

Social and Information Networks

University of Toronto CSC303
Winter/Spring 2019

Week 10: March 16-20 (2020)

Announcements

- As we all know, the University has cancelled classes until the end of term and the Faculty has cancelled Final Exams. In consultation with the Department of Computer Science, we have a revised grading scheme (next slide) which we will put to the class for a majority vote of approval.
- We will provide information during this period on the web page, on piazza and on these lecture slides. We will mainly respond to general questions as above and not by individual emails. However, we will try to make ourselves as available as possible by email and skype for questions that are more individual specific. My skype address is abborodin@gmail.com. In order to schedule meetings during the regular office hours, please request a 10-15 minute time slot by sending me an email with title "CSC303 - office hour meeting request" at the latest one hour before regular office hours. For meetings out of regular office hours entitle your email as "CSC303 - meeting request". I will try to meet all requests at some time.

Initial and proposed revised grading scheme

Original grading scheme:

- Assignments: two worth 15% each; Due dates: February 14 and March 16
- Critical review of a current article: worth 10% ; Due date: March 30
- 2 hour Term Test: worth 20% ; Date: March 4 and 6
- 3 hour Final exam (3 hours): worth 40%

New grading proposal:

- Two Assignemnts : two worth 20% each. Already completed
- Additional assignment: 15% Due April 15
- Critical review: 15% Due March 30
- Midterm 30% Already completed

Announcements continued

- Critical review is due March 30 by 11:59PM.
- So far (as of 11AM March 21), I only have 62 proposals for the critical review assignment. I repeat what you need to do.
 - ▶ Due date: March 30
 - ▶ You need to find a conference or journal article that has appeared in the last 3 years; to be precise lets say, has appeared since January 1, 2017.
 - ▶ The article can be about any topic in the coursei but I must approve your choice.
 - ▶ You are to provide a critical review of the article as if you were on program committee or a reviewer for a journal. I have elaborated on this assignment in a number of the lectures.
 - ▶ Please send me your suggestion for the paper you will critically review. If you sent me a suggestion and I did not reply then please resend. I wanted your suggestions by Wednesday, March 18. i
 - ▶ Almost all suggestions have been very appropriate. But there are a few that are either not appropriate or where the article may only appear say in the arXiv. I repeat that the article must appear (or has been accepted and will appear) in a conference or a journal. If you submit a review for a paper that is not appropriate, you will receive little or no credit.

Today's agenda and what will follow

- 1 We will complete chapter 21 with the discussion of genetic inheritance and “Mitochondrial Eve”
- 2 This week we mainly will discuss Chapter 12, Bargaining in a Network Exchange Model.
- 3 For the remainder of the course, we will discuss the Stable Marriage problem and then Congestion Networks and Braess' paradox.

Genetic inheritance and networks

Chapter 21 turns its attention to the issue of genetic inheritance, viewed as a random process taking place on a (directed acyclic) network of organisms (species, parts of a genome, etc).

Section 12.7 starts off with a very motivating example. In 1987, Cann, Stoneking and Wilson published a very striking and to many a very controversial paper. They asserted that if one traces their maternal lineage back in time, everyone's lineage traces back to a single woman (called Mitochondrial Eve), living sometime between 100,000 and 200,000 years ago and probably living Africa.

The chapter ignores the issue of the location of Mitochondrial Eve and focuses on the basis (i.e. a model based on various assumptions) for this bold assertion of a common ancestry.

Note: I suggest reading the text as to the caveats about the model.

Modeling the Mitochondrial Eve assertion

To understand the assertion, we have to make some simplifying biological assumptions. Later to understand the assertion more precisely (as in the advanced section 21.8 B) we also need to make some simplifying mathematical assumptions. These latter assumptions are easy to justify and do not change any of the conclusions.

The biological assumptions are beyond the scope of the course but we will accept them as they have been generally accepted in the sense that qualitatively they do not change the conclusions.

