An Experimental Study of Network Formation with Limited Observation*

Michael Caldara
Amazon

Michael McBride†
University of California, Irvine

July 2015

Abstract

Many social and economic networks emerge among actors that only partially observe the network when forming network ties. We ask: what types of network architectures form when actors have limited observation, and does limited observation lead to less efficient structures? We report numerous results from a laboratory experiment that varies both network observation and the cost of forming links. In particular, we find that limited network observation does not inevitably lead to highly inefficient networks but instead might actually inhibit inefficient, farsighted, positional jockeying among actors.

JEL Classifications: C92, D83, D85
Keywords: networks, limited observation, coordination

*This research was supported by funding from Air Force Office of Scientific Research Award No. FA9550-10-1-0569 and Army Research Office Award No. W911NF-11-1-0332. For valuable comments, we thank participants at the November 2013 *Trends in Social Networks Research* Conference at Academia Sinica, the May 2014 ICES End of Year Conference at George Mason University, and the June 2014 International Economic Science Association Meetings at University of Hawaii - Manoa. We also thank seminar participants at the Institute for Mathematical Behavior Sciences at UC Irvine and the Economic Science Institute at Chapman University. Finally, we thank the Experimental Social Science Laboratory at UC Irvine for use of its facilities.

†Corresponding author: Department of Economics and Experimental Social Science Laboratory, 3151 Social Science Plaza, Irvine, CA, 92697-5100, mcbride@uci.edu, 949-824-7417, 949-824-2182 (fax).

1 Introduction

Many economic activities share two important features. First, these activities occur on underlying social networks that emerge via decentralized forming and severing of relationships. Such is true for exchanges between buyers and sellers (Kranton & Minehart 2001) and the transmission of valuable information about production techniques (Conley & Udry 2001) and job openings (Calvó-Armengol & Jackson 2004). Recognition of the importance of social networks in economic life has led to a large theoretical literature on the formation of networks in various settings,[1] and, more recently, a growing experimental literature.[2]

The second feature is that an actor in a network generally observes only a limited and local portion of her social network, a fact established from examination of field data.[3] For example, a seller might not know the identities of a prospective buyer's other potential seller, and an individual engaged in a job search will not know with precision the sources of job information that comes to her via her social network. In the former case, the seller might not fully leverage her market power due to her misperception about the buyer's other exchange opportunities, while in the latter case, the job seeker might expend costly resources developing ties with another person already in her social network, whose information about jobs is already being transmitted via that network, but whose presence in the network is not known. Only a small theoretical literature examines how the presence of limited observation affects network formation,[4] and there is no experimental literature on the topic.

We conduct the first experimental study of network formation under limited observation with a focus on the following questions. Does the presence of this limited observation lead to

[1] See e.g., Bala & Goyal (2000) for a model of network formation in which ties are formed unilaterally and Jackson & Watts (2002) for a model of network formation in which ties are formed by mutual consent.

[2] Examples of the experimental literature include Falk & Kosfeld (2012), Callander & Plott (2005), Goeree, Riedl & Ule (2009), Pantz (2006), Corbae & Duffy (2008) Mantovani, Kirchsteiger, Mauleon & Vannetelbosch (2013), Dogan, van Assen & Potters (2013), Rong & Houser (2013), van Leeuwen, Offerman & Schram (2013), and Carrillo & Gaduh (2012).

[3] See Laumann (1969), Friedkin (1983), Kumbasar, Romney & Batchelder (1994), Bondonio (1998), and Casciaro (1998).

[4] See McBride (2006*a*), McBride (2006*b*), McBride (2008), Francetich & Troyan (2010), and Song & van der Schaar (2015). There is a separate theoretical literature on learning in networks (e.g., Acemoglu, Dahleh & Lobel 2011).

less-efficient network structures than would form with full observation? If so, what form do those inefficiencies take, and how large are the inefficiencies? Answers to these questions will provide a better understanding of the effectiveness of social networks in fostering or inhibiting economic activity and provide insights into potential efficiency gains. For example, would an intervention that allows an actor to increase the observational range of her network improve the efficiency of the network?[5]

We choose as our setting the Connections Model introduced by Jackson & Wolinsky (1996). Actors decide with whom to form costly ties, with each tie providing direct and indirect access to benefits. Mutual consent is required for a tie to form, but a tie can be terminated unilaterally. This model does not represent a specific real-life setting but rather is a flexible model that captures features of many real-life social networks. It has also been subsequently studied in later work (Jackson & Watts 2002, Watts 2001, Pantz 2006, Calvó-Armengol & Ilkiliç 2009, Mantovani et al. 2013, Carrillo & Gaduh 2012), including in a limited observation context (McBride 2006*b*, Francetich & Troyan 2010).

Our experiment randomly matches subjects into 12-person groups for 15 rounds of interaction. At the end of each round, a participant receives one point for each other participant with whom she is directly or indirectly connected and pays cost c for each of her direct ties. We exogenously vary the level of observation—under *full observation* the whole network is observed in each round, and under *neighbor observation* the actor only observes her own ties and the ties of her network neighbors—and the cost of forming ties—either low or high. This design differs from prior network formation experiments by allowing for limited observation and by using larger groups than typical in network formation experiments.[6] Having larger groups is necessary to allow for networks to emerge with large enough diameters for limited observation to matter.

We rely on the prior theoretical and experimental literature on network formation to

[5] The business development tool RelSci (www.relsci.com) intends to do precisely this by helping individuals to discover potential clients and investors that they are indirectly connected to through mutual acquaintances.

[6] Very few studies have used larger groups. Exceptions include Gallo & Yan (2015) and Candelo, Forbes, Martin, McBride & Allison (2014).

form testable hypotheses. Jackson & Wolinsky (1996) define and apply the *Pairwise Stability* (PS) concept to identify stable networks in the Connections Model, and McBride (2006*b*) generalizes (PS) as *Conjectural Pairwise Stability* (CPS) to identify stable networks when actors have limited observation. Analyses that more directly consider formation dynamics vary in what is assumed about the forward-lookingness of actors, though all prior work assumes full observation. In Bala & Goyal (2000) and Jackson & Watts (2002), for example, actors form ties myopically by looking only for best responses to immediately past or current period play. More recent work (Herings, Mauleon & Vannetelbosch 2009, Grandjean, Mauleon & Vince 2011 and Mantovani et al. 2013) defines the *Pairwise Far-sighted Stability* (PFS) in which actors are forward-looking in anticipating the future network structure. The static analysis and the myopic-actor dynamic analysis yield clearer predictions, yet prior experimental evidence reveals that subjects in network formation settings exhibit far-sighted behavior (Mantovani et al. 2013, Carrillo & Gaduh 2012, van Leeuwen et al. 2013), a fact we account for in our hypotheses.

We report a number of findings, and here mention four of them. (1) The likelihood that a network converges depends on the interaction of treatment conditions: convergence is higher in the Low Cost-Neighbor Observation and High Cost-Full Observation conditions. Far-sighted jockeying for position, whereby a subject adds or removes a link that reduces today's payoff in the hope of achieving a higher payoff in later rounds, hinders convergence in the other conditions. (2) Converged networks are near efficient and more efficient than non-converged networks. (3) Redundant links (structural cycles) occur more often under Neighbor Observation than Full Observation, consistent with the notion that limited observation prevents subjects from observing redundant links that could be removed. (4) Though convergence to inefficient networks is more common with Neighbor Observation, the lower incidence of jockeying under Neighbor Observation led, surprisingly, to higher overall efficiency under Neighbor Observation than under Full Observation.

We thus find that limited observation does not inevitably lead to inefficient networks as might be naïvely expected. To the contrary, limited observation can actually help efficiency when it reduces jockeying among actors attempting to gain advantages associated with certain positions in the network. This reduction in jockeying keeps the network more connected than otherwise which, in our setting, outweighs the inefficiencies in the form of redundant ties that we might expect under limited observation. Whether or not limited observation is a help or hindrance to network formation in another setting will likely depend on the relative importance of under-connections versus over-connections. We also learn that far-sighted behavior may be a source of stability, instability, efficiency or inefficiency, all depending on the context. Inefficient jockeying may destabilize otherwise stable networks in the short run. Alternatively, when it promotes efficiency, the anticipation of future jockeying may stabilize otherwise unstable networks in the short run. In settings where stable networks are highly inefficient, far-sighted behavior can even lead to efficient instability.

2 Theory

2.1 The Connections Model

Consider a set of actors $N = \{1, ..., n\}$, each of whom is a *node* in graph g. Let $ij \in g$ denote that there is a tie between i and j in g, while $ij \notin g$ means there is no tie. We assume that the ties are symmetric, i.e., $ij \in g \Leftrightarrow ji \in g$. Let $g + ij$ denote a graph that results from adding tie ij to g holding all else fixed, while $g - ij$ denotes the removal of the ij tie from g. We say that there is a *path* between i and j in g if either $ij \in g$ or there exists m distinct players i_1, i_2, ..., i_m, such that $\{ii_1, i_1i_2, ..., i_mj\} \subset g$. Let $N_i(g) \subseteq N$ denote i's component, which is the set of all $j \in N$ with a (finite) path to i. Let $L_i(g) \subseteq N$ denote the set of i's ties, i.e., the set of all $j \in N$, $j \neq i$, such that $ij \in g$. A network is *connected* if $N_i(g) = N$, and we say that a network is *minimally connected* if there is exactly one single path between any two i and j. Being minimally connected implies no cycles.

Figure 1 provides examples of networks and relevant network characteristics: (a) the empty network which has no ties; (b) a network that is not empty but also not connected;

Figure 1: Network Examples

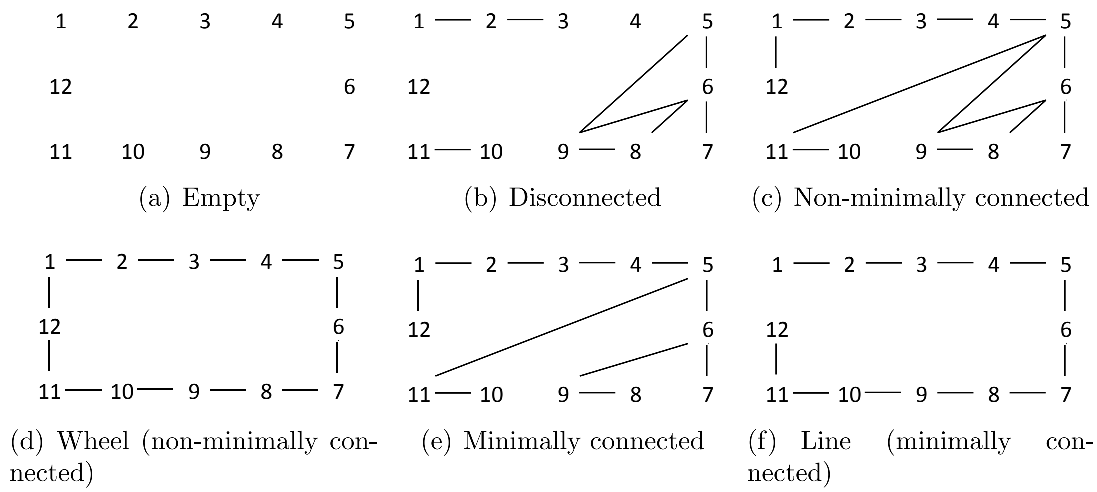

(a) Empty (b) Disconnected (c) Non-minimally connected

(d) Wheel (non-minimally connected) (e) Minimally connected (f) Line (minimally connected)

i.e., 4 and 12 are isolated; (c) a connected network that is not minimal, i.e., it has redundant ties that creates cycles of size three (5–6–9, 6–8–9) and four (5–6–8–9); (d) a wheel network that is connected; (e) a minimally connected network, created by removing the 5–9 and 6–8 ties from (c); and (f) another minimally connected network, called the line network.

