

A Theory of Market Efficiency

Anup Rao
University of Washington
anuprao@cs.washington.edu

March 28, 2017

Abstract

We introduce a mathematical theory called *market connectivity* that gives concrete ways to both measure the efficiency of markets and find inefficiencies in large markets. Our theory suggests metrics that can be used to find inefficiencies that may be profitable to exploit, and to evaluate the impact of policy and regulations on market efficiency.

A market's efficiency is tied to its ability to communicate information relevant to market participants. Market connectivity calculates the speed and reliability with which this communication is carried out via trade in the market. We model the market by a network called the *trade network*, which can be computed by recording transactions in the market over a fixed interval of time. The *nodes* of the network correspond to participants in the market. Every pair of nodes that trades in the market is connected by an *edge* that is weighted by the rate of trade, and associated with a vector that represents the type of item that is bought or sold.

We evaluate the ability of the market to communicate by considering the effect of *shocks* on the market. A shock is a change in the beliefs of market participants about the value of the products that they trade. We compute the effect of every *potential* significant shock on trade in the market. We give mathematical definitions for a few concepts:

- The *tension* and *energy* of the network are related concepts that measure the strength of the connections between sets of participants that trade similar items. They measure the amount of trade that is affected by significant shocks. They are high when there are many paths of high rate of trade that connect those with differing beliefs about the value of items. They are low when information from some large set of participants must take a long time to reach some other large set via trade.
- A *bottleneck* in the network is a collection of nodes that have undue power in controlling the dissipation of shocks. The nodes in the bottleneck have an incentive to set prices incorrectly and interfere with the fair transmission of information in the market.

We give explicit mathematical definitions that capture these concepts and allow for quantitative measurements of market inefficiency.

Contents

1	Introduction	3
1.1	Our Approach	3
1.2	Comparison with Related Work	8
2	The Trade Network for Single Item Markets	9
2.1	Motivating Examples	13
3	The Trade Network for Multiple Items	20
4	Rigorously Defining Market Connectivity	23
4.1	Single Item Markets	23
4.2	Markets with Multiple Items	25
5	Discussion and Future Work	27
6	Acknowledgements	28
A	Technical Proofs	29
B	Example Matching Fact 15	35

1 Introduction

Market-based economies have come to be the dominant system for the production and distribution of goods and services. The power of markets stems from their decentralized nature [Hay45]. A market is able to channel the effort of individuals in fruitful directions, even though no individual is aware of all the information necessary to decide how their effort ought to be expended. Complex decisions about how to react to scarcities, surpluses, and spikes in demand are handled without the intervention of a central entity. Over time, markets have proven themselves to be adept at taking advantage of economies of scale, and forming reliable and efficient distribution networks.

The forces of self-interest, competition, and the balancing of supply and demand help to transmit information and stabilize prices. An *efficient market* was defined by Fama [Fam70, Fam91] as a market in which prices always fully reflect available information. The famous *efficient markets hypothesis* asserts that financial markets are efficient, or close to being efficient. The hypothesis is supported by compelling intuition—if prices did not accurately reflect all information, then traders would have an opportunity to exploit the discrepancy to make a profit, which would lead to prices becoming more accurate. So, prices should not drift far from being correct. As Malkiel notes, “[...] [W]hen information arises, the news spreads very quickly and is incorporated into the prices of securities without delay.” [Mal03, p. 59]

However, much is left unanswered here. If a market is not efficient, how can one quantify its inefficiency? Exactly what information is *available* to the market and what information is *not available*? If a tree falls in a forest and no one is around to hear it, should the price of wood change? What if a woodcutter hears the tree falling, how much should the price of wood change then? What if a carpenter hears the tree falling, how much should the price change then? As Fama explains, one cannot hope to conclusively determine whether or not a market is efficient:

[...] [M]arket efficiency per se is not testable. It must be tested jointly with some model of equilibrium, an asset-pricing model. [...] [W]e can only test whether information is properly reflected in prices in the context of a pricing model that defines the meaning of “properly.” As a result, when we find anomalous evidence on the behavior of returns, the way it should be split between market inefficiency or a bad model of market equilibrium is ambiguous [Fam91, p. 1575].

In this work, we develop an alternative mathematical theory that provides new measures for the health of a market. Our definitions maintain the essential elements of the intuitions described above, yet allow for explicitly quantifying market inefficiencies. The result is a theory that is, in principle, testable. We hope that our theory leads to practical models that can be calibrated with data from the real world, and analyzed to find and quantify inefficiencies.

Such a model could have far reaching consequences. A useful mathematical model could provide an important parameter that indicates the health of the market. Regulators and policy makers could use it to objectively evaluate their decisions. Firms could use it to identify inefficiencies that can be exploited for profit, helping to eliminate the inefficiencies.

1.1 Our Approach

We introduce a mathematical theory that we call *market connectivity*. Market efficiency is closely tied to a market’s ability to communicate economic information. So, market connectivity measures

the speed and reliability with which the trade in the market is able to carry out this communication. Market participants do learn information from sources external to the market, but we believe that our theory automatically accounts for all significant external channels of information as well, because such channels will lead to parallel channels of trade. We elaborate on this point at the end of this section. Our theory focuses on the ability of the market itself to communicate information by trade; the advantage is that this can be measured with data. A market that is able to disseminate information quickly and reliably is deemed efficient, and conversely, a market that cannot disseminate information quickly and reliably is deemed inefficient. The concepts we develop allow us to quantify market inefficiencies—in our theory these are just obstructions to communication.

We represent the market by a network called the *trade network*. Every participant in the market corresponds to a node of the network—every firm, individual, or other entity that buys and sells in the market is represented by a node. Every pair of nodes is connected by an edge weighted by the rate of trade between the corresponding participants, and associated with a vector that represents the type of the item being traded. The rate of trade can be estimated by computing the ratio of the total monetary value of the trade between the two participants during an interval of time (say a year), divided by the length of the interval. The vectors that represent the type of item can be computed from the distribution of trade in the item (more on this in Section 3). These type vectors encode the similarity of different items traded in the network. The trade network itself evolves with time¹ as traders discover new trading partners and react to changes in prices. However, we restrict our attention to a fixed network that corresponds to a snapshot of the market during a particular interval of time.

The trade network of a market is likely to be sparse—most nodes will trade with a small fraction of the other nodes. Individuals tend to trade with local businesses more than they do with businesses located far away. Firms tend to trade with other firms with whom they have prior relationships, and large contracts are not often renegotiated. So, firms will tend to be sparsely connected in the market. Many trading relationships—like subscription to cell phone service or rental agreements—are maintained with contracts that deter participants from switching trading partners. Even markets that may seem well connected at first, like the stock market, may not turn out to be so, on closer inspection. Traders in the stock market trade only with other participants that either hold the same stocks, or are interested in the same stocks. If two publicly traded firms **A**, **B** have businesses that operate in the same space, yet traders that own stocks of **A** are not well connected to traders that own stocks of **B** in the trade network, the market may have an inefficiency. This is the kind of inefficiency that our model finds.

It is worth noting that one network can be better connected than another, even though both

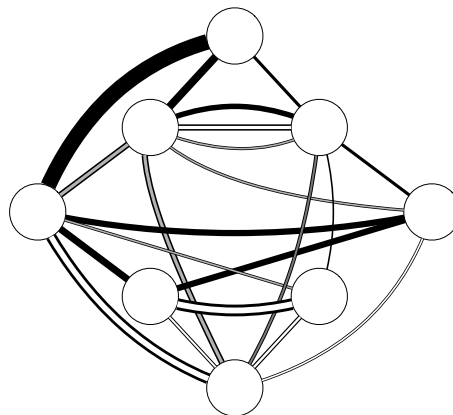
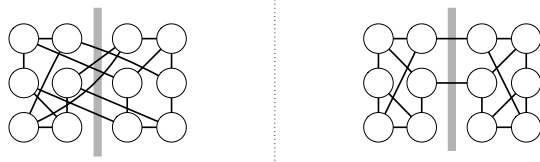
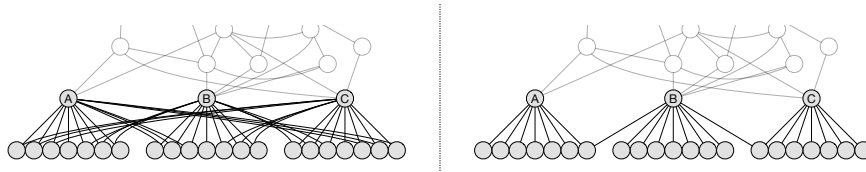


Figure 1: A trade network with 8 nodes, represented as circles. Edges connect nodes that trade. Edges are weighted by the rate of trade, and colored according to the type of item being traded.

¹The evolution of the trade network over time does encode some information about the efficiency of the market. It captures how quickly the market is able to form new trade connections in response to inefficiencies. We do not study this further here, but this seems to be a direction for future work that is worth pursuing. More on this in Section 5.



(a) Two trade networks with the same total volume of trade. Every node in both networks trades with 3 other nodes, and every edge has the same rate of trade, yet the network on the left is better connected. Information will take a long time to propagate across the gray divider in the network on the right.



(b) Two trade networks where **A,B,C** provide a service to a set of customers. **A,B,C** have the same share of the market in both networks, but the network on the left is more efficient, because most of **A**'s customers also trade with **B** and **C**. In the market on the right, if **B** and **C** reduce their prices, **A** has no incentive to reduce its price, since its customers do not trade with **B** and **C**.

Figure 2: The trade network can be used to find market inefficiencies.

have the same volume of trade. For example, Figure 2a shows two trade networks, each involving 12 participants. Every edge has the same rate of trade. In both cases, every participant is connected to 3 others, but the network on the left is better connected. There are only two edges crossing the gray divider on the right, so information that is known to the nodes on one side of the gray divider can take a long time to propagate to the other side. This illustrates the fact that measuring the connectivity of trade networks is not as straightforward as counting the total volume of trade, or estimating the number of participants involved in a lot of trade. The structure of the connections between the participants is important. The structure also reveals nodes that have undue control of prices because of their position. Figure 2b demonstrates this using two networks where the nodes **A,B,C** are service providers that use contracts to restrict the ability of customers to switch providers. **A,B** and **C** have the same share of the trade in both examples, so market share alone cannot discern a difference between the two networks. However, in the market on the right, few customers trade with more than one provider. So, if **B,C** reduce their prices, **A** does not have an incentive to follow suit. In the market on the left, customers do often trade with multiple providers, leading to a more efficient market. Observing the trade network can reveal such inefficiencies even though measuring market share does not reveal them.

A key difference between our approach and past work is that we do not attempt to describe the evolution of prices and do not attempt to estimate what prices ought to be. Instead, we measure the connectivity of the market by computing the effect of potential *shocks*. A shock is a change in the beliefs of market participants about the value of the products that they trade. For example, if many corn farmers experience a loss of their crop due to infestation, that event generates a shock in the market—some subset of market participants would then believe that corn is relatively more valuable. A *significant* shock is one that substantially changes the beliefs of many participants in the market. In an efficient market, a large volume of trade should be affected by any significant

shock, leading to pressure on the participants to readjust their beliefs to a new norm. We give 3 definitions to quantify market connectivity:

Tension and Energy Intuitively, the *tension* of the network is the total normalizing force that market participants experience after any significant shock, and the *energy* of the network is the minimum amount of work that needs to be done to overcome market forces and bring the market to the state of any significant shock. Suppose two participants trade an item at rate r , and after the shock, their beliefs for the amount of the item that is worth a unit of currency are x and y . Then the *force* generated by the shock is proportional to $r \cdot |x - y|$. The *energy* stored in the edge is equal to the amount of work that needs to be done to bring the market to this state: $(1/2) \cdot r \cdot (x - y)^2$. There is no force or energy if both nodes have the same beliefs. The force and energy are maximized when the difference in their beliefs is large and the rate of trade is high. The tension of the shock is the magnitude of the net force felt by all nodes. Up to normalization factors, the tension/energy of the network is the minimum over all significant shocks, of the tension/energy experienced by the network after the shock.

The tension and energy of a shock do have several qualitative differences. The tension is the total net force felt by all market participants, which corresponds to quantifying how quickly small changes in the beliefs of nodes helps to bring the nodes closer to having the same beliefs. The energy corresponds to quantifying the total potential energy stored in the edges when the nodes experience the shock. The energy is large when most nodes trade with partners that have conflicting beliefs from themselves. The tension is large when most nodes have beliefs that differ from the average of the beliefs of their trading partners.

Two large countries that do not trade correspond to a trade network with low tension/energy. Prices in one country can easily stray far from prices in the other country. These concepts can be used to distinguish the two examples shown in Figure 2a.

Bottlenecks A *bottleneck for tension/energy* is a small set of nodes that is responsible for a large fraction of the tension/energy of the network after some significant shock. Market participants that correspond to the bottleneck have an incentive to engage in anticompetitive practices, blocking the flow of information and interrupting the transmission of shocks.

A firm that has a monopoly on the production of an item that cannot be easily substituted with items available from competitors is a particular kind of bottleneck, but not the only one. Prices can be held at artificial levels by such a bottleneck. The concept of bottlenecks can be used to distinguish the two examples shown in Figure 2b.

A market with high *connectivity* is a market that has high energy, high tension, and no bottlenecks. The advantage of our theory over prior work is that it does not use an asset pricing model, so statements analogous to the efficient markets hypothesis are falsifiable in our theory. One can use our definitions to explicitly find and quantify inefficiencies if they exist. An inefficiency in our theory corresponds to a significant shock with low tension/energy, or a significant shock and a small set of nodes that form a bottleneck for the shock. Such an inefficiency identifies an obstacle for the flow of information that may be exploited for profit, reducing the inefficiency, or addressed by a change in policy or regulations.