The biological basis for the model is that “mitochondrial DNA (is to a first approximation) passed on to children entirely from their mothers”.

Mitochondrial Eve continued

Once we focus on mitochondrial DNA and accept that it is inherited only from ones mother, we are then able to consider a “single parent” ancestry model.

This model will conclude that our common mitochondrial DNA ancestry must have originated with with a single female Mitochondrial Eve and the mathematical analysis will give an estimate for the time period in which she lived.

This does not say that Mitochondrial Eve was the only female (or male) alive at this time but just that our mitochondrial DNA traces back to such a female. And of course our genomic makeup does come from both parents.

The Wright-Fisher single parent ancestry model

There are additional simplifying assumptions that need to be made to make the model more tractable. The model not only applies to mitochondrial lineage but also to reproduction in asexual reproduction and (with some further assumptions) to specific nucleotides in our genome.

We assume a fixed population of N individuals throughout the entire period of time. This is inconsistent with the fact that world population is growing. But we will argue that this does not change the nature of the conclusions or even the nature of the analysis.

In fact, once we accept that populations are growing, it is clear that certain individuals must be having multiple children which is also part of the model.

Single parent ancestry model continued

We assume that generations are completely synchronized, the generation of N individuals at some time t give rise to the next generation of N individuals at time $t + 1$.

Each individual at time $t+1$ has its “single parent” chosen uniformly at random from the previous generation, a significant assumption given geography, ethnicity, etc. To reconcile this (with respect to the assertion of a single Mitochondrial Eve), we need to understand the extent to which individual communities can be isolated. But the timing for when common ancestry would have taken place is not impacted by this assumption.

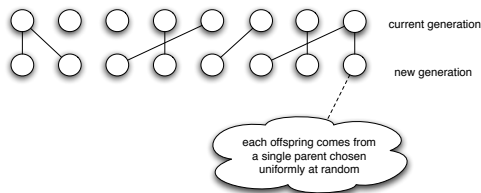
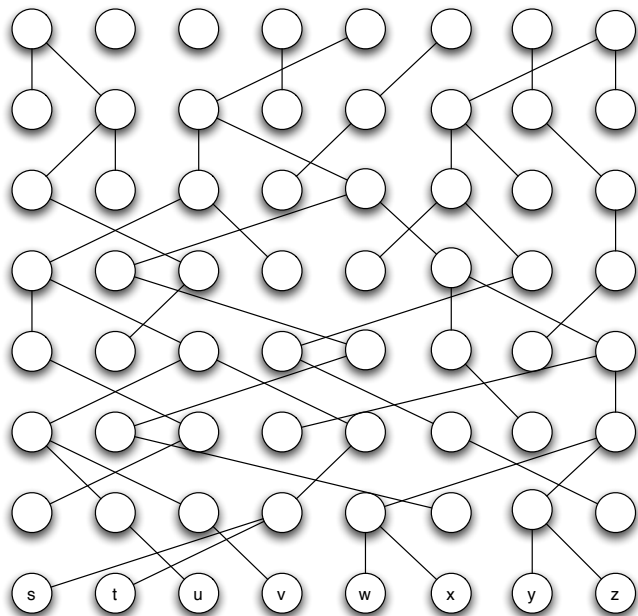
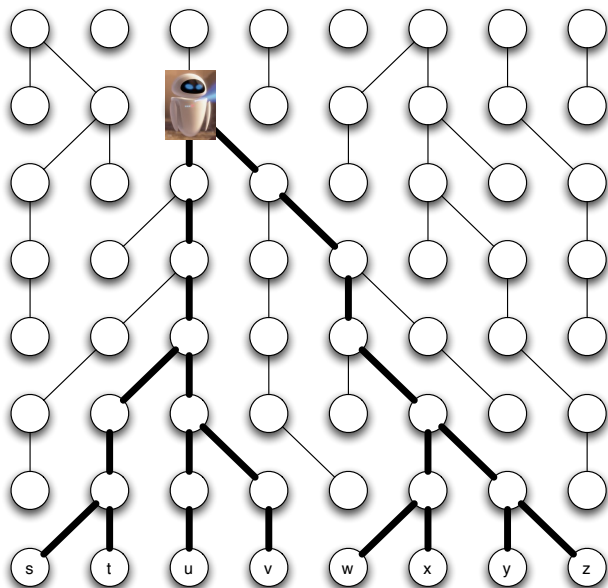


Figure: [Fig 21.11, E&K]

More generations of the model



Ancestry depicted.