The simplified version of the Connections Model examined here assumes each actor has payoff function

$$u_i\left(g\right) = \sum_{j \in N_i(g)} 1 - \sum_{j \in L_i(g)} c.$$

The value of being connected, either directly or indirectly, to another actor is 1, and the cost of forming a direct tie is c.

Our experiment uses $c \in \{0.7, 1.5\}$ and $n = 12$. It is straightforward to show that at each cost level, an efficient network, defined here as a network that generates the maximum sum of utilities, must be minimally connected, and any minimally connected network is efficient. The intuition comes from recognizing the positive externality in ties. With low cost, adding ties up to minimal connectedness adds consecutively more individual and social benefits than cost, while adding ties beyond minimal connectedness produces additional cost but no benefits. With high cost, minimally connected networks still produce the highest total net benefits (though, if the cost were sufficiently large or if n were sufficiently small, then the only efficient network is the empty network in which $L_i(g) = \varnothing$ for all $i \in N$).

2.2 Full Observation

Multiple stability concepts have been used to identify the types of network structures that may emerge in the Connections Model, two prominent ones being Pairwise Stability (PS) and Pairwise Farsighted Stability (PFS). The first examination of the Connections Model in Jackson & Wolinsky (1996) used the Pairwise Stability (PS) concept to identify stable networks. A graph g is *Pairwise Stable* (PS) if (i) for all $ij \in g$, it is true that $u_i(g) \geq u_i(g - ij)$ and $u_j(g) \geq u_j(g - ij)$ and (ii) for any $ij \neq g$, it is true that if $u_i(g + ij) > u_i(g)$ then $u_j(g) > u_j(g + ij)$. The first condition requires both parties to prefer keeping a tie to severing a tie, and the second implies that if one party strictly prefers a new tie then the other party to that tie must strictly prefer against that tie. Notice that the PS concept implicitly assumes that, when deciding to form or sever ties, each actor fully observes the entire network g, i.e., each actor perfectly knows the utility of $g + ij$ and $g - ij$.

It is straightforward to show that the set of PS networks is the set of minimally-connected networks with low cost $c = 0.7$, while the empty network is the unique PS network with high cost $c = 1.5$. With low cost, adding a tie to an actor not in your own component is individually rational, as is removing redundant ties. Only for minimally connected networks is it the case that no actor wants to add or remove ties. Under high cost, it is again true that removing redundant ties is individually rational, so any PS network must not have redundant ties. However, a minimal network that is not empty must necessarily have at least one stem, i.e., an actor has that exactly one tie (e.g., nodes 1 and 12 in Figure 1(f), and an actor directly connected to a stem is strictly better off removing the tie to that stem. The only network without redundant ties and stems is the empty network.

The Connections Model has been further examined in a dynamic framework by Watts (2001) and Jackson & Watts (2002), again implicitly assuming full observation. In each period, one (i, j)-pair is selected at random, and those actors decide myopically whether to keep the *status quo*, form a new tie, or sever an existing tie. As with the static PS concept, mutual consent is required to add a tie but ties can be removed unilaterally. Networks that this dynamic might enter but never leave are considered stable. Applying this dynamic in our setting yields a similar prediction as the PS concept: the system converges to a minimally connected network when the cost is low, and it converges to the empty network when the cost is high.

Other work considers network formation with far-sighted actors. One far-sighted concept used in the experimental literature on network formation is Pairwise Farsighted Stability (PFS) (Herings et al. 2009; Grandjean et al. 2011; Mantovani et al. 2013) which relies on the notion of a *far-sighted improving path*. As before, only one tie is changed at a time, but best responses are evaluated relative to the end network rather than the immediately adjacent network. A network g is PFS if all deviations from g to $g \pm ij$ are deterred by the threat of ending in some PFS network g' with no greater payoff for the deviator(s). The PFS concept refines the set of PS networks (Herings et al. 2009) although it is motivated by far-sighted rather than myopic actors. Here, the set of PFS networks is the same as the set of PS networks in both cost environments.[7] Far-sighted improving paths do exist between PS networks under the low cost (e.g., when an actor with multiple ties removes a tie and some other actor re-connects), but all such deviations could be credibly deterred by the reverse path, which would result in no improvement. The empty network is the unique PS, and thus PFS, network under high cost.

While PFS is a good descriptor of long-run group behavior in some settings (see, e.g., Carrillo & Gaduh 2012, Mantovani et al. 2013), we note two additional but relevant complexities in our environment that may generate different far-sighted behavior than implied by the prior theory. First, our paired inertia and fixed time horizon hinder the ability of the group to deter deviations. It is improbable that inertia will select the correct pairs to reverse

[7] Grandjean et al. (2011) and Mantovani et al. (2013) study environments in which some PS networks yield different aggregate payoffs. Thus certain network structures Pareto dominate the others. In our environment, all PS networks yield the same aggregate payoff.

any given deviation in the short-run, and thus these paths are not credible deterrents. An actor may still choose to deviate in the hope that it will lead to a network in which she obtains a higher individual payoff, which we refer to as *jockeying for position*.[8] Second, in our high cost environment, stability and efficiency are directly at odds. Far-sighted behavior could also lead the actors towards a recurrence class of (near) efficient networks and so far-sighted actors may still rationally achieve high levels of efficiency without ever converging to a PFS network.[9] Lastly, we note the main departure of our experiment from this previous theoretical work is that multiple ties can change at a time (see Section 3.3 for details). This aspect of our experimental design may introduce potential coordination problems, as situations could arise in which two distinct pairs collectively would like to add a single tie. However, this will only affect the path taken to the equilibrium networks rather than the set of equilibrium networks itself. Redundant links and missed connections that occur due to coordination failures will ultimately be remedied on the path to a PS network.

Bringing the theory and these issues together suggests the following predictions:

Claim 1 *Assume $n = 12$ and Full Observation.*

(a) With low cost $c = 0.7$, the process should converge to a minimally-connected network.

(b) With high cost $c = 1.5$ and myopic actors, the process should converge to the empty network.

(c) With high cost $c = 1.5$ and far-sighted actors, the process may or may not converge; it will either converge to the empty network or it will transition between several high efficiency networks.

[8] In a recent study, van Leeuwen et al. (2013) also find evidence of jockeying behavior in a setting where ties are formed unilaterally and investments in a public good determine the benefits.

[9] Due to the finite length of our experiment, not all improving paths will yield a positive expected payoff, especially in the later rounds. Thus we might observe some convergence to (near) efficient networks even with farsighted actors in the high cost setting.

Figure 2: Network Examples: Neighbor Observation

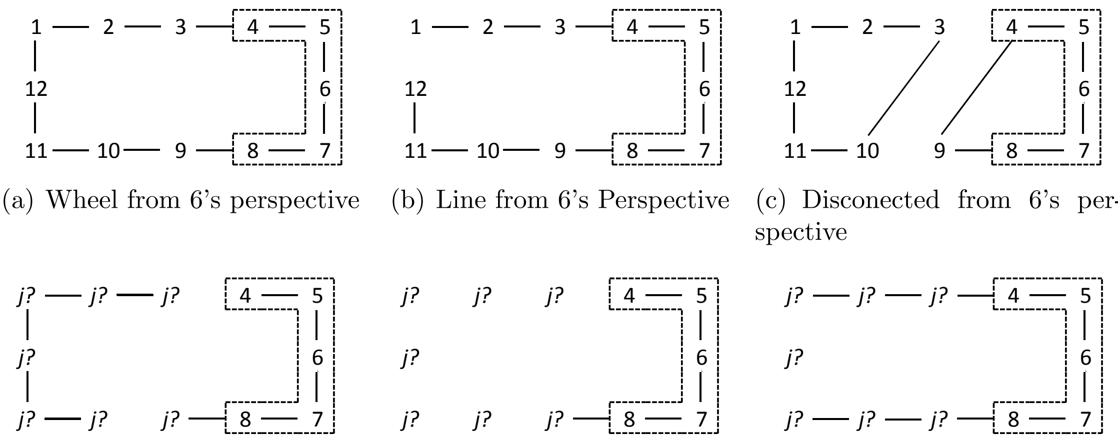

(a) Wheel from 6's perspective (b) Line from 6's Perspective (c) Disconected from 6's perspective

(d) Optimistic belief about (c) (e) Pessimistic belief about (c) (f) Uncertainty impedes connection

2.3 Neighbor Observation

There is less theoretical work on the Connections Model with Neighbor Observation. McBride (2006*b*) defines a Conjectural Pairwise Stable (CPS) network as a generalization of the PS concept to a limited observation setting. Each actor has a belief about the network, which takes the form of a probability distribution over possible networks. A network is CPS if each actor in that network believes she is better off keeping her existing ties and not forming any new ties and if her beliefs are not contradicted by what she observes in the network. The key difference between CPS and PS is that in CPS an actor may have incorrect beliefs about the network as long as she has no information to contradict those beliefs. With Neighbor Observation, this means that, in equilibrium, the actor must have correct beliefs about her own ties and her neighbors' ties; with common knowledge of the payoff structure and knowledge of her own payoff, each actor's beliefs must assign probability 0 to any network in which i's component differs in size from the actual network; but the actor's beliefs about the existence and location of other ties in the network may be incorrect.

Any PS network in which actors have correct beliefs is also CPS, but a few examples illustrate how some networks may be CPS but not PS. One particular case is to have a connected but not minimal network. Suppose the actual network is the wheel network in

Figure 2(a), but suppose actor 6 believes with probability 1 that the actual network is the line in Figure 2(b). This belief is consistent with what actor 6 observes and with her payoff, and she does not want to remove or add ties given those beliefs. If every other actor similarly believes that she is located at the middle of a line network, then this network is CPS. Notice that if an actor could see the cycle, then she would want to remove a tie. However, the cycle is outside of each actor's observational range, thereby allowing her beliefs to be incorrect about the network's exact structure. CPS networks can have smaller cycles, but the cycles must be of size 5 or more to be outside the observational range under Neighbor Observation.

Some disconnected networks may also be CPS networks. Figure 2(c) contains a disconnected network made of two 6-actor wheels. Redundant ties are again out of observational range. Yet, each actor knows from her payoff that her own component is of size 6. Whether or not an actor i wants to initiate a tie to some j not in her observational range depends on her belief about j's ties. If actor 6 believes that her network is that shown in Figure 2(d), with $j \in \{1, 2, 3, 9, 10, 11, 12\}$ equally likely to be connected to actor 8, then her expected net benefit from forming a tie is $\frac{6}{7}(6) + \frac{1}{7}(0) - c = \frac{36}{7} - c$, which is strictly greater than 0 for $c \in \{0.7, 1.5\}$. Thus, if she optimistically believes, without contradicting her observation, that unobserved parts of the network are highly connected, then she will form the tie. If instead she believes that the network is that shown in Figure 2(e), again with each unobserved actor equally likely to be connected with 8 and consistent with her observation, then her expected net benefit from forming a tie is $\frac{6}{7}(1) + \frac{1}{7}(0) - c = \frac{6}{7} - c$. So, if she pessimistically believes that unobserved parts of the network are highly disconnected, then she will want to form a tie if the cost is low $c = 0.7$ but not if the cost is high $c = 1.5$. Even when the size of all components is known, as in Figure 2(f), uncertainty about the position of each actor in the network may inhibit connection. Actor 6's expected net benefit from linking to j with $j \in \{1, 2, 3, 9, 10, 11, 12\}$ equally likely to be isolated is $\frac{1}{7}(1) + \frac{6}{7}(0) - c = \frac{1}{7} - c$ regardless of her level of optimism or pessimism. She will not want to form this connection in either cost setting, even though she would want to form this connection with low cost $c = 0.7$ if she knew the isolated actor's identity with certainty.

