Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 143 (2008) 36–58

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Informal insurance in social networks ☆

Francis Bloch [a], Garance Genicot [b,*], Debraj Ray [c,d]

[a] *GREQAM, Université de la Méditerranée and Warwick University, France*
[b] *Georgetown University, USA*
[c] *New York University, USA*
[d] *Instituto de Análisis Económico (CSIC), Spain*

Received 26 September 2006; final version received 28 January 2008; accepted 28 January 2008

Available online 24 May 2008

Abstract

This paper studies bilateral insurance schemes across networks of individuals. While transfers are based on social norms, individuals must have the incentive to comply. We investigate the structure of self-enforcing insurance networks. Network links play two distinct and possibly conflictual roles. They act as conduits for both transfers and information; affecting the scope for insurance and the severity of punishments upon noncompliance. Their interaction leads to a characterization of stable networks as suitably "sparse" networks. Thickly and thinly connected networks tend to be stable, whereas intermediate degrees of connectedness jeopardize stability. Finally, we discuss the effect of discounting on stability.
© 2008 Elsevier Inc. All rights reserved.

JEL classification: D85; D80; O12; Z13

Keywords: Social networks; Reciprocity network; Risk sharing; Norms; Informal insurance; Sparseness; Stable networks

☆ A first draft of this paper was written while Genicot and Ray were visiting the London School of Economics; we thank the LSE for their hospitality. Ray is grateful for funding from the National Science Foundation under grant No. 0241070. We are grateful for comments by seminar participants at CSIC (Barcelona), Cergy, CORE, Essex, Helsinki LSE/UCL, Oxford, Royal Holloway and University of Maryland and audiences at the NEUDC in Montréal, the AEA Meetings in Philadelphia, the SED in Budapest, conferences on networks and coalitions in Vaxholm and Guanajato, and the Conference in tribute to Jean-Jacques Laffont in Toulouse. Antoni Calvo-Armengol was extremely supportive of this research and we are grateful for his encouragement and constructive comments. This paper is dedicated to his memory.

* Corresponding author.

E-mail addresses: francis.bloch@univmed.fr (F. Bloch), gg58@georgetown.edu (G. Genicot), debraj.ray@nyu.edu (D. Ray).

0022-0531/$ – see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.jet.2008.01.008

1. Introduction

This paper studies networks of informal insurance. Such networks exist everywhere, but especially so in developing countries and in rural areas where credit and insurance markets are scarce and income fluctuations are endemic. Yet it is also true that everybody does not enter into reciprocal insurance arrangements with everybody else, even in relatively small village communities. A recent empirical literature (see for instance Fafchamps [12], Fafchamps and Lund [13], and Murgai, Winters, DeJanvry and Sadoulet [29]) shows that insurance schemes often takes place within subgroups in a community. One obvious reason for this is that everyone may not know one another at a level where such transactions become feasible, but—as Genicot and Ray [16,17] have argued—there may also be strategic reasons for limited group formation.[1] At the same time, even the supposition that insurance takes place in fully formed groups—small or large—may be open to question. A may insure with B, and B with C, but A and C may have nothing to do with each other. Empirically, such networks have attracted attention and have recently been mapped to some extent: for instance, Stack [32], Wellman [39], Dercon and de Weerdt [11], and Fafchamps and Gubert [14] reveal a complex architecture of risk-sharing networks. It is now widely recognized that social networks (based on kin, gender or occupation) play a dominant role in people's protection to risk in developing countries.

The very idea of an insurance *network* rather than a *group* suggests that our existing notion of insurance as taking place within an explicit "club" of several people may be misleading. Of course, such clubs may well exist, but a significant segment of informal insurance transactions is bilateral. A and B will have their very own history of kindness, reciprocity or betrayal. In these histories, either party may have been have cognizant of (and taken into account) her partner's obligations to (or receipts from) a third individual, but the fundamental relationship is nevertheless bilateral.

A principal aim of our paper is to build a model of risk-sharing networks which captures this feature. In the model studied here, only "directly linked" agents in some given[2] network make transfers to each other, though they are aware of the (aggregate) transfers each makes to others. Linked agents can observe only each other's commitments, but not necessarily the overall insurance scheme of the community.

Our approach takes an eclectic view to the extent of insurance between any linked pair. We view insurance as based on internalized *norms* regarding mutual help. A *bilateral insurance norm* between two linked agents specifies consumptions for every linked pair of individuals, as a function of various observables such as their identities, the network component they belong to, their income realizations and the transfers made to or received from other agents. These transfers are taken as given by the linked pair, but are obviously endogenous for society as a whole. We therefore introduce the notion of a *consistent* consumption allocation, one that allocates consumptions

[1] Genicot and Ray build on a large literature which studies insurance schemes with self-enforcement constraints; see, e.g., Posner [30], Kimball [22], Coate and Ravallion [9], Kocherlakota [24], Kletzer and Wright [23] and Ligon, Thomas and Worrall [27].

[2] In research that came to our attention after a first draft of this paper was written, Bramoullé and Kranton [6] study the formation of insurance networks under the assumption of equal division and perfect enforcement. By contrast, our paper studies a family of insurance schemes—including equal division—in an explicit context of self-enforcement, but assumes that the network is given for exogenous reasons such as friendship, family, or social contacts. In our model, links can be broken but new links can never be formed.

to everyone for each realization of the state, and which implicitly agrees with the bilateral norm for every linked pair.

With this setup as background, the paper then studies the stability of insurance networks, explicitly recognizing the possibility that the lack of commitment may destabilize insurance arrangements. Thus a consumption allocation cannot only be consistent (with the underlying norm); it must also be *self-enforcing*.

But precisely what does self-enforcement entail? In the "group-based" insurance paradigm, a natural supposition is that a deviating individual is thereafter excluded from the group, and that is what the bulk of the literature assumes. Yet if arrangements are fundamentally bilateral, this sort of exclusion needs to be looked at afresh. If A deviates from some arrangement with B, we take it as reasonable that B refuses to engage in future dealings with A.[3] The payoff consequences of this refusal may be taken to be the weakest punishment for A's misbehavior. But the punishment may conceivably be stronger: B might "complain" to third parties. If such parties are linked directly to A they, too, might break their links (such breakage would be sustained by the usual repeated-game style construction that zero interaction always constitutes an equilibrium). To go further, third parties might complain to fourth parties, who in turn might break with A if they are directly linked, and so on. Such complaints will travel along a "communication network" which in principle could be different from the network determining direct transfers, but in this paper we take the two networks to be the same.[4] If all agents are indirectly connected in this way, then the limiting case in which all news is passed on—and corresponding action taken—is the one of full exclusion typically assumed in the literature. We propose to examine the intermediate cases.

We focus our study on this informational effect, and initially study cases in which individual discount factors are close to unity. Our main result (Proposition 3) provides a full characterization of those insurance networks that satisfy the self-enforcement constraint for different "levels" of communication. By "level" we refer to the number of rounds q of communication (and consequent retribution) that occur following a deviation. For instance, if the immediate victim talks to no one else, $q = 0$, if she talks to her friends who talk to no one else, then $q = 1$, and so on. For any such q, we provide a characterization of those network architectures that are stable under the class of *monotone* insurance norms, those in which the addition of new individuals to a connected component by linking them to one member increases that member's payoff. The characterization involves a particular property of networks. As an implication, for any q, typically *both* thinly and thickly connected networks are most conducive to stability; intermediate degrees of connection are usually unstable.

To obtain an intuitive feel for this implication, fix a particular value of q. Now, even if this value is small, when the network is very thin the miscreant may still be effectively cut off (and thereby adequately punished) even though the accounts of his deviation do not echo fully through the network (he was tenuously connected anyway). On the other hand, if the network is fully connected a single round of complaints to third parties is *also* enough to punish the miscreant, because there will be many such "third parties" and they will all be connected to him. It is precisely in networks of intermediate density that the deviant may be able to escape adequate punishment.

---

[3] Of course, issues of renegotiation might motivate a reexamination of even this assumption, but we do not do that here.

[4] Certainly, it is perfectly reasonable to suppose that the communication network is a superset of the transfers network. From this perspective, the additional restriction that we impose is that all pairs who can talk can also make transfers.