In the restricted case where all trade in the market involves only one type of item, the formulas we use for the tension and energy of a shock have long been used to model many physical systems, and are closely related to concepts studied in spectral graph theory [Spi, Chu96]. The same

equations describe the behavior of electrical networks (where beliefs correspond to voltages, tension corresponds to the current flowing into the network, and energy corresponds to the power of the network), networks of springs, as well as the flow of heat in conducting materials. The fact that these equations have proven their worth in so many diverse settings is evidence that they capture something fundamental about the trade network. Viewed as a spring system, the rate of trade along an edge corresponds to the spring constant of the spring that it represents as in Hooke’s Law. Each node’s belief corresponds to a location on the number line. The placement of the nodes at these locations by the shock induces forces at the nodes, and the tension is the total magnitude of all these forces. The energy is the total potential energy stored in the system.

When multiple items are traded, the formulas we develop are natural multidimensional analogs of their single item counterparts—they loosely correspond to what happens when nodes are placed in a multidimensional space and the directions of the forces are twisted according to the type of item being traded. We associate each item with a vector that we call the *type vector* of the item. Similar items have nearly identical type vectors, and items that are different from each other have nearly orthogonal type vectors. We show how to use data from the trade network to generate the type vectors. Each edge applies a force to a node only in the direction of the type vector. These choices allow us to model how trade in one item can carry information about a shock in another item. To the best of our knowledge, the mathematics we have developed for this general setting is new.

We believe that the trade network is likely to encode *all* significant channels of information, even those that are external to the market. Let us make a few points to justify this belief. First, when we refer to the “market”, we mean the *whole* market. Prices on two different stock exchanges can be correlated even if stocks in one exchange cannot be traded on the other, because participants that trade on these exchanges are connected in the market via their trade in other items. Second, even if there are significant channels for the transfer of information that are truly external to the market, traders that act on this information will create a channel of trade that parallels the flow of information in the external channel, leaving a trace of that flow visible in the trade. For example, if traders in one country use the price of iron ore in another to make decisions about whether to sell or buy iron ore in their local market, then discrepancies in the price of iron ore between the two countries will give traders incentives to trade iron ore or products related to iron ore between the countries. Conversely, if trade between the two countries is impossible, then information about the price of iron ore in the other country is not actually a very useful source of information for the local trader. So, the extent to which trade connects different parts of the market is a good proxy for the extent to which those parts are able to communicate information relevant to prices. If a large number of people make trading decisions using information they obtain from the same public source, then there will be correlations in their trading decisions, which will lead to many of them being connected to each other via short routes of trade in the market. Consequently, this source of information is reflected in the trade of the market, and will be accounted for by the mathematics of the theory. Finally, the theory we develop is intended to quantify the macroeconomic health of the market, and the mathematics is tuned to finding large scale inefficiencies. A channel of information is important to our theory only if it affects decisions that lead to a significant volume of trade—all such channels are likely to be reflected in the trade network.

After briefly discussing related work, in Section 2 and Section 3, we explain the intuitions behind the choice of the model and develop the theory. We show how several ideas in economics can be better understood using the concepts we develop in our work. In Section 5, we discuss some open

questions, and possibilities for future work that we find worthwhile.

1.2 Comparison with Related Work

The efficient markets hypothesis is of fundamental importance, and there is a vast literature of prior work that has sought to understand it. This literature has proven to be illuminating and important, both theoretically and practically, and has had a major impact on the views and practices of professional investors. It is responsible for bringing data driven methods to bear on understanding finance. We refer the reader to the reviews of Fama [Fam70, Fam91] for details. Empirical economists have used different types of tests to try and understand if the market is efficient, which have been categorized in [Fam91] as follows. In *tests for return predictability*, historical data are used along with an asset pricing model to compute estimates for the correct returns on investments, and these are compared to actual returns realized by the market. In *event studies*, the speed at which prices adjust to significant new events is measured. In *tests for private information*, researchers attempt to understand whether or not market participants have private information that is not reflected in prices. To conclude that the market is *not* efficient using such tests, one must be sure that the information gathered by the test is at least as comprehensive as the information known to the market, and that the asset pricing model used is more trustworthy than the market itself. Perhaps not surprisingly, most of these tests have concluded that the financial market is efficient. That said, these tests do provide valuable insights into the inner workings of financial markets.

The theory we develop here is based on intuitions from this past work, especially the notions of event studies and tests for private information. However, the mathematics and data used in our theory are quite different from past work. We do not attempt to make any quantitative statements about the evolution of prices, nor do we attempt to understand what information the market should take into account. Instead, we evaluate markets as systems for communication. We give explicit ways to quantify obstructions for the the flow of information and find them if they exist, so market connectivity can be unambiguously determined. The primary motivation for our work is to strengthen the concept of market efficiency by giving it these features.

A different line of past work has tested a consequence of the efficient markets hypothesis: no trader should expect to make unusually large returns from trading in the financial market. Again, this is a vast area that we cannot hope to summarize here, so we refer the reader to the survey [Mal03]. This work is also extremely useful in practice, but the fact that traders cannot expect positive returns is not conclusive evidence of market efficiency. Traders may have difficulty generating positive returns even in an inefficient market due to high transaction costs, the difficulty of finding inefficiencies that they are capable of exploiting before they disappear, the lack of capital to compete with an entrenched system, or compliance with regulations.

Perhaps the biggest challenge to the efficient markets hypothesis has come from the field of behavioral economics [PS15]. Behavioral economics argues that the efficient markets hypothesis cannot hold because market participants are not rational, and are subject to human psychology. Our own model is built on the idea that market efficiency is limited by the bandwidth of information available to individual participants, so our work is inspired by similar considerations. However, unlike behavioral economics, our theory provides an alternative to the established theory of market efficiency that we believe has greater explanatory and predictive power.

There is a significant history of using networks to model social and economic phenomena. The book [Jac10] is a good reference. Several subfields here have developed very similar mathemat-

ics to the mathematics in our paper, but to the best of our knowledge, none are well suited to understanding market efficiency. One of the most related concepts from this past work is the *Bass* model of diffusion [Jac10, Chapter 7], which was invented to understand the spread of contagion in a social network. Contagion has been used to model the spread of financial crises [KRV03, AOTS15, GHM12, EGJ14]. Here the emphasis is on understanding how a crisis spreads over financial relationships, rather than how useful information spreads, so the network is defined differently from in our model. Another body of work with similar mathematics arises in the study of learning networks [Jac10, Chapter 8]. In the DeGroot model [DeG74], every node of the network holds an opinion, and each node updates its opinion according to the opinions of its neighbors in the network. More recently, Rabani and Schulman [RS16] studied Arrow-Debreu markets where the participants trade along a network, and used concepts similar to the energy of the network in our model to measure the speed of convergence to a new equilibrium.

Many of these past network models are mathematically and conceptually related to our own model for the single item case, but do not appear to take into account the subtleties that arise from trade in multiple types of items or the presence of insignificant shocks. The equations we use for the single item case have a much longer history, going back to the study of topics like electrical networks, networks of springs and the diffusion of heat. Another subtle, but important, difference is that much of this past work views the network as placing limits on the options available to participants. They study systems where the nodes are allowed to operate within the confines of the network. Our approach is different—we view the parameters of the trade network as data useful to evaluate market connectivity. The market is not limited by the network; the network encodes the limitations of the market.

DeGroot also develops a notion of *social influence* that seems superficially similar to our notion of bottlenecks, but is mathematically quite different. Social influence of a node in an *undirected* network is determined solely by the fraction of the total weight of edges that touch the node—the location of the node is not relevant. The concepts of *efficient networks*, and *graphical games* [Jac10, Chapter 6] are less related to the topic of this paper than they may seem from their names. They are concerned with understanding the space of strategies for playing games on networks.

2 The Trade Network for Single Item Markets

We use the notion of a trade network to understand the limits and power of markets to disseminate information via trade. The ideas we develop are simplest when there is only one item being produced and consumed in the market. We start by discussing this restricted case, because it is easier to understand. Later, we show how to model markets where an arbitrary number of items are traded, which is much more applicable to real markets.

Recall that we represent every market participant by a distinct node, and connect two nodes by an edge if the corresponding market participants directly trade with each other. We use V to denote the set of nodes and E to denote the set of edges. Every edge $e = \{u, v\}$ connects two nodes, u and v of the network. The edge is assigned a non-negative number called the *weight* of the edge, denoted $w(e)$, or $w(\{u, v\})$. The weight is the rate of the trade between the two participants. If $w(e) = 0$, that is equivalent to there being no trade along the edge e . To estimate the weight of each edge in real world markets, one can use past transactions to compute the monetary value of the trade between every pair of participants over a fixed time interval and then divide this number by the length of the interval.

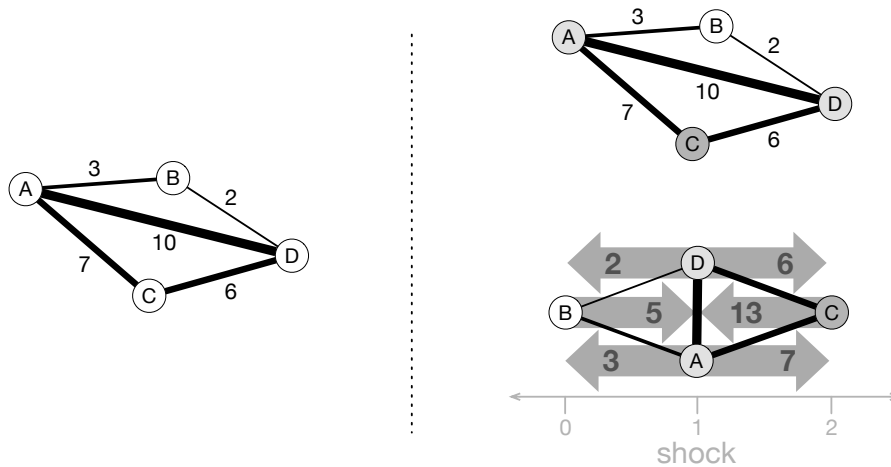


Figure 3: On the left: a single item trade network with 4 participants. The numbers on the edges denote the rate of trade. On the right: the network is shown with a mapping to the number line. The locations of the nodes correspond to a shock: a change to beliefs about the amount of the item that is worth a unit of currency. The arrows show the forces induced on each node by trade.

Since single item markets correspond to spring systems, we start by giving some intuition² about the correspondence between the two. Throughout the discussion we work with springs whose relaxed length is 0: namely the springs always pull nodes closer together and never push them apart.

It is helpful to think of placing each node at the location on the number line that corresponds to its belief for the amount of the item that is worth a unit of currency. In equilibrium, every node has the same belief, and there is no tension in the network. A *shock* is a change in the beliefs for the amount of the item that is worth a unit of currency. It can be represented by a n dimensional vector of numbers $\mathbf{x} \in \mathbb{R}^n$, where \mathbf{x}_v denotes the node v 's belief about the change. The shock induces forces in the trade network that helps to reduce the shock. The *force* on a node u along the edge $\{u, v\}$ is given by Hooke's Law:

$$w(\{u, v\}) \cdot (\mathbf{x}_u - \mathbf{x}_v).$$

The magnitude of the force is proportional to the product of the weight and the difference between the new beliefs of u, v about the value of the item. The *energy* stored in the edge $\{u, v\}$ is equal to

$$(1/2) \cdot w(\{u, v\}) \cdot (\mathbf{x}_u - \mathbf{x}_v)^2.$$

The energy is the amount of work that needs to be done to pull the two nodes of the spring apart. The larger the net forces felt by the nodes and the total energy of the system, the faster the shock is likely to dissipate in the market.

²This is just one possible interpretation for the model. The model can be interpreted in a number of different ways, each with its own merits.

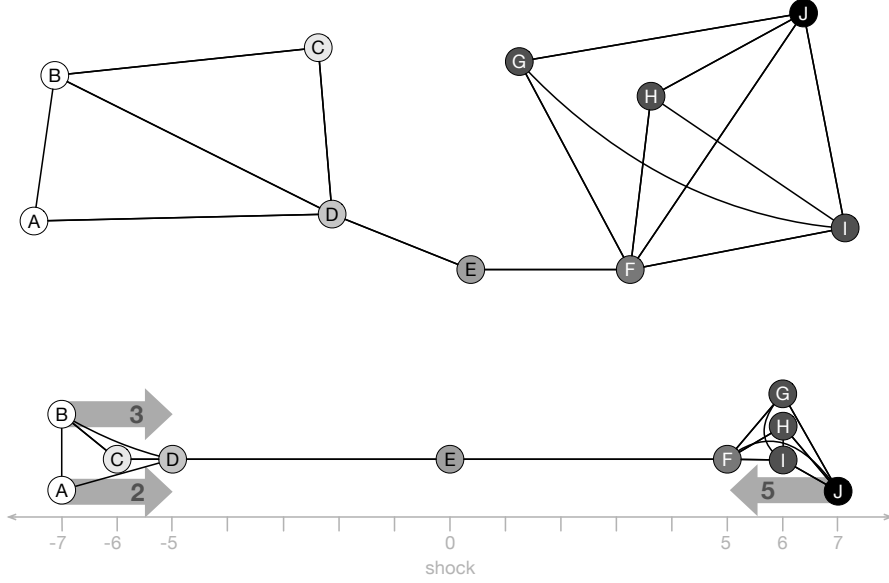


Figure 4: A mapping of a network to the number line. Every node is colored according to its location on the line, which is its belief about the amount of the item that a unit of currency is worth. The gray arrows show the forces experienced by nodes **A**, **B** and **J**. All other nodes experience no net force. The tension of the network for this shock is the sum of the magnitudes of the net force on each node, which is $T(\mathbf{x}) = 3 + 2 + 5 = 10$. By Fact 1, the energy for this shock is $E(\mathbf{x}) = 35$.