Despite the potential for many networks to be CPS, CPS networks share several commonalities that limit the scope of what is possible. First, in both cost environments, cycles of size less than 5 are not CPS. Each actor in such a cycle can see the whole cycle and thus will want to remove a tie. Second, with low cost $c = 0.7$, the network must have a component of at least size 8 before even the most pessimistic of actors would not benefit from connecting to an unobserved actor. For example, when all other actors are isolated, the expected benefit from connecting to an unobserved actor for the center actor in a line of size 7 is $\frac{5}{7}(1) + \frac{2}{7}(0) > 0.7$ and for the center actor in a line of size 8 is $\frac{4}{7}(1) + \frac{3}{7}(0) < 0.7$. Third, with high cost $c = 1.5$, any network with a stem is not CPS because the actor tied to this stem would prefer to remove this tie. However, connected networks with large cycles and no stems such as the wheel are CPS with high cost. An additional restriction on beliefs, the *rationalizability criterion*, which requires that all actors in the conjectured network are best responding to this network, further limits the scope of possible CPS networks. For example, for high cost, the line network shown in Figure 2(b) is not rationalizable under Neighbor Observation because actors 2 and 11 would benefit from removing the ties 1-2 and 11-12, respectively, thus implying that actor 6 cannot rationalize the network in Figure 2(b).

The CPS concept is static like the PS concept, and no theory to date has examined the Connections Model with limited observation in a dynamic framework like our experiment.[10] In order to adapt the findings of Watts (2001) and Jackson & Watts (2002) to the case of limited observation, we must specify the initial beliefs of all actors and the process by which these beliefs are updated as the network evolves. Here we consider two simple cases, at opposite extremes of the belief spectrum, to help pin down the set of possible CPS networks to which the system might converge. In both cases, we assume that beliefs satisfy *anonymity*,

[10]Song & van der Schaar (2015) consider dynamic formation in a different limited observation setting. The assume that upon being paired in a period, i observes j's component (not the entire network) and vice versa before deciding whether to form or sever a tie. Our information setting does not allow actors to have such information. A larger difference, however, is that, similar to McBride (2008), they allow actors to be of different type. This incomplete information setting introduces additional complexties not present in our setting.

stationarity, and *memorylessness*, in addition to the typical CPS requirement that beliefs do not contradict the observed information. Anonymity requires that all unobserved actors are assigned equal likelihood of being in each unobserved position in the conjectured network. Thus, decision makers only need to consider the structure of the conjectured network rather than all isomorphisms. Stationarity requires that beliefs be fixed unless the observed portion of the network changes. Thus decision makers do not attempt to forecast the dynamic path by which the network will evolve. Memorylessness requires an actor's belief about the existing network be formed independently of information observed in prior periods. Thus, once an actor leaves the decision maker's observational range, all information about this actor's position in the network is forgotten. These assumptions are admittedly strict, and we do not expect them to hold in all cases. However, they provide a simple and natural starting point for prediction and are consistent with extreme forms of myopia.

First consider the case of optimistic beliefs, e.g., Figure 2(d), in which actors assume that there are at most 2 disconnected components, both of known size. With low cost, actors will begin by connecting and continue connecting to unobserved actors until the network reaches at least size 10 ($\frac{3}{7}(3) + \frac{4}{7}(0) > 0.7$, but $\frac{2}{7}(2) + \frac{4}{7}(0) < 0.7$). Networks of size 10-12 are all possible depending on the order in which (i, j)-pairs are selected. Cycles of size 5 or greater are possible as (i, j)-pairs in the same component who cannot see each other may be randomly selected before the network reaches size 10+. With high cost, actors will begin by connecting and continue connecting to unobserved actors until the network reaches at least size 9 ($\frac{4}{7}(4) + \frac{4}{7}(0) > 1.5$, but $\frac{3}{7}(3) + \frac{4}{7}(0) < 1.5$). Networks of size 9-12 are possible and, as before, there may be cycles of size 5 or greater. However, whenever an actor is paired with a known stem she will remove the tie to that stem, and so the network would never converge in this case.

Now consider the case of pessimistic beliefs, e.g., Figure 2(e), in which actors assume all individuals outside of the component are isolated. With low cost, actors will begin by connecting and continue connecting to unobserved actors until the network reaches size 8

($\frac{5}{7}(1) + \frac{2}{7}(0) > 0.7$, but $\frac{4}{7}(1) + \frac{3}{7}(0) < 0.7$). Networks of size 8-12 are possible and there still may be cycles of size 5 or greater. With high cost, actors will never connect, and the network will remain empty. If instead, actors are assumed to be far-sighted, they may take myopically detrimental actions as part of a far-sighted improving path. Such actions might include adding ties to unobserved individuals beyond what is myopically rational in an effort to achieve a maximal component (accidental redundant ties could be removed later) or adding ties with pessimistic beliefs and high cost as part of a path to a more efficient network (the dynamic would still not converge in this case).

As before, allowing multiple ties to change at a time introduces potential coordination problems. This may lead to visible cycles or missed connections, but we expect these known inefficiencies to be remedied by the actors at the next opportunity. Coordination failures may also lead to unobserved cycles that are never removed, but most of these cycles could occur even when one tie changes at a time.

We summarize the above logic with this prediction:

Claim 2 *Assume $n = 12$ and Neighbor Observation.*

(a) With low cost $c = 0.7$, the process should converge to a network with no more than two components (one at least size 8) and no visible cycles. All CPS networks can be reached.

(b) With high cost $c = 1.5$, the process may or may not converge, but if it does, it will converge to the empty network.

2.4 Jockeying for Position and Positional Risk

Individual connection decisions become more complicated when actors jockey for position. Consider, for example, the stem position in a connected network (e.g., actor 1 is a stem in Figure 2(b)). The stem position is the best possible position in a component for an actor because she receives the maximal benefit at the minimum cost. This is no longer necessarily the case when actors jockey for position in the short-run. Recall that jockeying for position involves removing a tie with the expectation that some other actor will re-connect. The stem position is particularly vulnerable to this behavior because of the low cost of removal. Thus, with random paired inertia, the stem position is less attractive to a risk averse actor due to the small chance that her partner will remove this tie in subsequent periods. We expect risk averse actors to be more likely to avoid this position, when possible.

3 Experiment Design

3.1 Basics

We conducted eight experiment sessions at a large public university with a total of 240 subjects.[11] Our treatment variables are the cost of forming ties ($c = 0.7$ or $c = 1.5$) and the level of observation (neighbor or full). The treatment was chosen exogenously prior to the beginning of each session and remained fixed throughout the session. Two sessions were conducted for each of the four treatment conditions in our between-subjects 2x2 factorial design. Subjects were students who had previously registered to be in the laboratory subject pool, and were recruited at random via E-mail. Subjects were allowed to participate in at most one session. Each session had either 24 or 36 subjects who were randomly matched into groups of twelve. The sessions consisted of ten matches[12] and subjects were randomly assigned to a new group at the beginning of each match. Each match consisted of 15 periods and the groups were held constant across periods. The experiment was programmed and conducted using the z-Tree software package (Fischbacher 2007). Screenshots of the experiment can be seen in Figure 3. At the end of the session, each subject's risk preference

[11] We conducted a two-group, 24-subject pilot experiment prior to the first session to inform our choice of control variables. The two main conclusions from the pilot were that 20 periods per match was unnecessarily long (we reduced it to 15) and that assessing one relationship at a time was more manageable for the subjects than assessing 5 relationships at a time (we listened to this feedback from the questionnaire and restricted the decision to assessing one relationship at a time). We do not include the pilot data here due to the differences in the control variables and the absence of random rematching.

[12] One session of full observation with low cost was stopped a match early due to the cumulative effects of a slight, unanticipated, software lag at the end of each round that prevented all matches from being completed in the allotted time. We reduced the number of groups per full observation treatment to reduce the lag in future sessions. This adjustment fixed the problem.

Table 1: Session Information

	# of Subjects	% Male	Major						Avg Earn ($)
			% Bus/ Econ	% Psy/ Cog Sci	% Other Soc Sci	% Eng/ Phys Sci	% Bio/ Life Sci	% Hum/ Undec	
All	240	42.1	13.3	7.9	15.0	27.9	25.8	10.0	33.09
NO – LC	60	36.7	10.0	10.0	6.7	33.3	28.3	11.7	37.00
NO – HC	72	38.9	13.9	6.9	18.1	18.1	33.3	9.7	31.16
FO – LC	60	48.3	16.7	5.0	21.7	25.0	21.7	10.0	32.41
FO – HC	48	45.8	12.5	10.4	12.5	39.6	16.7	8.3	31.94

Demographic information obtained from questionnaire. Average earnings include $7 show-up payment.

was elicited using the Eckel & Grossman (2008) method and the subjects were asked to complete a short questionnaire.[13]

Subjects received $1.00 for every 50 units of experimental currency earned during the session and were given a lump sum payment ($2.00 for low cost and $4.50 for high cost) at the beginning of the session to prevent bankruptcy. Additionally, subjects received a $7.00 show-up payment and could earn between $0.00 and $14.00 in the risk preference elicitation, depending on that subject's choice of lottery and the outcome of the lottery. In total, subjects took home an average of approximately $33.00 for 2 hours of participation (Table 1).

3.2 A Single Match

At the beginning of a match, each participant is randomly assigned a group and a color and is shown an image of the initial (empty) network. The image (see Figure 3) contains 12 colored circles or "nodes;" one node corresponding to each participant in the group. The color represent's the subject's experimental identity and is used instead of assigning each

[13] A participant is given the choice between five lotteries and this choice maps to a range of possible risk coefficients under the assumption of constant relative risk aversion (the range of possible ρ is reported in parenthesis). The lotteries are: $4 with probability 100% ($2.63 < \rho$), $3 with probability 50% or $7 with probability 50% ($0.90 < \rho < 2.63$), $2 with probability 50% or $10 with probability 50% ($0.53 < \rho < 0.90$), $1 with probability 50% or $13 with probability 50% ($0.00 < \rho < 0.53$), and $0 with probability 50% or $14 with probability 50% ($\rho < 0.00$). The order of the lotteries was randomly permuted for each participant to avoid any order effect.