This suggests a *U-shaped relationship* between network density and stability for intermediate level of communication. To be sure, this is not a one-to-one relationship as networks of very different architecture have the same average density. However, by looking at the proportion of stable networks (of a given size) of different density levels, we can illustrate this U-shaped relationship. We show related results for networks with different degrees of clustering. This is particularly useful as the density and clustering coefficients (concepts that we will define precisely later) are basic characteristics of networks used in the social network literature (Wasserman and Faust [37]).

For values of the discount factor lower than unity, the short-run gains from deviation also play an important role. The architecture of the network also matters in calculating such gains. Consequently, the stable mutual insurance scheme networks are harder to characterize. We consider an example to illustrate the wide range of possibilities arising for different values of the discount factor, and the complexity of the relation between agents' degree of impatience and stable network architectures.

We believe that this paper represents a first step in the study of self-enforcing insurance schemes in networks. In taking this step, we combine methods from the basic theory of repeated games, which are commonly used for models of informal insurance, with the more recent theories of networks. It appears that this combination does yield some new insights, principal among them being our characterization of stable networks. However, it is only fair to add that we buy these insights at a price. For instance, it would be of great interest to study the case in which the aggregate of third-party transfers is not observable. This would introduce an entirely new set of incentive constraints, and is beyond the scope of the present exercise.

Qualifications notwithstanding, our findings contribute to a recent and growing literature on the influence of network structures in economics. See for instance, Calvó-Armengol and Jackson [7] on labor markets, Goyal and Joshi [18] on networks of cost-reducing alliances, Bramoullé and Kranton [5] on public goods, Tesfatsion [33,34] and Weisbuch, Kirman and Herreiner [40] on trading networks, Fafchamps and Lund [13] on insurance, Conley and Udry [10], Chatterjee and Xu [8] and Bandiera and Rasul [2] on technology adoption, and Kranton and Minehart [25,26] and Wang and Watts [36] on buyer-seller networks.

2. Transfer norms in insurance networks

2.1. Endowments and preferences

Consider a community of individuals occupying different positions in a social network (more on networks below). At each date, a state of nature θ (with probability $p(\theta)$) is drawn from some finite set Θ. The state determines a strictly positive endowment y_i for agent i. Denote by $\mathbf{y}(\theta)$ the vector of income realizations for all agents. Assume that every possible inter-individual combination of outputs has strictly positive probability. [This condition guarantees, in particular, that outputs are not perfectly correlated.]

Agent i is endowed with a smooth, increasing and strictly concave von Neumann–Morgenstern utility u_i defined over consumption, and a discount factor $\delta_i \in (0, 1)$. Individual consumption will not generally equal individual income as agents will make transfers to one another. However, we assume that the good is perishable and that the community as a whole has no access to outside credit, so aggregate consumption cannot exceed aggregate income at any date.

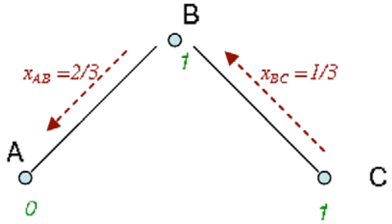

Fig. 1. Bilateral transfer scheme.

2.2. Networks

Agents interact in a social network. Formally, a network g is a set of agents (nodes) $N = \{1, \dots, n\}$ and a graph—a collection of pairs of agents—with the interpretation that the pair ij belongs to g if they are directly linked. In this paper, a bilateral link is a given: it comes from two individuals getting to know each other for reasons exogenous to the model. While such links may be destroyed (for instance, due to an unkept promise), no new links can be created.

Note that two individuals are *connected* in a network if they are directly or indirectly linked and a network is connected if *all* its members are connected. The *components* of a network are the largest subsets of connected individuals and their set of links. A *subnetwork* of g is a network g' among the same set of agents obtained by removing links from g.

For our purposes, a link between i and j means two things. First, it means that i and j can make transfers to each other. Second, it is a possible avenue for the transmission of information. We take these matters up in more detail below.

2.3. Bilateral norms

In contrast to existing literature, we adopt a decentralized view of insurance. Any two linked individuals may insure each other. This implies some degree of insurance for larger groups, but no deliberate scheme exists for such groups. To be sure, transfers from or to an individual must take into account what her partner is likely to receive from (or give to) third parties. In many situations this is easier said than done. Such transfers may not be verifiable, and in any case all transfers are made simultaneously. As a first approximation we assume that for every linked pair, and for every realization of the state, each partner's transfer receipts—positive or negative—from third parties are fully observed, as are partner incomes.

A *bilateral* (insurance) *norm* is a specification of consumptions—and therefore an implicit specification of transfers—across every linked pair $\{ij\}$ of individuals, as a function of observables: individual identities i and j, realized incomes y_i and y_j, and transfers from (or if negative, to) third-parties z_i and z_j. So a bilateral norm is just a function b such that

$$(c_i, c_j) = b(i, j, y_i, y_j, z_i, z_j)$$

for every vector of realizations, subject to the constraint

$$c_i + c_j = y_i + z_i + y_j + z_j.$$

For instance, Fig. 1 illustrates transfers that fully equate consumptions in a state in which B and C's income is 1 while A's income is 0.

Some norms could be derived from bilateral welfare functions. Such a welfare function would depend on state-contingent consumptions of the two agents, but could also depend on other variables, such as the ambient network component and the identity of the agents. Each pair—viewed provisionally as a single entity—would generate a bilateral norm by maximizing this function. For instance, if the bilateral welfare function consists of the sum of the utilities of the individuals, the resulting bilateral norm is *equal sharing*: the bilateral norm simply divides all available resources among the linked pair.

But bilateral norms also include a large class of sharing rules which are not easily amenable to a welfarist interpretation. For instance, suppose that each individual i has full, unqualified access to some fraction α_i of her income, and must only share the rest of her resources using, say, an even split. The resulting transfer norm would then look like this:

$$c_i = \alpha_i y_i + \frac{1}{2}\left[\sum_{k=i,j}(1-\alpha_k)y_k + z_k\right],$$

$$c_j = \alpha_j y_j + \frac{1}{2}\left[\sum_{k=i,j}(1-\alpha_k)y_k + z_k\right].$$

We may refer to these as *norms with a private domain*, the private domain in question being the inalienable quantity $\alpha_i y_i$ for each i.

Say that a bilateral norm *aggregates third-party obligations* if the consumption of each individual depends on z_i and z_j through their sum $z_i + z_j$ alone and is continuously increasing in this variable.

Bilateral norms that aggregate third-party obligations can still be asymmetric (i.e., depend on the index i and j), and they can also prescribe consumptions that are dependent on individual incomes in a variety of ways. Naturally, there is a large class of norms that satisfy this restriction. For instance, the equal sharing norm and the norms with a private domain do aggregate third-party obligations.

2.4. *Consistent consumption allocations*

Observe that all third-party transfers are endogenous functions of the state. Put another way, the society-wide operation of a bilateral norm not only generates transfers across every linked pair; it also generates the third-party receipts or obligations that the pair takes as "given." If we can cut through the implied circularity, a bilateral norm will yield a consumption allocation $\mathbf{c}(\theta)$ for everyone in the network, as a function of the realized state θ. Call this a *consumption allocation consistent with the norm*, or a consistent consumption allocation for short. Obviously, with all the interpersonal interactions in the network, there may be more than one consumption allocation consistent with any given bilateral norm. The following proposition describes when this cannot happen, so that a unique prediction is obtained.

Proposition 1. *Suppose that a bilateral norm aggregates third-party obligations. Then there is at most one consumption allocation consistent with that norm.*

Proof. Suppose the assertion is false. Then there are two consistent allocations—$\mathbf{c}(\theta)$ and $\mathbf{c}'(\theta)$—and a state θ such that the induced vectors of consumptions across individuals in that state are distinct. Then there must be some linked pair ij such that $c_i(\theta) \leqslant c'_i(\theta)$ and $c_j(\theta) > c'_j(\theta)$. But then at least one consumption is *not* continuously increasing in the sum of third-party obligations, a contradiction. □

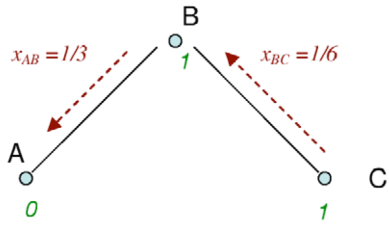

Fig. 2. Consistent allocation.

Recall the equal-sharing norm, in which every bilateral transfer is chosen to equalize consumption across a linked pair. It is easy to see that there is a unique consistent scheme associated with the equal-sharing norm, which entails "global" equal sharing of total output in any component of the network.

