateral, and other funds are channeled to the borrower for the purpose of their loan.

7 Subjective Risk Assessment

An assumption we have made up until this point is that all participants on a trust network are honest, repaying any and all borrowed funds in full with interest. However, this is naïve both in theory and practice: there will always be malicious actors in a trust-based network, and the default probabilities of participants cannot be assumed equal across the board: some people are naturally diligent with their finances, whilst others are profligate.

People from either camp - and everywhere in between - may opt to borrow for perfectly legitimate reasons, but in doing so, the risk they pose to the network must be quantified.

Without access to a participants' past financial history - as we are specifically aiming to provide credit access to those who have historically been financially underserved - we propose a method in which a potential lender can construct a subjective risk profile of a potential borrower, regardless of their distance.

This method is a novel stochastic variant of the EigenTrust reputation management algorithm (Kamvar, Schlosser, & Garcia-Molina, 2003), designed to aid in valuable, trustworthy community formation. Importantly, it is an endogenous mechanism (i.e. requires no external information).

At this point, we should make clear that the notion of 'trust' is becoming somewhat overloaded. Whilst we have defined trust as used thus far as the maximum credit capacity extended between a trustor and a trustee, the term is also associated with reliability, giving rise to the concepts of *default risk*, *recommender risk* and *subjective risk*. Respectively, these are: the risk that a lender will fail to make full repayment of a loan, the risk that a neighbour will include you in a chain of accountability with a high default risk, and the risk that you determine a particular agent to have given any/all information about it that has been provided to you by your neighbours. The word 'trust' is the most natural one to use when discussing these concepts, and going forward the word should be unambiguous through context.

We proceed by providing an overview of the notion of trust and reputation management in pseudonymous online frameworks, before discussing the EigenTrust algorithm and describing how its' usage gives rise to a robust risk metric which can be further refined over time by way of monitoring borrowing and repayment habits.

7.1 Reputation Systems & Their Limitations

In peer-to-peer networks such as those used for sales of goods (eBay) and file-sharing (BitTorrent), each user is pseudonymous - identified with a username or IP address - and associated with a 'reputation score' of sorts. For eBay, buyers and sellers can rate each other at the end of each transaction, whereas in BitTorrent, a user can give a positive or negative rating to a peer once they have finished downloading a file from them. These vary in scope between a simple 'good' or 'bad' decision, through to a five-star rating system.

These reputations become key to the success or failure of network participants, including financially leveragable advantages - given the choice between two retailers on eBay, a user is likely to pick the one with a 'better' reputation. As such, in situations where a user can freely rate others they have interacted with, malicious peers and collectives often arise in an attempt to inflate their own trustworthiness relative to others. Moreover, users new to a network do not necessarily know who they can trust, and may inadvertently trust a malicious actor. In systems such as the one we describe in this paper, there are multiple - albeit correlated - reputation metrics in play (i.e. recommender vs default risk), and when dealing directly with lending, the subjective views of participants asked to rate each other are likely to be so conservatively skewed as to inhibit the most efficient flow of capital.

We therefore propose that participants are *not* permitted to rate fellow network participants directly, with the only actionable method of expressing trustworthiness being that of the maximum credit capacity extended to them (which can be adjusted at will). We will see below how it is still possible to calculate risk profiles for participants even with this restriction in place.

7.2 EigenTrust: A Quick Primer

The EigenTrust algorithm proposes a way in which the global reputations and trust values of peers can be determined via the reputation-weighted aggregate of their own peers opinion of them, in a decentralised fashion. The nub of the algorithm is that global trust values can be calculated as the left principal eigenvector of a matrix of normalised trust values. We briefly define the notions that form the whole, and in the next subsection discuss how the algorithm is used within our framework. It should be highlighted here that what follows is *not* original research, but a condensed version of the main findings in (Kamvar et al., 2003).

7.2.1 Trust Values & Vectors

In EigenTrust, the concept of a *local trust value* is defined in two ways - either the 'transaction model' or the 'satisfaction model'. Within these, the trust peer i has in peer j is defined as -

$$\text{Transaction Model: } s_{ij} = \Sigma\, tr_{ij}$$
$$\text{Satisfaction Model: } s_{ij} = sat(i,j) - unsat(i,j)$$

- where tr_{ij} is a binary rating based on past transactions, whereas *sat* and *unsat* are predicates based on the number of satisfactory versus unsatisfactory transactions that have taken place. These models, however, do not account for transitive trust (determining how trustworthy a peer is that you have had no interaction with), which is crucial when determining whether or not to purchase novated debt via one of your trustees.