Figure 3: Software Screenshots

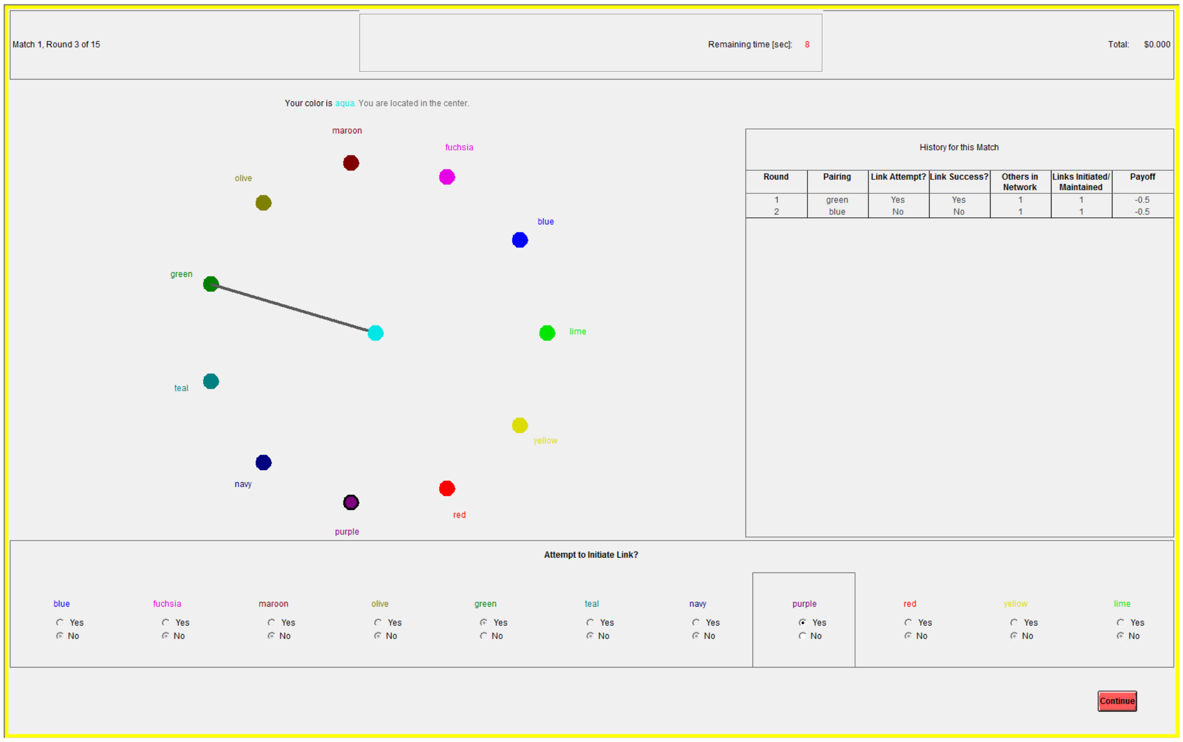

(a) Decision screen

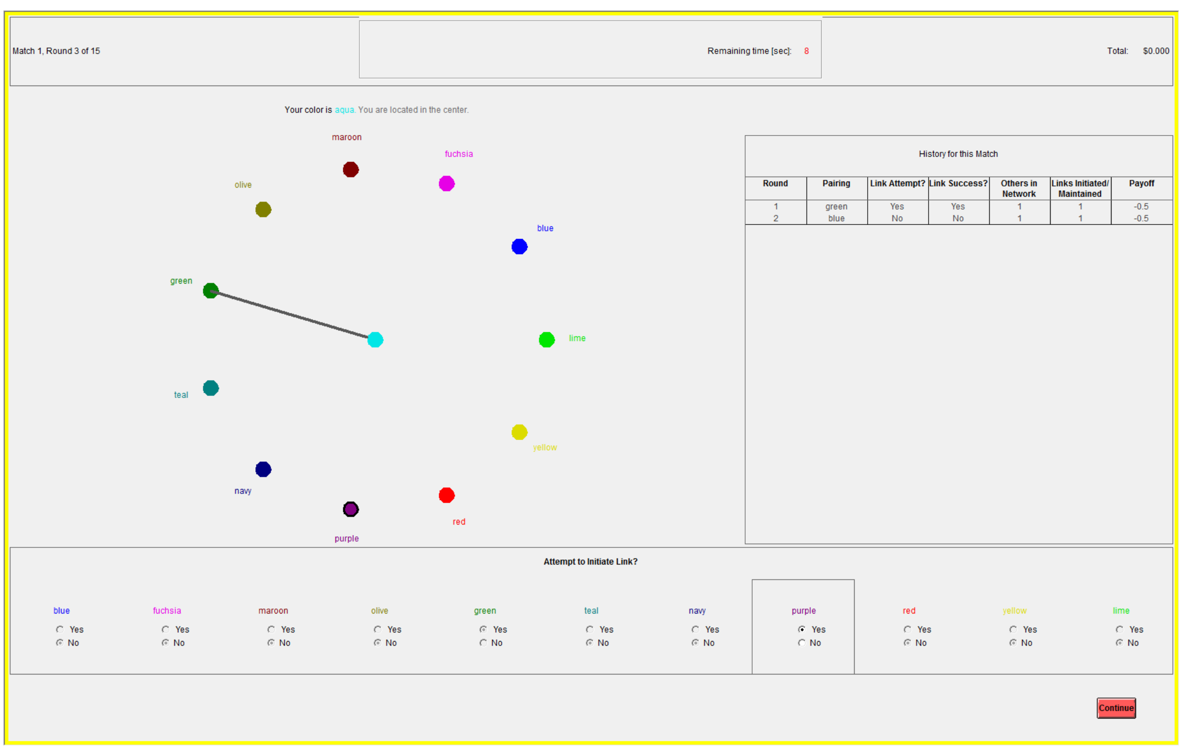

(b) Results screen

subject an ID number. A participant's own node is located at the center of her image surrounded by a ring consisting of the other 11 nodes. The order in which the nodes are patterned remains fixed throughout the match, but this order is determined randomly for each participant at the start of the match to avoid any position specific framing effect. Ties within a participant's observational range are depicted by a dark grey line connecting two nodes, whereas ties outside of a participant's observational range are not depicted. A match consists of 15 rounds and the image is updated at the end of every round to display all visible changes in the network. Participants are paid for every round, rather than the end network to mimic real-life network formation. The participants are shown the round number, the match number, their most recent earnings, their cumulative earnings, and a match specific history at all times throughout the match.

3.3 A Single Round

At the beginning of a round, each participant is randomly paired with one other participant in her group. These pairings are restricted to be non-overlapping and so there are six distinct pairs in each round.[14] Participants are first shown the Decision Screen (Figure 3 (a)) and have 15 seconds to decide whether or not to initiate (or maintain) a tie to their partner. The participant's partner is highlighted on the network image to aid in the decision-making process. A participant indicates her choice by selecting yes or no in the decision prompt and then pressing a submit button. The default option, which is initially selected and which is submitted on this participant's behalf if time runs out, is set so as to maintain the status quo. In the vast majority of cases (98%), the submit button was pressed before time ran out and so this design element did not substantively affect our results. After all decisions are submitted (or after time runs out), all participants are then shown the result screen (Figure

[14]The main difference between our intertia process and the inertia process of Jackson & Watts (2002) is that we allow six pairs to evaluate their relationship at a time rather than one pair at a time. Allowing six pairs to evaluate their relationship at a time substantially increases the speed of convergence, allowing us to collect more data in the allotted time. This decision comes at the cost of introducing potential coordination problems. However, this will only affect the path taken to the equilibrium networks rather than the set of equilibrium networks itself.

3 (b)), in which they learn the outcome of their decision, their payoff for the round, and are shown the visible portion of the resultant network. Participant i's payoff for network g is given by

$$u_i(g) = 1 \times (\#\text{ of direct ties}) - c \times (\#\text{ of direct ties} + \#\text{ of unsuccessful tie attempts})\,.$$

A participant pays the cost of any tie attempt she makes, whether or not it is successful, as well as the cost of any direct tie she formed in a prior round, whether or not she was able to change that tie in this round.

3.4 Treatment Variables

One first treatment variable is the level of observation. The subjects experience *Full Observation* or *Neighbor Observation* as described depending on the session. Full Observation has been used in prior experimental analyses, albeit with different control variables (e.g., Falk & Kosfeld 2012; Callander & Plott 2005; Berninghaus, Ehrhart & Ott 2006; Goeree et al. 2009; Pantz 2006; Mantovani et al. 2013; and Carrillo & Gaduh 2012). Under Neighbor Observation, each subject sees only her direct ties and the ties of her neighbors. Ties beyond this distance, including ties in other components, are not displayed. Individuals that are far away in one's own component are indistinguishable from individuals in a separate component and no information is provided as to the size of this component. Each subjects does, however, learn her own payoff at the end of a round and thus can infer how many individuals are in her own component.

Our other treatment variable is the cost of forming ties. The participants either experience *low cost* or *high cost* depending on the session, defined as $c = 0.7$ and $c = 1.5$, respectively.

3.5 Hypotheses

Given these parameters, we formulate the following testable hypotheses based upon the predictions from Section 2. The first two concern convergence, the next two concern efficiency, and the last three consider positional

1. Holding the observation level fixed, the frequency of convergence is higher under low cost than under high cost.

2. Holding the cost fixed, the frequency of convergence is higher under neighbor observation than under full observation.

3. Holding the observation level fixed, the average efficiency is higher under low cost than under high cost.

4. Holding the cost fixed, the frequency of cycles is higher under neighbor observation than under full observation.

5. Holding the cost fixed, stems are dropped more frequently under full observation than under neighbor observation.

6. Holding the observation level fixed, stems are dropped more frequently under high cost than under low cost.

7. A higher degree of risk aversion is associated with a lower likelihood of being in a stem position.

4 Results

Our unit of observation for hypothesis tests is a single match. When the relevant data is generated by period, we compute match-level averages. Because we match the same 24-36 participants repeatedly into random groups of 12 and provide regular feedback, our data could feasibly depend on the participants, group, and match. For all participant-level hypotheses, we correct the standard-errors for two-way clustering at the participant and group level (Cameron, Gelbach & Miller 2011, Thompson 2011; compared to other approaches in Petersen 2009) and include indicator variables to control for the match. For all group-level hypotheses, we assume in the main discussion that the observations do not depend on the participants or match. We then relax this assumption in Table A.1 and find that the p-values are practically identical, if not smaller. This is not surprising in our setting, because the large group size, random paired inertia, and required mutual consent to form ties limit the idiosyncratic effect of individual participants on group outcomes. Furthermore, while we do find significant trends in some group outcomes across matches (see Tables A.2 and A.3), these trends do not generally depend on the treatment. In the one case where the trend does depend on the treatment, our standard errors are biased upward, if anything.

For each match, we record whether the network converged, the efficiency of the convergent networks, and the end of match efficiency of the non-convergent networks. We also record network characteristics such as component size, whether there were cycles and the cycle size. As was first noted by **?**, measuring convergence in network formation experiments involves "trading off empirical certainty with experimental constraints" such as boredom and end-of-match effects (p. 1478). A stricter measure has the advantage of reducing the occurrence of false positives, where a non-convergent network is categorized as convergent, but it comes at the cost of increasing the likelihood of false negatives, where a convergent network is categorized as non-convergent due to non-strategic changes that may occur after a time of stationarity. For this reason, we consider a network to have converged if it remains unchanged for 5 consecutive rounds. This measure is relatively strict as 5 rounds represents 1/3 of the total match and requires between 30 and 60 independent decisions to not change the network. We compute the efficiency of a network by dividing the aggregate group payoff by the maximum possible group payoff (116.6 with low cost and 99 with high cost). This represents a measure of both absolute and relative efficiency as the aggregate group payoff of the empty network is 0. To allow for closer comparison to the convergent networks, we consider the efficiency of a non-convergent network to be its average efficiency over the last 5 rounds of the match.