Given a shock \mathbf{x} , we define the tension induced by \mathbf{x} to be the total magnitude of all the net forces felt by the nodes:

$$T(\mathbf{x}) = \sum_{u \in V} \left| \sum_{v \in V} w(\{u, v\}) \cdot (\mathbf{x}_u - \mathbf{x}_v) \right|.$$

The total *energy* of \mathbf{x} is the sum of the energies of each edge:

$$E(\mathbf{x}) = (1/2) \cdot \sum_{\{u, v\} \in E} w(\{u, v\}) \cdot (\mathbf{x}_v - \mathbf{x}_u)^2.$$

One can check that the gradient³ of the energy is the vector of forces that need to be applied to the nodes to hold them in place:

$$(\nabla E(\mathbf{x}))_u = \sum_{\{u, v\} \in E} w(\{u, v\}) \cdot (\mathbf{x}_u - \mathbf{x}_v), \quad (1)$$

so $-\nabla E(\mathbf{x})$ is the vector of net forces experienced by the nodes, and the tension is the L_1 norm of the gradient vector:

$$T(\mathbf{x}) = |\nabla E(\mathbf{x})|_1 = \sum_{u \in V} |(\nabla E(\mathbf{x}))_u|.$$

³The gradient is the vector of partial derivatives.

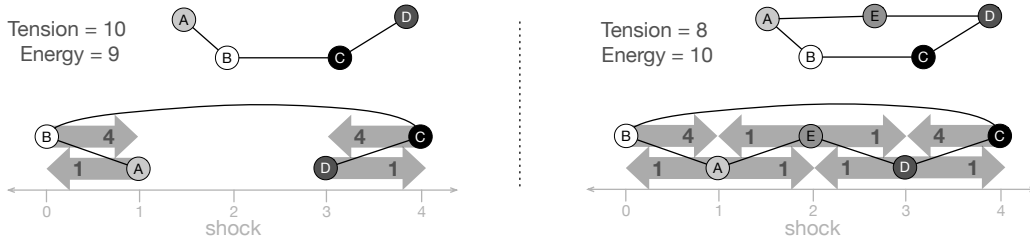
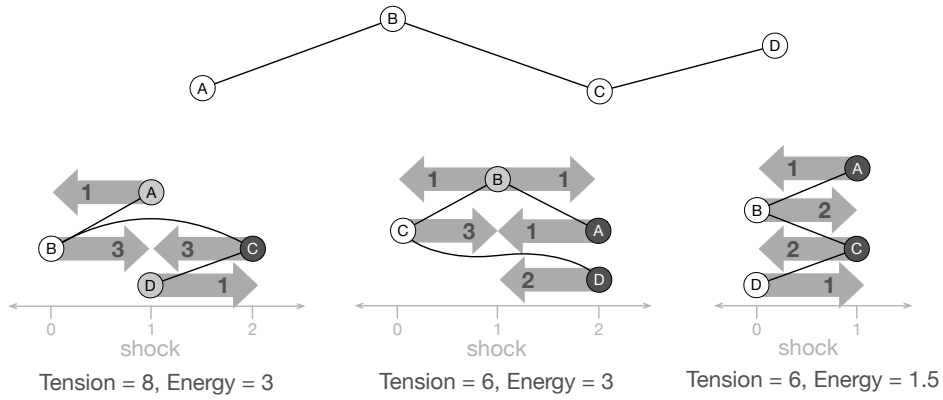
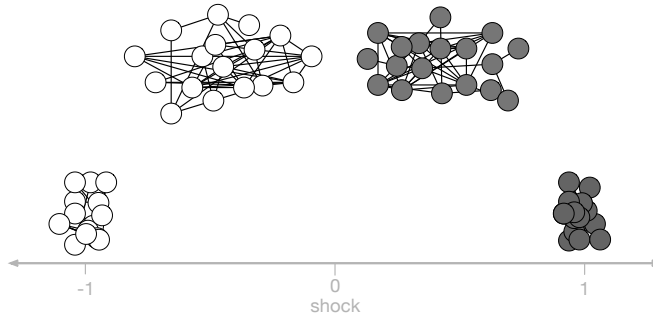


Figure 5: Examples illustrating the difference between the tension and energy of a shock.

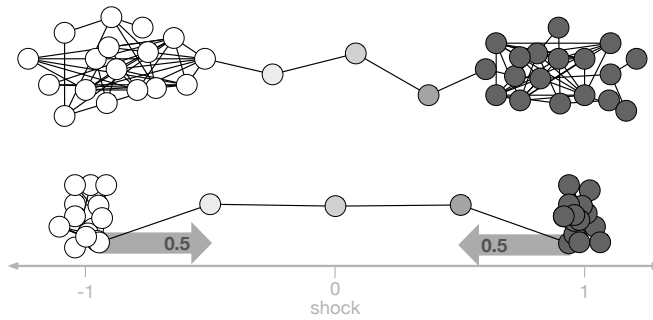
The energy of the network is equal to the amount of work that needs to be done to bring the network to its current position from a 0 energy state. The following fact is useful to keep in mind to understand the examples we discuss next:

Fact 1. *If some subset of the nodes are held at location a and some other subset are held at b , and the rest are left to settle at equilibrium, the energy is always equal to $|(a - b)t/4|$, where t is the tension.*

The tension and energy have qualitative differences that are explored in Figure 5. In Figure 5a, two of the shocks have the same energy, and two of them have the same tension. The first shock gives the highest tension because it is the easiest for the participants to reduce the energy locally. In the middle example, node **B** cannot help to reduce the energy. This is what leads to the second example having lower tension than the first. In the final example, the energy is much lower because the participants are much closer to having the same beliefs than in the other two examples. Loosely speaking, the tension corresponds to the ability of most participants to reduce the distance to equilibrium, while the energy corresponds to the total work that needs to be done to return the market to equilibrium. In Figure 5b, we see that the presence of a new node can actually decrease the tension of the network, though it can never decrease the energy. For simplicity of explanation, these subtleties are ignored in the examples we discuss next. The shocks in all of these examples have been chosen so that they do not distinguish between tension and energy.



(a) $T(x) = E(x) = 0$ when all the white nodes in one part are placed at -1 , and all the gray nodes from the other part are placed at 1 . The shock identifies the disconnection in the network.



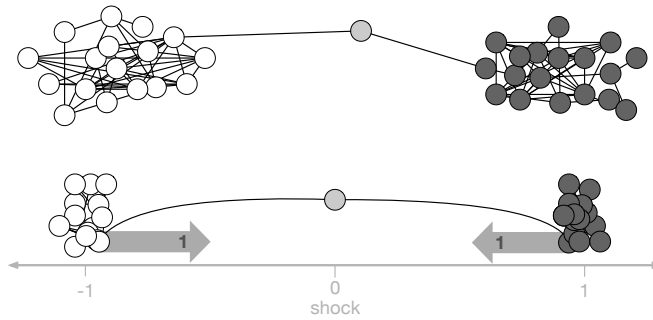
(b) Only the nodes at -1 and 1 experience a net force. The tension and energy of the shock has gone up to $T(x) = 1, E(x) = 0.5$ from Figure 6a.

Figure 6: All edges have unit weight in this sequence of examples showing that tension and energy increase with connectivity.

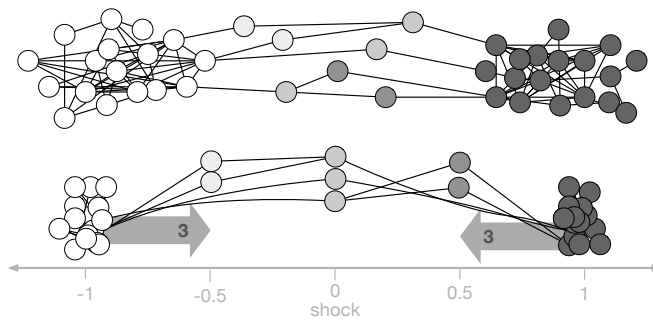
2.1 Motivating Examples

The model we define in this paper can be used to explain diverse phenomena that pertain to market efficiency. Here we consider several examples of well known ideas in economics that can be explained using our model. We stress that a proper treatment of actual markets should take into account that the items traded in the market are different from each other. We discuss the theory that handles this more realistic setting in Section 3. To compute the connectivity of the market, we find significant shocks that have low energy, low tension or reveal the presence of bottlenecks.

Suppose the market under consideration consists of n farmers that produce and consume corn. Every farmer buys and sells corn by trading in a common market place, and every farmer trades with every other farmer at a rate of 1. Then every pair of nodes in the trade network is connected by an edge of weight 1. If a single farmer produces less corn than expected, the effects of this shock are quickly felt by the whole market. The farmer will seek more corn from others, which increases the price of corn and encourages others to produce and sell more corn to compensate. If a large fraction of the farmers experience a drop in production, the effect on the network is even more pronounced. Viewed as spring network, where every edge has weight 1, we see that the



(a) A short trade route generates more tension than the trade route in Figure 6b. $T(\mathbf{x}) = 2, E(\mathbf{x}) = 1$. The node at 0 experiences no net force.



(b) Many parallel trade routes generate more tension than in the network of Figure 7a. $T(\mathbf{x}) = 6, E(\mathbf{x}) = 3$. Only the nodes at -1 and 1 experience a net force.

Figure 7: All edges have unit weight in this sequence of examples showing that tension and energy increase with connectivity.

tension between any two disjoint sets S, T that have different values in the shock is quite large. For example, suppose that S, T are each of size at least $n/3$, every node in S has a value of $+1$, and every node of T has a value of -1 . Then there are at least $n^2/9$ edges going directly from S to T , so the tension of any such shock is at least $4n^2/9$, and the energy is at least $2n^2/9$. So every shock gives high energy and tension, and the market has high connectivity.

It is a commonly held view that free trade increases market efficiency, a view that can be quantified with our model. Suppose that the market consists of farmers living in two adjacent towns, each with $n/2$ farmers that do not trade with each other. If the farmers in one of the towns experience an oversupply of corn, the price of corn, and so the economic activity, remains unaffected in the other town. Viewing the trade network as a network of springs, as in Figure 6a, we see that the tension and energy are minimized if every node in one town is placed at -1 , and every node in the other is placed at 1 . For this shock, we have $T(\mathbf{x}) = E(\mathbf{x}) = 0$. This matches the intuition that there is no way for information to flow between the two towns via trade. It will not be long before traders emerge between the two towns. Consider now the trade network depicted in Figure 6b, where there are three traders trading corn at a rate of 1 along a road connecting the two towns. This market is more efficient than the disconnected market that we had before, since trade

does communicate information between the two towns. In line with our intuition, if every node in one town is placed at -1 , and every node in the other is placed at 1 , and the rest of the nodes are allowed to relax at equilibrium, a force of 0.5 pulls the nodes on the left towards the nodes on the right. The tension is 1 and energy is 0.5 . Now suppose that there was just 1 trader between the two towns as in Figure 7a. Then the market would be better at communicating information from one town to the other, since only one trader is involved in communicating the shock from one town to the other. This is reflected in the fact that the tension of the shock has increased to 2 . The market is most efficient when there are multiple independent traders trading between the two towns, as in Figure 7b. The tension of the shock has increased to 6 .

Trade networks can be used to explain the success of many modern innovations in finance. Consider the role of *exchange traded funds* in the stock market. These are funds that hold assets like stocks or bonds, and allow investors to trade financial instruments that correspond to diversified portfolios. In recent years, the volume of trade in exchange traded funds has risen dramatically. The total annual trade in these funds exceeds the U.S. GDP, and the funds have a turnover rate that is more than 4 times larger than average stocks [Bal15]. Our model suggests that this high level of trade contributes tremendously to increasing market connectivity. It is quite hard for a typical investor to monitor and trade a truly diversified portfolio of stocks, so if these (or similar) funds were not available, investors would be trading only with others that are interested in the same stocks. Since the set of participants interested in trading a highly diversified exchange traded fund is much larger, trade in these funds makes the trade network much more connected than before. Investors that were formerly focused on disjoint portfolios that are nevertheless associated with the same exchange trade fund will now begin to trade with each other. This makes the trade network much more likely to have high energy and tension for a given shock⁴. Exchange traded funds offer a channel of communication that reaches a much broader set of participants than trade in the constituent stocks can provide.

A popular view in economics is that the availability of liquid assets is vital to the functioning of a healthy economy. For example, Chordia, Roll and Subrahmanyam write “Liquidity facilitates efficiency, in the sense that the market’s capacity to accommodate order flow is larger during periods when the market is more liquid.” [CRS08] This view is supported by our model as well. A market with high liquidity is one in which nodes are more likely to be connected to many trading partners via trade. Such a network is much more likely to have high energy and tension than a sparsely connected market.

A second kind of market inefficiency that we study corresponds to monopolies in the market. Suppose a small set of traders comes to control a very large fraction of the trade between the two towns, as in Figure 8. We have $T(\mathbf{x}) = 24$, $E(\mathbf{x}) = 12$, but the tension and would drop dramatically to $T(\mathbf{x}) = 4$, $E(\mathbf{x}) = 2$ if the highlighted nodes are removed. This is what we call a *bottleneck* in the trade network. It is a small set of nodes that contributes most of the tension/energy.

In some settings, the set of nodes involved in the bottleneck have an incentive to commit fraud, or engage in monopolistic practices, because they have a large market share. Even if these nodes behave in good faith, information flow in the market relies too heavily on the bottleneck, which can lead to the market being inefficient if these nodes make errors in their adjustments to prices.

