Circles - money for a multipolar world

Martin Köppelmann Paul Boes Friederike Ernst

Version 2.2.1

Abstract

We present Circles[1], a currency protocol that provides a completely decentralized and radically accessible complement or alternative to fiat money. It has been designed around the single goal to not put any participant or institution, current or future, at an undue and systematic advantage over another, while remaining attractive as a means of exchange. We believe that money which is freely adopted by people will enjoy lasting stability only through built-in fairness properties.

Circles achieves this by i) making every user an issuer of money, ii) providing a predictable bottom-up monetary policy that is based solely on the number of participants, while iii) allowing anybody to join without any centralized gatekeeping. The protocol relies on existing social relations between people and groups in place of incumbent political and financial institutions to back the currency, to safeguard against malicious parties and the undesired concentration of monetary control or extraction of profits from money issuance. By leveraging these social structures, Circles not only fosters local autonomy but also creates a viable path toward a global currency system.

Crucially, Circles represents an unaligned form of money—free from the influence of any single nation, corporation, or geopolitical bloc—and is thus ideally suited to thrive in a multipolar world.

This whitepaper provides a detailed account of this new currency protocol, its motivation, and inner workings, and also makes the case for its economic soundness, despite its radical design.

[1]Since its conception by Martin Köppelmann in 2014, Circles' vision has drawn diverse independent teams to contribute to its development - from early theoretical work through the 2020 technical launch by a cooperative and community implementations in Berlin and Bali to the work presented here. All code is open-sourced at https://github.com/aboutcircles. For early presentation of the idea, see for example this article and this article. There is also version of this whitepaper live as a decentralized doc, where we invite interested readers to provide feedback and participate in the discussion on how to improve the protocol and its presentation.

1 Introduction

A currency can be anything - gold coins, debt certificates, entries in a virtual ledger - that a group of people use for transacting with another. It lubricates the economy by facilitating exchange, serving as a store of value as well as as a unit of account (i.e. it allows the comparison of goods by value). Yet, currencies are also a source of political and economic power. The ability to issue money at a cost below face value gives the issuer the ability to impact the distribution of the money supply — who gets how much — and the conditions under which money is accessible, i.e. the monetary policy. Today, banks hold a monopoly over the creation of the national currency whose usage is, fundamentally, enforced through political authority. This money is issued, both by the central bank as well as commercial banks, mostly through the extension of loans. We believe that this monopoly over the issuance of currency is problematic, for the following reasons:[2]

- As with all concentration of power, the centralization of money issuance leads to exploitation by those in charge and creates an obvious target for those seeking undue influence. A transparent non-arbitrary mechanism for money genesis is inherently valuable and serves as a safeguard against such risks.
- Today, newly issued money primarily enters the economy through loans, disproportionately benefiting those with higher ability to take on credit — mainly big financial institutions and enterprises. These institutions benefit because being an upstream recipient of newly created money allows access to the market at conditions that have not yet incorporated the increased money supply — something that is known as *Cantillon* effect[3] — and moreover there are fewer middle men that charge them premiums than for downstream creditors. This mechanism is a historical artifact. A more equitable distribution of the value generated through money issuance, combined with a more radical democratization of its control, is not only technologically feasible but also economically and socially preferable.
- A government's and central bank's monetary policy should serve the economic interests of the people on whom money is imposed as legal tender. However, lack of visibility into the dynamics at the grassroots level and coarse tooling have led to many poor monetary decisions in the past.
- By construction, government-issued money puts those not served by that government at a disadvantage in a myriad of ways. For example, it can link economies without negotiation power to the domestic monetary policy of international powers, as is evident in the case of the US Dollar system. A global world should have an *impartial, unaligned* currency whose value is not tied to the interests or world views of any particular geopolitical bloc, nation or corporation.

The advent of blockchain technology has ushered in a wave of excitement and thinking

[2] In addition to the reasons given here, various authors have elaborated on the potential benefits of private and/or local currencies, such as smoothing of economic cycles. See, for example, Friedrich Hayek's *The Denationalization of Money* (1976) and Silvio Gesell's *The Natural Economic Order* (1929).

[3] See M. Blaug, *Economic Theory in Retrospect* (1997) as well as the original essay by R. Cantillon, *Essay on the nature of trade in general* (1755).

around how to overcome these problems. Cryptocurrencies like Bitcoin are minted in a decentralized manner and participation is open to everyone. However, no cryptocurrency, Bitcoin or otherwise, has yet succeeded in becoming a fully functional currency, due to what we believe are intrinsic design limitations. For instance, in a thought experiment in which Bitcoin becomes the new world currency, over 94% of the total money supply is concentrated in the hands of the 2-3% of the population that hold Bitcoin today. This is because Bitcoin has a finite supply, most of which has already been minted. Switching to such a monetary system is not in the economic interest of the 97+% who do not hold Bitcoin, as it would effectively result in a massive transfer of purchasing power to existing holders. Additionally, Bitcoin lacks a monetary mechanism to adjust its supply in response to the growth of an economy, making it unlikely to achieve the stability in value required to serve as unit of account and practical medium of exchange. Instead, Bitcoin today serves as a store of value, earning it the nickname "digital gold".[4] Still, the promise of the Web3 stack as enabling a world of decentralized governance and finance stands, and, with the ability to design sophisticated currency systems using smart contracts, the question remains: *What does a currency have to look like to be adopted by people in such a multipolar, decentralized world, not by virtue of necessity as today, but by virtue of its attractive properties alone?*

Circles is our answer to this question. It is founded on the principle that the most sustainable currency is one that, by design, ensures no party, present or future, is unfairly advantaged. For us, this implies the following requirements:

- **Universal access**: Participation is open to everyone.
- **Distributed issuance**: In a world in which money is not backed by a single authority, who should own the money creation process and benefit from its value? For us, the only reasonable answer is: everyone. At any given time, the benefits of money creation should be available on equal terms to all users.
- **Fair access to money supply:** A currency's value relies on trust that it will remain valuable in the future. To uphold this trust, every active and honest participant[5] must be able to accumulate significant money holdings under fair, non-coercive conditions, regardless of the distribution of supply at the time they join.
- **Possibility to organize in shared-interest groups:** People and communities know best the economic problems they struggle with and opportunities they have. A currency should allow these communities to organize themselves in order to effectively address these problems and opportunities using monetary mechanisms.

This whitepaper introduces Circles, shows that it meets the above requirements, and makes

[4]The recent paper The distributional consequences of Bitcoin by researchers at the European Central Bank makes a similar case. In a rebuttal paper, Bitcoin proponents state that it is not, in fact, the goal of Bitcoin (anymore) to become a global means of exchange, which to us indicates that they'd agree with our argument above.

[5]We here define an honest users as a user that is using the protocol as intended, and in particular one that uses a single account to create CRC. As the intention of the protocol might not always be clear, we are publishing a *Circles Etiquette* guide alongside this whitepaper in which a number of Do's and Don'ts for Circles are clarified.

the case for its economic soundness as a currency despite its radical departure from money as we know it. In the following, we use "Circles" to refer to the protocol and "CRC" to denote the currency created as part of the protocol.

2 Circles

The Circles protocol is completely defined by the following simple rules:

1. **Universal access:** *Anyone can create an account without any gatekeeping.*[6]
2. **Distributed issuance:** *Each account has the right to create their own, individual CRC at a rate of 1 CRC per hour.*[7]
3. **Demurrage:** *Every CRC gradually loses nominal value at a rate of 7% per year.*[8]
4. **Rule of Trust:** *Accounts can trust one another. If account A trusts account B, then anyone in the network holding CRC created by B can swap them, at any time, for any CRC held by A, at a rate 1:1.*
5. **Groups:** *Anyone can create a group and add accounts as its members, by trusting them. Groups support their own CRC, however these CRC are not issued continuously, instead they can be created from and redeemed against CRC that have been created by its members, at a rate 1:1.*

Before delving into the details, let us briefly describe how those rules map to the requirements above: The first two rules establish the eponymous requirements. The second part of rule 2 concerning the steady issuance rate, together with rule 3, creates a monetary policy that ensures fair access to money supply. It also ensures that early adopters don't receive an outsized benefit, such as with finitely supplied money. The fourth rule, the Rule of Trust, drives the mechanism by which we exclude dishonest parties from exploiting the system. It also enables the transition from the initial multitude of individual, non-fungible currencies created by the honest participants into a quasi-fungible state, leading effectively to the emergence of a single currency, CRC. Since all honest parties are allowed to create tokens of this shared currency at the same rate, rule 2 ensures an equitable distribution of seigniorage. Finally, Groups address our final requirement, to enable shared-interest groups to organize themselves economically. They form a powerful primitive in Circles.

In order to discuss these points in detail, it is helpful to distinguish between the *micro level* and the *macro level* of Circles. The micro level examines the interplay of all the many individual currencies that participants are creating, focusing on the possible states and transitions enabled by the rule of trust. The macro level, on the other hand, looks at the emergent

[6] A natural question is whether AI agents should be allowed to create an account. CRC is a currency that should only be created by people. As such, AI agents and applications are encouraged to exist on Circles, however they should be using the Organization accounts that we discuss below and that cannot create CRC.

[7] In the current implementation, a user has two weeks per CRC to make use of this right. CRC that have not been created for two weeks are forfeited.

[8] Demurrage as the property of money to lose value over time has been part of many alternative currency proposals. In the current implementation, the demurrage gets applied daily, starting from the time a user has the right to create: Every day (specifically at midnight UTC), every CRC balance in the system gets reduced by a factor of $0.93^{1/365.25} \approx 0.9998$.

properties of the system, in which most of the CRC created by honest participants become effectively fungible with another, allowing us to evaluate the "macroeconomic" properties of Circles.

We begin with examining the macro level to substantiate our claim that, assuming all participants are honest and the currency is fully fungible, CRC satisfies the outlined requirements. Subsequently, we address the micro level, where these assumptions are both dropped. We demonstrate how the rules of Circles ensure that, under mild conditions, a fungible currency will still emerge for the honest subset of accounts, excluding malicious accounts from economic activity. Finally, we introduce *Circles Groups* and *Organizations* in the last section.

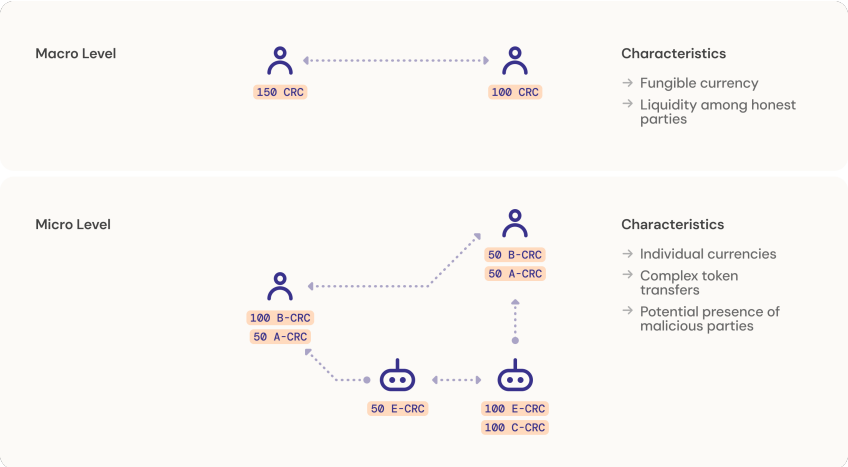

Fig.1: The micro and the macro levels of Circles. Arrows depict trust connections. At the micro-level, we distinguish between different types CRC depending on which account has created them. We also assume that accounts might be malicious and try to game the protocol to their unfair advantage. The rules of Circles are meant to ensure that, at the emergent macro level, from the point of view of honest users the difference between various types of CRC disappears or becomes largely negligible and malicious actors are effectively cut off from the network.

3 The Macro Level

For the purposes of this discussion of the macro level, we assume that all participating parties are honest and that all CRC issued are fungible and valued in the economy. We ask how an economy would look if Circles was to become generally adopted.

Circles in practice

Let us go through a simple example of Circles in practice: Say today you own 10000 CRC. Due to the 7% demurrage, 9300 of these CRC remain today in a year. Moreover, if you continuously create your CRC throughout the year at a rate of 1 CRC per hour, this yields an additional 8479 CRC (resulting from the roughly 8760 hours per year combined with demurrage), so that a year from today your total balance of CRC is 17779 CRC. If, on the other hand, you had started today with 200000 CRC,

then today in a year 186000 of these would remain, so that all together you'd end up with 194479 CRC.

3.1 Fair access to money supply

As stated above, we believe that for a currency to become adopted voluntarily, one important requirement is that every active user should be able to benefit in just the same way from the ability to issue money and accumulate significant money holdings over the course of their participation in the economy. Moreover, this ability should be independent of the existing distribution of money at the time a user joins the system. In particular, it should be impossible for any existing users, no matter how wealthy, to prevent new users from building up money holdings and purchasing power. This fair access is crucial for the currency to retain value over longer periods of time: *If a currency cannot ensure that it will be attractive to future adopters, then this anticipated lack of value tomorrow will undermine its value today.*

Circles satisfies this requirement in the following strong sense: In every period, newly created CRC constitute a significant portion of the total supply and are distributed equally among active users. This allows new-joiners to build a share of the money supply quickly, regardless of the distribution when they joined: Consider a user who enters the system at some time t and let $m_{t'}$ denote the percentage of the total money supply that they have issued at some later time $t' \geq t$. Then it is easy to show that, for a constant population N,

$$m_{t'} = \frac{1}{N}\left(\frac{1}{\gamma_{t,t'}+1}\right), \quad (1)$$

where $\gamma_{t,t'} \propto 0.93^{(t'-t)}$ is a factor that diminishes exponentially (albeit slowly) in the time difference $t'-t$ (in units of years). This shows that, regardless of the existing money supply and distribution at the time somebody joins or the actions of others, a user will contribute a percentage of the money supply that is roughly inversely proportional to the size of the population. This result also holds true for populations that vary over time, as we discuss in the Appendix. Hence, Circles provides fair access to the money supply in the sense of having all users benefit equally from the issuance of money

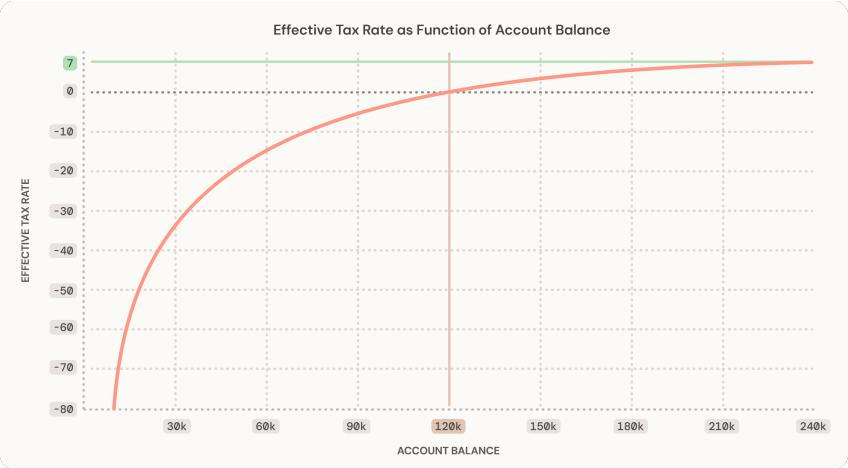

Fig. 2: The effective tax in Circles: This graph plots the inverse annual rate of relative change of the CRC balance, as a function of its initial balance, when a user is only creating new currency and not participating in any transactions : A user starting off with 30k CRC will see their balance grow by approximately 35%, equivalent to an effective tax of -35%. At the special point of 120804 CRC, the balance remains exactly unchanged, resulting in an effective tax of 0%. Balances above this threshold will overall lose in value. As the balances become larger, the effective tax converges to 7%, i.e. the demurrage completely dominates the negligible influx of newly created CRC. Note that the effective tax diverges as the balance goes to zero (which is not shown on this graph).