The equal-sharing norm, apart from its intrinsic interest, has the feature that there is some "multilateral norm" with which it is consistent; in this case, multilateral equal sharing. Bilateral transfer norms that allocate to each person a weighted share of consumption (depending perhaps on that person's identity or her income realization) also have this feature provided that the relative weights for every pair $\{ij\}$ equal the relative weights arrived at "indirectly" by compounding relative weights along any other path joining i to j. It is also possible to generate bilateral schemes by maximizing welfare functions that depend (for every linked pair ij) on $\mathbf{c}_i$ and $\mathbf{c}_j$. Such functions may also depend on the ambient network component and the identity of the agents, as also their income realizations, particularly if these serve as proxies for "outside options."

It is also easy to find the consistent consumption allocation associated with any norm with a private domain. For a given income realization, any individual i in a component d of size n will consume

$$c_i = \alpha_i y_i + \frac{1}{n}\sum_{j\in d}(1-\alpha_j)y_j.$$

As a particular instance, take the norm in which all individuals keep half their income and share the remainder of their resources ($\alpha_i = \frac{1}{2}$ for all i). Consider a three person network in which B is linked to both A and C. Fig. 2 illustrates the transfers associated with the consistent allocation $(1/3, 5/6, 5/6)$ in a state in which B and C's income is 1 while A's income is 0.

Now what about existence? Unfortunately, bilateral norms could be "incompatible enough" so that a consistent consumption allocation associated with that norm simply fails to exist. As an example, suppose there are three agents connected to each other in a circle. Assume that players 1 and 2 have a social norm that involves giving player 2 two-thirds of their joint endowment. Likewise, players 2 and 3 wish to give player 3 two-thirds of their joint endowment, and a symmetric circle is completed by players 3 and 1. Obviously, there is an incompatibility here, and it manifests itself in the nonexistence of a consistent consumption allocation.[5]

One way of seeing this incompatibility is to "follow" a natural fixed point mapping that would generate a consistent consumption allocation, were one to exist. If no consistent consumption allocation were to exist, such a mapping would prescribe larger and larger transfers as "best responses" in an ever-increasing spiral. If one proscribes unbounded transfers by assumption,

[5] The incompatibility does not arise from the weights alone: the linkage structure matters as well. For instance, there is no existence problem under this norm if the three agents are connected in a "line."

the problem goes away. We record this as a proposition, for it tells us that there are no *other* existence problems except for the one just described.

Proposition 2. *Suppose that for every linked pair ij, the bilateral transfer norm is continuous in $\mathbf{z}_i$ and $\mathbf{z}_j$, and that, in any state, the prescribed transfers cannot exceed some exogenous upper bound (say, the total output produced in society in that state). Suppose, moreover, that the norm never prescribes positive transfers from an individual with non-positive consumption to another with positive consumption.*

Then a consistent consumption allocation exists, and exhibits positive consumption for every individual at every state.

We relegate the proof to Appendix A.

Notice that every consistent consumption allocation $\mathbf{c}$ also implies an associated *transfer scheme* $\mathbf{x}$: a collection of payments $x_{ij}(\theta)$ from j to i (positive or negative) for every linked pair ij and every state θ. Observe that there may be several transfer schemes associated with the same consumption scheme: the uniqueness result of Proposition 1 does not apply to transfers.

We complete this section with a discussion of monotone norms.

2.5. Monotone norms

Say that a bilateral norm is "monotone" if, whenever more individuals are brought into a connected network by being connected to any particular individual, *that* individual's payoff increases. Intuitively, more individuals create better insurance possibilities, and a monotone norm should give some of the extra benefits to the individual serving as a "bridge."

More formally, consider a network g with m components—name them $1, 2, \ldots, m$—and an individual i in some component c. Now consider a network g' created by adding links from i to individuals j in other components ($j \notin c$). A norm is *monotone* if for every pair of associated consistent consumption allocations, one for g and one for g', the expected payoff to i under g' is higher than that under g.

Notice that monotonicity embodies more than a purely normative definition; it requires that a consistent solution not "misbehave" as we move across networks.[6]

It is easily seen that when agents are symmetric the equal division norm is monotone. It would be interesting to describe the full class of norms which satisfies monotonicity, though we do not know the answer to this question.

3. Enforcement constraints and stability

While a bilateral norm, as defined by us, comes from a fairly general class, it is time to emphasize a particular feature (already discussed in the Introduction). These norms *are* largely "normative" in that they take little or no account of self-enforcement constraints. But this is not to say that such constraints do not exist. Each individual may recognize that as a social being she is constrained to abide by the transfer norm in her dealings with j, *provided that she wants to maintain those dealings*. But she may not want to maintain them. It may be that (in some states)

[6] One possible source of "misbehavior" is nonuniqueness of consistent schemes for a given network. While this in itself is by no means inconsistent with monotonicity, it makes the concept less intuitive.

the transfer she is called upon to make outweighs the future benefits of maintaining a relationship with j under the bilateral norm. If that is the case, something must give, either the norm or the ij link. Our paper takes the point of view that the norm is more durable than the link, and that the link will ultimately fail.[7]

In a network setting, an agent could choose to renege on some (or all) transfers that she is required to make under a particular bilateral norm. In line with the bulk of the literature on risk-sharing without commitment (see, e.g., Coate and Ravallion [9], Fafchamps [12], Ligon, Thomas and Worrall [27] and Genicot and Ray [16,17]), individuals who are the *direct* victims of a deviation are presumed to impose sanctions on the deviant thereafter by not interacting with them.

This much may be clear, but nevertheless the extent of the punishment imposed on a deviant remains ambiguous. What about the rest of society, who were not directly harmed by the deviant? Do they, too, sever links with the deviant?

The answer to this question depends in part on what we are willing to assume about the extent of information flow in the society. In turn, this forces us to ask the question of just what the network links precisely mean. They certainly limit physical transfers, but do they also limit the flow of information? One possible interpretation is that the network represents a set of physical conduits and physical conduits alone, while information flows freely across all participants and is not constrained in any way by the architecture of the network. In this case the following notion of a punishment may be appropriate:

Strong punishment. Following a deviation, every agent severs its direct link (if any) with the deviant, so that the deviant is thereafter left in autarky.

In models of informal insurance in groups with self-enforcement constraints, this is the commonly adopted punishment structure. But in such scenarios, there are no networks, insurance is fully multilateral, and the event of a deviation is common knowledge among the group as a whole. In a situation in which network effects are under explicit consideration, the opposite presumption may seem more natural:

Weak punishment. Following a deviation, *only* those agents who have been directly mistreated by the deviant sever their links (with the deviant).

In our view, this concept is more appropriate to the case at hand than strong punishment. In the model that we study, insurance is bilateral, and linked agent pairs know relatively little about the particulars of other dealings (only the aggregate of transfers made to or received from third parties).[8] So it is entirely consistent to impose the restriction that while directly injured parties react, other agents do not, while strong punishments are more appropriate to a multilateral situation in which there are no restrictions on information flows and no network effects.

[7] Generally speaking, should we conceive of norms as restricted or unrestricted by incentives? This is an important open question that we do not pretend to address in any satisfactory way. Norms may range all the way from the fully idealistic (purely derived from ethical considerations, such as equal-sharing) to the purely pragmatic (wary of all enforcement and participation constraints, with ethical matters only invoked subject to the limits posed by such constraints). In this paper, we take the point of view that norms are not constrained by incentives, but of course we do use such incentive constraints to see if the resulting bilateral norms will or will not survive.

[8] However, for our model to have proper game-theoretic underpinnings—the formalities of which we do not explore here—it will be necessary to assume that the network itself is commonly known.

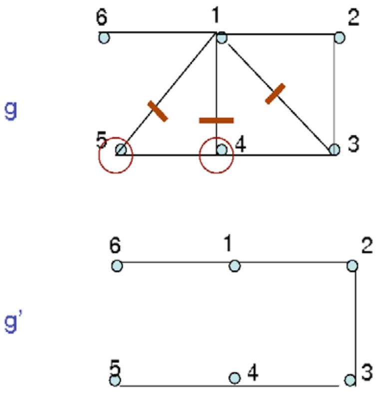

Fig. 3. Level-q punishment.