The analysis for estimating the time that the model coalesces on Mitochondrial Eve

Section 21.8B provides a mathematical analysis for estimating the time when a common ancestor (in the single parent model) will be reached. Along the way, some simplifying mathematical assumptions are made but these assumptions are easily defended and are not of the same nature as biological assumptions.

Suppose we have a total population of N and at some point of time $t + 1$ that we are down to k candidates (lineages) for a common ancestor. We want to consider the probability that two lineages will collide so that there be (at most) $k - 1$ candidates. .

The analysis

We want to determine an estimate for the number of generations for a lineage having k individuals will shrink down to $k - 1$ individuals.

Case: $k = 2$. Say the active lineage is individuals $\{a, b\}$. Let a trace back to some individual c (one of the N in the entire population). Then the probability that b will not trace back to the same c is $1 - \frac{1}{N}$.

Case: $k > 2$. Lets consider the probability that there will be no collapsing into $k - 1$ lineages. There will be no collapsing if the second node doesn't collide with the first, the third doesn't collide with the first two, etc, so this means that the probability of no collapsing is :

$$\left(1 - \frac{1}{N}\right)\left(1 - \frac{2}{N}\right) \cdots \left(1 - \frac{k-1}{N}\right)$$

The analysis continued

The previous product

$$\left(1 - \frac{1}{N}\right)\left(1 - \frac{2}{N}\right)\cdots\left(1 - \frac{k-1}{N}\right)$$

is at most:

$$1 - \left(\frac{1 + 2 + \cdots + k - 1}{N}\right) + \frac{g(k)}{N^2}$$

where $g(k)$ depends only on k and not on N . For any fixed k , the latter term is relatively negligible and we can say that the probability is

$$1 - \frac{k(k-1)}{2N}.$$

Moreover, we have ignored the probability that three or more lineages collide at the same node or that different paths collide, all of which would speed up the process.

The analysis continued

Fact: If we have a binary random variable Y_k (i.e., a heads coin flip) that is repeated independently each time with probability p , then the expected time $E[X_k]$ for X_k to occur is exactly $1/p$. Of course, if the probability is at least p , then the expected time can only be shorter.

Therefore, letting X_k denote the time to collapse from k to $k - 1$ lineages, then $E[X_k]$ is approximated by $\frac{2N}{k(k-1)}$

Note: Initially when k is large, the decrease is expected every generation going back. But when k is a small constant, then the expected number of generations to show a decrease is proportional to N .

Depiction of the lineages colliding

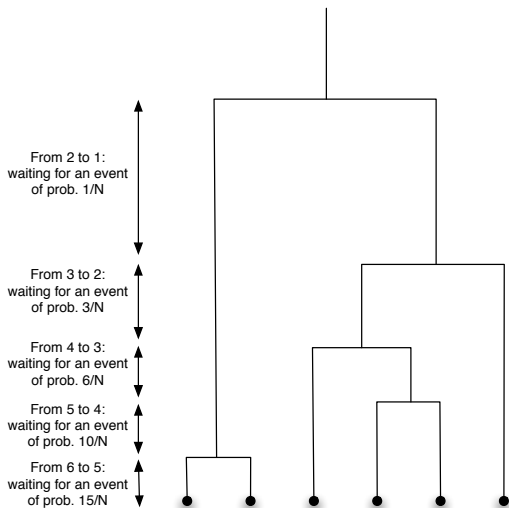


Figure: Assuming no three lineages collide simultaneously. [Fig 21.1(a), E&K]

Finishing the analysis

Let $X^k = X_k + X_{k-1} + \dots + X_2$ be the number of generation to reach a common ancestor starting from a lineage of k individuals.