Another way to look at this dynamic is shown in Figure 2, where the (inverse) rate of relative change in the CRC balance for a single account is plotted as a function of its initial balance. This rate can be understood as an effective tax resulting from the combined effects of continuous token issuance and the 7% annual depreciation of existing CRC.[9] At 120804 CRC, these opposing forces balance precisely. Accounts with balances below this threshold experience a net increase (negative tax), while those above see a net decrease (positive tax). Consequently, without transactions, each account's balance will naturally converge to this stable value. This means that, over time, the total share of the money created by any group will be proportional to its size.

Why 7% demurrage

The demurrage rate and the issuance rate are the only two free parameters in Circles. The challenge is to choose them such that CRC are sufficiently stable as an intermediate store of value to encourage people to accept and hold them, while the effects captured in Equation (1) also ensure the currency remains attractive for new participants. Our choice of 7% annual demurrage and 1 CRC/h issuance rate rests on two arguments, one empirical and one theoretical: Empirically, the long term average increase in the money supply of the US Dollar has been slightly below 7% since the 1960s (as measured in the so-called M2 supply). As discussed in the box on the inflationary view of Circles below, this is economically equivalent

[9] We should note, though, that this effective tax is not a money transfer in the sense that the CRC removed from one account are added to another account.

to Circles with a 7% demurrage rate. Theoretically, ensuring intergenerational fairness is a central part of Circles. The above choice of values means that accounts approximately reach the maximum level of 120804 CRC after creating CRC for 80 years, which is roughly the average life expectancy in industrialized countries (see Fig. 3). As such, the dynamics of the currency are gauged towards the natural unit of a user's average lifetime.

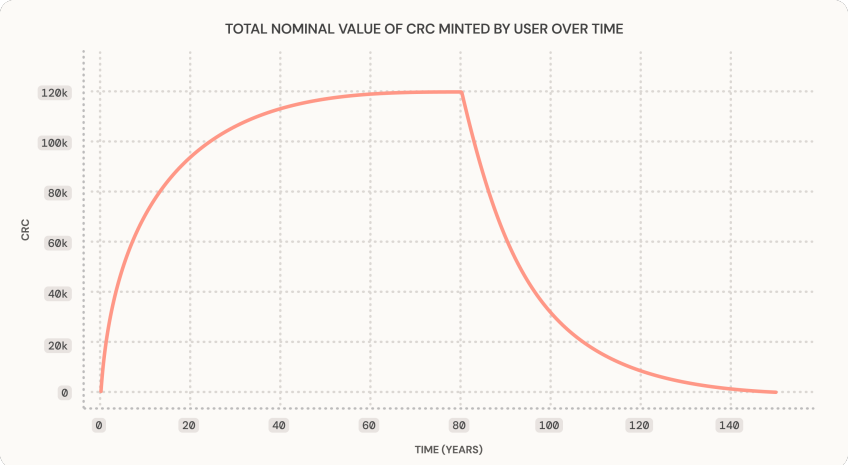

Fig. 3: This graph illustrates the total amount of CRC in circulation (i.e. the total nominal value) that have been created by a user over their lifetime (and beyond) over time, assuming they pass away at the age of 80. By this point, they have approximately reached the stable balance of 120804 CRC. After creation stops, the amount of their CRC in circulation decreases exponentially, with a half-life of approximately 10 years.

Finally, it may seem that this behavior of balances converging over time is some kind of radical feature embedded in the currency. However, this interpretation is misleading. Every currency exposed to inflation undergoes a similar wealth redistribution effect, a concept extensively analyzed by economists such as Keynes, Friedman, or Piketty.[10] In fact, there is an alternative way to understand Circles, the *inflationary view*, which reveals this dynamic within a more familiar economic framework (see below).[11]

A noteworthy implication of demurrage is that, despite the constant creation of new money, it is actually perfectly possible that prices in a Circles-based economy stabilize in the long run, even if the economy grows at a steady rate, or even go down. That is, while the distribution of purchasing power in the economy constantly shifts in favor of currently active

[10] See e.g. Ch. 6 of J.M. Keynes, *The Economic Consequences of the Peace* (London: Macmillan & Co., 1919); M. Friedman, *Money Mischief: Episodes in Monetary History* (New York: Harcourt Brace & Co., 1992). T. Piketty, *Capital in the Twenty-First Century* (Cambridge, MA: Harvard University Press, 2014); See also M. Doepke and M. Schneider, *Inflation and the Redistribution of Nominal Wealth* (Journal of Political Economy, 114:6, 2006)

[11] In an influential paper, Kocherlakota argues that money can be viewed as a mechanism for tracking past economic activity and achievements. He demonstrated that any state achieved by an economy using money could equally be achieved if, instead of money, there was a central memory ledger visible to all participants, recording past economic activities. In this analogy, we can interpret money with a wealth redistribution effect as a form of *leaking memory*: The wealth and accomplishments of former generations matter, but just like their monuments erode or are replaced over time, so the value of their money slowly fades.

members, just like in the case of money experiencing inflation, the prices of goods can remain unchanged, even over decades. To support this, in the Appendix we formulate a so-called *overlapping-generations model* — one of the foundational models in monetary economics — of a Circles based economy and show that any steady-state competitive equilibrium in this economy with constant per-capita output will have prices that stabilize over time.

The inflationary view of Circles

There is an alternative way of formulating rules 2 and 3 for Circles which introduces a purely nominal change into the currency. We refer to this as the *inflationary view* of Circles. In this perspective, no demurrage applies to coins, however, the amount of CRC that a user can create increases over time. On the surface, the two rule Groups appear different: For instance, the total money supply in the inflationary view grows unboundedly even in a finite economy, whereas Circles' standard rules imply a capped money supply in finite economies. However, in terms of real economic outcomes, the two views are equivalent. Specifically, equation (1) remains unchanged. The inflationary view of Circles resembles the behavior of money as we know it. For instance, inflation in fiat currencies reduces the real purchasing power of money earned in the past, similar to how demurrage affects older CRC holdings. In both cases, increasing wages or CRC creation allows current members to maintain stable purchasing power within the economy.

3.2 Distributed issuance

As discussed above, we consider the radical decentralization of money issuance essential for an equitable economy, as it would ensure that no subset of actors benefit exclusively from the advantages of controlling the money supply. These advantages are, on the one hand, deciding whom to give money to and under what conditions and, on the other hand, the ability to profit from the creation of a token at a cost below its face value. In the case of governments, this second advantage is actually known as *seigniorage*[12], which is defined as central bank revenue from the provision of the national currency, be that through the printing of cash or through interest payments that it receives on its loans. If we generalize the definition of seigniorage to mean simply revenue that an issuer of money receives from the process of issuance, then in Circles the distribution of issuance immediately implies a complete distribution of the seigniorage across the whole active user base. This is a central feature of Circles. An important question, then, is how much purchasing power this seigniorage actually provides. In other words: How much will people be able to buy from the CRC that they create? Although the answer depends on variables like the value of CRC in the economy, balance distribution, and demand patterns, a straightforward calculation

[12] See *Seigniorage* (2017) by J. Reich for an excellent overview of the history and current institutional relevance of the term.

suggests that the contribution of the seigniorage to the average money demand will be relatively small: Let z_t, denote the seigniorage per user (which is the same for all users under the assumption that all CRC are fungible) and $\bar{d}_t$ be the average demand for money per user, both over period t. Assuming a constant annual growth rate g in the number of participants, we generally expect, and show in the Appendix, that

$$\frac{z_t}{\bar{d}_t} \approx \frac{0.07 + g}{V(1+g)}.$$

Here, V denotes the *velocity* of money, which is a measure of how quickly a currency circulates within the economy. This fraction basically describes the percentage of an average user's expenses (say in a given month) that they can cover using the CRC that they created (during that month). Hence, for Circles, with an annual population growth rate of 2%, and at a velocity of around 1 (which is the both approximately the average velocity of major global currencies and also rougly that of the Euro), we estimate that seigniorage will contribute around 9% to a user's average money demand, in the long run. This insight also clarifies Circle's relationship to Universal Basic Income (UBI): Although Circles has previously positioned itself as a form of UBI, it becomes clear that the real economic contribution of created money to an individual's spending is notably lower than what is typically expected of a UBI, which would be closer to 25-30% or even higher. However, for individual users and in times of significant population growth, such as for example an initial phase of adoption, the real value of seigniorage can be much higher.

Moreover, we believe that distributed seigniorage could have a substantial compounding effect on the distribution of balances – and, by extension, on social inequality — in a society that adopts Circles. This is because the distribution of seigniorage dampens the capital compounding effect, whereby having more money makes it easier to accumulate even more.

Financing and lending in Circles

In a world that adopts Circles, we expect many existing financial products to continue to exist, however the business models of operators and the distribution of profits might differ significantly. Today, banks can create money by extending loans. The extent to which they make use of this ability depends on multiple factors, some of which are regulatory and some of which are market-driven. As such, money creation (and money destruction when loans are paid back) is endogenous. CRC creation and destruction, in contrast, is not endogenous. Consequently, in a Circles based economy, banks appear primarily as intermediaries that can only lend out deposits they hold and that charge for their services to assess the creditworthiness of borrowers. As a result, we believe that banks will stand in more competition for depositors' CRC, offering higher interest rates to depositors, which in turn reduces their spread. In effect, we can interpret this as the result of splitting the current profits into a "service"-component and a "rent"-component, where the first is a service fee and the second is a rent banks extract for their effective monopoly

over money creation. The reduced spread and increased deposit interests then essentially mean that the "rent" component is now passed on to depositors, while the "service" component remains with the bank.

As a simple example of how lending with CRC could work, consider Alice, who wants to start a business but is out of funds. Today she borrows 10000 CRC from Bob and over the course of a year makes revenue that results in her account holding exactly 11000 CRC today in a year (with ongoing demurrage factored in). Alice repays Bob his original 10000 CRC back and keeps a profit of 1000 CRC. For Bob, this loan is an investment with an effective annual interest of 7%, as his 10000 CRC would otherwise have depreciated to 9300 CRC.

4 The Micro Level

In the preceding section, we assume that all accounts belong to honest human agents and that all created CRC function as a fully fungible currency. These are expected emergent properties of the system as a whole. However, at the micro level—where the rules are directly enforced—these assumptions no longer apply. In this section, we examine the dynamics at this foundational level to understand the specific conditions under which the macro-level properties emerge.

4.1 Trust and trust-based swaps

Recall that, according to the second rule of the Circles protocol, CRC created by different accounts are inherently non-fungible across accounts (CRC created by the same account are fungible with one another). Thus, for two participants, Alice and Bob, we distinguish between Alice-CRC and Bob-CRC. The interaction between these individual currencies is governed by the trust relationships between accounts. By trusting Bob, Alice does two things:

1. She commits to accepting Bob-CRC as valid payment for any services and goods that she might offer for CRC.
2. She makes her funds effectively exchangeable with Bob-CRC, through the fourth rule, the rule of trust, which we here restate:

Accounts can trust one another. If account A trusts account B, then anyone in the network holding CRC created by B can swap them, at any time, for any CRC held by A, at a rate 1:1.

To understand this rule, let's consider a simple example (see Fig. 4). Suppose Alice trusts Bob, and Alice holds 5 Alice-CRC, while Carol holds 5 Bob-CRC. Carol can directly swap her Bob-CRC with Alices's Alice-CRC without needing additional permission from Alice, as this permission was given when Alice established trust with Bob. Furthermore, this exchange does not depend on whether Alice trusts Carol or Bob trusts Carol. The same would apply, at the same exchange rate, if Alice held any other type of CRC instead of their own.

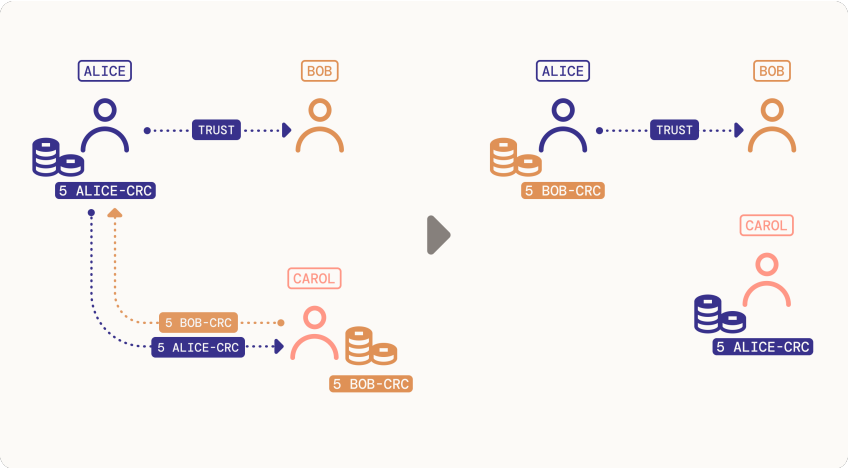

Fig. 4: A basic example of the rule of trust. If Alice trusts Bob, any party in the network holding Bob-CRC can exchange them for any CRC that Alice holds, both those she created and any others in her possession.

The rule of trust directly ties a user's wealth to their trust decisions: If Alice trusts Bob and Bob turns out to be dishonest leading to a loss of purchasing power of Bob-CRC, Alice risks her own CRC wealth. This is because a) she might have exchanged goods of real economic value against Bob-CRC, which are now useless and b) anybody in the system can exchange any CRC that Alice holds for Bob-CRC, depleting her purchasing power.

Trusting is therefore a serious responsibility and users should choose who they trust carefully, trusting — in their own best interest — only accounts that they are confident won't be unmasked as exploiting the system. We therefore expect trust relationships in the system to mirror people's personal social connections. It is the beauty of the Circles protocol that despite trusting only those they find most trustworthy, users end up being able to conduct business also with complete strangers, as we discuss now.