⁴If two publicly traded companies operate in the same space, but most shareholders of one company do not hold shares of the other, our model suggests that the introduction of an exchange traded fund giving exposure to both companies can improve market connectivity.

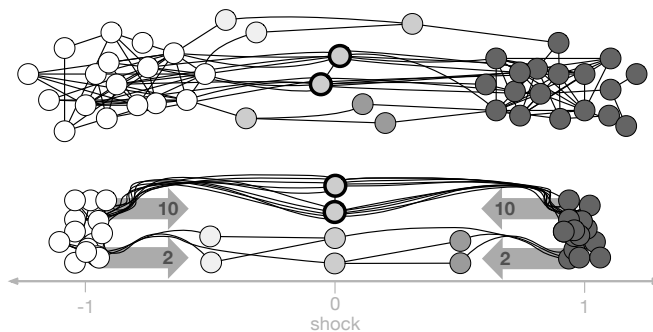


Figure 8: Two nodes that are responsible for most of the tension/energy form a bottleneck.

An extreme example of a bottleneck can be observed in the trade network of a centralized economy, shown in Figure 9. This network is well connected. In fact, from the perspective of the tension and energy, it is nearly as good as the fully-connected network, because there is a short path of length 2 connecting every two nodes in the network. However, every individual trades only with the central authority, depicted as the node in the middle. Even though the tension and energy of the network are quite high, the network clearly has a bottleneck. The burden for recognizing and transmitting shocks falls entirely onto a single entity, making it unlikely for information to be reliably communicated.

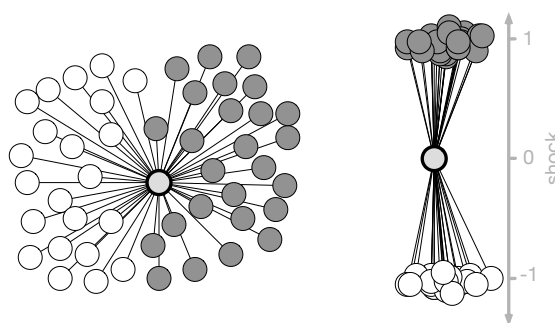


Figure 9: In a centralized economy, every node trades only with the central authority, which is a bottleneck.

It is widely recognized that the financial crisis of 2007-2009 began because of errors in accurately pricing mortgage backed securities (MBS). A crash in the value of these securities initiated a global recession. Data gathered about the organization of the market prior to the crisis suggest that a bottleneck was to blame. As Fligstein and Goldstein report:

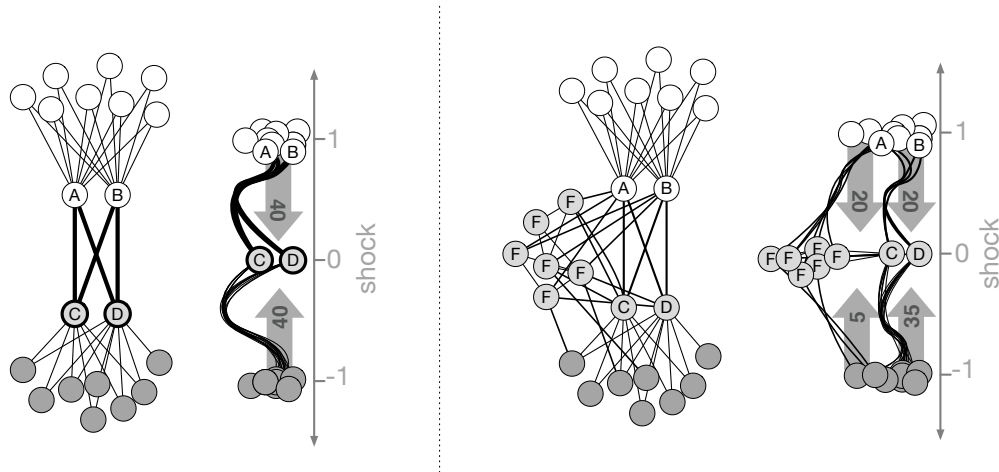
Another commonly voiced myth about the MBS market is that it was highly dispersed, with too many players to control any facet of the market. On the contrary, we show that over time all of the main markets connected to MBS, the originators, the packagers, the wholesalers, the servicers, and the rating companies became not only larger, but more concentrated. By the end, in every facet of the industry 5 firms controlled at least 40% of the market (and in some cases closer to 90%). Separate market niches also increasingly condensed around the same dominant firms. As a result, the mortgage field was not an anonymous market scattered across the country, but instead consisted of a few large firms. This concentration meant that firms very much collaborated and competed in these various markets. Firms would join together in MBS packages and assume different roles with each other. This meant that they had a great deal of knowledge of the market and what the others' moves were. [FG10, p. 33]

One of the major goals of our work is to develop tools that can be used to detect structural problems like bottlenecks before a crisis occurs. Economists and regulators could use the theory to identify bottlenecks in the trade network and investigate methods to tackle them before they lead to a crisis.

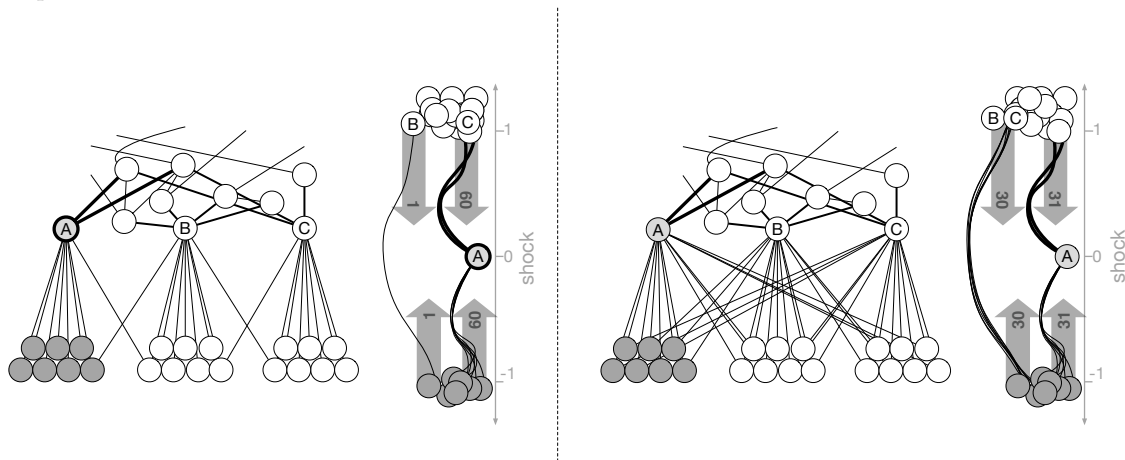
Market connectivity can also be used to assess financial solutions that mitigate the effects of bottlenecks. For instance, consider the role of *futures* in the commodities market. Futures contracts guarantee delivery of a commodity at a date in the future. Perhaps surprisingly, it is often the case that only a small proportion of contracts that are traded are settled by actual delivery. Our model suggests an explanation for this counterintuitive fact: futures contracts serve as a mechanism to break bottlenecks in the trade network, and they are most effective when the trade in the futures exceeds trade in the actual item. We illustrate this using the following anecdote by Hieronymus [Hei77] about the origins of futures contracts in corn:

A country merchant rides into town in January and proceeds to the place of business of the terminal merchant to whom he regularly sells. On offering 20,000 bushels for June delivery at the current price he hears, “I would like to bid that much but with the large stocks in Chicago and a large crop coming in I can only pay 15 cents below today’s price.” Our friend mentions the high price that he has already paid farmers, comments on the ancestry of terminal grain merchants in general, and takes himself to the nearest saloon to find solace. There he comments to all and sundry regarding the greed and cowardice of grain merchants in general and one in particular. On hearing this one stalwart soul says, “I know nothing of corn, being a builder of houses myself, but it occurs to me that the price of corn will be quite as high in June as it is now.” Being true to his occupation, just as we know country merchants today, our man asks, “Is that a firm offer?” “I shouldn’t want to go quite that far but I will bid five cents below today’s price. That will take you off of the hook and out of your cups and leave room for a bit of a profit for me.” “Done,” replies the country merchant and they sign a contract. Some weeks later the builder of houses, who has now become interested in these matters, notes that the price of corn for June delivery is five cents above the price that he has paid. He figures that \$1000 in his pocket is better than 20,000 bushels of corn in his lap in June so he peddles his contract to the nearest terminal merchant and wonders why he had not discovered this easy road to riches sooner. [Hei77, p. 74]

The transformation to the trade network described in this anecdote is shown in Figure 10a. The country merchants are represented by the nodes **A** and **B**, and the terminal merchants by the nodes **C** and **D**. If futures contracts were not available, the network would look like the one on the left, where **C** and **D** clearly form a bottleneck for the depicted shock. They are able to charge unfair prices from **A** and **B** because **A** and **B** cannot sell their corn to anyone else. The introduction of the futures market dramatically alters the trade network, as shown on the right. There are two reasons why **C** and **D** become less of a bottleneck according to our definitions. The first is more obvious: the tension and energy that they are responsible for has been reduced. For the given shock, **C** and **D** are responsible for only a small part of the tension, since a sizable fraction of the trade now involves the futures traders instead of **C** and **D**. Consequently their ability to extort higher prices has been greatly diminished. A second reason is that the significance of the depicted shock has been reduced, because a larger fraction of the traders have the same beliefs in the depicted shock. A shock is now significant only if the futures traders have significant disagreement about the values in the shock. **C** and **D** cannot manipulate prices very easily because the high volume of trade in futures contracts will reveal accurate prices for corn.



(a) A futures market can help to break a bottleneck. On the left, the merchants **A** and **B** can only trade with merchants **C** and **D**. **C** and **D** are a bottleneck for the depicted shock. The introduction of futures traders labeled **F** substantially changes the network to the figure on the right. **A**, **B** can now sell their contracts to the futures traders, which reduces the role of **C**, **D** in generating the tension in the network. **C** and **D** must now compete with the traders **F** for the contracts and cannot fix prices.



(b) High switching costs can give rise to trade networks with bottlenecks. In the network on the left, **A**, **B**, **C** provide service to a set of customers that correspond to the nodes on the bottom. Switching costs are high, so most customers trade with a single provider. **A** is a bottleneck for the depicted shock. If switching costs are low, we might see the network shown on the right, where **A** is no longer a strong bottleneck.

Figure 10: Bottlenecks explain the impact of high switching costs and the benefits of a futures market.

Next we consider the effects of switching costs on market efficiency. Markets where participants are locked-in to trade with particular nodes will tend to be less efficient than markets where participants can freely switch trading partners. As Farrell and Kemperer observe:

Switching costs [...] bind customers to vendors if products are incompatible, locking customers or even markets in to early choices. Lock-in hinders customers from changing suppliers in response to (predictable or unpredictable) changes in efficiency, and gives vendors lucrative ex post market power [...] [FK07, p. 1970]

The effects of two types of switching costs can be seen in Figure 10b. The figure shows three service providers, **A**, **B**, **C** that serve a group of customers, represented as the nodes on the bottom. If switching costs are prohibitively high, one might observe the network shown on the left. The typical customer trades with only one of the service providers. If a shock occurs where the customers of **A** have a different belief from the rest of the market, **A** becomes a bottleneck for the shock. The network on the right shows what the network might look like if switching costs are lower. Now most customers trade with multiple providers, and we see that **A** is less of a bottleneck for the same shock, because customers also experience tension and energy because of the edges that connect them to **B** and **C**. **A** can no longer fix prices because the competition with **B** and **C** is much more of a threat. Note that this example shows that trade networks are more effective at identifying monopolies than counts of market share. In this last example, **A** has only a 1/3rd share of the market. However, its ability to lock-in customers gives it a considerable advantage in the market.

Another interesting example of the explanatory power of trade networks comes from understanding *vertical integration*. Vertical integration refers to an arrangement where a firm owns several parts of the supply chain of a product. Does market efficiency improve with vertical integration? On the one hand, vertical integration does increase the energy and tension of most shocks; on the other hand, it has been observed that vertical integration is associated with increased rates of fraud. Fligstein and Roehrkasse [FR16] make the case that vertical integration was to blame for the elevated levels of fraud that contributed to the financial crisis of 2007-2009. They write:

[...] [T]he motivation toward fraud increases with vertical integration because the diffusion of fraud is a precondition of continuing normal operations. [...] [T]he transaction points that vertical integration eliminates are not only sites of potential fraud. They also represent key opportunities for quality control and due diligence. [FR16, p. 624]

The effect of vertical integration is illustrated in Figure 11. The network on the left shows a supply chain connecting two parts of the network. Each step of the supply chain involves 3 firms with equal market share. If there is a shock between the two ends of the chain, each of the firms on the boundary of the supply chain contributes equally to the tension and energy. The network on the right shows what happens if three of the firms integrate vertically. The new firm has not increased its market share or the rate of trade that it is involved in, but its share of the tension and energy have increased dramatically. Even though it is involved in only 1/3rd of the trade, it now contributes 1/2 of the tension and energy. This is because the edges associated with **V** have been stretched to be twice as long as the tension generating edges that **V** replaced.