Consider an illustrative example from the Neighbor Observation - High Cost condition

Figure 4: Example of Dynamic Network Formation by Period

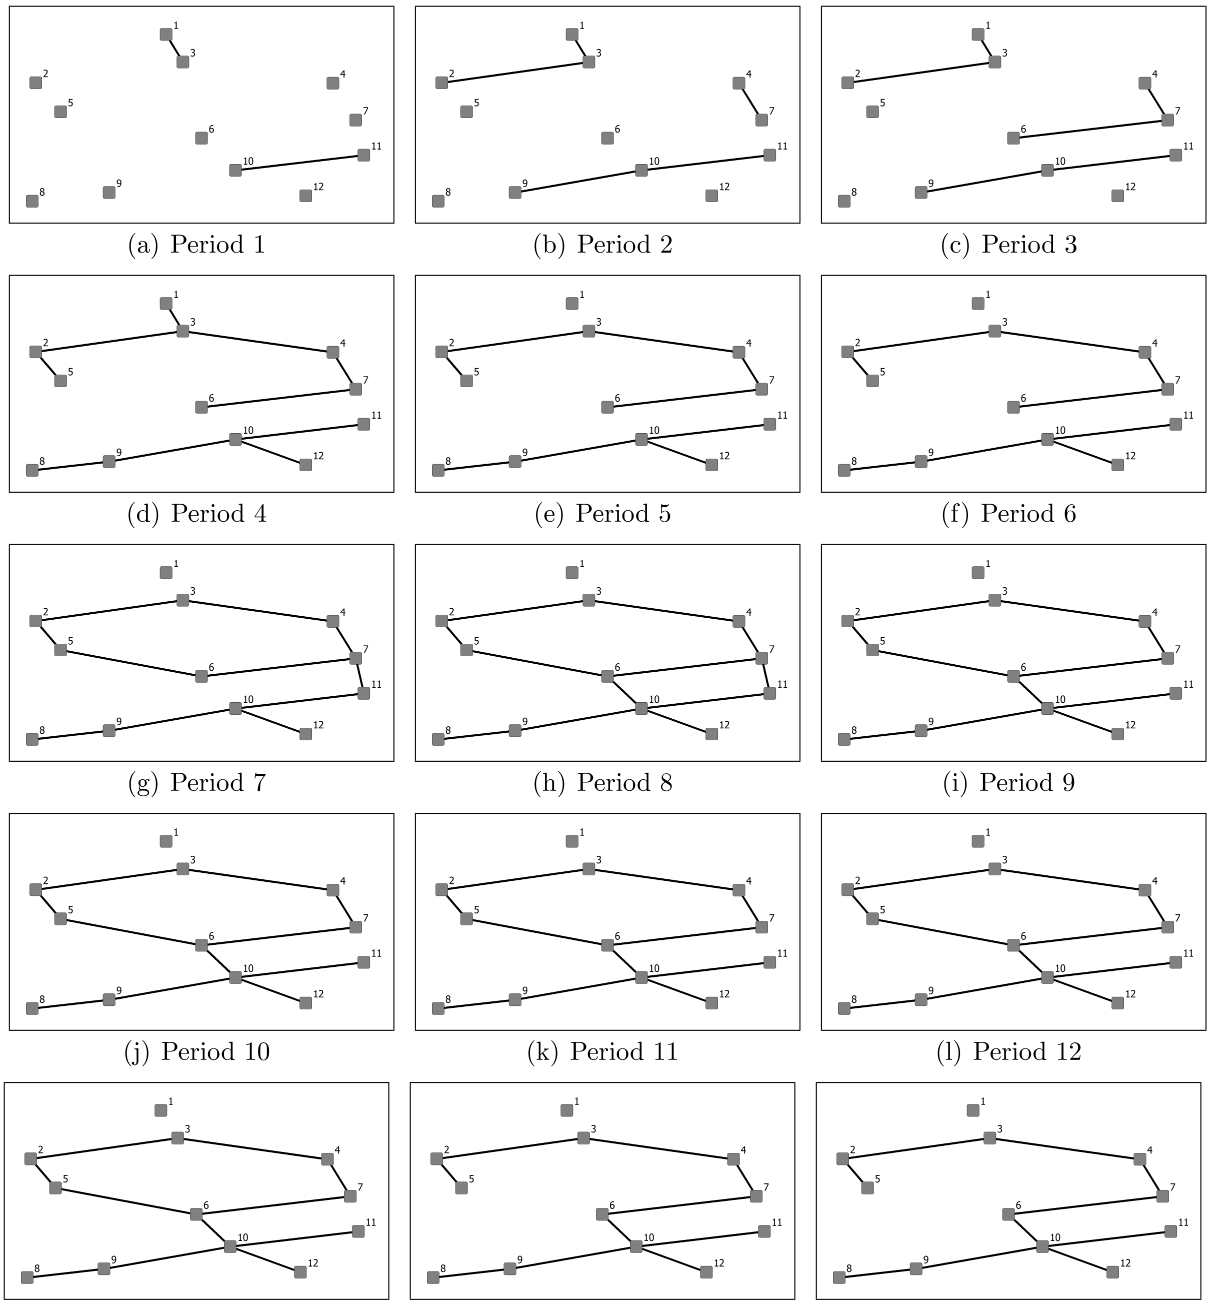

(a) Period 1 (b) Period 2 (c) Period 3 (d) Period 4 (e) Period 5 (f) Period 6 (g) Period 7 (h) Period 8 (i) Period 9 (j) Period 10 (k) Period 11 (l) Period 12 (m) Period 13 (n) Period 14 (o) Period 15

shown in Figure 4. The grey squares represent individual participants and the black lines represent ties between two participants. Ties are slowly added in periods 1–4. In period 5, participant 3 removes known stem 1. In period 7, cycle 2–3–4–5–6–7 forms, unbeknownst to the participants. In period 8, participants 6 and 10 add a tie and the visible cycle 6–7–

Table 2: Frequency of Convergence by Treatment

	Neighbor Observation	Full Observation
Low Cost	50.0% (25/50)	42.6% (20/47)
High Cost	23.3% (14/60)	60.0% (24/40)

10–11 forms. Participants 6 and 10 did not know in advance that they were in the same component. In period 9, participants 7 and 11 remove a visibly redundant tie, dissolving cycle 6–7–10–11. The network is unchanged from periods 9–13 and thus coded as having converged in period 9. It has components of size 11 and 1, one non-visible cycle of size 6, and an efficiency of 77.8%. The network is CPS. In period 14, participant 6 removes a tie to participant 5, dissolving non-visible cycle 2–3–4–5–6–7 in the process. Participant 6 did not know at the time whether this action would remove 0, 2, 3 or 4 individuals from her component. Depending on her belief about the network structure this may or may not have been jockeying for position. The network is unchanged in period 15.

We present our results, ordered by hypothesis, below. These results are summarized in Table A.1.

4.1 Convergence

In total, 83 out of 197 networks converged. Our convergence results, broken down by treatment, are shown in Table 2. Figure 5 shows the median number of network changes by period for each treatment. The typical network converges in period 7 for Neighbor Observation - Low Cost, in period 10 for Full Observation - Low Cost, and in period 5 for Full Observation - High Cost. The typical network for Neighbor Observation - Full Cost does not converge.

To test for a match-level trend in the frequency of convergence, we estimate a pooled logistic regression of the form

$$\Pr(\text{converge}_i = 1) = \text{logit}^{-1}(\mathbf{T}_i\vec{\alpha} + \beta\text{match}_i + \text{match}_i\mathbf{T}_i\vec{\gamma}),$$

where $\mathbf{T}_i$ is a vector of treatment dummy variables and match_i is a linear trend. We then

Figure 5: Median Number of Network Changes by Period

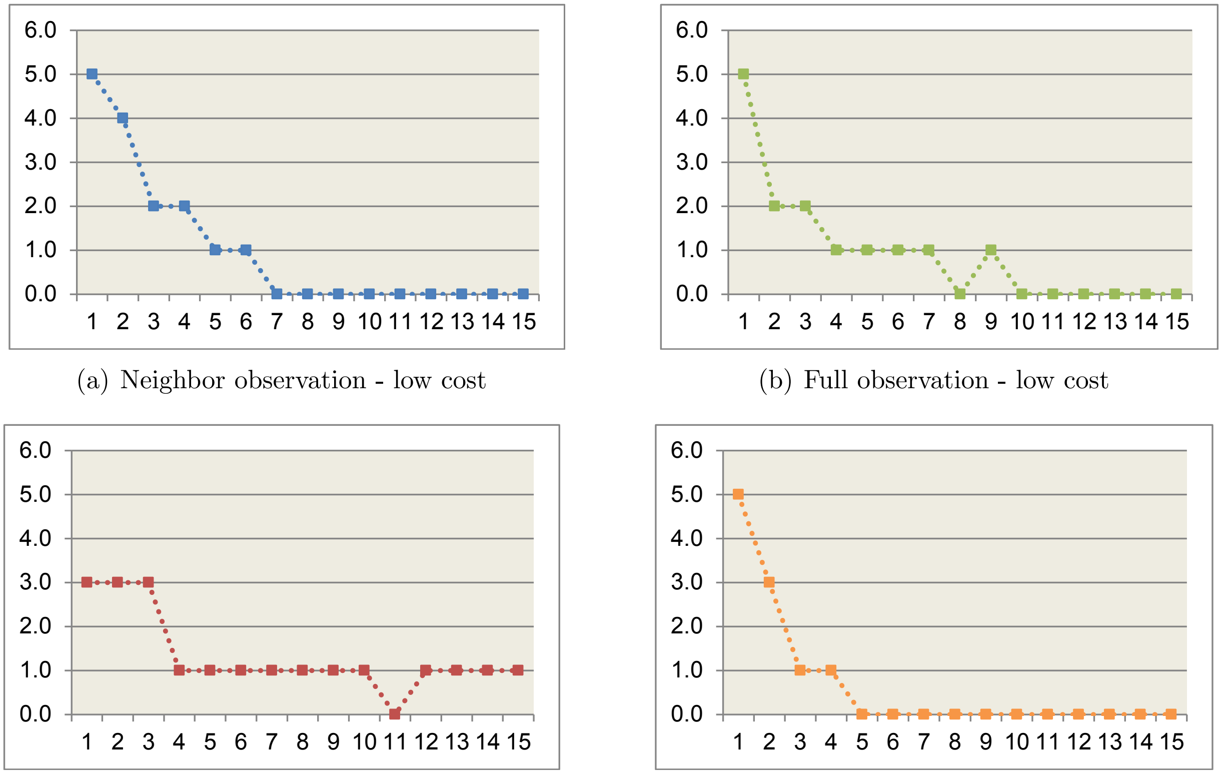

(a) Neighbor observation - low cost

(b) Full observation - low cost

(c) Neighbor observation - high cost

(d) Full observation - high cost

compute the average marginal effect for each treatment. On average there was a 5.5% increase in the frequency of convergence between matches in the Neighbor Observation - Low Cost treatment, a 3.2% increase in the Neighbor Observation - High Cost treatment, a 2.8% increase in the Full Observation - Low Cost treatment, and a 5.0% increase in the Full Observation - High Cost treatment. We cannot reject the null hypothesis that these trends are jointly equal (p-value 0.760). Overall, there was a 4.0% increase in the frequency of convergence between matches, significant at the 0.01 level.

We test the null hypothesis that the frequency of convergence under low cost is less than or equal to the frequency of convergence under high cost using Barnard's Exact Test. We reject this hypothesis at the 0.01 level for neighbor observation but cannot reject this hypothesis for full observation (p-value 0.938). For neighbor observation, networks were significantly more likely to converge with low cost (50.0%) than with high cost (23.3%) but for full observation

Table 3: Median Efficiency by Treatment

	Neighbor Observation	Full Observation
Low Cost	93.2% (50)	89.5% (38*)
C | NC	97.6% (25) | 88.4% (25)	99.4% (16) | 84.4% (22)
High Cost	80.4% (60)	75.8% (40)
C | NC	80.8% (14) | 78.0% (46)	88.9% (24) | 55.3% (16)

*Excludes 9 networks containing a subject who never connected.

there was no statistically significant difference in convergence. We interpret this as weakly confirming Hypothesis 1. The existence of (near) efficient (C)PS networks under low cost does seem to increase the likelihood of convergence, but this is clearly not capturing the full picture. We do not find the expected effect for full observation and, if anything, the evidence seems to point towards *less* convergence with low cost than high cost (42.6% vs. 60.0%).

We test the null hypothesis that the frequency of convergence under neighbor observation is less than or equal to the frequency of convergence under full observation using Barnard's Exact Test. We cannot reject this hypothesis under low cost (p-value 0.279) or high cost (p-value > 0.999). For high cost, networks were significantly less likely to converge with neighbor observation (23.3%) than with full observation (60.0%) and for low cost there was no statistically significant difference in convergence (50.0% vs. 42.6%). We interpret this as rejecting Hypothesis 2. The existence of CPS networks that are not PS networks is not sufficient to ensure a higher frequency of convergence under neighbor observation.

4.2 Efficiency

We record the final efficiency of each network pairing, separating converged networks (C) and non-converged (NC), we consider the final efficiency of a non-converged network to be its average efficiency over the last 5 periods. Our efficiency results, broken down by treatment, are shown in Table 3 (sample size in parentheses).