4.2 Transitive transfers

In the following section, we'll explore how payments between third parties function behind the scenes; however, keep in mind that all of these mechanics are fully abstracted away from the user, who only needs to curate their own trust connections.

As stated above, users in Circles commit to accepting CRC by accounts they trust in exchange for goods and services. As such, a key factor in ensuring agents can pay one another across the network is their ability to obtain trusted CRC from users that don't trust them directly. Fortunately, through iterative use of the rule of trust, users can access trusted CRC from across the entire network, by means of *transitive transfers*. Consider the following example to illustrate this process:

1. Bob trusts Alice
2. Alice trusts Carol

3. Alice has 5 Alice-CRC
4. Carol has 5 Carol-CRC
5. Carol wants to pay Bob 5 CRCCarol gives 5 Carol-CRC to Alice and takes 5 Alice-CRC from her.
6. Carol gives 5 Alice-CRC to Bob

In this example, Carol wants to pay Bob 5 CRC, perhaps as payment for some service. She can do this by obtaining 5 CRC that Bob trusts. Although Carol doesn't have these CRC herself, she can get them from Alice, who trusts her, by using the rule of trust. This sequence of swaps results in an effective flow of CRC from Carol to Bob, whose direction is opposite to the path of trust that connects Bob and Alice, as shown in the following figure.

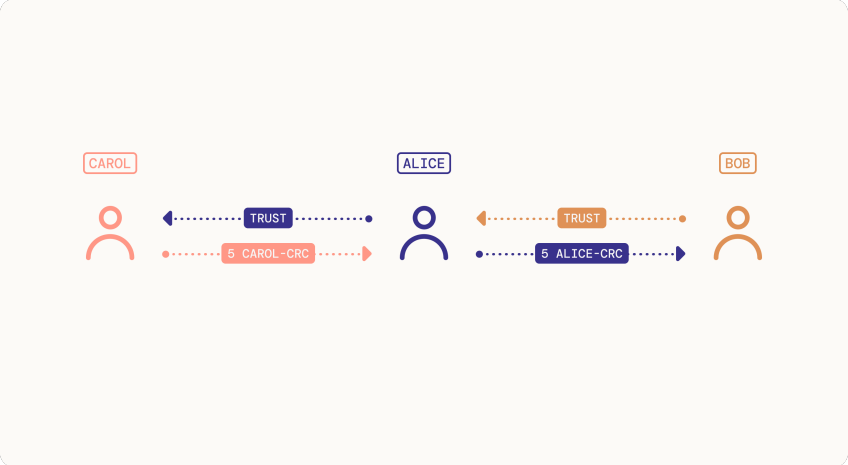

Transitive transfers can also involve multiple "hops" across trusted connections, as shown in the next example.

1. Bob trusts Alice
2. Alice trusts Carol
3. Carol trusts Dave
4. Alice has 5 Alice-CRC
5. Carol has 5 Carol-CRC
6. Dave has 5 Dave-CRC
7. Dave wants to pay Bob 5 CRC
8. Dave gives 5 Dave-CRC to Carol and takes 5 Carol-CRC from her
9. Dave gives 5 Carol-CRC to Alice and takes 5 Alice-CRC from her
10. Dave gives 5 Alice-CRC to Bob

This sequence of applications of the Rule of Trust results in the following effective flow:

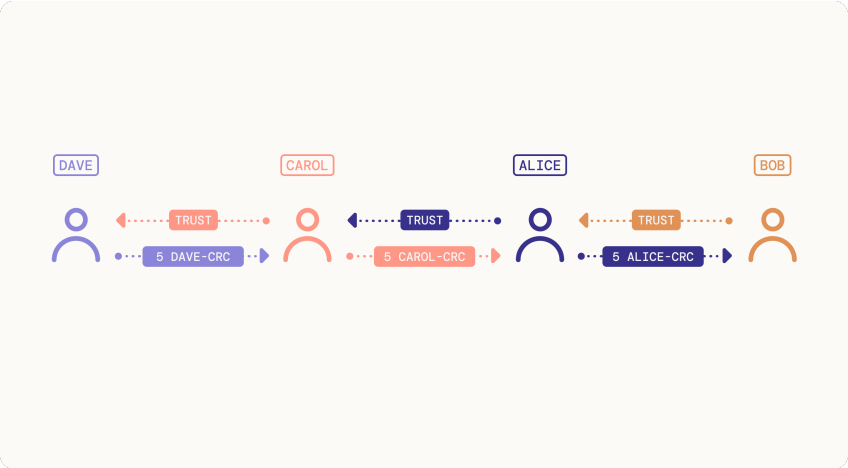

Note that, although these examples involve transferring different types of CRC along a path, the transaction appears at the macro level as the straightforward transfer of 5 CRC from sender to receiver. Additionally, the total amount of CRC holdings of each intermediate party remains constant throughout the process. This demonstrates an important general property of Circles:

Conservation of Trusted Balance

In Circles, each user's total balance of trusted CRC is protected and cannot be reduced by the actions of others.

The example transaction above was only possible because a chain of swaps existed, allowing the sender to obtain trusted CRC from intermediate parties who held sufficient CRC balances themselves.[13] As we will discuss in more detail below, we expect such chains of swaps to exist between any two honest agents in the network, based on the well-studied connectivity properties of social networks that are expressed in the famous "six degrees of separation". Thus, *the emergence of the effective global fungibility of CRC at the macro level fundamentally relies on the social network formed by its users*.

While we avoid technicalities in this text, it is helpful to introduce some basic notation around transitive transfers for the following sections: Let S denote the *state* of the Circles network at a given time, comprising the set of all existing accounts N, the trust relationships between them, and the balances B that each account holds of various individual CRC currencies.

For two states, we write $S \rightarrow_{\tilde{N}} S'$ if there exists a sequence of transitive transfers that can be achieved by accounts in some set $\tilde{N}$ that take an initial state S to a final state S'. Moreover, for any two subgroups of accounts $N_s, N_r \subseteq N$, we write $B(N_s \rightarrow N_r|S)$ to denote the total amount of CRC trusted by at least one account in N_r, that are held by accounts in

[13] Note also that this path depends not only on the presence of trust connections but also on the available balances of each participant. For instance, it is not sufficient that there exists a chain of trust from seller to buyer, since users might not hold sufficient amounts of their own currency.

N_s in state S.

We can then define

$$T(N_s \to N_r|S) := \max_{S':S\to_{N_s}S'} B(N_s \to N_r|S')$$

That is, $T(N_s \to N_r|S)$ is the maximal achievable amount of CRC, trusted by at least one account in N_r, that accounts in N_s can obtain by means of transitive transfers from an initial state S. We'll refer to this as the *transferrable trusted balance of N_s for N_r starting from S*. For example, if in state S there are only two accounts, Alice and Bob, that each hold 10 of their own CRC (and nothing else), and Alice trusts Bob but Bob does not trust Alice, then $T(A \to B|S) = 0$ while $T(B \to A|S) = 10$.

4.3 Resistance against malicious parties

At this point, some readers may wonder why Circles doesn't simply issue fungible CRC like a traditional currency, without relying on the trust rule. The answer is straightforward: This approach enables universal access to Circles without a centralized gatekeeper, while distributing issuance rights across the user base.

To see this, note that the absence of gatekeeping or KYC mechanisms in principle allows users to create several accounts. If all CRC created by these accounts would be fungible and have the same value, then in such a way dishonest users could secretly multiply the rate at which they create CRC compared to honest users — those that use only a single account. Such behavior would drive up prices and allow these users to increase their purchasing power at the expense of others.

The rule of trust is designed the mitigate this risk. It limits the influence of malicious users by restricting their economic participation to the degree that others in the network trust their accounts. To illustrate this, consider the following example:

1. Bob trusts Alice
2. Alice makes a fake account FakeAlice and trusts that account
3. Alice has 5 AliceCRC
4. FakeAlice has 5 FakeCRC
5. Alice wants to buy something worth 5 CRC from Bob using FakeCRC
6. FakeAlice gives 5 FakeCRC to Alice and takes 5 AliceCRC from her
7. FakeAlice gives 5 AliceCRC to Bob

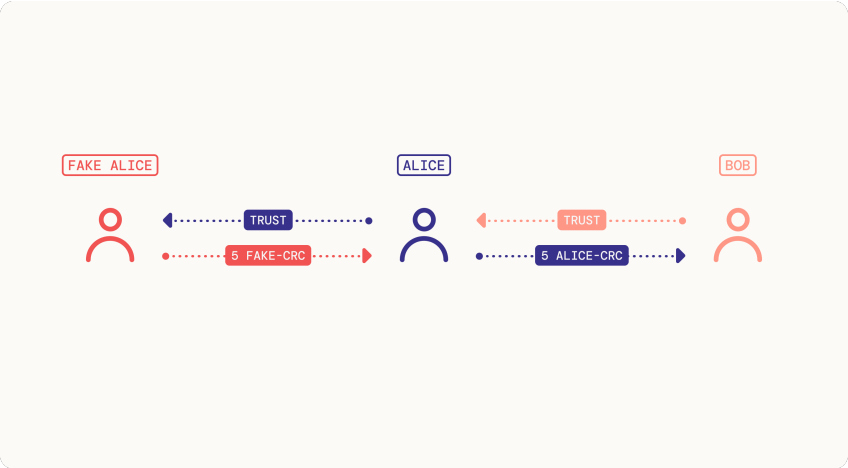

In this example, Alice holds a second, fake account, allowing her to create 2 CRC per hour in total. Bob, however, only trusts Alice's primary account, which gives him limited exposure to her network. Alice can therefore introduce 1 CRC per hour into the economy through Bob, while the coins created by her second, fake account, remain immobilized and effectively useless. Thus, Alice's "effective" issuance rate remains at 1 CRC/h.

This example illustrates a more general property of Circles:

Relative Sybil resistance

Let M be a set of accounts controlled by a malicious party, F be a set of accounts that trust at least one account in M (the "fooled" accounts) and R be the remainder of the network. Then, if accounts in M initially hold only their own funds, the transferrable trusted balance of M for R is limited by the amount of CRC trusted by accounts in R that are held by accounts in F,

$$T(M \to R|S) \leq B_T(F \to R|S).$$

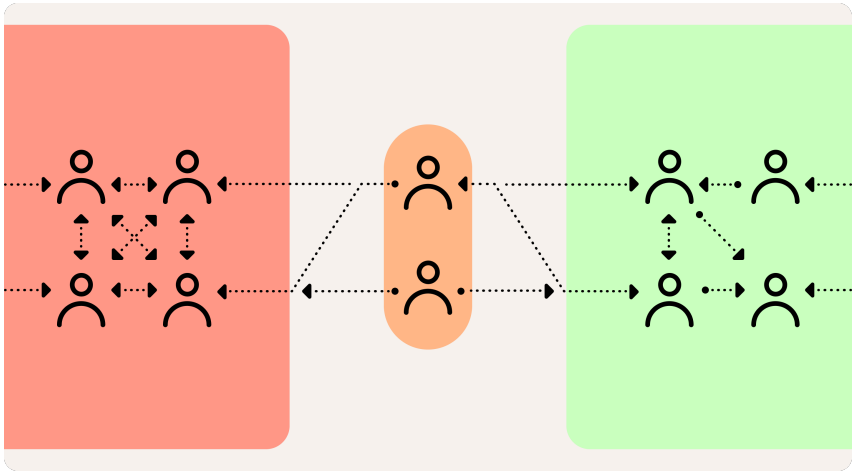

Fig. 5: Visual depiction of Relative Sybil resistance. The influence of a network of malicious nodes (red cluster) on the rest of the network (green cluster) is constrained by the total holdings of trusted CRC held by boundary nodes (orange cluster), which connect the green and red clusters.

Since nodes in M need trusted CRC to impact the economy of R, the above property limits their impact by the degree to which M is trusted by outside accounts that connect M to R. We call this property *Relative Sybil resistance* in reference a type of cyber attack known as *Sybil attack*, in which an attacker creates several pseudonymous identities.

Relative Sybil resistance is Circles' primary mechanism to safe-guard against users creating multiple accounts. Its driving force is the distribution of knowledge about others' integrity among its user base: If everybody makes their beliefs about the integrity of others public and if these beliefs are both accurate and sufficiently diversified among users, then Sybil resistance emerges as a manifestation of the *wisdom of the crowd*.

The above result also shows that the users most exposed to a malicious attack are those closest to the malicious nodes in the trust network. Moreover, 'leaks' in the form of honest agents trusting malicious nodes tend to resolve themselves, since a) once the malicious network depletes honest tokens, it can only acquire new tokens at a rate limited by the size of its trust boundary, which generally is negligible compared to the entire honest network and b) since accounts can always revoke trust, any detected malicious networks can quickly be isolated.

Of course, the upside of multiplying one's resources might still tempt users into trying to to maintain several accounts, even if that involves investing significant resources such as social engineering, etc. Since the lack of confidence of honest users into the protocol's ability to identify and exclude more sophisticated modes of exploitation would be detrimental to the value and adoption of CRC, it is a major concern for the Circles team to provide tools and intelligence to users, organizations and groups to help identify and untrust malicious accounts. Moreover, as we will discuss below Circles Groups provide a vehicle for additional layers of security if required.

4.4 Emergence of global fungibility

We now consider how fungibility between individual currencies emerges by the properties of the Circles trust graph. For our purposes, we here define *operational* fungibility as the ability of users to have their CRC accepted broadly within the parts of the economy they wish to transact with. A practical measure of this fungibility is the average fraction of their CRC that users can spend. We define this formally as the *average spendable fraction (ASF)* among a subset of nodes $\tilde{N} \subseteq N$,

$$ASF(\tilde{N}|S) = \frac{1}{|\tilde{N}|(|\tilde{N}|-1)} \sum_{n,n' \in \tilde{N}, n \neq n'} \frac{T(n \to n'|S)}{B(n|S)},$$

where $B(n|S)$ are the total CRC holdings of n in state S. ASF values range from 0 to 1, with 1 indicating complete operational fungibility within the group $\tilde{N}$, in state S. To illustrate, consider the following simple examples (see Fig.6):

- In a fully connected network where every account trusts every others, $ASF(\tilde{N}|S) = 1$ for any subset $\tilde{N}$ and any state S.
- In a network with trust paths between any two accounts in S and an equal balance of their own individual currencies, we have, for any subset, $ASF(\tilde{N}|S) = 1$.
- In a network, in which everybody only trusts a single account, $ASF(\tilde{N}|S) = 1$ if and only if S is such that only Circles of this single account are in circulation.