At the same time, if we take the network structure seriously, not just as a routeway for physical transfers but also for the flow of information, then we can define “intermediate” layers of punishment that are worth investigation in their own right. For instance, if I am an injured party and can communicate with those I am directly linked to, I can tell them about my experience. One might then adopt the equilibrium selection rule that all the individuals I talk to sever direct links (if any) with the deviant.

To be sure, once this door is opened, we might entertain notions in which the news of an individual’s mistreatment “radiates outwards” over paths of length that exceed a single link, and all those who hear about the news breaks off direct links (if any) with the original deviant. There are many ways to model such a scenario: we take the simplest route by indexing such punishments by the length of the required path.

Level-q punishments. Following a deviation, all agents who are connected to a victim by a path not exceeding length q (but not via the deviant) sever direct links (if any) with the deviant.

Fig. 3 illustrates this punishment scheme for $q = 1$. If in a given period individual 1 reneges on the transfers he owes to 4 and 5, then not only 4 and 5 but also 3 will sever their links to 1. This results in the subnetwork g'.

In this definition, q is to be viewed as a nonnegative integer, so that weak punishment may be thought of as a special case in which $q = 0$. In this sense, level-q punishments are quite general.[9] In Section 5.2, we discuss some other punishment structure.

Notice that we effectively treat q as an exogenous parameter to measure the extent of information flow. The reader may ask whether the passing-on of news about deviations may be detrimental to one’s own interests (see Lippert and Spagnolo [28] for more discussion on this point). This is an important question, but the issue is not critical here as our implicit punishment scheme is sustainable as a Bayesian equilibrium in a game in which network members send messages along their communication network (determined by q) to report whether the transfer was made or not. No individual will make a positive transfer to another if they do not expect reciprocity. It follows that if a deviant believes that one of his neighbors has been informed of his

[9] Of course, one might conceive of still more general punishment structures in which verifiable information decays—perhaps probabilistically—as it radiates along a path, but we avoid these for the sake of simplicity.

misbehavior, he will stop making transfers to her anyway. Given these beliefs, it is an equilibrium strategy for individuals along the path from the victim to the neighbor to transmit the information on the deviation.[10]

Now, given a level-q punishment structure in place and given a norm, we may define q-stable networks. (Sometimes, when there is no danger of ambiguity, we shall simply use the term "stable" instead of q-stable.) Consider a community of n individuals and a bilateral norm defined over all possible pairs of individuals. Because deviations result in the severance of links, we develop a recursive notion of stability. To this end, we begin with the empty graph in which all individuals live in isolation. The expected lifetime utility of an agent living in autarchy (normalized by the discount factor to a per-period equivalent) is

$$v_i^*(\emptyset) \equiv \sum_\theta p(\theta)u\big(y_i(\theta)\big).$$

The empty graph is obviously stable, as no transfers are called for and there is no question of a deviation. And there is only one payoff vector associated with that stable graph: $\mathbf{v}^*(\emptyset) \equiv \{v_i^*(\emptyset)\}$.

Proceeding recursively, consider a network g representing our n individuals and their links. Suppose that the set of stable *subnetworks* of g has been defined, along with collections of stable payoff vectors for each of those stable subnetworks. Now consider the network g, and pick a consistent consumption allocation $\mathbf{c}$ with expected payoff vector $\mathbf{v}$. Fix a transfer scheme $\mathbf{x}$ associated with it. Consider any individual i. For any realization θ, by abiding with the norm, i obtains a lifetime (normalized) expected payoff of

$$(1-\delta)u\big(c_i(\theta)\big) + \delta v_i. \tag{1}$$

In contrast, if i deviates by not honoring commitments to a set of neighbors S, a level-q punishment will set the new network to g', which is obtained by removing from g all direct links to i that are from individuals who are no more than q steps distant from some member of S (the connecting steps should not "pass through" i). Thus, the continuation payoff will depend on both the set of players who are her "victims," and the punishment level. Formally, the payoff to i of such a deviation is

$$(1-\delta)u\big(c_i(\theta,S)\big) + \delta v_i'(S,q). \tag{2}$$

The first term of the payoff in (2), $c_i(\theta,S)$, is the contemporary consumption from the deviation. At state θ, player i cheats on a subset S of links by selecting a transfer which is different from the transfer prescribed by the bilateral norm. In some cases, player i will be able to get away with no changes in transfers outside S, so that $c_i(\theta,S) = c_i(\theta) + \sum_{j\notin S} x_{ij}(\theta)$. Otherwise, just as in the literature on financial contagion and bank failures (Rochet and Tirole [31], Allen and Gale [1], Freixas, Parigi and Rochet [15]), the deviation by player i may trigger a chain reaction, and affect other transfers coming to him. Then $c_i(\theta,S)$ has a more complex expression that reflects his current consumption "at the end" of that process, but its exact form is unimportant.

How is the second term $v_i'(S,q)$ of the payoff in (2) determined? To describe the continuation payoffs following a deviation, we must adopt a convention that tells us the payoffs that accrue to player i when she finds herself at some network $g' \subset g$.

[10] Of course, there may be other equilibria in which badly treated partners resume dealings with the deviant, or fail to pass the information, or non-victims identify a deviant and also sever links with him, but we do not consider them here. Individuals who are direct victims or are informed of a deviation—and only these individuals—are presumed to transmit the information and impose sanctions according to the punishment scheme. We return to this point below.

If g' is stable, it is to be expected that i will enjoy a payoff of v_i', where this is the ith component of some stable payoff vector for g'. If g' is not stable, the resulting payoff will be presumably drawn from some stable subnetwork of g' itself. Two often-used devices to pin down the precise outcome in the face of potential multiplicity are "optimistic" and "pessimistic" beliefs (see, e.g., Greenberg [19]).[11] We assume that if g' is not stable then a subnetwork g'' would form where g'' is the or one of the largest stable subnetworks of g' to which i belongs. In this case, v_i' is the ith component of some stable payoff vector for g''. We do not insist on any particular selection rule at this conceptual stage, but we must take note of the "baseline" graph that player i induces on her deviation. This depends on two things: the set S of players who are her "victims," and the value of q that determines the punishment level.

Our formulation allows for a broad spectrum of interpretations, but it does require that the victims of a deviation be able to unambiguously identify the deviant, so that all the links specified by q-stability can indeed be severed. This requirement makes complete sense provided that insurance norms are bilateral and third-party transfers are fully observed. At the more abstract level of consumption allocations, absent any other information, the observation of a consumption allocation $c_i'(\theta) \neq c_i(\theta)$ does not necessarily allow a victim to identify the initial deviator. So the bilateral formulation helps us here.

Agents who were *not* the direct victims of a deviation may or may not be able to identify a deviant. The case most conducive to identification is that of three agents arranged along a line: $i - j - k$. Suppose that agent j reneges on his commitment to agent k. Depending on the bilateral norm in place and the distribution of income shocks, agent i may be able to infer that j has, in fact cheated.[12] Be that as it may, our formulation presumes that only the direct victims of a deviation, and those she informs, will break links with the deviant. We intend this formulation as a tractable shorthand to capture the effects of limited informational flow, but formally it may be viewed as a particular equilibrium selection device. It is presumed that the deviant will expect (and therefore reciprocate with) full noncooperation from all direct victims and their (q-step) neighbors, and from *only* this set: with the remainder, the bilateral norm continues to apply.

Comparing (1) and (2), we may therefore say that a network g is *q-stable* under a given bilateral norm if it has a consistent consumption allocation (with expected payoff vector $\mathbf{v}$) and associated transfer scheme such that for every player i, every state θ and every set of direct neighbors S of i,

$$u\big(c_i(\theta, S)\big) - u\big(c_i(\theta)\big) \leqslant \frac{\delta}{1-\delta}\big[v_i - v_i'(S, q)\big]. \tag{3}$$

This inequality requires that the short term deviation gain from not making the transfers be smaller than the long term gain from remaining in the original risk-sharing network.

If q is small, then punishments are "weak": not many individuals punish the miscreants. Indeed, for $q = 0$ we obtain precisely the notion of weak punishment introduced earlier. If q is large we approximate strong punishment and indeed that is what we get if the network is sufficiently connected to begin with. In any case, "strong stability" can always be defined by setting v_i' in (2) to $v_i^*(\emptyset)$; no recursion is needed.

[11] For instance, in the former case, v_i' would be the maximum value of v_i drawn from all stable payoff vectors drawn from any (stable) subgraph of g'.