Our primary hypotheses do not explicitly differentiate between convergent and non-convergent networks, but it is clear from Table 3 that convergent networks are more efficient. We reject the null hypothesis of equal medians (Mann-Whitney U Test) at the 0.05

Figure 6: Median Efficiency by Period

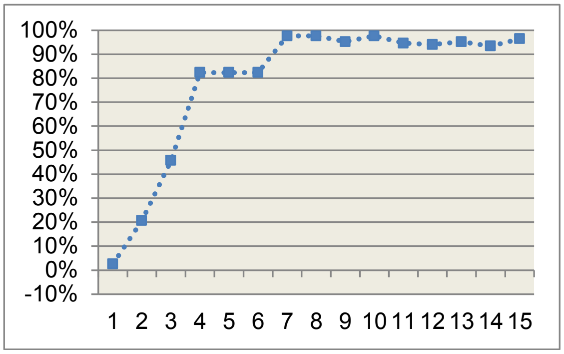

(a) Neighbor observation - low cost

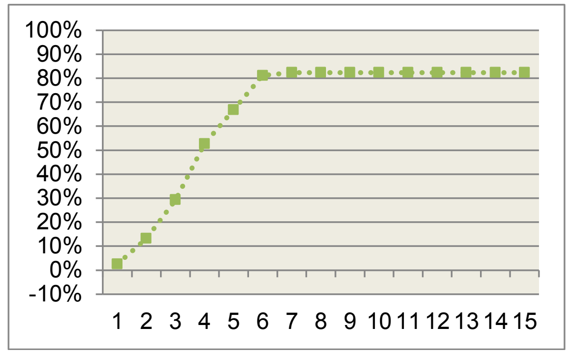

(b) Full observation - low cost

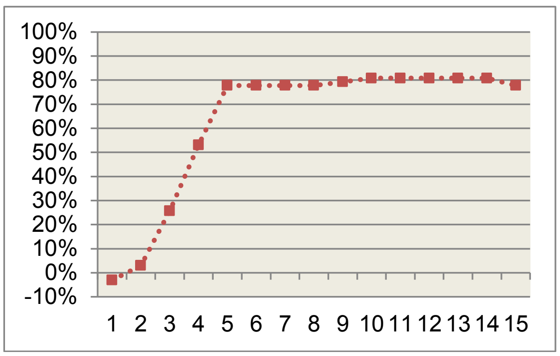

(c) Neighbor observation - high cost

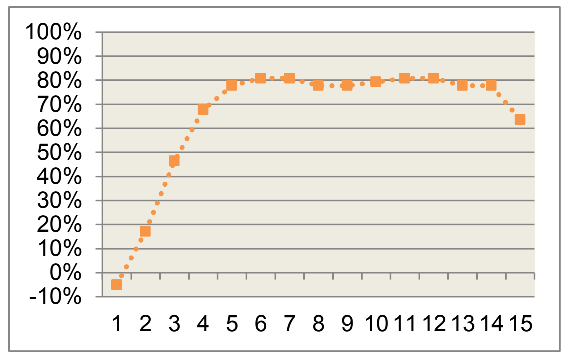

(d) Full observation - high cost

level for Neighbor Observation - High Cost and at the 0.01 level for the other three treatments. When we instead plot the median efficiency by period, as shown in Figure 6, we see that efficiency for the typical network increases over time and then eventually plateaus, with a slight efficiency drop-off at the end for the high cost treatments.

To test for a match level trend in the final efficiency of the network, we estimate a pooled regression of the form

$$\text{efficiency}_i = \mathbf{T}_i\vec{\alpha} + \beta\text{match}_i + \text{match}_i\mathbf{T}_i\vec{\gamma},$$

where $\mathbf{T}_i$ is a vector of treatment dummy variables and match_i is a linear trend. We then compute the average marginal effect for each treatment. On average there was a 0.6% decrease in efficiency between matches in the Neighbor Observation - Low Cost treatment, a 0.7% decrease in the Neighbor Observation - High Cost treatment, a 0.9% increase in the Full Observation - Low Cost treatment, and a 3.7% increase in the Full Observation - High

Table 4: Frequency of Cycles by Treatment

	Neighbor Observation	Full Observation
Low Cost	72.0% (50)	40.4% (47)
C | NC	60.0% (25) | 84.0% (25)	30.0% (20) | 48.1% (27)
High Cost	61.7% (60)	25.0% (40)
C | NC	57.1% (14) | 63.0% (46)	16.7% (24) | 37.5% (16)

Cost treatment. We reject the null hypothesis that these trends are jointly equal at the 0.01 level, the trend for Full Observation - High Cost is significantly different from the other treatments and zero. The overall trend for the remaining three treatments (0.2% decrease) is not significantly different from zero (p-value 0.537).

We test the null hypothesis that the efficiency under low cost is less than or equal to the efficiency under high cost using the Mann-Whitney U Test. We reject this hypothesis at the 0.01 level for both neighbor observation and full observation.[15] For both observation levels, median efficiency is greater with low cost than with high cost (93.2% vs. 80.4% for Neighbor Observation and 89.5% vs 75.8% for full observation). We interpret this as confirming Hypothesis 3. The stronger incentives to maintain connectivity with low cost do seem to yield higher average efficiencies.

In total, 33/83 convergent networks and 69/114 non-convergent networks had at least one cycle, and these are broken down by treatment condition in Table 4. To test for a match level trend in the frequency of cycles, we estimate a pooled logistic regression of the form

$$\Pr(\text{cycle}_i = 1) = \text{logit}^{-1}(\mathbf{T}_i\vec{\alpha} + \beta\text{match}_i + \text{match}_i\mathbf{T}_i\vec{\gamma}),$$

where $\mathbf{T}_i$ is a vector of treatment dummy variables and match_i is a linear trend. We then compute the average marginal effect for each treatment. On average there was a 5.6% decrease in the frequency of cycles between matches in the Neighbor Observation - Low Cost treatment, a 8.2% decrease in the Neighbor Observation - High Cost treatment, a 8.6% decrease in the Full Observation - Low Cost treatment, and a 4.2% decrease in the Full

[15] Controlling for the different match level trends under full observation and conducting a Wald test does not change the significance level. The p-value drops from 0.008 to < 0.001.

Table 5: Frequency of Stem Removal by Treatment

	Neighbor Observation	Full Observation
Low Cost	27.3% (65/238)	23.9% (62/259)
High Cost	47.5% (135/284)	34.6% (63/182)

Observation - High Cost treatment. We cannot reject the null hypothesis that these trends are jointly equal (p-value 0.189). Overall, there was a 6.9% decrease in the frequency of convergence between matches, significant at the 0.01 level.

We test the null hypothesis that the frequency of cycles under neighbor observation is less than or equal to the frequency of cycles under full observation using Barnard's Exact Test. We reject this hypothesis at the 0.01 level for both low cost and high cost. For both cost environments, cycles were significantly more likely to occur in convergent networks with Neighbor Observation than with Full Observation (60.0% vs. 30.0% for low cost and 57.1% vs 16.7% for high cost). We interpret this as confirming Hypothesis 4. Nearly 60.0% of convergent networks had cycles with Neighbor Observation, but only 15.4% (6/39) of these networks had *visible* cycles. Thus the limited observation does increase the likelihood with which inefficient cycles form.

Jockeying for position may result in periodic disconnection, particularly when a stem is paired with its neighbor. For each individual and match, we record the number of times the individual was paired with her own stem and the number of times the individual elected to remove the tie. From this we compute this individual's frequency of stem removal for the match. The overall frequency of stem removal by treatment is shown in Table 5.

To test the null hypothesis that the frequency of stem removal under neighbor observation is greater than or equal to the frequency of stem removal under full observation, we estimate a regression of the form

$$\text{droppedstems}_i = \mathbf{M}_i\vec{\alpha} + \beta\text{neighbor}_i$$

for each cost setting, where $\mathbf{M}_i$ is a vector of match dummy variables and neighbor_i equals 1 for Neighbor Observation, 0 otherwise. We then conduct a t-test of whether β is greater than or equal to 0, correcting for two-way clustering at the group and participant levels. We cannot reject this hypothesis under low cost (p-value 0.666) or high cost (p-value 0.914). For high cost, stems were significantly more likely to be dropped with neighbor observation (47.5%) than with full observation (34.6%) but for low cost there was no statistically significant difference in stem removal (27.3% vs 23.9%). We interpret this as rejecting Hypothesis 5. Stem actors are more likely to be dropped under neighbor observation even though it may be harder for these actors to reconnect.

In an effort to better understand this rejection, we conducted a follow-up analysis. Our hypothesis was premised on the notion that reconnection would be more difficult under Neighbor Observation than Full Observation due to uncertainty about one's own connectivity to an unobserved actor; however, an unobserved actor may also be viewed *more* favorably than an observed isolate when actors have optimistic beliefs (i.e., that the actors outside one's component are connected, not isolated). This effect is especially large in early periods, when components are much smaller than the maximum of 12, and with high cost, where connecting to an isolated actor is myopically detrimental.

We do not elicit the beliefs of our subjects, and so we cannot definitively show that optimistic beliefs are the cause of this rejection. However, there are several checks that we can conduct to assess the plausibility of this explanation. First, we note that the networks we observe under Neighbor Observation are more consistent with optimistic beliefs than pessimistic beliefs. Only 1/40 convergent networks have less than 10 actors in the maximal component. Further, only 2/71 non-convergent networks had less than 10 actors in the maximal component throughout the last 5 periods. If actors held pessimistic beliefs, we would expect more instances where the maximal component had 8 or 9 actors. Second, we note that under Neighbor Observation, a much larger percentage of stem removal occurs in the early rounds, before network efficiency plateaus (44.6% vs 27.4% for low cost, and 31.1% vs 14.3% for high cost). When we compare the distribution of stem removals by period, we reject the null hypothesis of equal medians at the 0.05 level for high cost. We cannot

Table 6: Fraction of Time Spent in Stem Position by Risk Category

$\rho < 0.00$	33.8%
$0.00 < \rho < 0.53$	33.0%
$0.53 < \rho < 0.90$	34.5%
$0.90 < \rho < 2.63$	29.0%
$2.63 < \rho$	30.2%

reject this hypothesis for low cost (p-value 0.24). This evidence points (weakly) towards a higher percentage of jockeying in the early rounds for Neighbor Observation than Full Observation, which is consistent with the incentives of far-sighted actors when beliefs are generally optimistic.

To test the null hypothesis that the frequency of stem removal under low cost is greater than or equal to the frequency of stem removal under high cost, we estimate a regression of the form

$$\text{droppedstems}_i = \mathbf{M}_i\vec{\alpha} + \beta\text{lowcost}_i$$

for each observation level, where $\mathbf{M}_i$ is a vector of match dummy variables and lowcost_i equals 1 for low cost, 0 otherwise. We then conduct a t-test of whether β is greater than or equal to 0, correcting for two-way clustering at the group and participant levels. We reject this hypothesis at the 0.01 level under Neighbor Observation and at the 0.05 level under full observation. For both observation levels, stems are more likely to be dropped with high cost than with low cost (47.5% vs. 27.3% for Neighbor Observation and 34.6% vs 23.9% for Full Observation). We interpret this as confirming Hypothesis 6. Actors do tend to remove ties whose cost exceeds the immediate benefit.

An individual's chosen position in the network may also depend on risk aversion. For each individual and match, we record the number of rounds the individual spent in the stem position. From this we compute this individual's fraction of time spent in the stem position for the match. For the purpose of this analysis, we exclude the first 5 periods, because networks do not typically begin to stabilize until round 6. The overall fraction of time spent in the stem position by risk category is shown in Table 6.