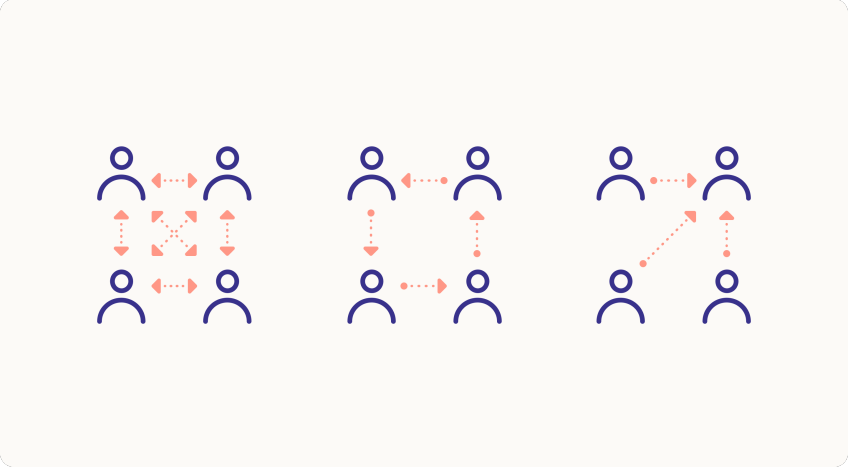

Fig. 6: Trust graphs with global fungibility: The complete graph in which everybody trusts another achieves complete fungibility regardless of the distribution of CRC (left). The "ring graph" in which every user is connected, achieves complete fungibility if all parties hold the same amount of their own CRC (middle). The "star graph" formation, in which everybody only trusts a single user, will achieve complete fungibility whenever only CRC of this user are in circulation (right).

Of course, these examples are highly simplified, and the actual state of the Circles network will be more complex. However, we expect global fungibility to reach a sufficient level for everyday economic exchanges within Circles, for the following reasons:

- *High connectivity*: The more independent paths that connect two parties, the greater

the volume of CRC that can flow between them (similar to water through a network of pipes). Since trust connections will reflect social connections of people, and social networks are known for their high connectivity levels, we expect the number of independent paths between any two users to grow with the size of the network.
- *Local economics:* Most users will interact regularly with only a small subset of the entire network. We expect the average connectivity between users and their core network to be significantly higher than with the overall network.
- *Unstable bottlenecks:* Transitive transfers require that all intermediate parties on a path hold sufficient CRC balance. As such, some paths may temporarily become unavailable, but the continuous CRC issuance for all active accounts helps quickly "unblock" these bottlenecks, maintaining network fluidity.
- *Diversified balances:* As users accumulate CRC from different accounts, they diversify their holdings. This diversity enables more paths for transitive transfers, since users are likely to hold CRC from a range of trusted sources, enhancing transaction flexibility.

4.5 Liquidity Clusters, Exchange Rates and Flow potentials

In the previous section, we explore fungibility from an *operational* perspective. Here, we shift our view to a *token* perspective. Instead of examining what accounts can pay one another, we analyze how one type of CRC can be converted into another. We then show that we can use this analysis to understand under which conditions the "value" of one user's CRC (in a sense made precise below) will be higher than that of another user and when they coincide. Our main result in this section is that value flows opposite to trust: Alice's CRC will, at equilibrium, be at least as valuable as those of all the people that trust her, and conversely her CRC will be at most as valuable as those of any person that she trusts.

4.5.1 Liquidity Clusters

For any two accounts $n, n' \in N$ and some state S, we write

$$n \succ_{l,S} n'$$

to indicate that in state S every account that initially holds some positive amount $l > 0$ of n-CRC can turn them into l n'-CRC. Moreover, for a given state S, we write $n \sim_{l,S} n'$ if $n \succ_{l,S} n'$ and $n' \succ_{l,S} n$. A set of currencies created by accounts $\tilde{N} \subseteq N$ forms a *l-liquidity cluster* in S if $n \sim_{l,S} n'$ for all $n, n' \in \tilde{N}$.

It follows that fungibility within a group as defined in the last section implies that all currencies form a liquidity cluster under mild additional assumptions. Consider a set $\tilde{N}$ of accounts and a state S such that $ASF(\tilde{N}|S) = 1$. Assume that each account holds l of its own currency. In this scenario, the currencies created by $\tilde{N}$ constitute a l-liquidity cluster. This holds because, by assumption, there exists a series of transfers enabling any two accounts in this set to exchange l of their own CRC for an equivalent amount of the other's currency.

Conversely, the existence of liquidity clusters provides a lower bound for the degree of operational fungibility of the system. Assume that $\tilde{N}$ forms a l-liquidity cluster in state S, and that every account in this set holds at least l of its own currency. Then one can show that

$$ASF(\tilde{N}|S) \geq \frac{l}{\bar{B}_S(\tilde{N})},$$

where $\bar{B}(\tilde{N}|S) = \frac{1}{|\tilde{N}|}\sum_{n\in\tilde{N}} B(n|S)$ are the average total token holdings of accounts in $\tilde{N}$.

4.5.2 Exchange Rates

The convertibility of currencies is also interesting from the point of view of the *exchange rates* between individual currencies that emerge in the presence of a public exchange. Essentially, stable exchange rates–where no party can extract arbitrage through transitive transfers–are expected to reflect the fungibility order introduced above.

Specifically, let $R(n \to n')$ denote an exchange rate between n-CRC and n'-CRC, given a public exchange for different variants of CRC. This rate reflects how many units of n'-CRC can be exchanged for one unit of n-CRC. We can show that, if R prevents arbitrage in state S, it must take the form

$$R(n \to n') = \frac{V(n)}{V(n')},$$

where V is a *value* function satisfying the property that for any $l > 0$,

$$n \succ_{l,S} n' \Rightarrow V(n) \geq V(n').$$

This characterization provides valuable insights. It implies that currencies within the same liquidity cluster will exchange at a 1:1 rate. In fact, it has the strong implication that, under the mild assumption that all accounts hold some of their own currency, arbitrage-free exchange rates will actually reflect trust:

Value at equilibrium flows opposite to trust

For any state S in which every account holds some of their own currency, then exchange rates in the absence of arbitrage reflect the trust connections between users: If n trusts n', then

$$V(n') \geq V(n).$$

This implies that, at equilibrium, the value of the CRC created by any user will be

- at least as much as each of the currencies of accounts that trust this user,
- at most as much as each of the currencies of accounts that this user trusts.

4.5.3 Flow potentials

Another implication of the above characterization of exchange rates is that the network can naturally divide into clusters. To visualize this, consider the graph shown in Figure 7, which depicts two "connected" clusters of users, referred to as Left and Right cluster respectively. Within each cluster, all users are connected, with a single one-sided trust connection going from one user in the left cluster to one user in the right cluster. In any state S where each user holds some amount of their own currency, the currencies created in the Left and Right cluster form two liquidity clusters. Under stable conditions, the currencies within each cluster have uniform value, denoted as V_L and V_R for the left and right cluster, respectively. Moreover, the trust structure ensures that $V_R \geq V_L$.

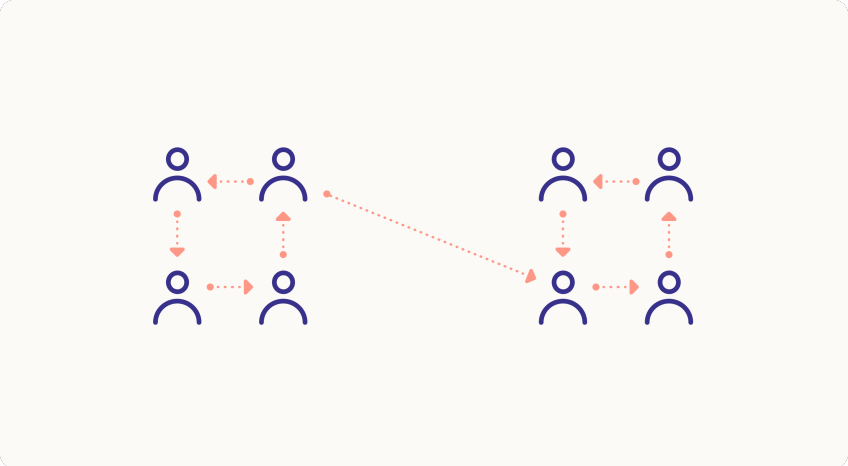

Fig. 7: A trust graph illustrating two connected clusters with a single one-way trust connection linking them. When each user holds some amount of their own currency, stable exchange rates ensure that the currency issued by the right cluster has a higher or equal value than that of the left one.

The existence of liquidity clusters and flow potentials imply the possibility that, instead of a single globally fungible currency, a diverse set of currencies could emerge, each with differing exchange rates. For the sake of conceptual simplicity, we have focused on the scenario where global fungibility between all CRC emerges. However, it is important to note that the liquidity cluster scenario is not just theoretically plausible, but also technically and economically feasible.

This concludes our examination of the micro-level dynamics of Circles. We have demonstrated how the simple rules of Circles generate a rich and adaptable structure. This structure achieves the dual objectives of providing universal access to the protocol while equitably distributing the rights to issue money among users. Ultimately, this system supports the emergence of one, or a few, globally fungible currencies at the macro level.

5 Organizations and Circles Groups

Above, we introduce the rules of Circles for personal accounts and demonstrate how these rules enable Circles to function as a currency issued by individuals and backed by their trust connections. However, economies are not made up solely of individuals. To extend Circles' functionality, support a broader range of use cases and satisfy our requirements of attractive money to allow shared-interest groups to organize themselves economically, we introduce two special kinds of accounts: *Organizations* and *Circles Groups*.

Organizations are accounts that function much like human accounts, but with one key difference: they don't issue their own currency. This reflects our conviction that Circles should be a currency *created exclusively by people*. In practice, shops, vendors, and other non-human entities will use Organization accounts to send and receive CRC.

Circles Groups (or simply Groups) are a more complex type of account with unique capabilities. Unlike Organizations, Groups *do* have their own currency. However, instead of being created, as with human accounts, this currency is issued in exchange for personal CRC from selected users—known as *members* of the Group.

Here is how it works:

1. Members' CRC can be exchanged for newly issued Group-CRC at a 1:1 rate.
2. These exchanged CRC are stored in the Group's *vault* as collateral, effectively removing them from circulation.
3. Holders of Group-CRC can redeem them at any time and at a 1:1 rate, against any CRC in the vault, which effectively burns the redeemed Group-CRC.

This issuance policy balances trust-based issuance with a robust mechanism that ensures the total circulating supply of CRC remains unchanged. See Fig. 8 for an illustration of this process.

While Organizations are relatively straightforward in scope, Groups introduce powerful capabilities to Circles, enabling a wide range of applications. We now turn to explore their primary use cases.

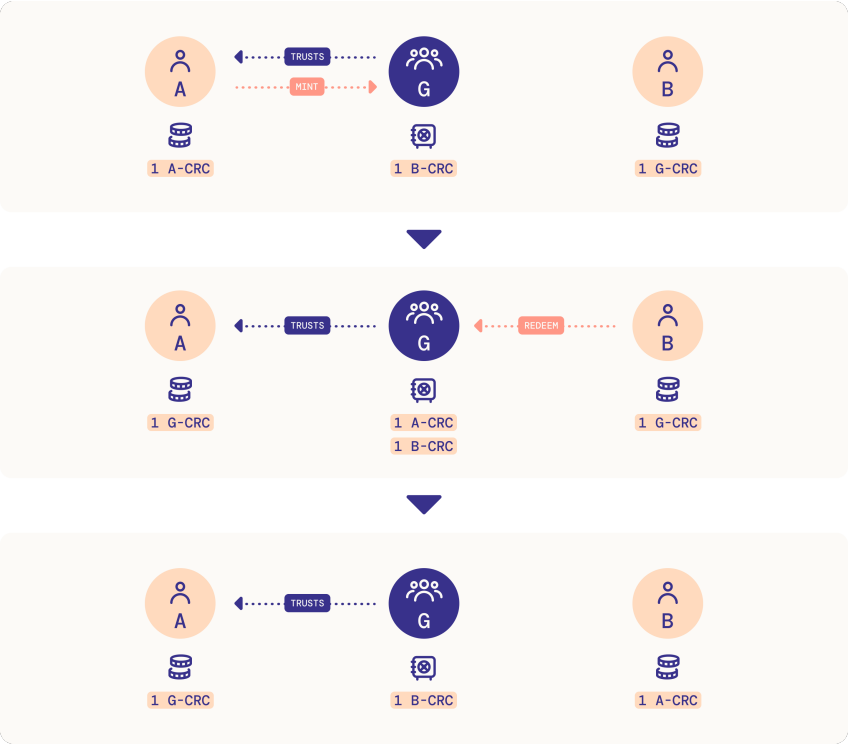

Fig. 8: The base Circles Groups minting policy in action: Alice is a member of the Circles Group G. Alice creates a Group-CRC by contributing one of her CRC as collateral, which is 'locked' into the Group's vault. Bob, who is not necessarily a member but happens to hold Group-CRC, then redeems it in exchange for Alice's CRC. Throughout this process, the total circulating supply of CRC not locked in a vault remains unchanged.

5.1 Currencies for shared interest groups

People naturally organize around shared interests, locations, economic needs, and other commonalities. Circles Groups provide a powerful tool for these communities to organize themselves economically, using CRC as the base currency. There are many potential types of Groups, including:

- Location-based (Humans of Berlin, ...)
- Requirements-based (Humans with proven identity, Active users, ...)
- Community-based (Humans of Urban Gardening Group XYZ, ...)
- Enterprise-based (Humans of Gnosis, ...)
- Event-based (Humans of DappCon 2024, ...)
- Interest-based (Humans who play Chess, ...)
- Education-based (Humans of Cambridge University, ...)
- Earth-based (Humans of Gaia, ...)

It's worth noting that multiple Circles Groups can target the same segment. For example, both Proof of Humanity and WorldID might use similar admission criteria but implement them differently. Since anyone can create a Circles Group and there is no central oversight or orchestration, this diversity is a natural consequence of the system's openness. Moreover, Groups can be trusted by both human accounts and other Group accounts. The

issuance mechanism for Groups then ensures that Group membership is effectively transitive: For instance, if A is a member of Group G and G is a member of Group H, then A can issue H-CRC as if they were a member of H. This transitivity enables Groups to organize into sub-Groups or larger federations.

What advantage do Circles Groups gain from issuing their own currency? First and foremost, it gives shared-interest groups access to monetary tools to pursue these interests. We foresee a plethora of different applications for Group currencies, such as

- local and complementary currencies
- loyalty programs that require minimum levels of activity to enjoy certain benefits
- voting and funding mechanisms for DAOs
- liquidity pooling
- safety-critical applications that require highly credible collateral.

While some of these applications can be expressed using just the default rules of Circles as discussed above, some require additional logic. For this reason, Groups can extend and build on the default issuance policy described above by calling external smart contracts at various points in the issuance and redemption process. Examples of such custom policies are

- limiting the total supply of Group-CRC
- dynamically adjusting supply based on network state or other inputs
- require exogenous collateral for issuance
- creating lock-ins to restrict redemption
- enforcing specific identification criteria for membership or issuance
- automatically granting or revoking membership based on predefined requirements.