[12] In other situations, non-victims may be unable to detect a deviation. Consider for example a star among four agents i, j, k, l with i as the hub. Suppose that the bilateral norm is the equal sharing norm, and that players j and k experienced a positive shock while players i and l experienced a bad shock and expect to receive a transfer. If player j deviates and fails to transfer money to the hub, agent l will be unable to detect whether j or k was the deviator.

Notice that q-stability exploits fully an awareness of repeated interaction between individuals and is therefore different from the stability concepts in Jackson [20,21] and Bala and Goyal [3].

4. Stable networks for high discount factors

In this section we characterize the set of q-stable networks for a given bilateral norm. The following notion of "sparse connectedness," related to the length of minimal cycles connecting any three agents in a network, will be central to the analysis.[13]

4.1. Sparseness

Consider a network g. For any triple of agents (i, j, k) such that (i, j) and (j, k) are directly linked to each other in the network ($ij, jk \in g$), compute the length of the *smallest cycle* connecting i, j and k, $\ell(i, j, k)$. By convention, if there is no cycle connecting those three points, we let $\ell(i, j, k) = 2$. For any integer $q \geqslant 0$, say that a network g is *q-sparse* if all minimal cycles connecting such triple of agents in the network have length smaller or equal to $q + 2$. If there is no connected triple of agents in the network, it is 0-sparse.

Observation 1 lists some of the properties of q-sparseness, but prior to stating it we need to introduce a couple of concepts.

We say that a network is composed of trees if it *only* has trees as its components. The *clusters* of a network g, $\{C_1, \ldots, C_M\}$, are the largest subsets of agents in g that are fully connected (have direct links to each other). Note that a cluster can consist in a singleton, so that every node belongs to at least one cluster. A *bridge-node* is a node such that removing the agent at that node and all her links increases the number of components in the remaining network. We say that a network g is composed of clusters connected by bridges if its clusters are connected only possibly through bridge-nodes.

Observation 1.

(a.) A network is q-sparse, then it is q'-sparse for all $q' \geqslant q$;
(b.) A network is 0-sparse iff it is composed of trees;
(c.) A network is 1-sparse if and only if it is composed of clusters connected by bridges.

As an illustration, Table 1 describes the minimal sparseness of the different network architectures for 8 connected agents illustrated in Fig. 4. The line is a tree so it is at least 0-sparse. Networks b and d are composed of clusters connected by bridges and therefore have a sparseness of 1 and higher. To be sure, the circle is the graph architecture for which the minimal sparseness is the largest. Indeed in a circle of size n, any three agents are connected by a cycle of length n, so that the circle is q-sparse only for values of $q \geqslant (n - 2)$.

4.2. Characterization of q-stability

When agents are sufficiently patient, the stability of a network has a simple characterization in terms of sparseness.

[13] We are grateful to Anna Bogomolnaia whose comments suggested this definition to us.

Table 1
Sparseness

	Graph	Minimal sparseness
a.	tree	0
b.	circle	6
c.	2 neighbors	2
d.	complete graph	1
e.	a bridge	1

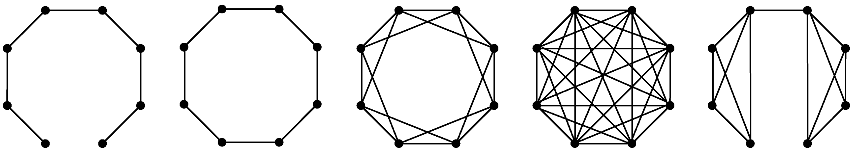

Fig. 4. q-sparseness & cycles.

Proposition 3. *Suppose that a bilateral norm is monotonic and aggregates third-party obligations. Then a network g is q-stable for all discount factors sufficiently close to* 1 *if and only if it is q-sparse.*

It is very important to note that q-sparseness is a purely graph-theoretic property. It requires no knowledge of utility functions, of the set of possible income realizations, nor of the stochastic law that governs such realizations. In contrast, q-stability is a more complex game-theoretic notion. Not only do we need knowledge of utilities and endowments to define the concept, we need the usual repeated game apparatus of deviations and punishments. Therefore, despite the linguistic similarity in terminology, q-stability and q-sparseness are very distinct concepts, and indeed this is why our characterization is potentially useful.

Proof. The proof of necessity uses two steps. The first step shows that for the bilateral norms under review, no punishment can be imposed on a deviant unless it is severe enough to break down the component of which the deviant is part.

Lemma 1. *Suppose that a bilateral norm aggregates third-party obligations. Then, assuming that a consistent consumption allocation exists for a connected network g, exactly the same allocation is consistent for every connected subnetwork of g.*

For a formal proof, see Appendix A. Bilateral norms that aggregate third-party obligations manage to impose the same consumption structure over every set of agents, provided that they are connected. The *particular* structure of connectedness does not matter. For instance, the equal division bilateral norm actually gives rise to *global* equal division over any set of agents, as long as they form a connected network.

The second main step is a restatement of q-sparseness.

Lemma 2. *A network g is q-*sparse *if and only if, for every linked pair $ij \in g$, the network formed by removing from g the links to i along all paths of length $m \leqslant q+1$ between i and j has strictly more components than g.*

Once again, the proof is postponed to Appendix A. Now connect Lemmas 1 and 2. Consider a network and a bilateral norm that aggregates third-party obligations. Assume that an agent i reneges on the transfers she owes to one or more partners. Then the immediate victims certainly sever their links to i, and so do all individuals (connected to i) who are connected to any of them (but not through i!) via a path of length q or less. If the graph g is *not* q-sparse, then, in the resulting subnetwork, i will *not* find herself in a smaller component, by Lemma 2. By Lemma 1, there can be no reduction in continuation utilities following the deviation, so q-stability is not possible.

The sufficiency direction is trivial. For monotone norms, q-sparse networks must cause a discrete loss in expected utility (Lemma 2 again). This is enough to yield q-stability for all discount factors close enough to 1. □

4.3. Strong and weak punishment

It is particularly interesting to contrast the stability of networks across different punishment schemes. For values of δ close to 1, *every* network is strongly stable. But this is no longer true as q comes down, weakening the flow of information. The following corollary summarizes the two extremes.

Corollary 1. *Suppose that the bilateral norm is monotonic and aggregates third-party obligations. Then a network is stable under weak punishment for discount factors close to unity if and only if it has only trees as components. In contrast, under strong punishment, every network is stable for high enough discount factors.*

4.4. Density and stability

Fig. 4 and Table 1 suggest that there is no monotonic relation between the sparseness of a network, and its density as measured by the number of links. Very thin networks such as trees and very dense networks such as the complete graph tend to have low sparseness indices. Intermediate graph structures on the other hand have higher sparseness indices, suggesting a U-shaped relation between sparseness and the number of links. For instance, for $q = 1$, only the line (a.), the complete graph (d.) and the bridge (e.) architectures are stable.

In order to test this hypothesis, we study the relationship between q-stability and density in networks of size six. The density is simply measured by the average number of links that individuals have (a number ranging from 0 to 5). We generate every possible network of six individuals. We compute the density of each possible network, as well as the lowest value of q for which that network is q-sparse. For any given value of $q \in \{0, \ldots, 6\}$, we can then calculate the proportion of networks of a given density (rounded up to the first decimal place) that are q-stable.[14] Fig. 5 plots the proportion of q-stable networks at each density level. Each curve corresponds to a different value of q, ranging from 0 to 6. Naturally, the higher is q the larger the proportion of q-sparse (and therefore q-stable) graphs, so higher curves correspond to larger q's. We also see that for intermediate values of q, the proportion of q-stable networks tends to first decrease with the average number of links that individuals have, and then increase with it. Hence, networks of intermediate density tend to be unstable.

[14] There are 32, 767 different networks (symmetric six-by-six matrices of ones and zeros with zero diagonal entries). The details of these calculations (made using Matlab) are available from the authors upon request.