Since local trust values can be easily exploited by a malicious actor or collective assigning trust in bad faith (i.e. artificially inflating a malicious peers trust value, or downvoting honest peers), these localised trust values must be *normalised*.

Regardless of the model which determined the local trust value, normalised local trust is calculated as -

$$c_{ij} = \frac{max(s_{ij},0)}{\Sigma_j(max(s_{ij},0))}, \text{ where } \Sigma_j(max(s_{ij},0)) \neq 0$$

- producing a value between 0 and 1. The predicate stating that the denominator must be non-zero is guaranteed by the notion of *pre-trusted peers*, which we will encounter shortly.

Due to the presence of potential malicious actors as well as the adage that you can never have too much information, it is natural for a peer i to ask its' peers how trustworthy they consider other peers to be, and aggregate their opinions, with each opinion weighted by the degree to which peer i trusts each respondent. As such, given a peer k that peer i has not interacted with in the past, the degree to which i trusts k is:

$$t_{ik} = \Sigma_j c_{ij} c_{jk}$$

One of the key observations in the original EigenTrust paper is that this has an elegant matrix notation: if C is the matrix $[c_{ij}]$ and $\vec{t_i}$ is a j-vector of values t_{ik}, then $\vec{t_i} = C^\intercal \vec{c_i}$. Crucially, this vector can be refined by asking increasingly wide circles of peers (i.e. the peers of peers), whereby $\vec{t_i} = (C^\intercal)^n \vec{c_i}$ for n iterations, and the vector $\vec{t_i}$ converged upon is identical for *every* peer i in the network - the left principal eigenvector of C - and represents a global trust vector which contains values corresponding to the trust in which the network as a whole places in each peer. If there are m peers in the network at the time of asking, the algorithm can be bootstrapped by taking $\vec{c_i}^0 = \vec{e}$, where $e_i = \frac{1}{m}$ is the uniform distribution over m peers.

7.2.2 Pre-Trusted Participants

In the absence of certainty as to who can be trusted (i.e. when a network is being initialised), it is helpful to designate early participants - those corresponding to the network developers, or any partnering institutions, for example - as being *pre-trusted*. A simple way of doing this is to define a m-vector $\vec{p}$ for a network of m peers, where $p_i = \frac{1}{|P|}$ if $i \in P$, and 0 otherwise.

Three more key observations within EigenTrust are that in the presence of malicious peers, $\vec{t} = (C^\intercal)^n \vec{p}$ converges faster than $\vec{t} = (C^\intercal)^n \vec{e}$. Moreover, in the case of a peer α who is brand new to the network (who has never rated anyone, and hence $c_{\alpha j} = 0 \ \forall j$), we can satisfy the predicate when calculating normalised local trust values by setting $c_{\alpha j} = p_j$. Finally, malicious peers can be subverted by incorporating the pre-trusted peers into the iterative step of computing $\vec{t}$:

$$\vec{t}^{(k+1)} = (1-a) C^\intercal \vec{t}^{(k)} + a\vec{p}, \text{ where } a < 1 \in \mathbb{R}$$

7.2.3 The EigenTrust Algorithm

With the above in hand, we can now make sense of the 'basic' (non-distributed) EigenTrust algorithm, which takes as input the vector of pre-trusted peers $\vec{p}$ and returns the global trust vector $\vec{t}$ within some margin of error ϵ -

$\vec{t}^{(0)} = \vec{p}$
`while` $\delta > \epsilon$:
  $\vec{t}^{(k+1)} = C^{\intercal} \vec{t}^{(k)}$
  $\vec{t}^{(k+1)} = (1-a)\vec{t}^{(k+1)} + a\vec{p}$
  $\delta = ||t^{(k+1)} - t^{(k)}||$

- where $||x||$ is the magnitude of the vector x.