To test the null hypothesis of equal proportions across risk categories against the alternative of different proportions, we first estimate a pooled regression of the form

$$\text{droppedstems}_i = \mathbf{M}_i\vec{\alpha} + \mathbf{R}_i\vec{\beta},$$

where $\mathbf{M}_i$ is a vector of match dummy variables and $\mathbf{R}_i$ is a vector of risk category dummy variables (one category omitted) corresponding to the choices in the risk elicitation task. We then conduct an F-test of whether the coefficients in $\vec{\beta}$ are jointly equal to 0, correcting for two-way clustering at the group and participant levels. We cannot reject this hypothesis (p-value 0.456). While time spent in the stem position does follow the predicted pattern across risk categories, the difference across categories is neither statistically or economically significant. Thus, the evidence does not support Hypothesis 7. The risk of being dropped from the network does not seem to be a major factor in each actor's decision about where to locate in the network.

4.3 Typical Networks

There is wide variance across realized networks within each treatment, yet it is instructive to identify what we can be considered typical for each treatment condition. Figures 5 and 6 provide one way to characterize the median behavior of the network formation process by treatment condition. We can infer from these analyses when the median network converges (if at all) and at what level of efficiency. These are just two of many dimensions on which networks can be compared; other dimensions include the number of redundant ties, number of cycles, size of cycles, and size of maximal component. To better understand the properties of the typical networks in each treatment, we compute the median values for each of these dimensions and then find a realized network from each treatment that most closely matches the median across dimensions. Visuals of these networks are shown in Figure 7. Each network converged within 1 period of the median for that treatment, has an efficiency within a few percentage points of the median, and exactly matches the median in all other dimensions. In these visuals, created using the software package UCINET (Borgatti, Everett & Johnson

Figure 7: Typical Network by Treatment

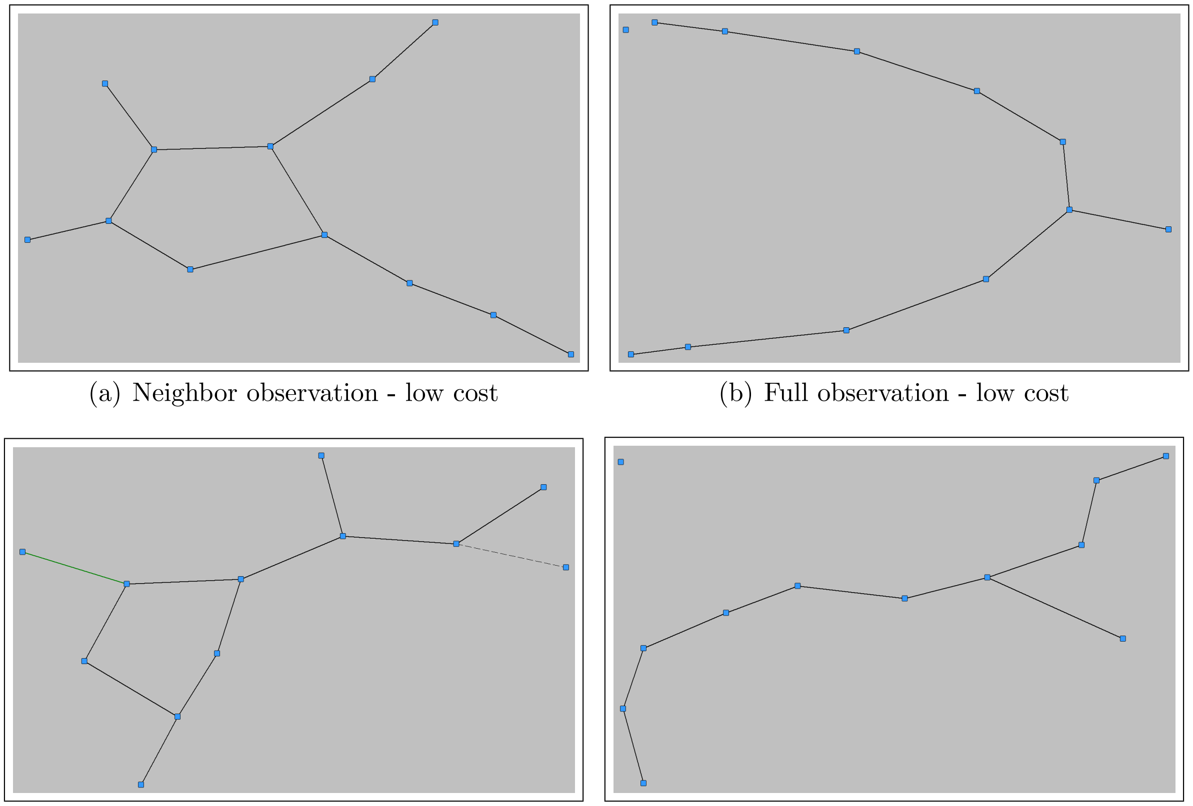

(a) Neighbor observation - low cost

(b) Full observation - low cost

(c) Neighbor observation - high cost

(d) Full observation - high cost

2002), actors are denoted by grey squares and ties are denoted by solid black lines. When the network did not converge, we present a visual of that network over the last 5 periods. Ties that were added after period 11 are colored grey, and ties that were removed before period 15 are denoted by dashed lines.

The typical network for Neighbor Observation - Low Cost converges to a connected network with one hidden cycle/redundant tie (Figure 7(a)), the typical network for Neighbor Observation - High Cost does not converge and has a component of size 11 (one isolate) with one cycle/redundant tie (Figure 7(c)), and the typical networks for both Full Observation treatments converge to a component of size 11 (one isolate) with no cycles/redundant ties (Figures 7(b) and 7(d)). That we see cycles as typical for the neighbor observation treatments but not with the full observation treatments nicely fits the motivation of our project, namely, that limited observation has implications for network structure and efficiency. The lack of convergence in the Neighbor Observation - High Cost is not surprising given that removing stems is myopically optimal. Perhaps surprising is that the typical networks with full observation are nearly identical, suggesting that the cost was not as important for determining network structure as expected. We remind the reader, however, that Figure 7 does not reveal the large variation observed across and within all conditions as our earlier data analysis reveals.

5 Conclusion

We present the first experimental study of network formation when actors have limited observation of the network. Our 2x2 experiment design exogenously varies the cost of forming links and the level of network observation. Observed behavior is sometimes, but not always, consistent with our predictions. Convergence is more likely with low link cost than high link cost, but it is also less likely with Neighbor Observation than with Full Observation. Efficiency is higher under low cost as expected, with more cycles observed under Neighbor Observation, and more stems dropped with Full Observation. Unexpectedly, stems are removed more often under full observation, and we do not find a significant relationship between network position and risk preference. Overall, we find that behavior under low cost is largely consistent with the theory.

We also find strong evidence of farsighted behavior, and this behavior plays a double-edged role in our setting. On the positive side, farsighted behavior leads to the rapid formation of near-efficient networks or near-efficient recurrence classes of networks. This is particularly beneficial in high cost settings where myopic behavior would lead to the severing of ties. However, on the negative side, the farsighted behavior leads to utility losses as subjects jockey for individually-advantageous network positions. The positives appear to largely outweigh the negative for the subjects in our setting, especially with high link cost. We further note that the PFS theory is not a good predictor of behavior because jockeying appears to deter movement into farsighted improving paths.

The frequency and characteristic of jockeying appears to be different under Neighbor Observation than Full Observation. While there is a higher frequency of inefficient, redundant links (cycles) under Neighbor Observation, there is also a lower frequency of inefficient jockeying, and this leads, surprisingly, to higher average utilities under Neighbor Observation. When jockeying occurs under Neighbor Observation, it appears relatively more often in earlier rounds, unlike under full observation, in which case it is later rounds. This may be evidence of optimistic beliefs.

Future work has many avenues to explore. The incentives to form and sever links vary widely across network settings, and additional work is needed to determine how many of our findings generalize. That the location of a subject in the network may depend on her risk preferences also deserves closer examination, especially in network settings where network position is more consequential for welfare than in our setting. Finally, future theoretical and experimental work can identify just how much network observation is enough to mimic Full Information, and can even examine the actors' incentives to collect information, perhaps at a cost, about others links. We expect this work to provide even more insights into the role limited observation plays in network formation.

References

Acemoglu, D., Dahleh, M. & Lobel, I. (2011), 'Bayesian learning in social networks', *Review of Economic Studies* **78**(4), 1201–1236.

Bala, V. & Goyal, S. (2000), 'A noncooperative model of network formation', *Econometrica* **68**, 1181–1229.

Berninghaus, S., Ehrhart, K.-M. & Ott, M. (2006), 'A network experiment in continuous time: The influence of link costs', *Experimental Economics* **9**, 237–251.

Bondonio, D. (1998), 'Predictors of accuracy in perceiving informal social networks', *Social Networks* **20**, 201–220.

Borgatti, S., Everett, M. & Johnson, J. (2002), *UCINET 6 for Windows: Software for Social Network Analysis*, Analytic Technologies.

Callander, S. & Plott, C. (2005), ‘Principles of network development and evolution: An experimental study’, *Journal of Public Economics* **89**, 1469Ű1495.

Calvó-Armengol, A. & Ilkiliç, R. (2009), ‘Pairwise-stability and nash equilibria in network formation’, *International Journal of Game Theory* **38**, 51–79.

Calvó-Armengol, A. & Jackson, M. (2004), ‘The effects of social networks on employment and inequality’, *American Economic Review* **94**, 426–454.

Cameron, A. C., Gelbach, J. & Miller, D. (2011), ‘Robust inference with multi-way clustering’, *Journal of Business and Economic Statistics* **29**, 238–249.

Candelo, N., Forbes, S., Martin, S., McBride, M. & Allison, B. (2014), ‘Endogenous formation of dark networks: Theory and evidence’, *Manuscript* .

Carrillo, J. & Gaduh, A. (2012), ‘The strategic formation of networks: Experimental evidence’, *Manuscript* . Working Paper.

Casciaro, T. (1998), ‘Seeing things clearly: Social structure, personality, and accuracy in social network perception’, *Social Networks* **20**, 331–351.

Conley, T. & Udry, C. (2001), ‘Social learning through networks: The adoption of new agricultural technologies in ghana’, *American Journal of Agricultural Economics* **83**, 668–673.

Corbae, D. & Duffy, J. (2008), ‘Experiments with network formation’, *Games and* **64**, 81–120.

Dogan, G., van Assen, M. & Potters, J. (2013), ‘The effect of link costs on simple buyer-seller networks’, *Games and Economic Behavior* **77**, 229–246.

Eckel, C. & Grossman, P. (2008), ‘Forecasting risk attitudes: An experimental study using actual and forecast gamble choices’, *Journal of Economic Behavior and Organization* **68**, 1–17.

Falk, A. & Kosfeld, M. (2012), ‘It’s all about connections. evidence on network formation’, *Review of Network Economics* **11**, 1–36.

Fischbacher (2007), ‘Z-tree: Zurich toolbox for ready-made economic experiments’, *Experimental Economics* **10**, 1710178.

Francetich, A. & Troyan, P. (2010), ‘A network formation game with private information’, *Manuscript* . Working Paper.

Friedkin, N. (1983), ‘Horizons of observability and limits of informal control in organizations’, *Social Forces* **62**, 54–77.

Gallo, E. & Yan, C. (2015), 'Efficiency and equilibrium in network games: An experiment', *Manu* .

Goeree, J., Riedl, A. & Ule, A. (2009), 'In search of stars: Network formation among heterogeneous agents', *Games and Economic Behavior* **67**, 445–466.

Grandjean, G., Mauleon, A. & Vince (2011), 'Connections among farsighted agents', *Journal of Public Economic Theory* **13**, 935–955.

Herings, J.-J., Mauleon, A. & Vannetelbosch, V. (2009), 'Farsightedly stable networks', *Games and Economic Behavior* **67**, 526–541.

Jackson, M. & Watts, A. (2002), 'The evolution of social and economic networks', *Journal of Economic Theory* **106**, 265–295.

Jackson, M. & Wolinsky, A. (1996), 'A strategic model of social and economic networks', *Journal of Economic Theory* **71**, 44–74.