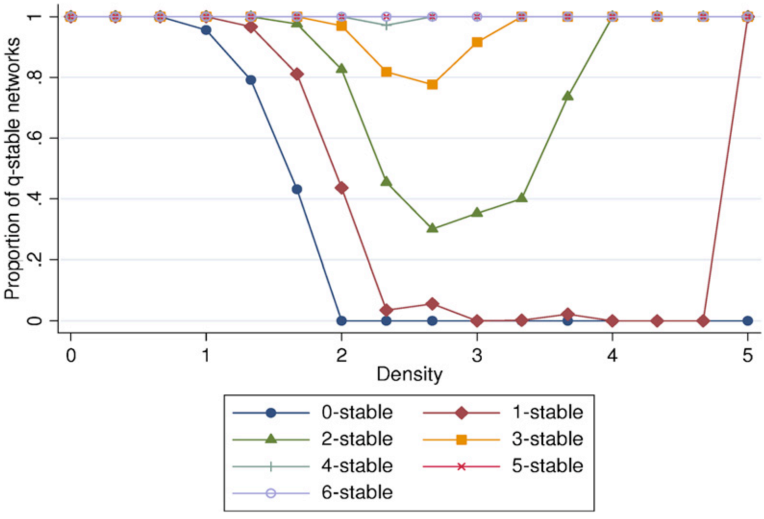

Fig. 5. Stability & density.

Consider a densely connected network. Our result suggests that if, following some exogenous shocks, some links are broken, it could destabilize the network and create additional unraveling of the existing relationships, resulting in a thinly connected graph. In this sense a relatively small shock (such as an increase in out-migration) can have a large effect on the social network.

4.5. *Stability and the clustering coefficient*

For intermediate values of q, stability is related to the *clustering coefficient* used in the literature on social networks that exhibit "small-world" properties (see Wasserman and Faust [37], and Watts and Strogatz [38]). Denote by $\aleph_i$ the set of direct neighbors of i. Presuming that $|\aleph_i| \geqslant 2$, the *clustering coefficient* for i in g is the actual number of links within $\aleph_i$ normalized by the maximum potential number of such links, which is $|\aleph_i| * |(\aleph_i - 1)|/2$. This coefficient lies between 0 and 1 and measures the propensity for i's neighbors to be linked to one another. The clustering coefficient of a graph is then simply the average clustering coefficient of its members.

Under strong or weak punishments, the clustering coefficient does not affect the stability of the network. However, at intermediate levels of punishment we expect highly clustered networks to be more stable. Intuitively, more clustered networks should be better able to ostracize deviants and therefore are more likely to be stable.

We can employ the database of networks of size 6 used in the previous section to look at the relationship between clustering and stability. For any given value of q, we calculate the proportion of networks with a given clustering that are q-stable and map their q-stability to their clustering coefficient. Fig. 6 reports the results. Each curve corresponds to a different level of q. As expected, stability and clustering are unrelated for weak and strong punishment. For intermediate values of q and strictly positive clustering coefficient, however, we see that the relationship between stability and clustering tends to be positive.

We note that in a different model, Vega-Redondo [35] uses numerical simulations to examine the importance of the architecture of the network—its density and cohesiveness—to transmit information about deviants in repeated Prisoner's Dilemma games.

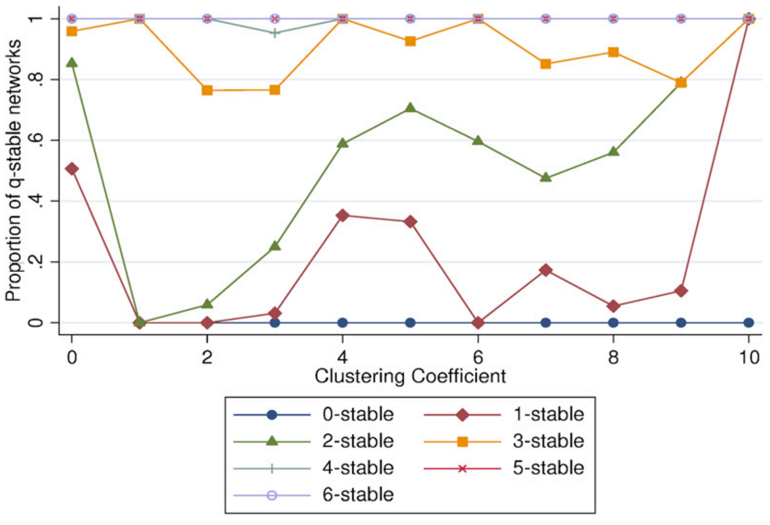

Fig. 6. Stability & clustering.

5. Discussion and possible extensions

5.1. Lower values of the discount factor

The analysis of the previous section was based on values of the discount factor close to unity, thereby emphasizing long run effects. In the short run, the current profits from a deviation, $u(c_i(\theta, S)) - u(c_i(\theta))$ depend on the architecture of the network. Assuming that agents are not affected by "financial contagion," an individual who is a *bottleneck* in the network, in the sense that all transfers must be channeled through him, also has the highest incentive to deviate in the short run. The stability of the informal insurance agreement thus depends on the short run incentives of this "bottleneck agent" to pocket all transfers and run. It is easy to see that the bottleneck agent's incentives to deviate will be strongest in some well-known network architectures, like the star or the line, where it is always possible to find an agent who receives transfers from k agents in order to distribute them to the remaining $n - k - 1$ agents. When the network becomes denser, and transfers can be routed along different channels, the incentives of bottleneck agents to deviate become weaker, and the network is more likely to be stable in the short run.

In general, then, the stability of a network will depend on the interplay between such long run and short run effects. For different values of the discount factor, different stable networks will emerge, as the short run effect dominates for low values of δ and the long run effect for high values of δ.

In order to gain intuition about the variation of stable networks with changes in the discount factor, we consider here a simple example of an equal sharing norm with weak punishment. For a more detailed investigation of the short run effects see our working paper [4].

Example 1. Endowments are i.i.d. across agents and take on just two values: h with probability p and ℓ with probability $(1 - p)$, with $h > \ell > 0$. The following parameters are set through the example: $\ell = 0$ and $h = 1$, and $u(c) = 2c^{1/2}$.

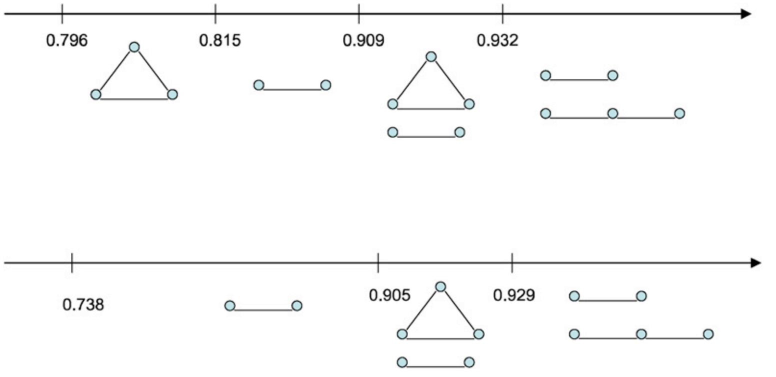

Fig. 7. An illustration of Example 1.

Fig. 7 depicts all stable networks for three agents, for two different values of p, as a function of the discount factor. In the upper panel, $p = 0.2$ and $p = 0.5$ in the lower panel.

Because we consider weak punishment (or 0-stability), trees emerge as the unique stable networks for high values of the discount factor. For lower values of δ, complete graphs emerge as stable network architectures. However, stable networks are not necessarily connected. Surprisingly, for $p = 0.2$, the complete connected graph is stable for low values of δ, then becomes unstable, and becomes stable again for higher values of the discount factor.

5.2. *Communication protocol and other punishment schemes*

The idea of information transmission over q steps captures two important features of communication within networks: it is generally limited and it depends on the graph itself.

One could imagine variations on our scheme. Consider, for instance, the following alternative communication protocol, in which information is transmitted only via individuals who know the person concerned. The idea is that before relaying to a person information on i, one would first ask to the person if she knows i. If she does, then the information is transmitted but otherwise not. In our model, knowing the person is taken to mean having a link with the person. We can define a new punishment structure using this particular communication protocol.

Alternative communication protocol. Information on a deviation is transmitted from an agent to another if and only if they are linked to each other *and* to the deviant. Following a deviation, all informed agents sever direct links with the deviant.

Formally, say that i is *strongly connected* to k via j if $ij \in g$ and there exists a sequence of individuals $\{i^0 = j, i^1, \dots, i^n\}$ such that $i^m i^{m+1} \in g$ for all $m = \{1, \dots, n-1\}$ and $i^m k \in g$ for all $m = 1, \dots, n$.

Using arguments similar to Proposition 3, it is easy to show that for monotone norms and a discount factor close to 1, a network g is stable under the alternative protocol if and only if for every linked pair $ij \in g$, the graph formed by removing from g the links between i and all individuals strongly connected to i via j has strictly more components than g.