8 Adapting EigenTrust Into Our Model

8.1 Axioms

1. Uncertainty is a Measure of Trust: The concept of trust describes the certainty of whether the agent will perform an action in the subject's point of view (Sun, Yu, Han, & Liu, 2006b).

$$\text{Probability: } p = P(subject : agent, action)$$
$$\text{Entropy: } H(p) = -plog_2(p) - (1-p)log_2(1-p)$$
$$\text{Trust: } T(subject : agent, action) = \begin{cases} 1 - H(p) & \text{for } 0.5 < p < 1 \\ H(p) - 1 & \text{for } 0 < p < 0.5 \end{cases}$$

2. Concatenation propagation of trust does not increase trust, rather uncertainty increases through propagation.

$$|T_{AC}| \leq \min(|R_{AB}, T_{BC}|) \text{ where } R \text{ is recommendation trust and A,B,C are agents}$$

3. Multipath propagation does not reduce trust, i.e. multiple recommendations of equal magnitude from different sources should be at least equivalent to the same magnitude recommendations from fewer sources.

$$T_{A_2C_2} \geq T_{A_1C_1} \geq 0 \text{ for } R_1 > 0 \text{ , } T_2 \geq 0$$
$$T_{A_2C_2} \leq T_{A_1C_1} \leq 0 \text{ for } R_1 > 0 \text{ , } T_2 \leq 0$$

4. Trust based on multiple recommendations from a single source should not be higher than from multiple indepen- dent sources; the same applies whenever multiple recommendations are highly correlated.

$$T_{A_2C_2} \geq T_{A_1C_1} \geq 0 \text{ if } T_{A_1C_1} \geq 0$$
$$T_{A_2C_2} \leq T_{A_1C_1} \leq 0 \text{ if } T_{A_1C_1} \leq 0$$

8.2 Bayesian Reputation via Beta Distributions

Given observations either through actual occurrence of actions or first/second-hand reports we use Bayes theorem to create a framework to generate a *belief* about the associated agent's reputation.

$$P(belief|obs) = \frac{P(obs|belief)*P(belief)}{P(obs)}$$

The beta distribution is the canonical probability distribution that describes the probability of the occurrence of probability or any ratio-valued random variable on the support [0,1]. In our case the underlying random variable is the probability that the agent will engage in the action (repay the loan). The beta-distribution is also the conjugate prior for the binomial distribution. The binomial distribution is a two parameter distribution ascribed by θ, the probability of success, and n, the number of trials. Given the observations n, number of trials, and m number of successes, estimated values for θ would follow a beta-distribution.

$$P_\beta(x; \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{(\alpha-1)} 1 - x)^{(\beta-1)}$$

where x is the probability the desired action occurs
α, β are parameters
Γ is the gamma function

To accommodate the edge-case where an agent has no prior observations and no connections on the network we modify P_β to be $P_\beta(x; \alpha+1, \beta+1)$ because $P_\beta(x; \alpha = 1, \beta = 1) = P_{uniform}(0, 1)$.

There are three data-sources from which we will perform Bayesian updates to P_β

- direct observations (in terms of successful vs unsuccessful actions) of the agent
- a direct risk profiling of the agent provided from a central source
- recommendations others on the network make of the agent

We describe the mathematical formalism of each kind of Bayesian update in the following sections (Ganeriwal, Balzano, & Srivastava, 2003).

8.3 Bayesian Update From Direct Observation

Suppose we observe s successes and r failures (the number of trials is $s + d$) for an agent. Then we can update their existing reputation P_β as follows:

$$\alpha_{new} = \alpha + s$$
$$\beta_{new} = \beta + r$$

8.3.1 Bayesian Update From Central Risk Profiling

We assume that

1. the central source creates a risk profile for the agent that also follows beta distribution $\hat{P}_\beta(x; \hat{\alpha}+1, \hat{\beta}+1)$

2. the central source itself has reputation that all nodes observe $\tilde{P}_\beta(x; \tilde{\alpha}+1, \tilde{\beta}+1)$ that is almost a delta function at 1, e.g. $\tilde{\alpha} = 500, \tilde{\beta} = 1$.