Despite this flexibility, the base rules of Circles always apply and connect the Circles Groups and their economics.

5.1.1 Active and Passive Circles Groups

Membership in Circles Groups is defined by the Group (whose actions are carried out by one or several Group Admins) by extending trust to an account, rather than the other way around. Since Circles imposes no restrictions on who an account can trust, we expect that users will be members of various Circles Groups without actively seeking membership. For example, a location-based Group might automatically trust all users in a specific region. We call such Groups *passive* Circles Groups because membership is extended without requiring the user's involvement. These Groups may have little practical impact on users who don't engage with them, as the trust is unilateral and poses no risk. In contrast, *active* Circles Groups require users to actively participate, apply, or meet certain conditions to gain and retain membership. Examples include cooperatives, DAOs, or other enterprise-based Groups.

5.2 The resilience-efficiency tradeoff

An important purpose of Groups is that they provide a way for members of a Circles-based economy to navigate the tradeoff between monetary efficiency and resilience. By efficiency, we mean minimizing overhead, frictions, and transaction costs when using Circles. By resilience, we refer to the system's ability to resist exploitation by malicious parties and free-riders, as well as the capability to shield parts of the economy from shocks or crises in another.

The tradeoff is a well-known concept in mainstream economics, especially in the macroeconomics of international finance: Using a single currency across multiple markets (or to a lesser extent, by linking different markets via fixing the exchange rates of their respective domestic currencies, as in the gold standard) can reduce frictions in trade and labour mobility between these markets, increasing their overall efficiency. However, this comes at the cost of resilience: without independent monetary policy tools, individual markets are more vulnerable to crises spilling over from others, which might have been mitigated or avoided with localized currencies.

Circles, at its core, is designed with resilience in mind. Each member has their own personal currency, which is, by default, non-fungible with others. Trust connections between accounts can be revoked or altered at any time, offering individuals robust tools to safeguard the value of their personal CRC against shocks such as the discovery of a Sybil network.

The resilience-first design of Circles comes with a tradeoff in efficiency. Seemingly simple transactions at the macro level can require significant computational effort at the micro level to identify the paths for transitive transfers. Moreover, newcomers may face challenges in quickly having their CRC accepted by others, creating barriers to seamless participation.

Groups address these challenges by offering a 'fast lane' to efficiency. For example, consider a Group that automatically grants membership to any account that is not obviously a bot. In this scenario, users could exchange their individual CRC for the Group's Group-CRC, creating a single, fully fungible currency with significantly lower transaction costs (even at the micro level) and easier onboarding for new users. However, this efficiency comes at the cost of reduced resilience. A low barrier to entry would make the Group an easy target for malicious parties. While this represents an extreme example, in practice, Groups are likely to span a spectrum of efficiency and resilience. Users can then choose to hold a mixture of Group-CRC and individual CRC based on their preferences and risk tolerance.

6 Conclusion

With Circles, we have set out to propose an answer to the question: *What does a currency have to look like in order to be adopted by people in a multipolar, decentralized world, not by virtue of necessity, but by virtue of its attractive properties alone?* We have argued that a currency should satisfy three necessary conditions to qualify as a good answer to this question: It should be universally accessible, its issuance mechanism should be decentral-

ized, and it should prevent any group from maintaining persistent control over the money supply. We have demonstrated that Circles satisfies these criteria by leveraging the existing social trust connections between people. Additionally, we introduced complementary structures, Circles Groups and Organizations, to facilitate the adoption of Circles among communities with shared interests and goals.

History has seen many proposals for currencies, mainstream and alternative, come and go. Unsurprisingly, the currencies most likely to succeed have typically been those backed and supported by the institutions with the most political and military power. Against this backdrop, any new currency, including Circles, faces an uphill battle. Yet, the 21st century has shown us glimpses of profound transformation in the nature of money, driven by digitization and decentralization. Circles represents the next step in the evolution of this movement, reinventing money for a decentralized world: Truly global, inherently resilient, free from control by any one entity, and built to stand the test of time.

7 Acknowledgments

We'd like to thank the many people whose work and feedback and commitment have been absolutely critical to this project and to the ideas presented in this text. There are too many to name individually, so we choose here to thank you collectively.

8 Disclaimer

This whitepaper is for informational purposes only and is not intended to constitute, nor should it be construed as, any form of financial, legal or investment advice. This whitepaper is not an offer or solicitation to purchase tokens. The project is open source and each user is responsible for creating their own CRC. While we aim to provide accurate and up-to-date information, we do not guarantee the accuracy, completeness or suitability of the content. This whitepaper has not been reviewed or approved by any regulatory authority and readers should do their own research and ensure they inform themselves of, and comply with, relevant laws, regulations and tax requirements that apply in connection with their use of the Circles Protocol or CRC.

9 Appendix

In this appendix, we formally derive the results presented in the paper.

9.1 Macro level basics

We begin by setting up the basic notation. We consider a finite set N of all users that have been active at some point in the history of the network. We model time as a non-negative discrete parameter $t \in \mathbb{N}_+$. Let $\theta : N \times \mathbb{N}_+ \to \{0, 1\}$ be the indicator function whether a given

user creates at a given time. Under the assumption that all active users in the economy create while they are active, we denote as $N_t = \{n \in N : \theta(n,t) = 1\}$ the population at time t and as $s_t = |N_t|$ its size. Moreover, for any subset of accounts $\tilde{N} \subseteq N$, we denote $N_t(\tilde{N}) = |N_t \cap \tilde{N}|$ as the active minting subset of these accounts at time t.

Let $r > 0$ be demurrage rate (in units of t) and a the amount of CRC that an individual member of the economy can create per time period (i.e. in our case r is 7% p.a. and a is 1 CRC/h). The total value of *money created* at some time t' between (and including) two periods t and t', by any set of accounts $\tilde{N}$, is

$$M_{t,t'}(\tilde{N}) = \sum_{\tau=t}^{t'} a s_\tau(\tilde{N}) \rho^{(t'-\tau)},$$

where we introduced $\rho = (1-r)$ for notational convenience. We write $M_t = M_{0,t}(N)$ for the *total money supply* at period t (we assume that money is always created as the first thing in every period, meaning the money created at t is always at users' disposal in that very period). Thus, $M_0 = a s_0$, $M_1 = a s_0 \rho + a s_1$ etc. For any time and set of accounts, the money supply obeys the recurrence relation

$$M_{t,t'}(\tilde{N}) = M_{t,t'-1}(\tilde{N})\rho + a s_{t'}(\tilde{N}).$$

A special case of the above is given by the situation in which all members of the set create at every time, i.e. $N_t(\tilde{N}) = \tilde{N}$ for all relevant times. In this case we have

$$M_{t,t'}(\tilde{N}) = a|\tilde{N}|\eta(0, t'-t, \rho),$$

where we defined

$$\eta(a,b,c) := \sum_{k=a}^{b} c^k.$$

This special case yields a geometric series and is therefore easy to evaluate analytically. In particular, using standard results on geometric series, we have for $\rho \neq 1$,

$$\eta(0,b,\rho) = \frac{1-\rho^{b+1}}{1-\rho}.$$

This implies that for any single account $n \in N$ that creates continuously once they join the economy at some time t,

$$M_{t,t'}(n) = \frac{a(1-\rho^{t'-t+1})}{r} = \frac{a}{r} + O(e^{-(t'-t)})$$

In the case of Circles, which applies the demurrage in 24 hour intervals, $a = 24$ and $r = 1 - 0.93^{1/365.25}$ which results in the limit of 120804 CRC mentioned in the main text.

9.2 Fair access to supply

We derive the a generalized form of equation (1) in the main text. For this, let $B(n)_t$ denote the total CRC balances of some account $n \in N$ at time t. If this user creates throughout, then for any later time $t' \geq t$, we have

$$B_{t'}(n) = B_t(n)\rho^{(t'-t)} + a\eta(0, t'-t, \rho)$$

Consequently, writing $\Delta t = t' - t$,

$$\begin{aligned}\frac{B_{t'}(n)}{M_{t'}} &= \frac{B_t(n)\rho^{\Delta t} + a\eta(0,\Delta t,\rho)}{M_t\rho^{\Delta t} + a\sum_{\tau=t}^{t'} s_\tau \rho^{t'-\tau}} \\ &= \frac{\eta(0,\Delta t,\rho)}{\sum_{\tau=t}^{t'} s_\tau \rho^{t'-\tau}} \left(\frac{\alpha_{t,t'}+1}{\gamma_{t,t'}+1}\right) \\ &= \frac{1}{\bar{N}_{t,t'}(\rho)} \left(\frac{\alpha_{t,t'}+1}{\gamma_{t,t'}+1}\right),\end{aligned}$$

where we defined

$$\alpha_{t,t'} = \frac{B_t(n)\rho^{\Delta(t)}}{a\eta(0,\Delta t,\rho)} \leq \frac{B_t(n)\rho^{\Delta(t)}}{a},$$

$$\gamma_{t,t'} = \frac{M_t(n)\rho^{\Delta(t)}}{a\sum_{\tau=t}^{t'} s_\tau \rho^{t'-\tau}} \leq \frac{M_t(n)\rho^{\Delta(t)}}{a\,\mathsf{min}_{\tau:t\leq\tau\leq t'}\, s_\tau},$$

as well as the *time-discounted average population size*

$$\bar{N}_{t,t'}(\beta) = \frac{\sum_{\tau=t}^{t'} s_\tau \beta^{t'-\tau}}{\sum_{\tau=t}^{t'} \beta^{t'-\tau}} = \sum_{\tau=t}^{t'} w_\tau s_\tau, \quad w_\tau = \frac{\beta^{t'-\tau}}{\sum_{\tau=0}^{\Delta t}\beta^\tau}$$

for some discount-factor β. This is a form of weighted average for the actively minting population between times t and t', in which population sizes farther in the past are weighted less. For $\beta = 1$, this is just the unweighted average population over the time window. To interpret the above equality, we note that both $\alpha_{t,t'}$ and $\gamma_{t,t'}$ are upper bounded by exponentially decaying terms and hence will vanish over longer time windows. This implies that any member's share in the total money supply converges, over time, to the time-discounted average population size (which is itself a moving target of course) with discount rate $\beta = \rho$, at least if that member only creates CRC and does not participate in further economic transactions.

Equation (1) presented in the main text then emerges as a special case in which i) the population is constant and ii) the user starts without any initial balance, $B_t(n) = 0$, so that $\alpha_{t,t'} = 0$.

9.3 The inflationary view

In this section, we formally introduce the inflationary view and discuss how it differs. For this, let us introduce another currency, ICRC, whose issuance follows the same rules as that of CRC, with the following changes: ICRC is not subject to demurrage and accounts can create an increasing amount of ICRC. Formally, the issuance rate of ICRC at time t is $a_t = a(1+i)^t$ for some *inflation rate* i. The total amount of ICRC created by some set $\tilde{N}$ up until time t then is

$$\tilde{M}_t(\tilde{N}) = \sum_{\tau=0}^{t-1} s_\tau(\tilde{N}) a_\tau$$

By construction, the total supply of ICRC diverges over time, under the assumption that $s_t > 1$ at all times. However, to support our claims the real economics of the inflationary view coincide, in the sense that principle A and principle B are still satisfied, note that under the assumption that the inflation rate is chosen such that $(1+i) = \rho^{-1}$

$$\rho^{(t-1)} \tilde{M}_t(\tilde{N}) = a\rho^{(t-1)} \sum_{\tau=0}^{t-1} s_\tau(\tilde{N}) \rho^{-\tau} = M_t(\tilde{N})$$

This is true regardless of t or the minting schedule of the accounts in $\tilde{N}$. This relates the money supply of two economies, one using ICRC and one using CRC, that follow the same minting schedule. An immediate implication is that the relative contribution to the total money supply by any two Groups would be the same: That is, for any Groups $N_1, N_2 \subseteq N$, any minting schedule and any time t,

$$\frac{M_t(N_1)}{M_t(N_2)} = \frac{\tilde{M}_t(N_1)}{\tilde{M}_t(N_2)}$$

Moreover, a similar statement holds true for the development of balances in the absence of transactions, in that a user will see their share in the total money supply develop in exactly the same way, regardless of which of the two views we're using: To see this, consider again two economies whose user base is the same and follows the same minting schedule but uses ICRC and CRC respectively. Let $\tilde{B}_t(n)$ denote the ICRC balance of an account $n \in N$ in the first economy, at some initial time t and let $\tilde{B}_{t'}(n)$ denote the ICRC balance of n if in the time window (t, t') they only create ICRC but do not transact with other accounts. Similarly, let $B_t(m)$ and $B_{t'}(m)$ denote the initial and final balances of an account $m \in N$ for the other economy under the same conditions. Then we have that,

$$\frac{\tilde{B}_t(n)}{\tilde{M}_t} = \frac{B_t(m)}{M_t} \Rightarrow \frac{\tilde{B}_{t'}(n)}{\tilde{M}'_t} = \frac{B_{t'}(m)}{M'_t}$$

that if the two accounts hold the same percentage of the total money supply in these economies, then the percentage of their holdings will evolve in just the same way. This follows from

$$\begin{aligned}\frac{\tilde{B}_{t'}(n)}{\tilde{M}'_t} &= \frac{\tilde{B}_t(n) + a\eta(t, t', 1+i)}{\tilde{M}_t + a|N|\eta(t, t', 1+i)} \\ &= \frac{B_t(n)\rho^{-t} + a\eta(-t', t, \rho)}{M_t\rho^{-t} + a|N|\eta(-t', -t, \rho)} \\ &= \frac{B_t(n)\rho^{\Delta t} + a\eta(0, \Delta t, \rho)}{M_t\rho^{\Delta t} + a|N|\eta(0, \Delta t, \rho)} \\ &= \frac{B_{t'}(n)}{M'_t},\end{aligned}$$

where we used that $\rho^t \tilde{M}_t = M_t$ due to the same minting schedules up until time t, that, by the above assumption, we consequently have $\rho^t \tilde{B}(n)_t = B_t(n)$ and also that $(1+i) = \rho^{-1}$. Since the real economic relevance of the balance is not dependent on the absolute number of tokens, ICRC or CRC, held but by their share in the total money supply, this means that the purchasing power of accounts in economies that differ only with respect to whether they use ICRC or CRC is the same, as claimed. Moreover, of course, this also means that the bounds we established in the last section about the convergence to the average, directly apply to the inflationary picture as well.

9.4 The value of seigniorage

In this section we derive the claims we make in the main text about the value of the seigniorage in equilibrium. For simplicity, we assume an economy with a single perishable good (meaning all of the available good for a single period must be produced and then consumed in that period) and that CRC is the only available currency. We denote as $p_t > 0$ the price of this good in units of CRC and as $c_t(n) \geq 0$ denote the amount of the good that an account $n \in N_t$ consumes at time t. We say that this simple economy is *in equilibrium* at time t if both goods and money markets clear, that is, all of the available good for this period gets consumed and the price is such that the total demand for money equals the total supply (for a Walrasian equilibrium we'd also assume that users choose the consumption that maximizes their utility given the price but here we don't consider utilities explicitly).