The clustering coefficient discussed in Section 4.5 will be particularly important for stability under this new protocol. We illustrate this using a specific example and simulations.

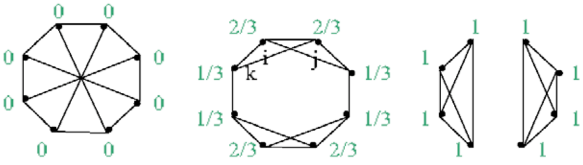

Fig. 8.

Fig. 8 pictures three different graphs with 8 individuals having 3 links each. The clustering coefficient of each individual is indicated next to her node. In the first network, the *wheel*, each individual has a clustering coefficient of 0. It follows that, for high values of δ, this graph will not be stable. Next, consider the *neighbors* network. If individual i in the Figure defects on a transfer to individual j, where both have clustering coefficients of $2/3$, then our information protocol would isolate and effectively punish i. In contrast, if it is individual k (whose clustering coefficient is only $1/3$) who defects on i, we see that our information protocol will be unable to punish him. Hence, graph b is also unstable. Finally, study the *cliques*. Here individuals are clustered in two cliques of four individuals. Their clustering coefficient is 1 and therefore it is possible to fully isolate any deviant. Hence, the graph is stable for monotone norms.

6. Conclusion

In this paper, we develop a model of *risk-sharing in social networks*. Only individuals who are linked can make transfers to each other. There are two important features of our model. First, a risk-sharing arrangement at the level of the entire network results from a collection of bilateral arrangements among linked individuals. Second, for each linked pair of individuals, bilateral *insurance norms* prescribe consumptions and transfers as a function of the individuals' income, identities and their net obligations to others in the network. These generate consumption allocations and associated transfer schemes that satisfy a society-wide consistency requirement: a consistent transfer scheme is a *fixed point* of these bilateral state-contingent transfers.

After describing conditions for the existence of consistent consumption allocation schemes, we assess the stability of self-enforcing insurance networks. We look at discount factors close to one. We show that for monotone norms the stability of a network can be fully characterized by a measure of its sparseness. It is of interest that the sparseness of a network is purely a function of its architecture, while its stability depends, of course, on behavioral postulates.

This result implies, in particular, that for discount factors close to unity, all networks are stable under strong punishment, while only networks of trees are stable under weak punishment. For intermediate levels of punishment both very dense graphs and minimally connected graphs are stable while graphs of intermediate density tend to be unstable. For lower discount factors, the architecture of the network also affects short run incentives, and the stable insurance schemes are harder to characterize.

By introducing a model of risk-sharing networks, this paper opens the door to much future research. There are a number of possible extensions of our analysis. First, different social norms have different stability properties. Comparing them would give us insights on the endogenous emergence of insurance norms. Second, we have explored only one source of instability of insurance networks: the lack of commitment. There are other potential sources of instability, such as moral hazard or the imperfect observability of incomes. Finally, allowing for link formation and analyzing the endogenous formation of risk-sharing networks is an important issue for future work.

Appendix A

A.1. Proof of Proposition 2. [Existence]

Fix a network g. Denote by $M(\theta)$ the size of the bound on transfers in state θ. Let $\mathbf{X}$ be a Euclidean cube, where $\mathbf{x} \in \mathbf{X}$ is a vector of transfers x_{ij}, one for every linked pair, such that $x_{ij} \in [-M(\theta), M(\theta)]$.

For each $\mathbf{x} \in \mathbf{X}$, and for every linked pair $ij \in g$ and state θ, construct the agent's net transfer from others $(z_i(\theta), z_j(\theta))$ as follows:

$$z_i(\theta) \equiv \sum_{k \neq j: ik \in g} x_{ik}(\theta) \quad \text{and} \quad z_j(\theta) \equiv \sum_{k \neq i: jk \in g} x_{jk}(\theta). \tag{A.1}$$

Now for each linked pair $ij \in g$ the bilateral norm allocates consumptions $c_i(\theta), c_j(\theta)$. Build $x'_{ij}(\theta)$ by setting

$$x'_{ij}(\theta) = c_i(\theta) - y_i(\theta) - z_i(\theta)$$

provided that the absolute value does not exceed $M(\theta)$. Otherwise, set x'_{ij} equal to $M(\theta)$ or to $-M(\theta)$ as the case may be. Denote by $\mathbf{x}'$ the vector so generated; by construction, $\mathbf{x}'$ also lies in $\mathbf{X}$. Let ϕ be the mapping that takes $\mathbf{x}$ to $\mathbf{x}'$; obviously, ϕ continuously maps $\mathbf{X}$ into itself. By Brouwer's fixed point theorem, there is $\mathbf{x}^* \in \mathbf{X}$ such that $\phi(\mathbf{x}^*) = \mathbf{x}^*$, and the proof of existence is complete.

It remains to prove that any fixed point yields strictly positive consumption to all individuals in all states. Suppose this is false for some fixed point $\mathbf{x}^*$. We know that in any component d all consumptions cannot be negative, as aggregate income is positive. Hence, there exists a linked pair of agents $ij \in d$ and a state θ such that $c_i(\theta) \leqslant 0$ and $c_j(\theta) > 0$. By our assumption, it must be that i is making no transfer to j, and her output $y_i(\theta) > 0$. So she must be making a transfer to some other set of individuals, say K. By the same logic, $c_k(\theta) \leqslant 0$ for all $k \in K$. Repeating this argument for all individuals in K (and continuing on to other linked individuals if necessary), we see that ultimately we must encounter a pair of linked individuals—say ℓm—such that ℓ makes a positive transfer to m, $c_\ell(\theta) \leqslant 0$, and $c_m(\theta) > 0$. This is a contradiction. □

A.2. q-sparse networks

Proof of Observation 1. Part (a.) is trivial. Part (b.) follows from the fact that there are no cycles connecting any three points in a tree, so networks composed of trees are 0-stable, and any network that is not composed of trees has at least a cycle of length 3. Now consider part (c.).

Take a network g composed of clusters and bridges and denote as $C_1, \ldots, C_M$ its clusters. If there is no triple of agents (i, j, k) such that ij and jk belongs to g, the network is 0-sparse by definition and therefore 1-sparse. Hence, assume that there is a connected triple and take any triple of agents (i, j, k) such that ij and jk belongs to g.

If i and k belong to the same cluster C, then clearly j does too. By definition of a cluster $ik \in g$ and $\ell(i, j, k) = 3$. Now assume that i and k do not belong to the same clusters. In this case, it must be that j is a node-bridge: removing it would increase the number of components of the network. It follows that there is no other path from i to k, $\ell(i, j, k) = 2$. As either way $\ell(i, j, k) \leqslant 3$, g is 1-sparse.

Now, assume that g is not composed of clusters connected by bridges. This implies that there exists at least two clusters who are connected. Select a triple of agents (i, j, k) and two connected clusters C and C' such that $ij, jk \in g$, i is in cluster C and not in C' and k is in cluster C' and not in C. First, notice that i and k cannot be directly linked to each other as otherwise they would belong to the same cluster. Second, observe that there must be a path between i and k that does not go through j as otherwise j would be a bridge-node (removing it would disconnect i and k). Hence, $\ell(i, j, k) > 3$ and therefore g's sparseness is at least 2. □

A.3. Lemmas for the characterization of q-stability

Proof of Lemma 1. Let the bilateral norm generate a consistent consumption allocation for the connected network g. Now take any connected subnetwork g' obtained by removing *one* link, say ij, from g ($g' = g \setminus ij$). Since g' is connected, there is a path of direct links connecting i and j via some ordered set of individuals $1, \ldots, p$, where $i = 1$, $j = p$, and $(k, k+1) \in g'$ for all $k = 1, \ldots, p-1$.

Consider some transfer scheme $\mathbf{x}$ associated with the consistent consumption allocation on g. We construct a transfer scheme $\mathbf{x}'$ on g' as follows. Fix some state θ. For each $k \in \{1, \ldots, p-1\}$, define $x' \equiv x_{k,k+1}(\theta) + x_{ij}(\theta)$. Set $x'_{k\ell}(\theta) = x_{k\ell}(\theta)$ for every other pair of linked agents $k\ell$ in g'. It is easy to see that for *every* linked pair $k\ell$, the sum of third-party obligations is unchanged over $\mathbf{x}$ and $\mathbf{x}'$: $z'_k(\theta) + z'_\ell(\theta) = z_k(\theta) + z_\ell(\theta)$. Moreover, this scheme yields exactly the same consumption vector as before. Because the bilateral norm aggregates third-party obligations, it will indeed recommend the same consumption allocation for g' as it did for g. Therefore the earlier consumption allocation is still consistent.