Ergo we can perform a Bayesian update of the agent's existing reputation to be P_β as follows:

$$\alpha_{new} = \alpha + \frac{2 * \hat{\alpha} * \tilde{\alpha}}{(\tilde{\beta}+2)(\hat{\alpha}+\hat{\beta}+2)+(2\tilde{\alpha})}$$
$$\beta_{new} = \beta + \frac{2 * \hat{\beta} * \tilde{\alpha}}{(\tilde{\beta}+2)(\hat{\alpha}+\hat{\beta}+2)+(2\tilde{\alpha})}$$

8.3.2 Recommendations From Others

In a similar manner as we update P_β from a central source we can do so from multiple sources where those sources are other agents on the network. Suppose we have a set $K = [1...k]$ of other agents from whom we can receive recommendations about the agent of interest. Then

1. each agent k possesses a risk profile for the agent of interest that follows beta distribution $\hat{P}^k_\beta(x; \hat{\alpha}^k+1, \hat{\beta}^k+1)$

2. we maintain a reputation profile for each agent k, $\tilde{P}^k_\beta(x; \tilde{\alpha}^k+1, \tilde{\beta}^k+1)$

Ergo we can perform a Bayesian update of the agent's existing reputation to be P_β as follows:

$$\alpha_{new} = \alpha + \sum_{k=1}^{K} \frac{2 * \hat{\alpha^k} * \tilde{\alpha^k}}{(\tilde{\beta^k}+2)(\hat{\alpha^k}+\hat{\beta^k}+2)+(2\tilde{\alpha^k})}$$
$$\beta_{new} = \beta + \sum_{k=1}^{K} \frac{2 * \hat{\beta^k} * \tilde{\alpha^k}}{(\tilde{\beta^k}+2)(\hat{\alpha^k}+\hat{\beta^k}+2)+(2\tilde{\alpha^k})}$$

Obviously, there will arise situations where an agent will have both their own direct assessment of another agent's reputation but also would like to amend that information with reputation assessments other agent's may have about the agent of interest. In that situation we simply amend the above equations to include a weighting term $w > 0.5$ as follows:

$$\alpha_{new} = w * \alpha + (1-w) \sum_{k=1}^{K} \frac{2 * \hat{\alpha^k} * \tilde{\alpha^k}}{(\tilde{\beta^k}+2)(\hat{\alpha^k}+\hat{\beta^k}+2)+(2\tilde{\alpha^k})}$$
$$\beta_{new} = w * \beta + (1-w) \sum_{k=1}^{K} \frac{2 * \hat{\beta^k} * \tilde{\alpha^k}}{(\tilde{\beta^k}+2)(\hat{\alpha^k}+\hat{\beta^k}+2)+(2\tilde{\alpha^k})}$$

8.4 Applied to Credit Networks

The above presents a highly parallelisable, recursive Bayesian updating system built entirely on beta distributions that can allow any agent on a network to ascertain the inherent risk of another node on the network. However there are some aspects to this 'Beta Reputation System' that are a bit more subtle:

1. The risk or reputation profile is subjective, i.e. every agent will have their own opinion of another agent's risk. There are actually two kinds of *risk*:
   - Recommendation Risk - this is equivalent to $\tilde{P}_\beta(x; \tilde{\alpha}+1, \tilde{\beta}+1)$, the probability distribution that the signal actually provided by the recommender can be considered trustworthy. From henceforth we will use the term R^β with parameters α^R, β^R to describe this.
   - Default Risk - this is the equivalent to $\hat{P}_\beta(x; \hat{\alpha}+1, \hat{\beta}+1)$, the actual probability that the agent will perform the desired action (repay). From henceforth we will use the term P^β with parameters α^P, β^P to only describe this.
2. For the beta distribution to cleanly map to risk we must vary its parameters dependent on the loan size; furthermore the network is always dynamic, and any risk assessment will quickly age. Therefore risk must be computed at the time the loan is requested.
3. If allowed to iterate, via Expectation Maximization we should arrive on a converged set of values for either P^β or R^β that are consistent across all subjects. However there is no compelling reason at the moment to do this for the Default Risk, P^β but we should do this for R_β if possible using the same recursive procedure outlined below. We would first simply update R^β until we have global values for R_i^β and then plug into recursively solve for subjective values P_{ij}^β. Ergo P^β is objective vs R^β is objective.