Since $\mathbb{E}[x_j] = \frac{2}{j(j-1)}$ and $\frac{1}{j(j-1)} = \frac{1}{j-1} - \frac{1}{j}$, by linearity of expectations we have:

$$\begin{aligned}\mathbb{E}[X^k] &= \sum_{j=2}^k \frac{2N}{j(j-1)} \\ &= 2N \left(\left[\frac{1}{1} - \frac{1}{2} \right] + \left[\frac{1}{2} - \frac{1}{3} \right] + \dots + \left[\frac{1}{k-1} - \frac{1}{k} \right] \right) \\ &= 2N \left(1 - \frac{1}{k} \right)\end{aligned}$$

Note: Further more detailed analysis is consistent with the basic analysis that was presented in the text.

Chapter 12: Bargaining and Power in Networks

We begin a subtle and fascinating topic, namely how individuals in a network come to agreement on an outcome. This chapter is part of a larger subject called cooperative game theory and to some extent touches on behavioural game theory. As previously discussed, we have a course (CSC304) which covers game theory and in our course we will only present what is necessary regarding game theory. What we need is rather minimal (e.g., as when we were discussing network coordination in chapter 19).

But perhaps here is a good place to mention some basic game theory concepts to keep in mind (and again we have at least implicitly seen these concepts in our discussions to date). The following is a very brief set of informal comments.

- Individuals (agents) have strategies or actions and employ a (pure or mixed/randomized) strategy so as to act in self interest, always trying to maximize benefit or minimize cost.

A few more comments on game theory concepts

- **Note:** There is a lot of subtlety in how one understands benefits and costs as it often cannot be explained simply in monetary terms (or at least one has to have some way to assign monetary values to more subjective values, such as fairness, pride, reputation, etc.)
- The fact that agents are acting in self interest implies that their actions are decentralized. Mechanism design concerns how a central agent can introduce incentives so as to influence the actions of agents.
- A central concept in game theory is that of an equilibrium, which are states in which no agent has an incentive to change their strategy assuming no one else is changing. And again, we have also seen this concept, for example when considering the Schelling segregation model in Chapter 4, structural balance in Chapter 5. Equilibria will again be an important concept in Chapter 12 and the study of relative power. In the next two weeks we will see two more examples of equilibria, namely stable “marriages” (matchings) in a bipartite graph, and traffic equilibria in a congestion network.

Power as a relative relation between people

The chapter deals with individuals in a *bargaining network* and the relative power between any two people in this network.

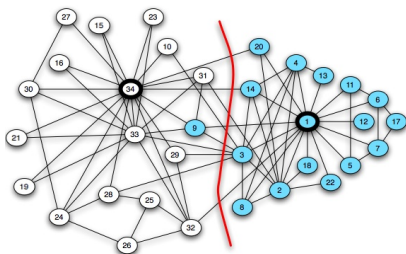
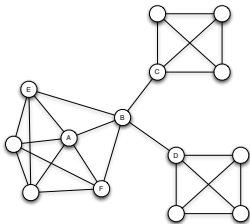
Power between individuals can come from two distinct sources:

- The relative reputation, status, official position, exceptional attributes (intelligence, finances), etc.
- The pivotal position of the person in the network.

In the first week of the term we mentioned the network of Florentine marriages and the *centrality* of the Medici family that was said to have conferred power to the Medicis. In the second week of the course we discussed the *bridging capital* of a node, such as node *B* in Figure 3.11 of the text, as well as the bonding capital and centrality of a node such as node *A* in Figure 13.1 or nodes 34 and 1 in the Karate Club in Figure 3.13.