The seigniorage z_t at time t (which is the same across users), is simply

$$z_t = \frac{a}{p_t}$$

since this is the value created (in units of goods) by creating a units of CRC (assuming negligible issuance costs).[14] Then the average contribution of the seigniorage to consumption across users at t is, in equilibrium

[14] In this section, we model the issuance of CRC as a fiat currency as the only source of seigniorage. In principle, users can also provide loans using their CRC and derive seigniorage from interests on these loans. Hence, a more complete model would take into account both of these sources of seigniorage. We leave such an extension to future work.

$$\begin{aligned}\frac{z_t}{\bar{c}_t} &= \frac{s_t a}{p_t \sum_{n\in N_t} c_t(n)} \\ &= \frac{(M_t - M_{t-1}\rho)}{VM_t}\end{aligned}$$

where for the nominator we used the recursion relation for the total money supply and for the denominator we used the fact that, in equilibrium, $\sum_{n\in N} c_t(n) = VD_t/p_t = VM_t/p_t$, with D_t the total demand for money and V its velocity, which we here assume, for simplicity, to be constant over time (at least within the considered time window). The first equality links the consumption to the total money demand through the equation of exchange, under the assumption that all consumed goods are purchased by its consumers (a standard assumption in economic models that the fraction of goods consumed by the producers of that good is negligible), while the second equality equates money demand and money supply by the assumption of equilibrium. This shows that, in this very simplified model, the contribution of seignorage to average consumption equals the relative increase of the money supply, after applying demurrage.

To further evaluate this, we now introduce the assumption that the active population grows at a constant rate $g \geq 0$, so that $s_t = (1+g)^t s_0$ for some initial population $s_0 > 0$. Under this assumption,

$$\begin{aligned}M_t &= a s_0 \rho^t \eta(0, t, \zeta) \\ &= a s_0 \rho^t \frac{(\zeta^{t+1} - 1)}{(\zeta - 1)}\end{aligned}$$

for $\zeta = \frac{(1+g)}{\rho} > 1$. Plugging this into the above equation yields

$$\begin{aligned}\frac{z_t}{\bar{c}_t} &= \frac{(\zeta^{t+1} - \zeta^t)}{V(\zeta^{t+1} - 1)} \\ &= \frac{(1 - \zeta^{-1})}{V(1 - \zeta^{-(t+1)})} \\ &= \frac{1}{V(1 - \zeta^{-(t+1)})} \frac{(r+g)}{(1+g)}.\end{aligned}$$

Therefore, under the various simplifying assumptions we make, the relative contribution of the seigniorage to the average consumption converges to the claimed ratio at an exponential rate.

9.5 On the price stability of CRC - An OLG model

In this section, we provide a simple overlapping-generations (OLG) model of a Circles based economy. One noteworthy feature is that the presence of demurrage leads to a stabilization

of the price level in a competitive equilibrium, despite a positive growth of the economy.

The OLG model, introduced by Paul Samuelson, is one of the most celebrated and foundational models of an economy in the presence of money in that it captures a variety of monetary phenomena while being analytically tractable. While a full-blown exposition of this model is beyond the scope of this whitepaper, we here define all the essential structure: An OLG economy consists of generations of people, each of which lives for two periods and each of which overlaps for one period with the previous generation (when they are young and the previous generation is old) and for one period with the next generation (when they are old the next generation is young). When young, every user receives some fixed quantity y of a perishable consumption good (which can be interpreted as food or working time depending on the context). Every agent's utility depends on the amount of the consumption good that they consume in each of the two periods for which they are alive. Since the good by assumption perishes within a single period, agents cannot consume any of their own endowed good when they are old. The introduction of money solves this problem as it allows agents to sell a portion of their endowment when young and then purchase some of the good endowed to others when old.

Let us introduce the necessary quantities formally. *Note that the details of the model differ from those of the macro level above, since the assumption of agents that only live for two periods requires a slight adaption.* Let Y_t and O_t denote the number of young and old people in period t, respectively. By construction, $O_t = Y_{t-1}$. Let $C_Y(t)$ denote the aggregate amount of the good that is consumed by the young in period t and similarly $C_O(t)$ for the old. A *consumption schedule* $(C_Y(1), C_O(1), C_Y(2), \dots)$ for this economy is a specification of these amounts for the whole of history. A consumption schedule is *feasible* if for all times,

$$C_Y(t) + C_O(t) \leq yY_t.$$

This expresses the necessity that the old and young together cannot consume more than the total endowment. Moreover, let P_t denote the price for one unit of the good in terms of units of money and let $M_O(t)$ denote the amount of money held by the old at t. In terms of these quantities, a consumption schedule is said to be *affordable given prices and distribution* if, for all times,

$$P_t C_O(t) \leq M_O(t).$$

This "cash-in-advance" constraint simply states that the old can only purchase as much of the good as their budget allows. Finally, a consumption schedule is said to be *optimal at given prices and distribution* if it maximizes the utility of all agents across the affordable consumption schedules. A crucial concept for an OLG is that of a *competitive equilibrium (CE)*. This is defined as a tuple of

- consumption schedule
- a sequence of prices $(P_t)_t$

- a sequence of money distributions $(M_O(t))_t$

such that the consumption schedule is both feasible and optimal at given prices and distribution and that markets clear (i.e. that all of the available good gets consumed or sold and that the demand for money equals its supply). A CE is moreover a *steady-state* CE if C_Y and C_O are constant over time. Competitive equilibria are relevant insofar as they describe stable trajectories of the economy that are consistent with the incentives of the individual agents if understood as utility maximizing rational agents.

The following is an OLG model for a Circles-based economy with a steady growth-rate of the economy: We assume that $Y_t = (1+g)^t Y_0$ for some initial set of young people Y_0. We moreover assume that newly created money gets distributed in a lump-sum fashion across all old users at some rate a. That is, the total money supply available at the beginning of time t is

$$\bar{M}_t = \bar{M}_{t-1}\rho + O_t a$$

with $\bar{M}_0 = 0$.[15] We then have the following fact.

Existence of steady-state CE with stable prices

For the OLG model of the Circles economy and under the additional assumption that all agents, across periods, share the same utility function that is monotonically increasing in the amount of goods consumed, any steady-state CE needs to have prices that satisfy

$$\frac{P_{t+1}}{P_t} = \frac{(1-\zeta^{-(t+1)})}{(1-\zeta^{-t})},$$

where $\zeta = \frac{(1+g)}{\rho} > 1$.

Proof Sketch: By the market clearing condition and the monotonicity of the utility functions, any consumption schedule that is part of a CE will have to saturate the feasibility and affordability constraints as well as $\bar{M}_t = M_O(t)$. The feasibility constraint is saturated by definition of the market clearing condition, the affordability constraint is saturated since agents can never maximize their utility by leaving money unspent and, as a consequence of this fact paired with the fact that new money is distributed to the old, the totality of the money supply is in the hands of the old. Moreover, by the steady-state condition and the identity of the utility functions, we have that consumption amounts are time-independent, i.e. $C_Y(t) = c_Y Y_t$

[15] Note that this is different from the total money supply M_t presented above. This is because here, we only provide money to the old. The reason is that in an OLG context, we consider the whole life of a user in only two periods. As such, we assume that users start with very little CRC and for them to be able to accumulate significant CRC after several decades.

and similarly for the young. Together, this results in the identity

$$P_t = \frac{\bar{M}_t}{Y_t(y - c_Y)}.$$

Moreover, using reasoning exactly analogous to the last section, we find that

$$\begin{aligned}\bar{M}_t &= aY_0\rho^{t-1}\eta(0, t-1, \zeta) \\ &= aY_0\rho^{t-1}\frac{(\zeta^t - 1)}{(\zeta - 1)}\end{aligned}$$

for $\zeta = \frac{(1+g)}{\rho} > 1$. Using this and that $Y_t = (1+g)^t Y_0$, we obtain

$$\begin{aligned}\frac{P_{t+1}}{P_t} &= \frac{Y_t\bar{M}_{t+1}}{Y_{t+1}\bar{M}_t} \\ &= \frac{\zeta^{t+1} - 1}{\zeta(\zeta^t - 1)} \\ &= \frac{(1 - \zeta^{-(t+1)})}{(1 - \zeta^{-t})},\end{aligned}$$

which converges to 1 at an exponential rate in t.

Note that the above OLG model assumes constant per capita output, through fixing the constant y. As such, it is assuming that output, defined as $Y_t y$, grows exactly proportional to the size of the population. While dropping this assumption would often lead to the non-existence of a steady-state CE, it is intuitively clear how the results would change: Longer term trends in per-capita output would be reflected in longer term trends in price levels, with increasing per-capita output leading lower prices and decreasing per-capita output leading to higher prices.

9.6 Micro level basics

We now turn to the micro level description of Circles. We first set up the basic notation. Abstracting, for the moment, from the fact that individual nodes create their own currency, we consider a world consisting of accounts N and a set of non-fungible currencies Kthat are in circulation. We describe the instantaneous state of the system at any time as the tuple

$$S = (G, B),$$

where $G \in \{0,1\}^{|N|\times|K|}$ is the *adjacency matrix* of the *trust graph*, with G_{nk} indicating whether user n trusts currency k at time t, and $B \in \mathbb{R}_+^{|N|\times|K|}$ is the *balance matrix*, with B_{nk}

indicating the amount of currency k that account n holds in this state. The state space $\mathcal{S}$ of Circles is then given by $\{0,1\}^{|N|\times|K|} \times R_+^{|N|\times|K|}$. We also denote as

$$K_G(n) = \{k \in K | G_{nk} = 1\}$$

the set of trusted coins for a given account and $K_G(\tilde{N}) = \cup_{n\in\tilde{N}} K_G(n)$.

9.6.1 State transitions – Definition

Let us now turn to state transitions that are allowed by the Circles protocol. For any two states $S, S' \in \mathcal{S}$ and any $n \in N$ we write $S \to_n S'$, if a transition from S to S' is can be brought about by account n in Circles, using a combination of the following four possible types of actions. In the following E_{ij} denotes a matrix (of appropriate size) with all zero entries except 1 in entry ij:

1. **Trusting**: Any account can trust any of their untrusted accounts.

   $$S \to_n (G', B)$$

   if $G' = G + E_{nk}$ for some $k \in K$.

2. **Untrusting**: Any account can untrust any of their trusted accounts.

   $$S \to_n (G', B)$$

   if $G' = G - E_{nk}$ for some $k \in K$.

3. **Direct transfer**: Any account can send any of their balance to any other account as a direct transfer.

   $$S \to_n (G, B')$$

   if $B' = B + b(E_{n'k} - E_{nk})$ for some $b \geq 0$, $k \in K$ and $n' \in N$.

4. **Trust swap**: Accounts can swap tokens using the rule of trust.

   $$S \to_n (G, B')$$

   if $B' = B + E_{nk} - E_{n'k'}$ for some $k, k' \in K$ and $n' \in N$ such that $G_{n'k'} = 1$.

The above also exhaustively specify the relation $\to_n$. That is, by definition $S \to_n S'$ *only if* the state transition can be brought about by n by means of a sequence of (any combination of) the above operations.

Definition: Achievable transition

For any set of accounts $\tilde{N} \subseteq N$ we write $S \to_{\tilde{N}} S'$ and say that *a transition from S to S' can be achieved by accounts $\tilde{N}$* if there exists an $I \in \mathbb{N}$, a sequence of states $(S_i)_{i=0}^{I}$ and a sequence of accounts $(n_i)_{i=0}^{I-1}$ such that $n_i \in \tilde{N}$ for all i, $S_0 = S, S_I = S'$ and

$$S_i \to_{n_i} S_{i+1}, \quad i = 0, \ldots, I-1$$

The following decomposition result makes it easier to think about the complex transitions that are possible and will be frequently invoked in proofs in later sections.

Decomposition of transitions

Consider any two states $S = (G, B), S' = (G', B') \in \mathcal{S}$ and a subset $\tilde{N} \subseteq N$. If $S \to_{\tilde{N}} S'$, then there exist states $S_1, S_2, S_3 \in \mathcal{S}$ such that

$$S \to_{\tilde{N}} S_1 \to_{\tilde{N}} S_2 \to_{\tilde{N}} S_3 \to_{\tilde{N}} S',$$

where the first transition is achieved by means of trusting actions only, the second transition is achieved by means of trust swaps only, the third transition is achieved by means of direct transfers only and the fourth by means of untrusting only.

Proof: Let $(S_i)_{i=0}^{I}$ be the sequence of states that exist by assumption such that $S_0 = S, S_I = S'$ and

$$S_i \to_{n_i} S_{i+1}, \quad i = 0, \ldots, I-1$$

and where each transition is achieved by means of one of the four basic kinds of actions describe above. We first note that the set of possible state transitions that can be achieved by accounts in $\tilde{N}$ by means of trusted swaps and direct transfers can only be enlarged by those accounts trusting one another and decreased by untrusting one another and that more over this set is independent of the trust connections of accounts in $\tilde{N}$ into the remainder of the network. As such, we can assume, without loss of generality, that $S_1 = (\bar{G}, B)$, where $\bar{G}$ has a trust connection for all members in $\tilde{N}$ and otherwise equals G, while $S_3 = (\bar{G}, B')$. In other words, we can assume that in the first step, all accounts in $\tilde{N}$ trust one another and then, in the last step, they untrust one another to achieve the trust graph G' of the final state S'. We now turn to the intermediate transitions. For this, assume that there exists some $i < I - 1$ such that the transition $S_i \to_{n_i} S_{i+1}$ corresponds to a direct transfer of some amount b_i of k_i-tokens from n_i to some account n_i' and $S_{i+1} \to_{n_{i+1}} S_{i+2}$ corresponds to a trust swap in which n_{i+1} swaps some amount b_{i+1}

of k_{i+1}-tokens with another party n'_{i+1} in exchange for the same amount of k'_{i+1} tokens, with $k_{i+1} \neq k'_{i+1}$. Then in the following we show that that the transition $S_i \to_n S_{i+2}$ could also have been achieved by either first applying a trust swap and then a two direct transfers or by means of three direct transfers. To see this, we distinguish two cases:

1. $n'_i \neq n'_{i+1}$ or $k_i \neq k'_{i+1}$: In this case we can simply invert the the two operations.
2. $k_i = k'_{i+1}$ and $n'_i = n'_{i+1}$: We denote $n' = n'_i = n'_{i+1}$ and $k = k_i = k'_{i+1}$ and $k' = k_{i+1}$. In this case n_i sends k-tokens to the other party $n'$and then n_{i+1}exchanges some of that currency back later. The reason we cannot simply exchange the operations in this case is that n' might not initially have sufficient balance of the k token for the swap. However, we can construct alternative operations that achieve the same effect: If $b_i \leq b_{i+1}$, then n_{i+1} can first swap an amount $b_{i+1} - b_i$ of their k' tokens for k tokens and then conduct a direct transfer of an amount b_i of k' tokens to n', while n_i transfers b_i of k tokens to n_{i+1}. This results in the same state as the original operations above and we are guaranteed the possibility of these transfers, since by assumption of the existence of the two operations above, $B^{(i)}_{n'k} + b_i \geq b_{i+1}$ and $B^{(i)}_{n_{i+1}k'} \geq b_{i+1}$ and $B^{(i)}_{n_ik} \geq b_i$. If $b_i > b_{i+1}$, then we instead replace the above operations by means of three direct transfers: n_i transfers an transfers an amount b_{i+1} of k tokens to n_{i+1} as well as an amount $b_i - b_{i+1}$ of k tokens to n', while n_{i+1} transfers an amount b_{i+1} of k' tokens to n'. We are guaranteed the possibility of these transfers by $B^{(i)}_{n_ik} \geq b_i$ and $B^{(i)}_{n_{i+1}k'} \geq b_{i+1}$.