8.4.1 Inputs to Algorithm

We will use the following terminology in the proceeding sections:

- i - agent assessing default risk
- j - agent for which the default risk is being assessed
- $k \in \{K\}$ - agents trusted by i to provide a report of j's default risk
- $h \in \{H\}$ - agents directly connected to node j who can directly observe j's default risk

The following variables and functions will be given:

- w - global variable trading off direct vs indirect reports of risk
- $loanSize_j$ - loan size requested by the borrower
- $ptflSize_{\{i,j,k,h\}}$ - portfolio size of all agents
- $f_{scale}(loanSize_j, ptflSize_h, crdLine_{hj}) = [\alpha_{hj}^P; \beta_{hj}^P]$ - global function that - for j's direct trustors $h \in \{H\}$ - will scale their assessment of j's default risk dependent on the loan request size. This will only be called at the time the borrower requests a loan and provides the initial priors for the borrower's immediate trustors.
- $P_{*j}^\beta(\alpha_{*j}^P, \beta_{*j}^P)$ - risk assessment of the central authority for the borrower's loan.

- Every agent maintains their own table of the Recommendation Risk R^β of their connections as follows:

Node ID	Trustee/Trustor	Credit Line	$R^\beta(\alpha^R, \beta^R)$
1	Trustee	$crdLine_1$	$[\alpha_1^R, \beta_1^R]$
...	Trustor	$crdLine_{...}$	$[\alpha_{...}^R, \beta_{...}^R]$
m	Trustor	$crdLine_m$	$[\alpha_m^R, \beta_m^R]$

- Every agent maintains their own table of the repayments of any of their connections with outstanding loans as follows:

Node ID	Trustee/Trustor	Loan Amount	Amount Repaid	Amount with any Default
1	Trustee	$loanAmt_1$	$repaid_1$	$default_1$
...	Trustor	$loanAmt_{...}$	$repaid_{...}$	$default_{...}$
m	Trustor	$loanAmt_m$	$repaid_m$	$default_m$

8.4.2 Algorithm Start

Step 1: Agents update their Recommendation Risk R^β given the successes/failures of the recommendations of their connections.

1. A scaling parameter κ is set by the network

2. Each agent computes

$$s = \tfrac{1}{\kappa} \textstyle\sum_{m \subseteq Trustors} repaid_m$$
$$r = \tfrac{1}{\kappa} \textstyle\sum_{m \subseteq Trustors} default_m$$

3. Each agent broadcasts (r, s) to their connections

4. Each receiving agent updates the R^β of their connections as follows:

Node ID	Trustee/Trustor	Credit Line	$R^\beta(\alpha^R, \beta^R)$
1	Trustee	$crdLine_1$	$[\alpha_1^R + s_1, \beta_1^R + r_1]$
...	Trustor	$crdLine_{...}$	$[\alpha_{...}^R + s_{...}, \beta_{...}^R + r_{...}]$
m	Trustor	$crdLine_m$	$[\alpha_m^R + s_m, \beta_m^R + r_m]$

Step 2: Agents h directly connected to the borrower j compute their subjective view of j's default risk P^β:

$$[\alpha_{hj}^P; \beta_{hj}^P] = f_{scale}(loanSize_j, ptflSize_h, crdLine_{hj})$$

8.4.3 Recursion

We have a set of agents $\{i, k, h\} \in V \subseteq \{All\ Agents\}$ who need to assess j's risk. Each node i recursively solves via depth-first search as follows:

1. retrieve $V_k = \{v_1, ..., v_K\}$ the set of all neighbouring nodes in propagation direction (trustors)
2. update V_k to remove all nodes previously visited $\longrightarrow$ break if V_k now empty
3. iterate over each V_k : we know $R^{\beta}_{i,k}$ but need to know $P^{\beta}_{k,j}$
   (a) if agent amongst trusted seeds set α^P and β^P to trusted values
   (b) retrieve P^{β}_{kj} if computed by k and store in table
   (c) else recurse from Step 1 if not computed
   (d) record node k as visited

   When the for-loop ends, we have a table as follows:

Node ID	$P^{\beta}_{.j}(\alpha^P_{.j}, \beta^P_{.j})$	$R^{\beta}_{i.}(\alpha^R_{i.}, \beta^R_{i.})$
$k = 1$	$[\alpha^P_{1,j}, \beta^P_{1,j}]$	$[\alpha^R_{i,1}, \beta^R_{i,1}]$
...	$[\alpha^P_{...} + s_1, \beta^P_{...}]$	$[\alpha^R_{...}, \beta^R_{...}]$
$k = K$	$[\alpha^P_{K,j}, \beta^P_{K,j}]$	$[\alpha^R_{i,K}, \beta^R_{i,K}]$

4. Now the node i itself must compute and store it's own value for $P^{\beta}_{i,j}(\alpha^P_{i,j}, \beta^P_{i,j})$:

$$\alpha_{i,j} = \sum_{k=1}^{K} \frac{2 * \alpha^R_{i,k} * \alpha^P_{k,j}}{(\beta^R_{i,k}+2)(\alpha^P_{k,j}+\beta^P_{k,j}+2)+(2\alpha^R_{i,k})}$$
$$\beta_{i,j} = \sum_{k=1}^{K} \frac{2 * \alpha^R_{i,k} * \beta^P_{k,j}}{(\beta^R_{i,k}+2)(\alpha^P_{k,j}+\beta^P_{k,j}+2)+(2\alpha^R_{i,k})}$$

8.4.4 Final Computations

We also store the result in its entropy form -

$$H(P^{\beta}_{i,j}) = ln(\mathbf{B}(\alpha_{i,j}, \beta_{i,j})) - (\alpha_{i,j} - 1)\Psi(\alpha_{i,j}) - (\beta_{i,j} - 1)\Psi(\beta_{i,j}) + (\alpha_{i,j} + \beta_{i,j} - 2)\Psi(\alpha_{i,j} + \beta_{i,j})$$

- where $\mathbf{B}$ is the beta function (not beta distribution) and Ψ is the digamma function.

9 Future Research

Whilst the authors have created a proof of concept that demonstrates that the underlying concept is sound, there are a number of areas in which further rigour is required, particularly regarding the underlying mathematics and computational optimisation thereof. In order of importance, the following topics must be addressed to move to the next stage of development:

1. Incorporating a ‘patience’ term within the optimisation, factoring in the opportunity cost of lending funds at the first chance, rather than waiting for superior opportunities.

2. A more robust formalisation of the optimisation problem itself: ideally in terms of a modified Weibull distribution (Weibull, 1951) in order to describe the distribution of portfolio returns inclusive of time.

3. A neural-network based formalization of the problem to incorporate memory and learning to produce ideal agent goal-states

4. We are open to alternative ways to solve the problem at hand. There is a significant body of literature on credit networks, but in almost all cases interest rates and collateral are treated as exogenous: they do not arise as a consequence of an agents’ perceived risk. Further, credit networks as typically studied are more concerned with contagion, as collateral is usually proffered in the form of other risky assets. In our formulation we do not allow for this, eliminating contagion risk at the expense of capital efficiency.

References

Aleem, I. (1990, September). Imperfect Information, Screening, and the Costs of Informal Lending: A Study of a Rural Credit Market in Pakistan. *World Bank Economic Review*, *4*(3), 329-349. Retrieved from https://ideas.repec.org/a/oup/wbecrv/v4y1990i3p329-49.html

Bachmann, A., Becker, A., Buerckner, D., Hilker, M., Kock, F., Lehmann, M., ... Funk, B. (2011, 08). Online peer-to-peer lending – a literature review. *Journal of Internet Banking and Commerce*, *16*.

Carmichael, D. (2014). Modeling Default for Peer-to-Peer Loans. Retrieved from http://dx.doi.org/10.2139/ssrn.2529240

Dehejia, R., Montgomery, H., & Morduch, J. (2012). Do Interest Rates Matter? Credit Demand In The Dhaka Slums. *Journal of Development Economics*, *97*(2), 437 - 449. Retrieved from http://www.sciencedirect.com/science/article/pii/S0304387811000587

Emekter, R., Tu, Jirasakuldech, & Lu. (2015, 01). Evaluating Credit Risk And Loan Performance In Online Peer-to-Peer (P2P) Lending. *Applied Economics*, *47*. doi: 10.1080/00036846.2014.962222

Fenwick, M., Vermeulen, E., & McCahery, J. (2017). *Fintech And The Financing Of Entrepreneurs: From Crowdfunding To Marketplace Lending* (Vol. Working Paper No. 369; Tech. Rep.).