Bridging and bonding capital of nodes

The early chapters of the text provided some insights about the importance of centrality and bonding and bridging capital with regard to the *flow of information*.



Power arising from asymmetries in pairwise relations

In contrast to these earlier discussions as to the pivotal relation of certain nodes in a network, Chapter 12 considers **power** in terms of the pivotal relation between two individuals that results in different values being conferred upon them corresponding to the imbalance in their relative power.

Note: In this context, centrality can sometimes be misleading.

Clearly, the above paragraph is vague and does not give us a definition of power. But, in fact, the study of power in this context (of imbalance) is a well studied concept that has led to applicable precise definitions. Here we emphasize that we are isolating power in terms of the position in a network and ignoring status aspects.


For motivation and following the text, we first present some simple but illustrative network examples. This will be followed by a social experiment that will provide insight as to the power imbalances in these simple networks. This in turn will lead to precise definitions.

Some illustrative examples

We will soon carefully describe the social experiment but briefly think a \$1 being placed on each edge of the network, and then each node trying to reach an agreement (within a fixed amount of time) with at most one other adjacent node as to how to split the dollar. (This pairing up of nodes is called a *matching* in graph theoretic terms. Formally, a matching M in a graph is a subset of edges such that no node is adjacent to more than one edge in the matching.)

Who will have relative power (i.e., receive more than half a dollar in the following networks)?

Let's start with the simplest possible network:



(a) 2-Node Path

Does either party have an advantage?

Some illustrative examples

We will soon carefully describe the social experiment but briefly think a \$1 being placed on each edge of the network, and then each node trying to reach an agreement (within a fixed amount of time) with at most one other adjacent node as to how to split the dollar. (This pairing up of nodes is called a *matching* in graph theoretic terms. Formally, a matching M in a graph is a subset of edges such that no node is adjacent to more than one edge in the matching.)

Who will have relative power (i.e., receive more than half a dollar in the following networks)?

Let's start with the simplest possible network:

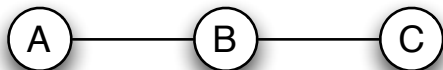


(a) 2-Node Path

Does either party have an advantage?

No; a $\frac{1}{2} - \frac{1}{2}$ split is a reasonable predicted split that is observed in the experiments.

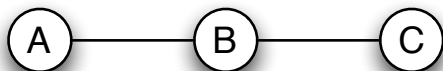
A three node path



(b) *3-Node Path*

What matching might occur and who each holds power?

A three node path



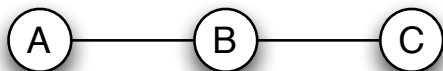
(b) *3-Node Path*

What matching might occur and who each holds power?

Clearly since we need a matching, either A and C will have to be left out. Intuitively then, node B holds much more power than A or C . The basic theory and experiments support this intuition.

What fraction of the \$ would you expect B to obtain in negotiating between A and C ?

A three node path



(b) 3-Node Path

What matching might occur and who each holds power?

Clearly since we need a matching, either A and C will have to be left out. Intuitively then, node B holds much more power than A or C . The basic theory and experiments support this intuition.

What fraction of the \$ would you expect B to obtain in negotiating between A and C ?


There is a difference between the basic theory and the social experiments. In the experiments, B gets a $(\frac{5}{6})^{th}$ fraction of the \$. The basic theory would predict that B gets almost all of the \$. Why the difference?

End of Wednesday, March 18 lecture

Today's agenda

- We ended the last lecture as we were beginning to consider a four node path and that is there we will start today.
- Complete Chapter 12
 - ▶ Some more suggestive examples
 - ▶ The social exchange experiment
 - ▶ The Nash Bargaining solution
 - ▶ Stable outcomes
 - ▶ The Ultimatum game experiment
 - ▶ Balanced outcomes
- If we have time we will start the *Stable marriage problem* and the Gale Shapley algorithm


A four node path



(c) *4-Node Path*

What matching might occur and how might the money be split? Would *B* get more or less in this four node network than in the previous three node path?