The above establishes that for any sequence of transfers that involves a direct transfer followed by a trust swap, there exists either a trust swap followed by two direct transfers or three direct transfers that achieves the same transition. By iteratively applying this insight, we can always generate a sequence of states $(\tilde{S}_i)_{i=0}^I$ such that $\tilde{S}_0 = S, \tilde{S}_I = S'$ and in which there exists some i^* such that for all $i \leq i^*$, all corresponding state transitions are generated by transitive transfers and for all $i > i^*$, all corresponding state transition are generated by direct transfers, as claimed.

9.6.2 State transitions - Characterization

We now study the structure of the relation $\to_{\tilde{N}}$. To this end, we first define

$$\mathcal{S}_{\tilde{N}}(S) := \{S' \in \mathcal{S} : S \to_{\tilde{N}} S\}$$

as the *reachable state space* for a given set of accounts and initial state S. We also introduce some additional notation: We denote, for any state $S = (G, B) \in \mathcal{S}$, any set of nodes $\tilde{N} \subseteq N$ and any set of currencies $\tilde{K} \subseteq K$,

$$B(\tilde{N}, \tilde{K}|S) := \sum_{n\in\tilde{N}, k\in\tilde{K}} B_{nk}$$

as the total balance of $\tilde{K}$ coins held by members in $\tilde{N}$ (we drop the explicit mention of the state if it's clear from the context). For instance, $B(n, K)$ equals the total token holdings of a single account, $B(N, k)$ equals the total supply of currency k in state S, while $B(N\backslash\{n\}, K_G(n))$ is the total supply of tokens trusted by some account n that are not already held by this account. Finally, for two sets $N_1, N_2 \subseteq N$, we denote as

$$B_T(N_1 \to N_2|S) := B(N_1, K_G(N_2)|S)$$

the *trusted balance* of N_1 *towards* N_2, that is, the total amount of currency held by accounts in N_1 that are trusted by at least one account in N_2 (in the main text, we dropped the subscript T for sake of simplicity). We write $B_T(\tilde{N}|S) \equiv B_T(\tilde{N} \to \tilde{N}|S)$.

Necessary conditions

Using this notation, we can establish some necessary conditions for state transitions from observation:

1. The total amount of coins in circulation remain constant:

$$S \to_{\tilde{N}} S' \Rightarrow B(N, k|S) = B(N, k|S') \quad \forall\, k \in K$$

1. The balance of the "signers" cannot increase and that of non-signers in the network not decrease:

$$S \to_{\tilde{N}} S' \Rightarrow B(\tilde{N}, K|S) \geq B(\tilde{N}, K|S')$$

$$S \to_{\tilde{N}} S' \Rightarrow B(N\backslash\tilde{N}, K|S) \leq B(N\backslash\tilde{N}, K|S')$$

1. Trusting and untrusting alone cannot affect the trusted wealth of the non-signers:

$$(G, B) \to_{\tilde{N}} (G', B) \Rightarrow B_T(\bar{N}|S) = B_T(\bar{N}|S'), \quad \forall\, \bar{N} \subseteq N\backslash\tilde{N}$$

1. The total trusted balance of the non-signers cannot decrease:

$$S \to_{\tilde{N}} S' \Rightarrow B_T(\bar{N}|S) \leq B_T(\bar{N}|S'), \quad \forall\, \bar{N} \subseteq N\backslash\tilde{N}$$

This last property is what we refer to as the *Conservation of trusted balance* in the main text.

Sufficient conditions

An obvious question is whether these necessary conditions are also sufficient. That is, can any state that satisfies the above necessary conditions also be reached? For the case that $\tilde{N} = N$, the answer is affirmative, as the following argument shows.

Fact: Reachable state set for whole network

For any state $S \in \mathcal{S}$,

$$\mathcal{S}_N(S) = \{S' \in \mathcal{S} : \sum_n B'_{nk} = \sum_n B_{nk}, \forall k \in K\}.$$

Proof Sketch: Let B' be any final distribution of the coins that satisfies the above conservation property. Then this state can be reached as follows from the initial state S: First, all parties send their whole balance to a single account via direct transfer. Then, this account transfers the amounts B'_{nk} to every party n via direct transfer. Since these direct transfers are independent of the trust graph, any combination of trusting and untrusting is possible.

For the case that the whole network is involved in the state transition, the reachable state set is therefore simple to characterize. However, for the case of a generic set of signers $\tilde{N}$, we are not aware of a similarly simple characterization and believe that it does not exist without imposing further structure. It turns out, though, that the absence of a simple analytical characterization of the reachable state set is in practice not problematic, as we care primarily about the ability of parts of the network to pay other parts. It turns out that we don't need to fully characterize the relation $\rightarrow_{\tilde{N}}$ in order to understand this ability.

9.7 Transferrable trusted balance

As part of trusting an account, users commit to accept tokens they trust as viable payment for goods they price in CRC. As such a key measure of liquidity in Circles in any given state S is the total amount of tokens that a sender set of accounts can manage to transfer over to a (disjoint) receiver set of accounts, such that the for each token sent, at least one account in the receiving set trusts that token. We call this quantity the *transferrable trusted balance* and define as follows: For any $N_s, N_r \subseteq N$ such that $N_s \cap N_r = \emptyset$,

$$T_G(N_s \rightarrow N_r|S) := \max_{S':S\rightarrow_{N_s}S'} B_T(N_r|S') - B_T(N_r|S)$$

In the following we're going to study the properties of this key quantity and in fact will be able to completely characterize its value as a function of the balances held by accounts in state S. First, we note that T_G is equivalent to the total amount of tokens that the senders can obtain that are trusted by the receivers:

Equivalence of obtainable and transferrable trusted balances

For any disjoint sets $N_s, N_r \subseteq N$, the transferrable trusted balance is given as

$$T_G(N_s \to N_r|S) = \max_{S':S\to_{N_s}S'} B_T(N_s \to N_r|S').$$

Proof Sketch: We first show LHS $\geq$ RHS: Let $\tilde{S} \in \arg\max_{S':S\to_{N_s}S'} B_T(N_s \to N_r|S')$ and let S' be the state that results when we start in $\tilde{S}$ and accounts in N_s transfer all the trusted amount $B_T(N_s \to N_r|\tilde{S})$ to some (arbitrary) account in N_r (via direct transfer. Then $S \to_{N_s} S'$ by assumption so that

$$\begin{aligned} T_G(N_s \to N_r|S) &\geq B_T(N_r|S') - B_T(N_r|S) \\ &= B_T(N_s \to N_r|\tilde{S}) + B_T(N_r|\tilde{S}) - B_T(N_r|S) \\ &\geq B_T(N_s \to N_r|\tilde{S}), \end{aligned}$$

where we the conservation of trusted balance property in the last step.

We now show LHS $\leq$ RHS: Let $\tilde{S} \in \arg\max_{S':S\to_{N_s}S'} B_T(N_r|S')$. By the decomposition of transitions that we showed above, there exist states S_1 and S_2 such that

$$S \to_{N_s} S_1 \to_{N_s} S_2 \to_{N_s} \tilde{S}$$

and such the three transitions occur by means of only trusting and untrusting, trust swaps and direct transfers respectively. Note that the trusting and untrusting transition cannot change trusted balance of N_r, since N_s and N_r are disjoint, so that $B_T(N_r|S_1) = B_T(N_r|S)$. For the second transition that consists of a sequence of trusted swaps, assume that this sequence includes swaps that decrease the amount of tokens trusted by accounts in N_r that are held by accounts in N_s. By skipping all such swaps and all subsequent swaps that "rely" on this swap (in the sense that some of the incoming tokens are later sent away using other swaps), we obtain a final state S_2' with the property that

$$B_T(N_s \to N_r|S_2') \geq B_T(N_s \to N_r|S_2) + B_T(N_r|S_2) - B_T(N_r|S_1),$$

i.e. all trusted balance that would be transferred to N_r under the swaps leading to S_2 is now retained by N_s. Using this fact, together with the fact that $B_T(N_r|\tilde{S}) \leq B_T(N_s \to N_r|S_2) + B_T(N_r|S_2)$ by virtue of the direct transfer, we obtain

$$\begin{aligned}\max_{S':S\to_{N_s}S'} B_T(N_s \to N_r|S') &\geq B_T(N_s \to N_r|S_2') \\ &\geq B_T(N_s \to N_r|S_2) + B_T(N_r|S_2) - B_T(N_s|S_1) \\ &\geq B_T(N_r|\tilde{S}) - B_T(N_r|S_1) \\ &= B_T(N_r|\tilde{S}) - B_T(N_r|S) \\ &= T_G(N_s \to N_r|S).\end{aligned}$$

9.7.1 Transferrable trusted balance as maximum flow problem

In this section, we show that the problem of determining the transferrable trusted balance can be phrased as a maximum flow problem. This is great because this problem is one of the best-understood optimization problems and there exist efficient algorithms to solve them in practice. The problem involves finding the maximum amount of flow that can be sent from a source node to a sink node in a directed graph, subject to capacity constraints on the edges.

Definition: Maximum flow problem

An instance of a max flow problem is defined by the following:

- A directed graph $\mathcal{G} = (V, E)$, where V is the set of nodes and $E \subseteq V \times V$ is the set of directed edges.
- A capacity function $c : E \to \mathbb{R}_+$, where $c(u, v)$ is the maximum flow allowed on edge (u, v).
- A source node $s \in V$ and a target node $t \in V$.

A flow function $f : E \to \mathbb{R}_+$ is called *feasible* if it satisfies:

1. Capacity constraints:

$$f(u, v) \leq c(u, v), \quad \forall\, (u, v) \in E$$

1. Flow conservation:

$$\sum_{v\in V} f(v, u) = \sum_{v\in V} f(u, v), \quad \forall\, u \in V \setminus \{s, t\}$$

The goal of the max flow problem is to find a feasible flow that maximizes the *value* of the flow, defined as

$$V(f) := \sum_{v \in V} f(s, v)$$

The maximum flow problem admits the following conceptually helpful decomposition theorem:

Fact: Flow decomposition

Consider an instance $(\mathcal{G}, c, s, t)$ of the max-flow problem and any feasible flow f. Then there exists a collection of feasible flows $f_1, \ldots, f_I$ for some $I \leq |E|$ and a collection of paths $p_1, \ldots, p_I$ along $\mathcal{G}$, each of which starts in s and ends in t, such that

- $\sum_i V(f_i) = V(f)$
- $\sum_i f_i(u, v) \leq f(u, v)$, for all $(u, v) \in E$.

Using the flow decomposition theorem, one can show the following deep connection between evaluating the transferrable trusted balance and the maximum flow problem:

Fact: Transferrable trusted balance as maximum flow problem

Let $S \in \mathcal{S}$ be some state and $N_s, N_r \subseteq N$ be two disjoint subsets. We define the following instance of a maximum flow problem:

- $V = \{s, t\} \cup \{v_n\}_{n \in N} \cup \{V_{nk}\}_{n \in N, k \in K}$ is a set of consisting of $|N|(|K| + 1) + 2$ different nodes: the source and the target, one node for each account in N and one node for every combination of account and currency.
- $E = \{(s, v_n)\}_{n \in N_s} \cup \{(v_{n'}, t)\}_{n' \in N_r} \cup \{(v_n, v_{nk}\}_{n \in N, k \in K} \cup \{(v_{nk}, v_{n'})\}_{n, n' \in N, k \in G(n')}$ consists of edges going from the source to the senders, from the receivers to the target, from each account node to the corresponding combination nodes and from each combination node to all those accounts nodes that trust the currency in the combination.
- The capacity function is as follows:
  - $c(s, v_n) = c(v_{n'}, t) = \infty$
  - $c(v_n, v_{nk}) = c(v_{nk}, v_{n'}) = B_{nk}$
  - $c(v_{nk}, v_{n'}) = B_{nk}$

Let f denote the optimal flow for this instance. Then

$$V(f) = T_G(N_s \to N_r|S)$$

While this is by no means obvious , we here skip a formal proof of this statement. To state the main result of this section, the final ingredient we need is the max-flow min-cut theorem, which is a well-established and foundational theorem in graph theory

Max-flow min-cut theorem

Let $(\mathcal{G}, c, s, t)$ be an instance of a max-flow problem. An **s-t cut** is defined as an ordered partition (S, T) of the vertex set V into two disjoint subsets $S, T \subseteq V$ such that $S \cup T = V$, $S \cap T = \emptyset$, $s \in S$, and $t \in T$. The **capacity of the s-t cut** (S, T) is given by

$$c(S, T) := \sum_{(u,v)\in E \ u\in S, v\in T} c(u, v).$$

The Max-flow min-cut theorem states that for every instance of the maximum flow problem, the maximum value of a feasible flow from s to t is equal to the minimum capacity of any s-t cut. Formally:

$$\max_f V(f) = \min_{(S,T):s\in S, t\in T} c(S, T)$$

Using the equivalence between the max-flow mapping and the problem of obtaining the transferrable trusted balance, we can now show the following characterization of the transferrable trusted balance:

Characterization of transferrable trusted balance

For any disjoint sets $N_s, N_r \subseteq N$ and any state $S \in \mathcal{S}$,

$$T_G(N_s \to N_r|S) = \min_{\tilde{N}\subseteq N: N_s\subseteq\tilde{N}, N_r\subseteq N\setminus\tilde{N}} B_T(\tilde{N} \to N\setminus\tilde{N}|S).$$

Proof Sketch: The idea for the proof is simple. From the max-flow min-cut theorem and the equivalence above, it follows that the transferrable trusted balance is given by the minimal capacity over all N_s - N_r cuts for the instance defined above. Now, we can further partition the set of all these cuts into subsets such that all cuts in each partition element share the same account nodes v_n in both parts of the cut. Then, every partition element corresponds to a set of nodes $\tilde{N}$ that contains all sender nodes and no receiver node. In a second step, a simple argument shows

that the minimal capacity over all cuts in the partition element corresponding to $\tilde{N}$ is given by $B_T(\tilde{N} \to N \backslash \tilde{N}|S)$. As such, the transferrable trusted balance is then simply equal to the minimal of this quantity over all elements of the partition.