Ganeriwal, S., Balzano, L., & Srivastava, M. (2003, 01). Reputation-based framework for high integrity sensor networks. *TOSN*, *4*. doi: 10.1145/1029102.1029115

Jones, M. (2009, 01). Kumaraswamy's Distribution: A Beta-Type Distribution With Some Tractability Advantages. *Statistical Methodology*, *6*, 70-81.

Kamvar, S. D., Schlosser, M. T., & Garcia-Molina, H. (2003). The Eigentrust Algorithm for Reputation Management in P2P Networks. In *Proceedings of the 12th International Conference on World Wide Web* (p. 640–651). New York, NY, USA: Association for Computing Machinery. Retrieved from https://doi.org/10.1145/775152.775242

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, *7*(1), 77–91. Retrieved from http://www.jstor.org/stable/2975974

Mendoza, R. U. (2011, January). Why Do The Poor Pay More? Exploring The Poverty Penalty Concept. *Journal of International Development*, *23*(1), 1-28. Retrieved from https://ideas.repec.org/a/wly/jintdv/v23y2011i1p1-28.html

Mills, K., & McCarthy, B. (2014). The State of Small Business Lending: Credit Access During the Recovery and How Technology May Change the Game. *Harvard Business School General Management Unit Working Paper No. 15-004*. Retrieved from http://dx.doi.org/10.2139/ssrn.2470523

Mock, W. B. T. (2011). Pareto Optimality. In D. K. Chatterjee (Ed.), *Encyclopedia of Global Justice* (pp. 808–809). Dordrecht: Springer Netherlands. Retrieved from https://doi.org/10.1007/978-1-4020-9160-5_341 doi: 10.1007/978-1-4020-9160-5_341

Morduch, J. (2000, April). The Microfinance Schism. *World Development*, *28*(4), 617-629. Retrieved from https://ideas.repec.org/a/eee/wdevel/v28y2000i4p617-629.html

Morse, A. (2015, January). *Peer-to-Peer Crowdfunding: Information And The Potential For Disruption In Consumer Lending* (Tech. Rep. No. 20899). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w20899 doi: 10.3386/w20899

Serrano-Cinca, C., & Gutiérrez-Nieto, B. (2014, 02). Microfinance, the long tail and mission drift. *International Business Review*, *23*. doi: 10.1016/j.ibusrev.2013.03.006

Serrano-Cinca, C., Gutiérrez-Nieto, B., & López-Palacios, L. (2015, 10). Determinants of Default in P2P Lending. *PLOS ONE*, *10*. doi: 10.1371/journal.pone.0139427

Stiglitz, J. (1969). Distribution of income and wealth among individuals. *Econometrica*, *37*(3), 382-97. Retrieved from https://EconPapers.repec.org/RePEc:ecm:emetrp:v:37:y:1969:i:3:p:382-97

Sun, Y., Yu, W., Han, Z., & Liu, K. J. R. (2006a, 02). Information Theoretic Framework Of Trust Modeling And Evaluation For Ad Hoc Networks. *IEEE Journal on Selected Areas in Communications*, *24*, 305-317. doi: 10.1109/JSAC.2005.861389

Sun, Y., Yu, W., Han, Z., & Liu, K. J. R. (2006b, 02). Information theoretic framework of trust modeling and evaluation for ad hoc networks. *IEEE Journal on Selected Areas in Communications*, *24*, 305-317. doi: 10.1109/JSAC.2005.861389

Thyfronitis Litos, O. S., & Zindros, D. (2017, Jan). Trust Is Risk: A Decentralized Financial Trust Platform. In (p. 340-356). doi: 10.1007/978-3-319-70972-7_19

Weibull, W. (1951, Jan). A Statistical Distribution Function Of Wide Applicability. *Journal of Applied Mechanics*, *18*, 293-297.

Weston, J., & Strahan, P. E. (1996, March). Small Business Lending and Bank Consolidation: Is There Cause for Concern? *Current Issues in Economics and Finance*, *2*(3).

Yunus, M., & Jolis, A. (2003). *Banker To The Poor: Micro-Lending And The Battle Against World Poverty.* New York, NY: PublicAffairs.