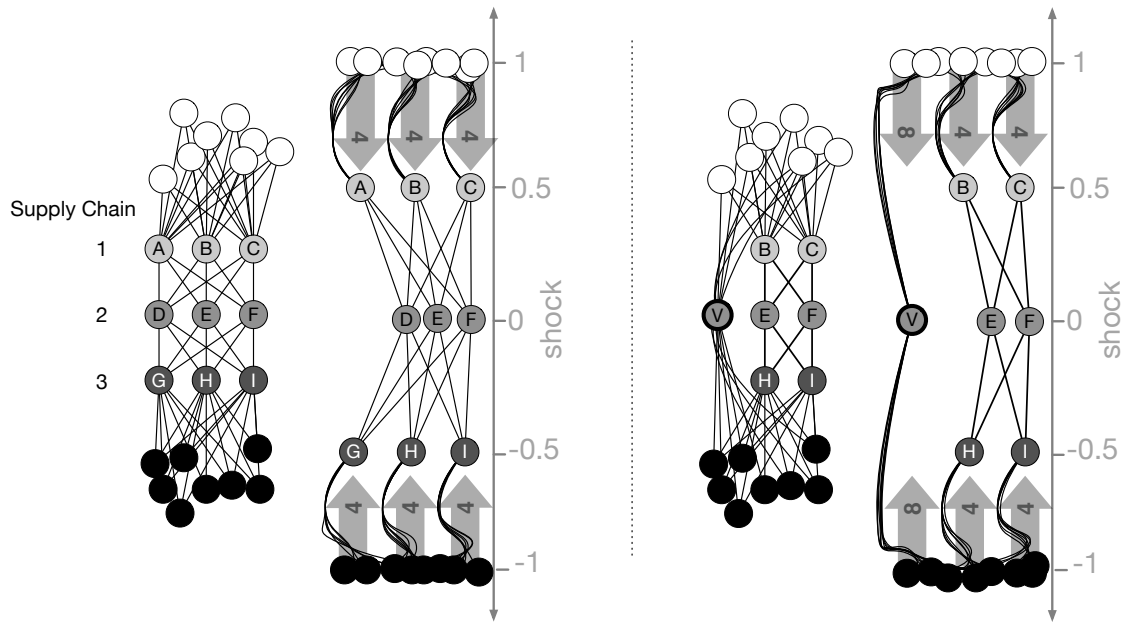


Figure 11: Vertical integration can lead to bottlenecks. The network on the left shows a supply chain involving the firms A, B, C, D, E, F, G, H, I. Each firm trades at an equal rate with the adjacent levels of the supply chain. The network on the right shows the same shock applied after firms A, D, G have been combined into the firm V. The tension and energy have increased, but V now contributes 1/2 of the tension and energy, even though it is involved in only 1/3'rd of the trade.

3 The Trade Network for Multiple Items

The real power of markets emerges when one considers trade networks involving many different kinds of items. It is a mistake to view such a market as a collection of distinct markets trading each item, because shocks in one item can be dissipated via trade in another item. This makes multi-item networks more efficient than any of the subnetworks obtained by restricting attention to individual items.

Consider the market shown in Figure 12, which is a market of farmers that grow both corn and wheat. Viewing the market as a network for trading corn and a separate network for trading wheat suggests that the corn network has 0 tension and energy for a significant shock, since no corn is traded between the left and the right. But this is misleading. If the farmers on the left experience a sharp drop in the

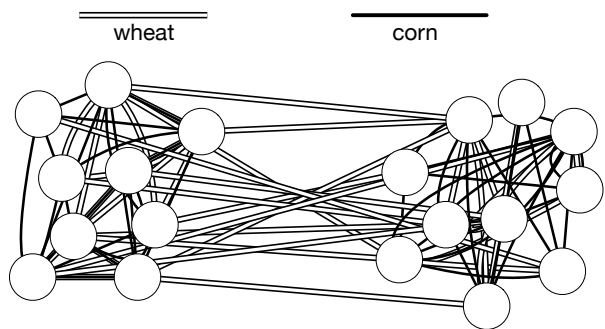


Figure 12: Trade in wheat can carry information about corn, so one should not view the trade network as two different trade networks for each item.

production of corn, the price of corn will rise. The consumption of corn will decrease on the left, and the consumption of wheat (another grain) will increase. This will increase the demand for wheat on the left, making the price of wheat rise on the right, and eventually leading to a rise in the price of corn on the right—the trade in wheat transmits shocks in the corn market. However, if the trade between the left and right is in a commodity, say iron ore, that is not easily substitutable for corn, then the market is much less effective at communicating shocks in the corn network. These considerations motivate the choices we make next.

We associate every item traded in the market with a unit vector $\boldsymbol{\tau} \in \mathbb{R}^d$, for some number d . We refer to $\boldsymbol{\tau}$ as the *type* vector of the item. The type vectors are chosen in such a way that the inner product $\langle \boldsymbol{\tau}, \boldsymbol{\eta} \rangle$ captures the correlation between the items represented by the vectors $\boldsymbol{\tau}, \boldsymbol{\eta}$. So two items that are completely different from each other are represented by nearly orthogonal vectors— $\langle \boldsymbol{\tau}, \boldsymbol{\eta} \rangle \approx 0$ —and two items that are very similar in terms of their role in the market are represented by vectors that are almost the same— $\langle \boldsymbol{\tau}, \boldsymbol{\eta} \rangle \approx 1$. Every edge e of the trade network corresponds to an item, and we let $\boldsymbol{\tau}(e)$ denote the corresponding type vector⁵. To allow for the fact that two participants may trade multiple types of items, we allow the set of edges E to contain multiple edges between two vertices u, v . The mathematics we develop has the feature that the values of the inner products between all pairs of items is the *only* relevant feature of the type vectors to the model. For example, rotating all the type vectors by the same degree will preserve all the quantities we define, because doing so does not change any of the inner products.

One can treat the vectors that represent items in the market as parameters of the model, or use the trade network itself to generate them. If the market has n participants, and E' denotes the subset of the edges that trade item z , let $\boldsymbol{\tau}'$ be the $d = n$ dimensional vector such that $\boldsymbol{\tau}'_v = \sum_{e \in E'} w(e)$ is equal to the total amount of trade of item z that node v engages in. Then represent item z using the unit vector $\boldsymbol{\tau} = \frac{\boldsymbol{\tau}'}{\|\boldsymbol{\tau}'\|}$. Two items that are largely traded by disjoint parts of the market will correspond to nearly orthogonal vectors, and two items that are often traded by the same nodes will correspond to vectors that have a large inner product. Two items that are substitutable will likely be sold and traded by the same nodes in the market, and two items that are unrelated to each other will likely be traded by different nodes. This choice corresponds somewhat to our intuitions, but may not always be ideal. One may need to come up with these vectors using a combination of data driven methods and information about the underlying items.

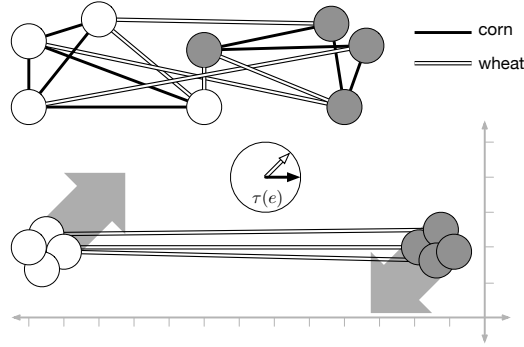
Note that if I types of items are traded in the market, it is always possible to set $d \leq I$, since we can always project the type vectors down to an I dimensional space while preserving the inner products. Moreover, by the Johnson-Lindenstrauss lemma [JL84] one can always project the type vectors to a space whose dimension is proportional to $\log I$, while introducing tiny errors in the inner products. So, one can ensure that d is much smaller than n .

A shock in the system is an nd dimensional vector $\boldsymbol{x} \in \mathbb{R}^{nd}$, that can be interpreted as n vectors of dimension d . The coordinates $\boldsymbol{x}_u \in \mathbb{R}^d$ correspond to the beliefs of node u about the amounts of all items that are worth a unit of currency, and the inner product $\langle \boldsymbol{\tau}, \boldsymbol{x}_u \rangle$ corresponds to u 's belief about the value of the item whose type vector is $\boldsymbol{\tau}$.

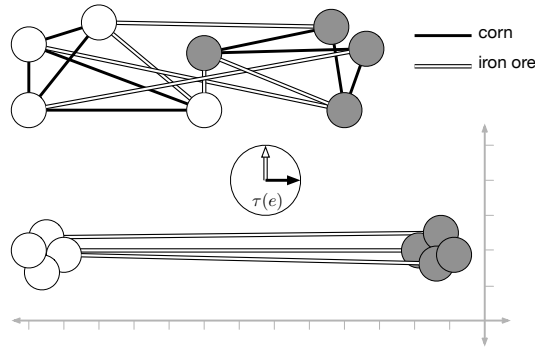
Trade along an edge $e = \{u, v\}$ induces a normalizing force on v of

$$w(e) \cdot \langle \boldsymbol{\tau}(e), \boldsymbol{x}_u - \boldsymbol{x}_v \rangle \cdot \boldsymbol{\tau}(e).$$

⁵These type vectors can be used to encode the geography of the market even in single item markets. If we want to take into account the fact that items that are nearby are likely to be more correlated, we can associate every edge with a unit vector that is proportional to the vector encoding the location of the trade.



(a) If the white nodes are placed at the point $(-1, 0)$, and gray nodes are placed at the point $(1, 0)$, the market experiences some tension along the direction corresponding to wheat, because the type vector for wheat is not orthogonal to the shock.



(b) If the white nodes are placed at the point $(-1, 0)$, and the gray nodes are placed at the point $(1, 0)$, the market experiences no tension, because the type vector for iron ore is orthogonal to the shock.

Figure 13: A shock on a multi-item network, showing the effect of choosing different type vectors.

The energy of the edge is

$$(1/2) \cdot w(e) \cdot \langle \boldsymbol{\tau}(e), \mathbf{x}_u - \mathbf{x}_v \rangle^2.$$

Thus, if $\mathbf{x}_u - \mathbf{x}_v$ is orthogonal to the type vector, then no force is induced and the energy is 0, since the parties agree on their beliefs for the item they trade. Note that these quantities remain identical if the type vector is replaced with its negation.

Given a shock \mathbf{x} , we define the tension of the shock to be the sum of the magnitudes of all forces felt by the nodes

$$\mathbf{T}(\mathbf{x}) = \sum_{u \in V} \left\| \sum_{\{u,v\} \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \mathbf{x}_v - \mathbf{x}_u \rangle \cdot \boldsymbol{\tau}(e) \right\|,$$

where here $\|\mathbf{y}\| = \sqrt{\sum_{j=1}^d \mathbf{y}_j^2}$ denotes the length of the vector \mathbf{y} . Figure 13a shows the tension a

network experiences when the shock is along a direction in one of the items, and the other item has a type vector that is not orthogonal. Figure 13b shows a similar situation when the type vectors are orthogonal. We define the energy of the shock to be the sum of the energies on all edges:

$$E(\mathbf{x}) = (1/2) \sum_{\{u,v\}=e \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \mathbf{x}_u - \mathbf{x}_v \rangle^2.$$

As in the single item case, the gradient of the energy is the amount of force needed at each node to hold the network in place:

$$(\nabla(E(\mathbf{x})))_u = \sum_{u \in e = \{u,v\}} w(e) \cdot \langle \boldsymbol{\tau}(e), \mathbf{x}_u - \mathbf{x}_v \rangle \cdot \boldsymbol{\tau}(e),$$

So adjusting beliefs in the directions of the forces felt at each node is the quickest way to reduce the energy.

4 Rigorously Defining Market Connectivity

In this section, we give mathematical definitions for the tension and energy of trade networks and show how to identify bottlenecks mathematically. We aim to give numbers that capture the behavior of the network for all significant shocks. To ease comprehension, we explain the definition for the single item case first, and then generalize the definitions to handle multiple items.

4.1 Single Item Markets

Key to the definition is what kinds of shocks should be considered significant. A small, completely disconnected part of the network can give rise to 0 tension/energy. Suppose the network consists of two disjoint sets \mathbf{A} , \mathbf{B} that have no edges between and the sum of all weights of edges in \mathbf{A} is ϵ and the sum of all weights in \mathbf{B} is $w - \epsilon$. If ϵ is very small, and \mathbf{B} is well connected, this should be considered an efficient market. But placing all the nodes of \mathbf{A} at $w - \epsilon$ and all the nodes of \mathbf{B} at $-\epsilon$ gives $E(\mathbf{x}) = T(\mathbf{x}) = 0$, even though $\sum_{u \in V} w(u) \cdot \mathbf{x}_u = 0$. Such shocks focus too much attention on small anomalies in the network and should not be considered significant.