The same argument can be extended to all connected subnetworks of g by removing one link at a time. □

Proof of Lemma 2. *Sufficiency*. Suppose that there exists a triple of agents i, j and k with a shortest cycle of length greater than $q + 2$. Let j delete his link with i. Then the length of the path between i and k in the cycle is greater than q, and because the cycle is minimal, there is no other path between i and k of smaller size. But this implies that if one deletes all links to j along paths from j to i of size smaller or equal to $q + 1$, link jk remains. In particular, this implies that j remains linked to i in the resulting graph, and remains linked (through i) to all agents who have severed their direct link to j. Hence, the number of components of the graph has not changed.

Necessity. Conversely, consider a link ij for which the deletion of all links to i along the paths to j of size smaller than $q + 1$ does not increase the number of components in the graph. This implies that i and j remain connected in the graph, so that there exists an agent k directly connected to i, such that all paths from k to j have size greater than q. Then, for the tripe k, j, i, all cycles connecting the three agents must have length greater than $q + 2$. □

Proof of Corollary 1. Networks that have only trees as components are the only network that are 0-sparse. Since stability under weak punishment corresponds to q-stability for $q = 0$, the proof of the first part of the corollary is a direct application of Proposition 3.

Take any network g. There exists $\bar{q}$ such that g is q-sparse for $q \geqslant \bar{q}$. It follows from Proposition 3 that g is q-stable under monotone transfer norms for $q \geqslant \bar{q}$. That is, the inequality (3) where $v_i(g')$ is determined by level q punishment for $q \geqslant \bar{q}$ holds for every player i, every state

θ and every set of direct neighbors S of i. Under strong punishment, any individual receives a payoff $v_i^*(\emptyset)$ following a deviation. To be sure, $v_i(g') \geqslant v_i^*(\emptyset)$ for all i and g'. Hence, (3) holds under strong punishment too and g is stable. □

References

[1] F. Allen, D. Gale, Financial contagion, J. Polit. Economy 108 (2000) 1–33.

[2] O. Bandiera, I. Rasul, Social networks and technology adoption in northern Mozambique, Econ. J. (2006) 869–902.

[3] V. Bala, S. Goyal, A non-cooperative model of network formation, Econometrica 68 (2000) 1181–1229.

[4] F. Bloch, G. Genicot, D. Ray, Informal insurance in social networks, BREAD Working Paper No. 145, 2007.

[5] Y. Bramoullé, R. Kranton, Public goods in networks: How social networks can shape social learning and innovation, J. Econ. Theory 135 (2007) 478–494.

[6] Y. Bramoullé, R. Kranton, Risk-sharing networks, J. Econ. Behav. Organ. 64 (2007) 275–294.

[7] T. Calvó-Armengol, M. Jackson, The effects of social networks on employment and inequality, Amer. Econ. Rev. 94 (2004) 426–454.

[8] K. Chatterjee, S. Xu, Technology diffusion by learning from neighbors, Adv. in Appl. Probab. 36 (2004) 355–376.

[9] S. Coate, M. Ravallion, Reciprocity without commitment: Characterization and performance of informal insurance arrangements, J. Devel. Econ. 40 (1993) 1–24.

[10] T.G. Conley, C.R. Udry, Social learning through networks: The adoption of new agricultural technologies in Ghana, Amer. J. Agr. Econ. 83 (2001) 668–673.

[11] S. Dercon, J. de Weerdt, Risk-sharing networks and insurance against illness, J. Devel. Econ. 81 (2006) 337–356.

[12] M. Fafchamps, Solidarity networks in preindustrial societies: Rational peasants with a moral economy, Econ. Devel. Cult. Change 41 (1992) 147–174.

[13] M. Fafchamps, S. Lund, Risk sharing networks in rural Philippines, J. Devel. Econ. 71 (2003) 261–287.

[14] M. Fafchamps, F. Gubert, The formation of risk-sharing networks, J. Devel. Econ. 83 (2007) 326–350.

[15] X. Freixas, B. Parigi, J.C. Rochet, Systemic risk, interbank relations, and liquidity provision by the Central Bank, J. Money Credit Banking 32 (2000) 611–638.

[16] G. Genicot, D. Ray, Group formation in risk sharing arrangements, Rev. Econ. Stud. 70 (2003) 87–113.

[17] G. Genicot, D. Ray, Informal insurance, enforcement constraints, and group formation, in: G. Demange, M. Wooders (Eds.), Group Formation in Economics: Networks, Clubs and Coalitions, Cambridge University Press, Cambridge, UK, 2005, pp. 430–446.

[18] S. Goyal, S. Joshi, Networks of collaboration in oligopoly, Games Econ. Behav. 43 (2003) 57–85.

[19] J. Greenberg, The Theory of Social Situations, Cambridge University Press, Cambridge, MA, 1990.

[20] M. Jackson, The stability and efficiency of economic and social networks, in: B. Dutta, M. Jackson (Eds.), Networks and Groups: Models of Strategic Formation, Springer-Verlag, Heidelberg, 2003.

[21] M. Jackson, A survey of models of network formation: Stability and efficiency, in: G. Demange, M. Wooders (Eds.), Group Formation in Economics: Networks, Clubs and Coalitions, Cambridge University Press, Cambridge, UK, 2005, pp. 11–57.

[22] M. Kimball, Farmer cooperatives as behavior towards risk, Amer. Econ. Rev. 78 (1988) 224–232.

[23] K. Kletzer, B. Wright, Sovereign debt as intertemporal barter, Amer. Econ. Rev. 90 (2000) 621–639.

[24] N. Kocherlakota, Implications of efficient risk sharing without commitment, Rev. Econ. Stud. 63 (1996) 595–609.

[25] R. Kranton, D. Minehart, Competition for goods in buyer–seller networks, Rev. Econ. Design 5 (2000) 301–332.

[26] R. Kranton, D. Minehart, A theory of buyer–seller networks, Amer. Econ. Rev. 91 (2001) 485–508.

[27] E. Ligon, J. Thomas, T. Worrall, Mutual insurance and limited commitment: Theory and evidence in village economies, Rev. Econ. Stud. 69 (2002) 115–139.

[28] S. Lippert, G. Spagnolo, Networks of relations and social capital, CEPR Discussion Paper 3522, 2005.

[29] R. Murgai, P. Winters, E. Sadoulet, A. de Janvry, Localized and incomplete mutual insurance, J. Devel. Econ. 67 (2002) 245–274.

[30] R. Posner, A theory of primitive society, with special reference to law, J. Law Econ. 23 (1980) 1–53.

[31] J.C. Rochet, J. Tirole, Interbank lending and systemic risk, J. Money Credit Banking 28 (1996) 733–762.

[32] C.B. Stack, All Our Kin: Strategies for Survival in Black Community, Harper & Row, New York, 1974.

[33] L. Tesfatsion, A trade network game with endogenous partner selection, in: H. Amman, et al. (Eds.), Computational Approaches to Economic Problems, Kluwer Academic Publishers, 1997, pp. 249–269.

[34] L. Tesfatsion, Gale–Shapley matching in an evolutionary trade network game, Iowa State University Economic Report 43 (1998).

[35] V. Vega-Redondo, Building up social capital in a changing world, J. Econ. Dynam. Control 30 (2006) 2305–2338.
[36] P. Wang, A. Watts, Formation of buyer–seller trade networks in a quality-differentiated product market, Can. J. Econ. 39 (2006) 971–1004.
[37] S. Wasserman, K. Faust, Social Network Analysis: Methods and Applications, Cambridge University Press, New York, 1994.
[38] D.J. Watts, S.H. Strogatz, Collective dynamics of 'small-world' networks, Nature 393 (1998) 440–442.
[39] B. Wellman, Which types of ties and networks give what kinds of social support? Advances in Group Processes 9 (1992) 207–235.
[40] G. Weisbuch, A.P. Kirman, D.K. Herreiner, Market organisation and trading relationships, Econ. J. 110 (2000) 411–436.