Kranton, R. & Minehart, D. (2001), 'A theory of buyer-seller networks', *American Economic Review* **91**, 485–508.

Kumbasar, E., Romney, K. & Batchelder, W. (1994), 'Systematic biases in social perception', *American Sociological Review* **100**, 477–505.

Laumann, E. (1969), 'Friends of urban men: An assessment of accuracy in reporting their socioeconomic attributes, moral choice, and attitude agreement', *Sociometry* **32**, 54–69.

Mantovani, M., Kirchsteiger, G., Mauleon, A. & Vannetelbosch, V. (2013), 'Limited farsightedness in network formation', *Manuscript* . Working Paper.

McBride, M. (2006*a*), 'Imperfect monitoring in communication networks', *Journal of Economic Theory* **126**, 97–119.

McBride, M. (2006*b*), 'Limited observation in mutual consent networks', *Advances in Theoretical Economics* **6**, 1–29.

McBride, M. (2008), 'Position-specific information in social networks: Are you connected?', *Mathematical Social Sciences* **56**, 283–295.

Pantz, K. (2006), The Strategic Formation of Social Networks: Experimental Evidence, PhD thesis, Humbolt Universität.

Petersen, M. (2009), 'Estimating standard errors in finance panel data sets: Comparing approaches', *Review of Financial Studies* **22**(1), 435–480.

Rong, R. & Houser, D. (2013), 'Emergent star networks with ex ante homogeneous agents', *Manuscript* . Working Paper.

Song, Y. & van der Schaar, M. (2015), 'Dynamic network formation with incomplete information', *Economic Theory* **59**(2), 301–331.

Thompson, S. (2011), 'Simple forumlas for standard errors that cluster by both firm and time', *Journal of Financial Economics* **99**, 1–10.

van Leeuwen, B., Offerman, T. & Schram, A. (2013), 'Superstars need social benefits: An experiment on network formation', *Manuscript* . Working Paper.

Watts, A. (2001), 'A dynamic model of network formation', *Games and Economic Behavior* **34**, 331–341.

Appendix

A Supplementary Tables

Table A.1: Primary Hypothesis Tests

Hypothesis Test	Statistic (p-value)		
H_1: Convergence$_{LC} \leq$ Convergence$_{HC}$		Observation Level	
H_{1a}: Convergence$_{LC} >$ Convergence$_{HC}$		Neighbor	Full
	Barnard's Exact	0.002	0.938
Logistic Regression: Group Convergence	t-test	0.001	0.939
H_2: Convergence$_{NO} \leq$ Convergence$_{FO}$		Tie Cost	
H_{2a}: Convergence$_{NO} >$ Convergence$_{FO}$		Low	High
	Barnard's Exact	0.279	> 0.999
Logistic Regression: Group Convergence	t-test	0.264	> 0.999
H_3: Efficiency$_{LC} \leq$ Efficiency$_{HC}$		Observation Level	
H_{3a}: Efficiency$_{LC} >$ Efficiency$_{HC}$		Neighbor	Full*
	Mann-Whitney U	< 0.001	0.008
Regression: Individual Share of Group Efficiency	t-test	< 0.001	< 0.001
H_4: Cycles$_{NO} \leq$ Cycles$_{FO}$		Tie Cost	
H_{4a}: Cycles$_{NO} >$ Cycles$_{FO}$		Low	High
	Barnard's Exact	0.001	< 0.001
Logistic Regression: Group Cycles	t-test	< 0.001	< 0.001
H_5: Drop Stem$_{NO} \geq$ Drop Stem$_{FO}$		Tie Cost	
H_{5a}: Drop Stem$_{NO} <$ Drop Stem$_{FO}$		Low	High
Regression: Dropped Stems (Match Proportion)	t-test	0.666	0.914
H_6: Drop Stem$_{LC} \geq$ Drop Stem$_{HC}$		Observation Level	
H_{6a}: Drop Stem$_{LC} <$ Drop Stem$_{HC}$		Neighbor	Full
Regression: Dropped Stems (Match Proportion)	t-test	0.001	0.014
H_7: Time as Stem Equal by Risk Category		Pooled Data	
H_{7a}: Different Proportions			
Regression: Time as Stem (Match Proportion)	F-test	0.456	

Note: Bold denotes hypotheses predicted to be true. All regressions utilize subject level data with standard errors two-way clustered at the group and subject level. Regressions also include controls for the match number.

*Excludes 9 networks containing a subject who never connected. Including these networks changes p-values to 0.017 and 0.003. The treatment of this subject does not qualitatively affect the result of any hypothesis test.

Table A.2: Trend in Group Outcomes (Match)

Treatment	Dependent Variable		
	Convergence (%)	Efficiency (%)	Cycles (%)
Neighbor Observation - Low Cost			
Average Marginal Effect	5.5	-0.6	-5.6
Standard Error	(1.9)	(0.8)	(1.8)
p-value	0.005	0.437	0.002
Neighbor Observation - High Cost			
Average Marginal Effect	3.2	-0.7	-8.2
Standard Error	(1.8)	(0.7)	(1.1)
p-value	0.075	0.308	< 0.001
Full Observation - Low Cost			
Average Marginal Effect	2.8	0.9	-8.6
Standard Error	(2.5)	(0.9)	(1.1)
p-value	0.271	0.311	< 0.001
Full Observation - High Cost			
Average Marginal Effect	5.0	3.7	-4.2
Standard Error	(2.2)	(0.8)	(2.2)
p-value	0.026	< 0.001	0.055
Hypothesis Test	Wald statistic (p-value)		
$\mathbf{H}_{\mathrm{L}}$: Trends Jointly Equal			
H_{La}: Different Trends	0.756	< 0.001	0.189

Note: Bold denotes hypothesis predicted to be true.

Table A.3: Paired Comparisons: Trend in Group Outcomes (Match)

Comparison	Bonferroni Adjusted p-value
$\text{Efficiency}_{\text{NO-LC}} = \text{Efficiency}_{\text{NO-HC}}$	> 0.999
$\text{Efficiency}_{\text{NO-LC}} = \text{Efficiency}_{\text{FO-LC}}$	> 0.999
$\text{Efficiency}_{\text{NO-LC}} = \text{Efficiency}_{\text{FO-HC}}$	< 0.001
$\text{Efficiency}_{\text{NO-HC}} = \text{Efficiency}_{\text{FO-LC}}$	0.925
$\text{Efficiency}_{\text{NO-HC}} = \text{Efficiency}_{\text{FO-HC}}$	< 0.001
$\text{Efficiency}_{\text{FO-LC}} = \text{Efficiency}_{\text{FO-HC}}$	0.141

B Experiment Instructions

Text enclosed in square brackets, e.g., [string1/string2], depends on the treatment.

SCREEN 1:

Welcome to this experiment at UC Irvine. Thank you for participating.

You are about to participate in a study of decision-making, and you will be paid for your participation in cash, privately at the end of this session. What you earn depends partly on your decisions and partly on chance.

Please turn off your cell phone.

The first part of this experiment consists of 10 matches with 15 rounds per match. You will be paid according to the outcome of *each* round and you will receive the sum of all your individual round earnings, in addition to the show up fee and your earnings in the second part, at the end of the experiment. You will receive further instruction when the second part of the experiment begins.

All rounds will take place through the computer terminals. It is important that you do not talk with any other participants during the session.

When you are ready, please click "Continue" to go to the instructions.

SCREEN 2:

In each match, you will be randomly grouped with 11 other participants. The participants that you are grouped with will stay the same for the 15 rounds in this match, but you will be randomly grouped with a new set of participants at the start of the next match. Each participant will be assigned a color and will retain this color throughout the match.

During each match, you and the other participants will form a network. In each round, you will be presented with a [partial view of the network/complete view of the network]. At the start of the first round, the network has no connections. In all following rounds, you will be shown [part/all] of the network formed at the end of the prior period.

The network consists of 12 colored circles, each representing a participant in your group, which may or may not be connected via lines (we refer to these lines as "links"). Participants are either directly linked, denoted by a line connecting two participants, or not directly linked, denoted by the absence of a line.

Please press "Continue" to proceed.

SCREEN 3:

After observing [part/all] of the network, each participant is then randomly paired with another participant in the network and given the opportunity to change his or her link status with that participant. If you already have a link with that participant, you can maintain that link or remove that link. If you do not have a link with that participant, you can initiate a link to that participant or not initiate a link. Your link status with all other participants stays the same. Links that have already been initiated are maintained.

NOTE: Links are initiated and maintained bilaterally, meaning that a link is initiated or maintained only if BOTH participants choose to do so. If one of the participants does not want the link then it will not form.

You will have 15 seconds per round to make your decision, and you must press the "Continue" button before time runs out for your decision to be recorded. If you do not press "Continue," then the computer will assume you do not want to make any changes to the network that formed last period. Please press "Continue" to proceed.

SCREEN 4:

We say that two participants are "indirectly linked" if these participants are not directly linked but are connected through a series of links with other participants.

In each round you will [only see your own links and the links of participants that you are directly linked to/be shown the links of every participant].

Your payoff for each round will be 1 point for every participant you are either directly or indirectly linked to minus [0.7/1.5] points for every participant that you have attempted to initiate or maintain a link to. Note that you will pay this cost even *if your link attempt is not successful.*

EXAMPLE 1: If A and B are both directly linked to each other and no-one else, each receives 1 point for being linked to one other participant and pays a cost of [0.7/1.5] for having one direct link. The resulting payoff is [0.3/-0.5].

EXAMPLE 2: If A and B are both directly linked to C and no-one else, then A and B both receive 2 points for being linked to two other participants and pay a cost of [0.7/1.5] for having one direct link. The resulting payoff is [1.3/0.5].

NOTE: You will receive a payoff and pay the cost of all your direct links (including those links to participants with whom you were not paired) *in every round.* [You receive 1 point for every participant you are linked to even if you cannot see how you are linked to that participant./] Please press "Continue" to view an example.

SCREEN 5:

Below is a sample network. You are red and have exactly one link. This link is to yellow, and you can see that yellow is linked to lime. [There may be other links in the network that you cannot see./] At the bottom of the screen, there is a box around yellow denoting that in this round you must choose whether to change your link status with yellow. If you were paired with blue, then the box would be around blue.

In this example, you can maintain the link with yellow or remove it. Please select "No" to remove the link and then press "Continue."

SCREEN 6:

You are no longer linked to yellow because you chose not to maintain this link. [Assuming no other participant has changed their link status and that there were no other links you couldn't see in the network, the new network is shown to the left below. Your view of the new network is shown to the right below. Notice that you no longer see the link between yellow and lime because you are no longer linked to yellow./Assuming no other participant has changed their link status, the new network is shown below.]

Please press "Continue" to proceed.

SCREEN 7:

This is an example of what your decision screen will look like at the start of the first round. In this example, you are red and you are paired with yellow. You will ordinarily have 15 seconds to make a decision on this screen.

SCREEN 8:

The first of the 10 matches will now begin. Each match has 15 rounds and you will be paid for *each* round. **You will be paid $0.02 for each point you earn during the session**. You will also receive [$2.00/$4.50] at the beginning of the first match.

REMINDER: You will only have 15 seconds each round on the decision screen. You must press "Continue" before time runs out for your link decision to be recorded.

You will receive 1 point for each participant you are directly or indirectly linked to and you will pay a cost of [0.7/1.5] points for each link attempt you initiate or maintain. You pay the cost to maintain a link even if you were not able to change the status of that link in this round. You do not pay the cost for link attempts that failed in a prior round.

You will start with no links in the first round of each match. In future rounds, you will start with the links you had in the prior round.

In each round you will [only see your own links and your neighbors' links, but you will still receive points if you are indirectly linked to a participant that you cannot see./be shown the links of every participant, regardless of whether or not you are linked to that participant.]