Observe that the tension and energy of a shock remain the same if every entry of the shock is shifted by the same number. Namely, if $\mathbf{1}$ denotes the all 1's vector, then $E(\mathbf{x}) = E(\mathbf{x} + \gamma \mathbf{1})$ and $T(\mathbf{x}) = T(\mathbf{x} + \gamma \mathbf{1})$, for every number γ . Keeping this in mind, it will be convenient to shift each shock so that the average belief of the nodes (weighted by the volume of trade) is 0. Let

$$w(u) = \sum_{v \in V} w(\{u, v\}),$$

denote the sum of the weights of the edges that touch node u , and let

$$w = \sum_{u \in V} w(u).$$

Then for any shock \mathbf{x} , define

$$\bar{\mathbf{x}} = \mathbf{x} - \frac{\sum_{u \in V} w(u) \cdot \mathbf{x}_u}{w} \cdot \mathbf{1}.$$

$\bar{\mathbf{x}}$ is the shock obtained by shifting \mathbf{x} so that the average shock value (weighted by the weight of the nodes) is 0:

$$\sum_{u \in V} w(u) \cdot \bar{\mathbf{x}}_u = \sum_{u \in V} w(u) \cdot \mathbf{x}_u - \sum_{u \in V} w(u) \cdot \mathbf{x}_u = 0.$$

For any node u , $\bar{\mathbf{x}}_u$ is a measure of how far node u 's belief is from the average belief in the network. In order to preclude shocks that put undue weight on small parts of the network, we define the set of significant shocks of magnitude α , for $0 \leq \alpha \leq 1$, to be:

$$\mathcal{V}_\alpha = \left\{ \mathbf{x} : \text{for every } v \in V, \sum_{u \in V} \frac{w(u) \cdot |\bar{\mathbf{x}}_u|}{w} \geq \alpha \cdot |\bar{\mathbf{x}}_v| \right\}.$$

Significant shocks are those where the magnitude of the shock is well spread among the nodes in the sense that the shock affects at least an α (weighted) fraction of the trade in the market. Our measure of connectivity should be scale invariant. One can always increase or decrease the energy and tension by multiplying the shock by a scalar. To ignore such scaling factors, we measure the ratio of the tension and energy to a normalization factor:

Definition 2. *The α -tension of the trade network is*

$$T_\alpha = \min_{\mathbf{x} \in \mathcal{V}_\alpha} \frac{T(\mathbf{x})}{2 \sum_{u \in V} w(u) \cdot |\bar{\mathbf{x}}_u|}.$$

Next we define the energy⁶ of the trade network. Significant shocks in this context are given by:

$$\mathcal{U}_\alpha = \left\{ \mathbf{x} : \text{for every } v \in V, \sum_{u \in V} \frac{w(u) \cdot \bar{\mathbf{x}}_u^2}{w} \geq \alpha \cdot \bar{\mathbf{x}}_v^2 \right\}.$$

Definition 3. *The α -energy of the trade network is*

$$E_\alpha = \min_{\mathbf{x} \in \mathcal{U}_\alpha} \frac{E(\mathbf{x})}{\sum_{u \in V} w(u) \cdot \bar{\mathbf{x}}_u^2}.$$

The differences in the choices of the normalization factors and the choices of significant shocks for tension and energy can be explained by the properties they ensure. The proofs of these facts are elementary, and can be found in Appendix A.

Fact 4. $1 \geq T_\alpha \geq 0$ and $1 \geq E_\alpha \geq 0$.

Both quantities are relatively large when the network is well connected:

Fact 5. *For every α , the fully connected network on n nodes with unit edge weights has*

$$T_\alpha = \frac{1}{2} + \frac{1}{2(n-1)} = E_\alpha.$$

⁶The energy of the network is closely related to the concept of *conductance* in spectral graph theory. The expression for the energy is the same as the expression for the conductance of the normalized Laplacian of a graph.

When the number of participants is $n = 2$, Fact 5 shows that $T_\alpha = 1 = E_\alpha$. Since the set of significant shocks only gets bigger as α gets smaller, we have:

Fact 6. For every $\alpha \leq \beta$, $T_\alpha \leq T_\beta$, $E_\alpha \leq E_\beta$.

Although the set of significant shocks for tension and energy are different, they are closely related:

Lemma 7. $\mathcal{U}_\alpha \subseteq \mathcal{V}_\alpha \subseteq \mathcal{U}_{\alpha^2}$.

Finally, we have:

Fact 8. $E_\alpha = 0 \Rightarrow T_\alpha = 0 \Rightarrow E_{\alpha^2} = 0$.

Let B be a subset of k nodes. Define $T^B(\mathbf{x})$ to be the tension of the shock after deleting the edges of the network that do not touch B .

Definition 9. A (k, α, β) bottleneck for tension is a set of k nodes B such that there is a vector $\mathbf{x} \in \mathcal{V}_\alpha$ with $T(\mathbf{x}) > 0$ and

$$\frac{T^B(\mathbf{x})}{T(\mathbf{x})} \geq \beta.$$

Similarly, let $E^B(\mathbf{x})$ denote the energy due to all the edges that touch the set B .

Definition 10. A (k, α, β) bottleneck for energy is a set of k nodes B such that there is a vector $\mathbf{x} \in \mathcal{U}_\alpha$ with $E(\mathbf{x}) > 0$ and

$$\frac{E^B(\mathbf{x})}{E(\mathbf{x})} \geq \beta.$$

A (k, α, β) bottleneck is a serious obstruction to efficiency when k is small and α, β are relatively large. Intuitively, when the network has a (k, α, β) bottleneck, there is a small set of k nodes B that is responsible for an unusually large share β of the tension/energy of the network when the network experience a large shock of magnitude α .

4.2 Markets with Multiple Items

As in the single item case, we define:

$$w(u) = \sum_{v \in V} w(\{u, v\}),$$

to be the total weight of all edges that touch u , and

$$w = \sum_{u \in V} w(u)$$

to be sum of all the weights of each node.

We need a technical lemma whose proof is deferred to Appendix A:

Lemma 11. For every shock \mathbf{x} , there is a vector $\mathbf{z} \in \mathbb{R}^d$ such that

$$\sum_{u \in e \in E} w(e) \langle \boldsymbol{\tau}(e), \mathbf{x}_u \rangle \cdot \boldsymbol{\tau}(e) = \sum_{u \in e \in E} w(e) \langle \boldsymbol{\tau}(e), \mathbf{z} \rangle \cdot \boldsymbol{\tau}(e).$$

Let \mathbf{y} be the vector such that for every $u \in V$, $\mathbf{y}_u = \mathbf{z}$, and define

$$\bar{\mathbf{x}} = \mathbf{x} - \mathbf{y}.$$

The point of this definition is that the energy and tension of $\bar{\mathbf{x}}$ remain exactly the same as those of \mathbf{x} ,

$$\mathbf{E}(\mathbf{x}) = \mathbf{E}(\bar{\mathbf{x}}), \mathbf{T}(\mathbf{x}) = \mathbf{T}(\bar{\mathbf{x}}),$$

but the weighted average belief of $\bar{\mathbf{x}}$ is 0:

$$\sum_{u \in e \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle \cdot \boldsymbol{\tau}(e) = 0,$$

and $\bar{\mathbf{x}}$ still captures the inefficiencies described in \mathbf{x} .

Towards defining the α -tension, for each $\alpha > 0$, we define the set of significant shocks to be

$$\mathcal{V}_\alpha = \left\{ \mathbf{x} : \text{for every } v \in V, \sum_{u \in e \in E} \frac{w(e) \cdot |\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle|}{w} \geq \alpha \cdot \|\bar{\mathbf{x}}_v\| \right\}.$$

Definition 12. For $\alpha > 0$, the α -tension of the trade network is

$$\mathbf{T}_\alpha = \min_{\mathbf{x} \in \mathcal{V}_\alpha} \frac{\mathbf{T}(\mathbf{x})}{2 \cdot \sum_{u \in e \in E} w(e) \cdot |\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle|}.$$

For the energy, we define significant shocks with the set:

$$\mathcal{U}_\alpha = \left\{ \mathbf{x} : \text{for every } v \in V, \sum_{u \in e \in E} \frac{w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle^2}{w} \geq \alpha \cdot \|\bar{\mathbf{x}}_v\|^2 \right\}.$$

Definition 13. For $\alpha > 0$, the α -energy of the trade network is

$$\mathbf{E}_\alpha = \min_{\mathbf{x} \in \mathcal{U}_\alpha} \frac{\mathbf{E}(\mathbf{x})}{\sum_{u \in e \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle^2}.$$

In the case that $\boldsymbol{\tau}(e)$ is the same vector for every edge e , these definitions are identical to those defined for the single item case. To understand the definitions, we prove several facts whose proofs can be found in Appendix A. Continuing the analogy with the case of single item markets, we have:

Fact 14. $1 \geq \mathbf{T}_\alpha \geq 0$ and $1 \geq \mathbf{E}_\alpha \geq 0$.

Fact 15. For any choice of type vectors, if the network has n nodes such that every node trades every item with every other node at unit rate, we have $\mathbf{T}_\alpha \geq \frac{\alpha}{2}$, and $\mathbf{E}_\alpha = \frac{1}{2} + \frac{1}{2(n-1)}$.

Note that when the number of participants is $n = 2$, Fact 15 shows that $\mathbf{E}_\alpha = 1$. For large n , the energy of this completely connected network is close to $1/2$. To contrast, the tension can actually be small. In Appendix B, we give an example showing that one can have fully connected networks where the tension is close to $\alpha/2$.

Since the set of significant shocks only gets bigger as α gets smaller, we have:

Fact 16. For every $\alpha \leq \beta$, $T_\alpha \leq T_\beta$ and $E_\alpha \leq E_\beta$.

Lemma 17. $\mathcal{U}_\alpha \subseteq \mathcal{V}_\alpha \subseteq \mathcal{U}_{\alpha^2}$

Finally, we have:

Fact 18. $E_\alpha = 0 \Rightarrow T_\alpha = 0 \Rightarrow E_{\alpha^2} = 0$.

Given the changes we have already made, the definitions that capture bottlenecks are identical to the corresponding definitions in the single item case. Define $T^B(\mathbf{x})$ to be the tension of the shock after deleting the edges of the network that do not touch B .

Definition 19. A (k, α, β) bottleneck for tension is a set of k nodes B such that there is a vector $\mathbf{x} \in \mathcal{V}_\alpha$ with $T(\mathbf{x}) > 0$ and

$$\frac{T^B(\mathbf{x})}{T(\mathbf{x})} \geq \beta.$$

Similarly, let $E^B(\mathbf{x})$ denote the energy due to all the edges that touch the set B .

Definition 20. A (k, α, β) bottleneck for energy is a set of k nodes B such that there is a vector $\mathbf{x} \in \mathcal{U}_\alpha$ with $E(\mathbf{x}) > 0$ and

$$\frac{E^B(\mathbf{x})}{E(\mathbf{x})} \geq \beta.$$

A (k, α, β) bottleneck is a serious obstruction to efficiency when k is small and α, β are relatively large. Intuitively, when the network has a (k, α, β) bottleneck, there is a small set of k nodes B that is responsible for an unusually large share β of the tension/energy of the network when the network experience a large shock of magnitude α .

5 Discussion and Future Work

The principal contribution of this work is the definition of the trade network and the concepts of tension, energy and bottlenecks, which give new metrics for measuring the health of a market. The trade network provides a bridge from observable data to concepts like market efficiency, without making any assumptions about the evolution or correctness of prices. The trade network encodes data about the information that is visible to market participants and the information participants are able to act on.

We believe that the mathematics of our theory is useful to measure efficiency because the single item case is similar to equations used to model analogous concepts in physics. Experiments have confirmed that these equations accurately describe electrical networks, networks of springs, and the diffusion of heat. That said, one can only know if the choices we have made in the formulas for tension, energy and bottlenecks are well-founded by experimenting with real data. Experiments can help to validate the model, or find issues that suggest refinements to the model. Gathering data to approximate the trade network would be extremely useful.

It would be very useful to find sampling based techniques to estimate the energy/tension of the network, and use them to compute the statistics we have defined in this work. Can we sample small amounts of data about the trade network and use it to compute the tension and energy, and find bottlenecks?

There are several ways to extend our model that we believe could be fruitful:

- One could study the evolution of the trade network over time. An efficient market should have a rapidly evolving trade network, with traders switching trading partners to take advantage of new information and better prices. If nodes do not switch trading partners in response to shocks, that is an inefficiency. One could look for a model that rigorously captures this kind of market efficiency using data.
- One could study the relationship between the rates of trade and prices. If nodes do not often trade with partners offering them the best price, the market is likely to be inefficient. Again, it remains open to rigorously define this kind of efficiency.

6 Acknowledgements

Thanks to Paul Beame, Morgan Dixon, Ran Duchin, Abe Friesen, Shayan Oveis Gharan, Kevin Minter, Shay Moran, Alex Jaffe, Anna Karlin, Kayur Patel, Sivaramakrishnan Natarajan Ramamoorthy, Darcy Rao, Tim Roughgarden, Brad Wagenaar, Mark Westerfield, Alex White and Amir Yehudayoff for useful comments. Thanks to Kayur Patel for many suggestions that improved the figures in this paper.

References

- [DeG74] Reaching a consensus. *Journal of the American Statistical Association*, 69:118–121, 1974.
- [AOTS15] Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, February 2015.
- [Bal15] Eric Balchunas. The amount of ETF shares being traded has eclipsed U.S. GDP. *BloombergMarkets*, July 2015.
- [Bol98] Bela Bollobas. *Modern Graph Theory*. Springer, 1998.
- [CRS08] Tarun Chordia, Richard Roll, and Avanidhar Subrahmanyam. Liquidity and market efficiency. *Journal of Financial Economics*, 87:249–268, February 2008.
- [Chu96] Fan Chung. *Spectral Graph Theory*, volume 92 of *Regional Conference Series in Mathematics*. AMS, 1996.
- [EGJ14] Matthew Elliott, Benjamin Golub, and Matthew O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–53, October 2014.
- [Fam70] Eugene F. Fama. Efficient capital markets. *The Journal of Finance*, 25(2):383–417, 1970.
- [Fam91] Eugene F. Fama. Efficient capital markets: II. *The Journal of Finance*, 46(5):1575–1617, 1991.
- [FK07] Joseph Farrell and Paul Klemperer. Coordination and lock-in: competition with switching costs and network effects. In *Handbook of Industrial Organization*, chapter 31, pages 1967–2072. Elsevier, 2007.