9.8 Relative Sybil resistance

In this section, we derive sybil resistance that is stated in the main text. Let us first restate it using notation that is adapted to the notation in this appendix:

Relative Sybil resistance

Let M be a set of accounts controlled by a malicious party, F a set of accounts not in M that are trust at least one account in M (i.e. the ingoing outer boundary of M) and let $R = N \backslash M \cup F$ be the remainder of the network. Then, if accounts in M initially hold no funds trusted by users in R, for any given state S,

$$T(M \to R|S) \leq B_T(F \to R|S).$$

Proof: This follows immediately from the characterization of the transferrable trusted balance in the last section, by setting $N_s = M, N_r = R$ and $\tilde{N} = M \cup F$:

$$\begin{aligned} T(M \to R|S) &\leq B_T(M \cup F \to R|S) \\ &= B_T(M \backslash \delta M \to R|S) + B_T(F \to R|S) \\ &= B_T(F \to R|S) \end{aligned}$$

9.9 Liquidity clusters and ASF

In this section, we provide the formal background to our discussion of liquidity clusters.

Definition: Fungibility order and liquidity clusters

For any two distinct currencies $k, k' \in K$, a state $S = (G, B)$ and some $l \geq 0$, we write

$$k \succ_{l,S} k'$$

if any set of l tokens of currency k can be translated by its holders into the same amount of currency k'. Formally, we require that for every account n such that

$B(n, k|S) \geq l$, there exists a state $S' = (G', B')$ such that $S' \in \mathcal{S}_n^T(S)$ and

$$B(n, k|S') = B(n, k|S) - l$$

$$B(n, k'|S') = B(n, k'|S) + l.$$

Moreover, to avoid vacuous truths, we require that there is at least $n \in N$ with sufficient funds of k and we define $k \succ_{l,S} k$ for any k. We call the family of orders $(\succ_{l,S})_{l \in \mathbb{R}_+}$ the *fungibility orders with respect to state* S. Finally, we define

$$k \sim_{l,S} k' :\Leftrightarrow k \succ_{l,S} k' \wedge k' \succ_{l,S} k.$$

We say that a set of currencies $\tilde{K} \subseteq K$ *forms a* l*-liquidity cluster in state* S if

$$k \sim_{l,S} k' \quad \forall\, k, k' \in \tilde{K}.$$

Note that we restrict the possible state transitions in the definition of the fungibility order to trust swaps. This is because we are interested in the ability of accounts in the network to interchange amounts of different currencies and if we included direct transfers, then our definition would not capture this intention (the same does not apply to trusting and untrusting, but they make no difference in this context). The following provides a conceptually simple sufficient criterion for the fungibility between two currencies:

Fact: Sufficient criterion for fungibility

For any two currencies $k, k' \in K$ and a state S, if there exists an integer $I > 0$, a sequence of currencies $(k_i)_{i=0}^I$ and accounts $(n_i)_{i=0}^I$ such that, $k_0 = k, k_I = k'$ and it holds that,

$$\begin{aligned} k_i &\in K_G(n_{i+1}), \quad i = 0, \ldots, I-1 \\ B(n_i, k_i|S) &\geq l, \quad i = 0, \ldots, I \end{aligned}$$

then $k \succ_{l,S} k'$.

Proof: Since fungibility holds reflexively by definition, we can assume w.l.o.g. that $k \neq k'$. Next, we construct a new sequence from the original sequence in which every pair (n_i, k_i) appears only once, by iterating through the pairs and truncating the original sequence between any pairs that appear more than once. Let us then assume w.l.o.g. that the sequence above already satisfies this property. Now, consider a single party $\tilde{n}$ such that $B(\tilde{n}, k) \geq l$. We consider the state transition

induced by a sequence of I trust swaps, where in the ith swap $\tilde{n}$ swaps an amount l of currency k_i against currency $k_{i+1}$with n_i. This trust transfer is possible by the assumptions above and the fact that every pair (n_i, k_i), so that previous swaps involving the same party cannot have affected the relevant holdings of the party n_i. At the end of the swaps, $\tilde{n}$ has transformed l k-tokens into l k'-tokens. Moreover, by assuymption we know that there exists at least one account, n_0 that holds sufficient funds initially, as required.

Average spendable fraction

In this section, we prove the stated relationships between average spendable fraction (ASF) and k-liquidity clusters from the main text. Recall that ASF is defined, using the notation of the appendix, as

$$ASF(\tilde{N}|S) = \frac{1}{|\tilde{N}|(|\tilde{N}|-1)} \sum_{n,n' \in \tilde{N}, n \neq n'} \frac{T_G(n \to n'|S)}{B(n, K|S)},$$

for any set of accounts $\tilde{N}$ and state S. We now formally define the fungibility order and l-liquidity clusters.

We are now in a position to prove the following:

Fact:

For any state $S \in \mathcal{S}$, let $\tilde{N} \subseteq N$ and $\tilde{K} \subseteq K$ be two sets such that every account in $\tilde{N}$ trusts at least one currency in $\tilde{K}$. Then, if $\tilde{K}$ forms a l-liquidity cluster in state S,

$$T(n \to n'|S) \geq \min\{\max_{k \in \tilde{K}} B(n, k|S), l\}, \quad \forall\, n, n \in \tilde{N} : n \neq n'$$

Proof: Consider any distinct $n, n' \in \tilde{N}$. Let $l^* = \max_{k \in \tilde{K}} B(n, k|S)$ and $k^* = \arg\max_{k \in \tilde{K}} B(n, k|S)$. Moreover, let $k' \in \tilde{K}$ be a currency that n' trusts. This currency exists by assumption. Since $\tilde{K}$ forms a l-liquidity cluster, n can achieve a final state S' in which they hold an amount of at least $\max\{l, l^*\}$ of currency k' from their holdings of k^*, using trust swaps. Hence,

$$T(n \to n'|S) \geq B(n, k'|S') \geq \max\{l, l^*\}.$$

Corollary:

For any state $S \in \mathcal{S}$, $\tilde{N} \subseteq N$ and $\tilde{K} \subseteq K$, the following hold:

If

- *every account in $\tilde{N}$ t$\tilde{N}\tilde{K}$,*
- *$\tilde{K}$ is a l-liquidity cluster,*
- *every account in $\tilde{N}$ holds at least l of some $k \in \tilde{K}$,*

then

$$ASF(\tilde{N}|S) \geq \frac{l|\tilde{N}|}{B(\tilde{N}, K|S)}.$$

Proof: We have

$$\begin{aligned} ASF(\tilde{N}|S) &= \frac{1}{|\tilde{N}|(|\tilde{N}|-1)} \sum_{n,n' \in \tilde{N}, n \neq n'} \frac{T(n \to n'|S)}{B(n, K|S)}, \\ &\geq \frac{l}{|\tilde{N}|} \sum_{n} \frac{1}{B(n, K|S)} \\ &\geq \frac{l|\tilde{N}|}{\sum_n B(n, K|S)} = \frac{l|\tilde{N}|}{B(\tilde{N}, K|S)}, \end{aligned}$$

where in the first step we used that $T(n \to n'|S) \geq l$, and in the second step we used the AM-HM inequality, which states that for any $n \geq 1$ and positive real numbers $a_1, a_2, \ldots, a_n$,

$$\frac{1}{a_1} + \frac{1}{a_2} + \cdots + \frac{1}{a_n} \geq \frac{n^2}{a_1 + a_2 + \cdots + a_n}.$$

9.9.1 Exchange rates

So far we have been concerned only with the possible state transitions under the rules of Circles. We now introduce the possibility of exchanging different kinds of currencies. We distinguish two scenarios:

"Internal" exchange for different currencies

In the first scenario, accounts have the possibility to exchange some of the currencies at a given rate $R : K \times K \to \mathbb{R}_+$, where $R(k \to k')$ denotes the number of k' tokens received in exchange for a single token of currency k. We write $S \to_{\tilde{N},R} S'$ for two states S, S' and and a set of accounts $\tilde{N}$ such that S' can be reached from S by means of a sequence of state transitions and exchanges that are initiated by accounts in $\tilde{N}$. We assume that these rates are provided by market makers and for now, we make the simplifying assumption that exchange rates don't change in the course of the exchanges and we also place no limit on the amount of tokens that can be exchanged at the rate. We expect that both of

these assumptions approximately hold when $\tilde{N}$ makes up only a small fraction of the total account base and their exchanges involve only a small subset of the total supply of the involved currencies.

For a given exchange rate function R, we can define the *wealth function* W via the mapping

$$W(\tilde{N}, k|S) := \sum_{k' \in K} B(\tilde{N}, k'|S) R(k' \to k)$$

so that $W(\tilde{N}, k|S)$ indicates the total amount of k tokens that accounts $\tilde{N}$ could obtain if they exchanged all their holdings into that currency. Using this wealth function, we can define the notion that a given set of exchange rates allows no arbitrage:

Definition: Absence of arbitrage

We say that an exchange rate function R *allows no arbitrage in state* $S = (G, B)$, if the following two conditions are satisfied:

1. *Account-facing:* It is impossible for any set of users to increase their wealth in any currency using any combination of state transitions, exchanges and borrowing. Formally, we require that for all $\tilde{N} \subseteq N, k \in K$, and every $B^C \in \mathbb{R}_+^{N \times K}$ that has support only over $\tilde{N}$,

$$W(\tilde{N}, k|S) = \max_{(G, B+B^C) \to_{\tilde{N},R} (G', B'+B^C)} W(\tilde{N}, k|(G', B')),$$

where as before B^C corresponds to tokens that members of $\tilde{N}$ have borrowed and that they need to return at the end of the transition.

1. *Market-maker facing:* Rates are *competitive* in the sense that it is impossible for market makers to make a profit from cyclic conversions: Let $(k_i)_{i=0}^I$ be any cyclic sequence of currencies such that $k_0 = k_I$, then we require that

$$\prod_{i=0}^{I-1} R(k_i \to k_{i+1}) \geq 1.$$

The absence of arbitrage implies that internal exchange rates have to reflect the presence of liquidity clusters, as we now show:

Internal exchange rates without arbitrage respect liquidity clusters

For any state $S = (G, B)$ with B not all zero, and any exchange rate function R. If R allows no arbitrage in S, then there exists a value function $V : K \to \mathbb{R}$ such that

$$R(k \to k') = \frac{V(k)}{V(k')}$$

and with the property that, for any $l > 0$,

$$k \succ_{l,S} k' \Rightarrow V(k) \geq V(k').$$

Proof: Assume that R allows no arbitrage in S. By the fact that that rates are competitive, R is strictly positive. Let n be a user that holds a positive balance. This user exists by the assumption that B is not all zero. Since n holds a positive balance and R is strictly positive, $W(n, k)$ is also strictly positive for any $k \in K$. Now, let $k_1, k_2 \in K$ be an arbitrary pair of currencies. Consider the sequence of transitions, in which n exchanges all their holdings for k_1, then for k_2 and then back to k_1. By definition of the wealth function and the account facing arbitrage property, it follows that

$$W(n, k_1) \geq W(n, k_1)R(k_1 \to k_2)R(k_2 \to k_1).$$

Using the fact that $W(n, k_1) > 0$ and combining this with the market-maker facing condition implies that

$$R(k_1 \to k_2) = \frac{1}{R(k_2 \to k_1)}$$

Now, finally, let $k_3 \in K$ be a third currency and consider the same cyclic sequence of transitions as above but now using k_3 as an additional intermediate currency. Then by the same argument, and using the above invertibility property,

$$\begin{aligned} R(k_2 \to k_3) &= \frac{R(k_2 \to k_1)}{R(k_3 \to k_1)} \\ &= \frac{V(k_2)}{V(k_3)}, \end{aligned}$$

where we defined $V(k) := R(k \to k_1)$ for any $k \in K$. The second property easily follows by contradiction: Assume that there exist k, k' and an $l > 0$ such that $k \succ_{l,S} k'$ but that $V(k) < V(k')$. Then, by definition of the fungibility ordering there exists an account n that holds some amount l of k tokens that they can convert to the same amount of k' tokens. Now, this user can perform the following sequence of transitions: First they turn l k tokens into k' tokens, then exchange them for k_1 tokens and then then exchange the k_1 tokens back to k tokens. This sequence of

transitions has strictly increased their holdings of k tokens, since

$$l(R(k' \to k_1)R(k_1 \to k) - 1) = l\left(\frac{V(k')}{V(k)} - 1\right) > l,$$

while it has left the remaining holdings of n unchanged. As such, by the positivity of R, it has also increased the wealth of n with respect to any currency, in contradiction with the assumed absence of arbitrage.

"External" exchange rate

We now change the scenario and assume that there exists an external asset z and that some market maker offers exchange rates $R(z \to k)$ and $R(k \to z)$ for all currencies $k \in K$. The exchange rate function R naturally induces an extension to "internal" exchange rates, via

$$R(k \to k') \equiv R(k \to z)R(z \to k')$$

Using this extended definition, we define the wealth function in the absence of arbitrage property exactly like in the previous section, only for the set $K \cup \{z\}$. As such, the characterization of the exchange rates through fungibility order respecting value functions also carries through. In particular, by setting, as we can without loss of generality, $V(k) = R(k \to z)$, the characterization implies that, in the absence of arbitrage, prices in the external assets will have to reflect the fungibility between tokens, in a given state.

Implications for accounts that create their own currency

In the previous sections, we have developed the theory of Circles, by treating accounts and currencies separately and without any special relationship between individual accounts and individual currencies. This was for conceptual clarity, but things simplify once we explicitly account for the fact that accounts create their own currency: In particular, let us now add the additional assumptions that $|N| = |K|$ and that there exists a bijective mapping $k : N \to K$ so that $k(n)$ is the currency that is created by account n. In this case, the balance matrix B and trust matrix G are $n \times n$ matrices and we sometimes write $G_{nn'}$ to indicate the trust relation between two accounts. Moreover, we assume that $k(n) \in K_G(n)$ for all $n \in N$, that is, all accounts trust their own currency. Then, we get the following simple implication of the above results

Fact: Value flows opposite to trust

For any state S in which every account holds some of their own currency, then exchange rates in the absence of arbitrage reflect the trust connections between users:

$$G_{nn'} = 1 \Rightarrow V(k(n')) \geq V(k(n)), \quad \forall n, n'$$

Proof: By virtue of the path-based sufficient criterion for fungibility, the assumption of trust and that $B(n, k(n)|S) > 0$ for all n, implies that $k(n') \succ_{l,S} k(n)$ for some amount $l > 0$, since n' can simply exchange their own tokens for that of n. Then, by the above statement about exchange rate implications of fungibility, the result follows.