- [FG10] Neil Fligstein and Adam Goldstein. The anatomy of the mortgage securitization crisis. *Research in the Sociology of Organization*, 30A:29–70, 2010.
- [FR16] Neil Fligstein and Alexander F. Roehrkasse. The causes of fraud in the financial crisis of 2007 to 2009. Evidence from the mortgage-backed securities industry. *American Sociological Review*, 81(4):617–643, 2016.
- [GHM12] Christian Gouriroux, Jean-Cyprien Ham, and Alain Monfort. Bilateral exposures and systemic solvency risk. *Canadian Journal of Economics/Revue canadienne d'économie*, 45(4):1273–1309, 2012.
- [Hay45] Friedrich Hayek. The use of knowledge in society. *The American Economic Review*, 35(4):519–530, 1945.
- [Hei77] Thomas A. Heironymus. *The Economics of Futures Trading*. 1977.
- [Jac10] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, 2010.
- [JL84] W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Conference in Modern Analysis and Probability*, 26:189–206, 1984.
- [KRV03] Graciela L. Kaminsky, Carmen M. Reinhart, and Carlos A. Végh. The unholy trinity of financial contagion. *Journal of Economic Perspectives*, 17:51–74, Fall 2003.
- [Mal03] Burton G. Malkiel. The efficient market hypothesis and its critics. *Journal of Economic Perspectives*, 17(1):59–82, 2003.
- [PS15] Devin G. Pope and Justin R. Sydnor. Behavioral economics: Economics as a psychological discipline. In *The Wiley Blackwell Handbook of Judgment and Decision Making*, chapter 28. 2015.
- [RS16] Yuval Rabani and Leonard J. Schulman. The invisible hand of laplace: the role of market structure in price convergence and oscillation, February 24 2016.
- [Spi] Dan Spielman. Course notes on spectral graph theory. <http://www.cs.yale.edu/homes/spielman/561/>.

A Technical Proofs

Proof of Fact 4. The tension is clearly non-negative. We have:

$$\begin{aligned}
T(\mathbf{x}) = T(\bar{\mathbf{x}}) &= \sum_{u \in V} \left| \sum_{v \in V} w(\{u, v\}) \cdot (\bar{\mathbf{x}}_u - \bar{\mathbf{x}}_v) \right| \\
&\leq \sum_{u \in V} \sum_{v \in V} w(\{u, v\}) \cdot (|\bar{\mathbf{x}}_u| + |\bar{\mathbf{x}}_v|) \\
&= 2 \sum_{u \in V} w(u) \cdot |\bar{\mathbf{x}}_u|,
\end{aligned}$$

proving that $T_\alpha \leq 1$, for $\mathbf{x} \in \mathcal{V}_\alpha$, since such \mathbf{x} must have $2 \sum_{u \in V} w(u) \cdot |\bar{\mathbf{x}}_u| > 0$.

Similarly, the energy is clearly non-negative. We have

$$\begin{aligned}
\mathbf{E}(\mathbf{x}) &= \mathbf{E}(\bar{\mathbf{x}}) = (1/2) \sum_{\{u,v\} \in E} w(\{u,v\}) \cdot |\bar{\mathbf{x}}_u - \bar{\mathbf{x}}_v|^2 \\
&\leq \sum_{\{u,v\} \in E} w(\{u,v\}) \cdot (|\bar{\mathbf{x}}_u|^2 + |\bar{\mathbf{x}}_v|^2) \quad \text{since } 2a^2 + 2b^2 \geq (a+b)^2 \text{ for all } a, b. \\
&= \sum_{u \in V} w(u) \cdot |\bar{\mathbf{x}}_u|^2,
\end{aligned}$$

proving that $\mathbf{E}_\alpha \leq 1$, for $\mathbf{x} \in \mathcal{U}_\alpha$. □

Proof of Fact 5. For any shock $\mathbf{x} \in \mathcal{V}_\alpha$, we have

$$\begin{aligned}
\mathbf{T}(\mathbf{x}) &= \mathbf{T}(\bar{\mathbf{x}}) = \sum_{u \in V} \left| \sum_{v \in V} (\bar{\mathbf{x}}_u - \bar{\mathbf{x}}_v) \right| \\
&= \sum_{u \in V} \left| n \cdot \bar{\mathbf{x}}_u - \sum_{v \in V} \bar{\mathbf{x}}_v \right| \\
&= \sum_{u \in V} n \cdot |\bar{\mathbf{x}}_u|. \quad \text{by the definition of } \bar{\mathbf{x}}
\end{aligned}$$

Thus we have

$$\mathbf{T}_\alpha = \min_{\mathbf{x} \in \mathcal{V}_\alpha} \frac{\mathbf{T}(\mathbf{x})}{2 \sum_{u \in V} (n-1) |\bar{\mathbf{x}}_u|} = \frac{n}{2(n-1)} = \frac{1}{2} + \frac{1}{2(n-1)}.$$

For any $\mathbf{x} \in \mathcal{U}_\alpha$, we have

$$\begin{aligned}
\mathbf{E}(\mathbf{x}) &= \mathbf{E}(\bar{\mathbf{x}}) = (1/2) \sum_{\{u,v\} \in E} (\bar{\mathbf{x}}_u - \bar{\mathbf{x}}_v)^2 \\
&= (1/2) \sum_{\{u,v\} \in E} \bar{\mathbf{x}}_u^2 + \bar{\mathbf{x}}_v^2 - 2\bar{\mathbf{x}}_u\bar{\mathbf{x}}_v \\
&= (1/2) \sum_{u \in V} (n-1)\bar{\mathbf{x}}_u^2 - 2 \sum_{\{u,v\} \in E} \bar{\mathbf{x}}_u\bar{\mathbf{x}}_v.
\end{aligned}$$

Now since $\sum_u \bar{\mathbf{x}}_u = 0$, we have that for every u , $\sum_{u \neq v} \bar{\mathbf{x}}_v = -\bar{\mathbf{x}}_u$. So

$$\begin{aligned}
\mathbf{E}(\mathbf{x}) &= (1/2) \sum_{u \in V} (n-1)\bar{\mathbf{x}}_u^2 - 2 \sum_{u \in V} \bar{\mathbf{x}}_u(-\bar{\mathbf{x}}_u) \\
&= (1/2) \sum_{u \in V} (n-1)\bar{\mathbf{x}}_u^2 + 2 \sum_{u \in V} \bar{\mathbf{x}}_u^2 \\
&= (1/2) \sum_{u \in V} (n+1)\bar{\mathbf{x}}_u^2.
\end{aligned}$$

Thus

$$\mathbf{E}_\alpha = \min_{\mathbf{x} \in \mathcal{U}_\alpha} \frac{\sum_{u \in V} (n+1)\bar{\mathbf{x}}_u^2}{2 \sum_{u \in V} (n-1)\bar{\mathbf{x}}_u^2} = \frac{n+1}{2(n-1)} = \frac{1}{2} + \frac{2}{2(n-1)}.$$

□

Proof of Lemma 7. Suppose $\mathbf{x} \in \mathcal{V}_\alpha$. Then we have

$$\begin{aligned}
\alpha \cdot |\bar{\mathbf{x}}_v| &\leq \sum_{u \in V} \frac{w(u) \cdot |\bar{\mathbf{x}}_u|}{w} \\
&\leq \sqrt{\sum_{u \in V} \frac{w(u)}{w}} \cdot \sqrt{\sum_{u \in V} \frac{w(u)}{w} \cdot \bar{\mathbf{x}}_u^2} && \text{by the Cauchy-Schwartz inequality} \\
&= \sqrt{\sum_{u \in V} \frac{w(u)}{w} \cdot \bar{\mathbf{x}}_u^2},
\end{aligned}$$

proving that

$$\sum_{u \in V} \frac{w(u) \cdot \bar{\mathbf{x}}_u^2}{w} \geq \alpha^2 \cdot \bar{\mathbf{x}}_v^2,$$

and so $\mathbf{x} \in \mathcal{U}_{\alpha^2}$. For the other containment, suppose that $\mathbf{x} \in \mathcal{U}_\alpha$, and $\bar{\mathbf{x}}_{v'}$ is the coordinate that has maximum magnitude. Then for every $v \in V$,

$$\begin{aligned}
\alpha \cdot |\bar{\mathbf{x}}_v| &\leq \alpha \cdot |\bar{\mathbf{x}}_{v'}| = \alpha \cdot \frac{|\bar{\mathbf{x}}_{v'}|^2}{|\bar{\mathbf{x}}_{v'}|} \leq \frac{1}{|\bar{\mathbf{x}}_{v'}|} \sum_{u \in V} \frac{w(u) \cdot \bar{\mathbf{x}}_u^2}{w} && \text{since } \bar{\mathbf{x}} \in \mathcal{U}_\alpha \\
&= |\bar{\mathbf{x}}_{v'}| \sum_{u \in V} \frac{w(u) \cdot \bar{\mathbf{x}}_u^2}{w \cdot \bar{\mathbf{x}}_{v'}^2} \\
&\leq |\bar{\mathbf{x}}_{v'}| \sum_{u \in V} \frac{w(u) \cdot |\bar{\mathbf{x}}_u|}{w \cdot |\bar{\mathbf{x}}_{v'}|} && \text{since } \frac{|\bar{\mathbf{x}}_u|}{|\bar{\mathbf{x}}_{v'}|} \leq 1 \\
&= \sum_{u \in V} \frac{w(u) \cdot |\bar{\mathbf{x}}_u|}{w},
\end{aligned}$$

so $\mathbf{x} \in \mathcal{V}_\alpha$. □

Proof of Fact 8. Suppose $\mathbf{x} \in \mathcal{U}_\alpha$ is a shock such that $\mathbf{E}(\mathbf{x}) = 0$. Then by lemma 7, $\mathbf{x} \in \mathcal{V}_\alpha$. Moreover, for every edge $\{u, v\}$ in the network, we have $\mathbf{x}_u = \mathbf{x}_v$, or else the energy would be positive. So we must have $\mathbf{T}(\mathbf{x}) = 0$, proving that $\mathbf{T}_\alpha = 0$.

Suppose $\mathbf{x} \in \mathcal{V}_\alpha$ and $\mathbf{T}(\mathbf{x}) = 0$. Then by lemma 7, $\mathbf{x} \in \mathcal{U}_{\alpha^2}$. Let $f(t) \in \mathbb{R}^n$ be the vector of forces felt by the nodes during the shock $t\mathbf{x}$. Then we see that $f(t)$ is t times the vector of forces felt when the shock is \mathbf{x} , so $f(t) = \mathbf{0}$ for every t . Since this vector is the gradient of the energy at $t\mathbf{x}$, we must have $\mathbf{E}(1 \cdot \mathbf{x}) = \mathbf{E}(0 \cdot \mathbf{x}) = 0$. □

Proof of Lemma 11. This proof requires some knowledge about positive semidefinite matrices. View $\boldsymbol{\tau}(e)$ as a d dimensional column vector, and consider the $d \times d$ matrix

$$A = \sum_{e \in E} w(e) \cdot \boldsymbol{\tau}(e) \boldsymbol{\tau}(e)^\top.$$

We have

$$A\mathbf{z} = \sum_{e \in E} w(e) \langle \boldsymbol{\tau}(e), \mathbf{z} \rangle \cdot \boldsymbol{\tau}(e),$$

and A is positive semidefinite, since for every \mathbf{z} , we have

$$\mathbf{z}^\top A \mathbf{z} = \sum_{u \in e \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \mathbf{z} \rangle^2 \geq 0.$$

Thus there is an orthonormal basis $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_d$ for the space \mathbb{R}^d and non-negative numbers $\gamma_1, \dots, \gamma_d$ such that

$$A = \sum_{i=1}^d \gamma_i \cdot \mathbf{e}_i \mathbf{e}_i^\top.$$

Moreover, every type vector is contained in the span of vectors \mathbf{e}_i for which $\gamma_i > 0$, because if, for example there is some edge for which $\boldsymbol{\tau}(e)$ is not contained in this span, then we can express $\boldsymbol{\tau}(e) = \mathbf{a} + \mathbf{b}$, where \mathbf{a} is in the span, and $\mathbf{b} \neq \mathbf{0}$ is orthogonal to it. Then we have

$$0 = \mathbf{b}^\top \left(\sum_{i=1}^d \gamma_i \cdot \mathbf{e}_i \mathbf{e}_i^\top \right) \mathbf{b} = \mathbf{b}^\top A \mathbf{b} = \sum_{u \in e' \in E} w(e') \langle \boldsymbol{\tau}(e'), \mathbf{b} \rangle^2 \geq w(e) \langle \boldsymbol{\tau}(e), \mathbf{b} \rangle^2 = w(e) \cdot \|\mathbf{b}\|^2 > 0,$$

which is a contradiction. So, for $\mathbf{q} = \sum_{u \in e \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \mathbf{x}_u \rangle \cdot \boldsymbol{\tau}(e)$, we can set

$$\mathbf{z} = \sum_{i: \gamma_i > 0} (1/\gamma_i) \langle \mathbf{q}, \mathbf{e}_i \rangle \cdot \mathbf{e}_i,$$

and then we get

$$\sum_{u \in e \in E} w(e) \langle \boldsymbol{\tau}(e), \mathbf{z} \rangle \cdot \boldsymbol{\tau}(e) = A \mathbf{z} = \sum_{i=1}^d \langle \mathbf{q}, \mathbf{e}_i \rangle \cdot \mathbf{e}_i = \mathbf{q},$$

as required. \square

Proof of Fact 14. The tension is clearly non-negative. We have:

$$\begin{aligned} T(\mathbf{x}) &= T(\bar{\mathbf{x}}) = \sum_{u \in V} \left\| \sum_{\{u,v\}=e \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_v - \bar{\mathbf{x}}_u \rangle \cdot \boldsymbol{\tau}(e) \right\|^2 \\ &\leq \sum_{u \in V} \sum_{\{u,v\}=e \in E} w(e) \cdot \|\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_v - \bar{\mathbf{x}}_u \rangle \cdot \boldsymbol{\tau}(e)\|^2 \\ &\leq \sum_{u \in V} \sum_{\{u,v\}=e \in E} w(e) \cdot (|\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_v \rangle| + |\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle|)^2 \\ &= 2 \sum_{u \in e \in E} w(e) \cdot |\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle|^2, \end{aligned}$$

proving that $T_\alpha \leq 1$, for $\mathbf{x} \in \mathcal{V}_\alpha$.

Similarly, the energy is clearly non-negative. We have

$$\begin{aligned} E(\mathbf{x}) &= E(\bar{\mathbf{x}}) = (1/2) \sum_{\{u,v\}=e \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u - \bar{\mathbf{x}}_v \rangle^2 \\ &\leq \sum_{\{u,v\}=e \in E} w(e) \cdot (\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle^2 + \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_v \rangle^2) \quad \text{since } a^2 + b^2 \geq (a+b)^2/2 \text{ for all } a, b. \\ &= \sum_{u \in e \in E} w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle^2, \end{aligned}$$

proving that $E_\alpha \leq 1$, for $\mathbf{x} \in \mathcal{U}_\alpha$. \square

Proof of Fact 15. Let I denote the set of items, and let $\boldsymbol{\tau}(i)$ denote the type vector of item i . Then for any shock $\boldsymbol{x} \in \mathcal{V}_\alpha$, we have

$$\begin{aligned}
\mathbb{T}(\boldsymbol{x}) &= \mathbb{T}(\bar{\boldsymbol{x}}) = \sum_{u \in V} \left\| \sum_{v \in V, i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_v - \bar{\boldsymbol{x}}_u \rangle \cdot \boldsymbol{\tau}(i) \right\| \\
&= \sum_{u \in V} \left\| (n-1) \sum_{i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle \cdot \boldsymbol{\tau}(i) - \sum_{v \in V, i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_v \rangle \cdot \boldsymbol{\tau}(i) \right\| \\
&= \sum_{u \in V} \left\| (n-1) \sum_{i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle \cdot \boldsymbol{\tau}(i) \right\|. \quad \text{by the definition of } \bar{\boldsymbol{x}}
\end{aligned}$$

To bound this expression, we use the Cauchy-Schwartz inequality twice:

$$\begin{aligned}
&(n-1) \sum_{u \in V} \left\| \sum_{i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle \cdot \boldsymbol{\tau}(i) \right\| \\
&\geq (n-1) \sum_{u \in V} \frac{\langle \sum_{i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle \cdot \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle}{\|\bar{\boldsymbol{x}}_u\|} \quad \text{by Cauchy-Schwartz} \\
&\geq (n-1) \sum_{u \in V, i \in I} \frac{\langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle^2}{\|\bar{\boldsymbol{x}}_u\|}.
\end{aligned}$$

Now since $\bar{\boldsymbol{x}} \in \mathcal{V}_\alpha$, we have

$$\begin{aligned}
&\geq (n-1) \sum_{u \in V, i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle^2 \cdot \frac{\alpha n |I|}{\sum_{u \in V, i \in I} |\langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle|} \quad \text{since } \bar{\boldsymbol{x}} \in \mathcal{V}_\alpha \\
&\geq (n-1) \frac{\left(\sum_{u \in V, i \in I} |\langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle| \right)^2}{n |I|} \cdot \frac{\alpha n |I|}{\sum_{u \in V, i \in I} |\langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle|} \quad \text{by Cauchy-Schwartz} \\
&= \alpha (n-1) \sum_{u \in V, i \in I} |\langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle|.
\end{aligned}$$

So we get

$$\mathbb{T}_\alpha = \min_{\boldsymbol{x} \in \mathcal{V}_\alpha} \frac{\mathbb{T}(\boldsymbol{x})}{2 \sum_{u \in V, i \in I} (n-1) |\langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle|} \geq \frac{\alpha}{2}.$$

Similarly, we have that for any shock $\boldsymbol{x} \in \mathcal{U}_\alpha$,

$$\begin{aligned}
\mathbb{E}(\boldsymbol{x}) &= \mathbb{E}(\bar{\boldsymbol{x}}) = (1/2) \sum_{u \neq v, i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u - \bar{\boldsymbol{x}}_v \rangle^2 \\
&= (1/2) \sum_{u \neq v, i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle^2 + \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_v \rangle^2 - 2 \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_v \rangle \\
&= (1/2) \sum_{u, i \in I} (n-1) \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle^2 - \sum_{u \in \mathcal{U}, i \in I} \langle \boldsymbol{\tau}(i), \bar{\boldsymbol{x}}_u \rangle \cdot \left\langle \boldsymbol{\tau}(i), \sum_{v \neq u} \bar{\boldsymbol{x}}_v \right\rangle. \quad (2)
\end{aligned}$$

Now by the definition of $\bar{\mathbf{x}}$,

$$\mathbf{0} = \sum_{u \in V, i \in I} \langle \boldsymbol{\tau}(i), \bar{\mathbf{x}}_u \rangle \boldsymbol{\tau}(i) = \sum_{i \in I} \left\langle \boldsymbol{\tau}(i), \sum_{u \in V} \bar{\mathbf{x}}_u \right\rangle \cdot \boldsymbol{\tau}(i),$$

so

$$\begin{aligned} 0 &= \left\langle \sum_{i \in I} \left\langle \boldsymbol{\tau}(i), \sum_{u \in V} \bar{\mathbf{x}}_u \right\rangle \cdot \boldsymbol{\tau}(i), \sum_{u \in V} \bar{\mathbf{x}}_u \right\rangle \\ &= \sum_{i \in I} \left\langle \boldsymbol{\tau}(i), \sum_{u \in V} \bar{\mathbf{x}}_u \right\rangle^2, \end{aligned}$$

proving that for every i , $\langle \boldsymbol{\tau}(i), \sum_{u \in V} \bar{\mathbf{x}}_u \rangle = 0$, and $\langle \boldsymbol{\tau}(i), \sum_{v \neq u} \bar{\mathbf{x}}_v \rangle = -\langle \boldsymbol{\tau}(i), \bar{\mathbf{x}}_u \rangle$. Returning to (2), we have

$$\begin{aligned} \mathbb{E}(\bar{\mathbf{x}}) &= (1/2) \sum_{u, i \in I} (n-1) \langle \boldsymbol{\tau}(i), \bar{\mathbf{x}}_u \rangle^2 - \sum_{u \in u, i \in I} \langle \boldsymbol{\tau}(i), \bar{\mathbf{x}}_u \rangle \left\langle \boldsymbol{\tau}(i), \sum_{v \neq u} \bar{\mathbf{x}}_v \right\rangle \\ &= (1/2) \sum_{u, i \in I} (n-1) \langle \boldsymbol{\tau}(i), \bar{\mathbf{x}}_u \rangle^2 + \sum_{u \in u, i \in I} \langle \boldsymbol{\tau}(i), \bar{\mathbf{x}}_u \rangle^2 \\ &= (1/2) \sum_{u, i \in I} (n+1) \langle \boldsymbol{\tau}(i), \bar{\mathbf{x}}_u \rangle^2. \end{aligned}$$

This proves that

$$\mathbb{E}_\alpha = \min_{\mathbf{x} \in \mathcal{U}_\alpha} \frac{\mathbb{E}(\mathbf{x})}{\sum_{u \in V, i \in I} (n-1) \langle \boldsymbol{\tau}(i), \bar{\mathbf{x}}_u \rangle^2} = \frac{n+1}{2(n-1)} = \frac{1}{2} + \frac{1}{2(n-1)}.$$

□

Proof of Lemma 17. If $\mathbf{x} \in \mathcal{U}_\alpha$, let v' be the node for which $\|\bar{\mathbf{x}}_{v'}\|$ is maximized. Then for every $v \in V$,

$$\begin{aligned} \alpha \cdot \|\bar{\mathbf{x}}_v\| &\leq \frac{\alpha \cdot \|\bar{\mathbf{x}}_{v'}\|^2}{\|\bar{\mathbf{x}}_{v'}\|} \leq \sum_{u \in e \in E} \frac{w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle^2}{w \cdot \|\bar{\mathbf{x}}_{v'}\|} = \|\bar{\mathbf{x}}_{v'}\| \sum_{u \in e \in E} \frac{w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle^2}{w \cdot \|\bar{\mathbf{x}}_{v'}\|^2} \\ &\leq \|\bar{\mathbf{x}}_{v'}\| \sum_{u \in e \in E} \frac{w(e) \cdot |\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle|}{w \cdot \|\bar{\mathbf{x}}_{v'}\|} \\ &= \sum_{u \in e \in E} \frac{w(e) \cdot |\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle|}{w}, \end{aligned}$$

so $\mathbf{x} \in \mathcal{V}_\alpha$. If $\mathbf{x} \in \mathcal{V}_\alpha$, then for every $v \in V$,

$$\begin{aligned} \alpha \cdot \|\bar{\mathbf{x}}_v\| &\leq \sum_{u \in e \in E} \frac{w(e) \cdot |\langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle|}{w} \\ &\leq \sqrt{\sum_{u \in e \in E} \frac{w(e)}{w}} \cdot \sqrt{\sum_{u \in e \in E} \frac{w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle^2}{w}} \quad \text{by the Cauchy-Schwartz inequality} \\ &= \sqrt{\sum_{u \in e \in E} \frac{w(e) \cdot \langle \boldsymbol{\tau}(e), \bar{\mathbf{x}}_u \rangle^2}{w}}, \end{aligned}$$

proving that $\mathbf{x} \in \mathcal{U}_{\alpha^2}$. □

Proof of Fact 18. Suppose $\mathbf{x} \in \mathcal{U}_\alpha$ is a shock such that $E(\mathbf{x}) = 0$. Then by Lemma 17, $\mathbf{x} \in \mathcal{V}_\alpha$. Moreover, for every edge $e = \{u, v\}$ in the network, we have $\langle \boldsymbol{\tau}(e), \mathbf{x}_u - \mathbf{x}_v \rangle = 0$, or else the energy would be positive. So we must have $T(\mathbf{x}) = 0$, proving that $T_\alpha = 0$.

Suppose $\mathbf{x} \in \mathcal{V}_\alpha$ and $T(\mathbf{x}) = 0$. Then by Lemma 17, $\mathbf{x} \in \mathcal{U}_{\alpha^2}$. Let $f(t) \in \mathbb{R}^n$ be the vector of forces felt by the nodes during the shock $t\mathbf{x}$. Then we see that $f(t)$ is t times the vector of forces felt when the shock is \mathbf{x} , so $f(t) = \mathbf{0}$ for every t . Since this vector is the gradient of the energy at $t\mathbf{x}$, we must have $E(1 \cdot \mathbf{x}) = E(0 \cdot \mathbf{x}) = 0$. □

B Example Matching Fact 15

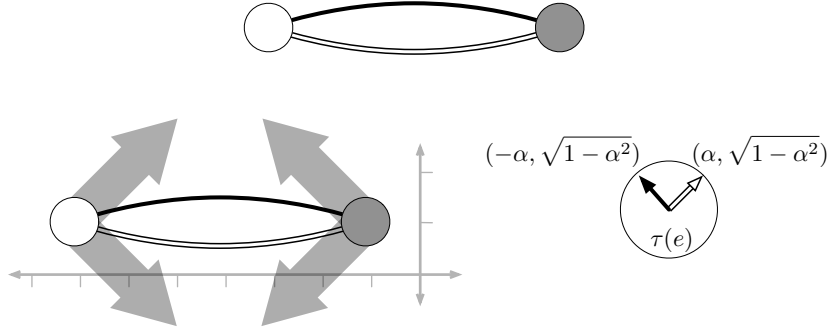


Figure 14: Even the fully connected network with 2 participants has $T_\alpha = \alpha$.

Here we show that one can have a fully connected network whose tension is proportional to α , matching the lower bound proved in Fact 15 upto a factor of 2. Similar ideas can be used to give a large network that asymptotically matches the bounds given by Fact 15.

Consider the network shown in Figure 14, which has just two participants, 1, 2. Suppose the first edge has type vector $\boldsymbol{\tau}(1) = (\alpha, \sqrt{1 - \alpha^2})$, and the second has type vector $\boldsymbol{\tau}(2) = (-\alpha, \sqrt{1 - \alpha^2})$, and both have weight 1. Then if we set $\mathbf{x}_1 = (1, 0)$ and $\mathbf{x}_2 = (-1, 0)$, we get that \mathbf{x} is a significant

shock with $\mathbf{x} = \bar{\mathbf{x}}$. Indeed, $\|\mathbf{x}_1\|, \|\mathbf{x}_2\| \leq 1$, and

$$\sum_{i=1,2} \sum_{u=1,2} |\langle \boldsymbol{\tau}(i), \mathbf{x}_u \rangle| = 4\alpha = \alpha w,$$

so $\mathbf{x} \in \mathcal{U}_\alpha$. But the tension of the shock is:

$$\begin{aligned} \mathbb{T}(\mathbf{x}) &= \left\| \sum_{i=1,2} \langle \boldsymbol{\tau}(i), \mathbf{x}_1 - \mathbf{x}_2 \rangle \boldsymbol{\tau}(i) \right\| + \left\| \sum_{i=1,2} \langle \boldsymbol{\tau}(i), \mathbf{x}_2 - \mathbf{x}_1 \rangle \boldsymbol{\tau}(i) \right\| \\ &= 2 \left\| \sum_{i=1,2} \langle \boldsymbol{\tau}(i), \mathbf{x}_1 - \mathbf{x}_2 \rangle \boldsymbol{\tau}(i) \right\| \\ &= 2 \left\| 2\alpha \cdot (\alpha, \sqrt{1-\alpha^2}) - 2\alpha \cdot (-\alpha, \sqrt{1-\alpha^2}) \right\| = 8\alpha^2, \end{aligned}$$

so we get $\mathbb{T}_\alpha \leq \frac{\mathbb{T}(\mathbf{x})}{2 \sum_{i=1,2} \sum_{u=1,2} |\langle \boldsymbol{\tau}(i), \mathbf{x}_u \rangle|} = \frac{8\alpha^2}{8\alpha} = \alpha$.