A four node path

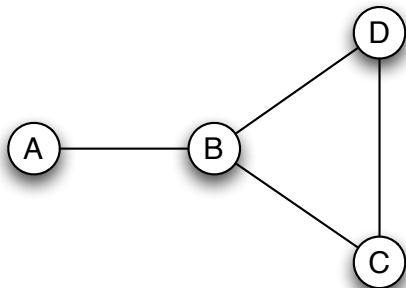


(c) *4-Node Path*

What matching might occur and how might the money be split? Would *B* get more or less in this four node network than in the previous three node path?

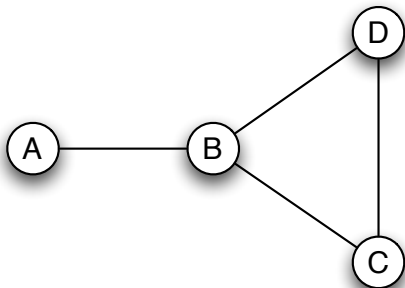
Here the experiments show that *B* gets a fraction of between $\frac{7}{12}^{th}$ and $\frac{2}{3}^{rd}$ of the \$, less than what we obtained in the three node network. **Why?**

The stem graph in figure 12.3



What matching might occur and how might the money be split? Would *B* get more or less in this stem network than in the previous three and four node paths?

The stem graph in figure 12.3



What matching might occur and how might the money be split? Would B get more or less in this stem network than in the previous three and four node paths?

Experiments show that B in the stem graph makes slightly more money than B in the four node path (but less than in the three node path). **Why?**

A five node path



(d) *5-Node Path*

Does *C* have any power (i.e. fraction of money obtained) compared to other nodes?

A five node path

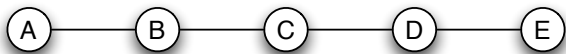


(d) *5-Node Path*

Does *C* have any power (i.e. fraction of money obtained) compared to other nodes?

Intuitively *B* and *D* have most of the power in the five node path network. The text states that in experiments, *C* has slightly more power than *A* or *E*.

A five node path



(d) *5-Node Path*

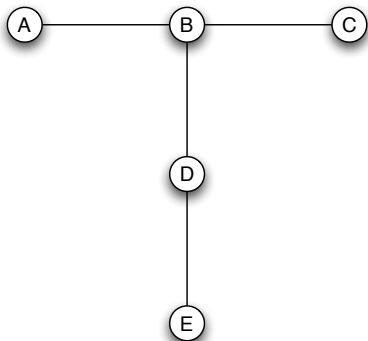
Does *C* have any power (i.e. fraction of money obtained) compared to other nodes?

Intuitively *B* and *D* have most of the power in the five node path network. The text states that in experiments, *C* has slightly more power than *A* or *E*.

Note that *C* is the most central node in terms of being on all shortest paths. However, this has not translated into substantial batgaining power.

Another graph to consider

The previous examples may help us reason about the following example from the text.



The *network exchange* social experiment

The following network exchange social experiment (and variants) is repeated a number of *rounds* so that some form of learning is taking place. There are many variants and the text presents one particular setting.

- Individuals (not knowing each other since we want to focus on the network aspects and not on the status, etc. of individuals) are placed at computer terminals and can interact with certain other individuals.
- In a complete information setting, one might see the entire network. The text considers the setting where an individual only knows and negotiates with their neighbouring nodes.
- For some known duration of time for a given round, negotiations take place for sharing say one \$ on each edge. (We could allow larger and different sums for each edge). Once a pair have decided how to share the \$, they leave the game.
- There is one more important condition on the experiment; namely in any given round, the outcome has to be a matching. This is called the *1-exchange rule*.

How much do these experimental findings depend on the exact setting.

We would, of course, like to have results that are robust and do not differ that much in the exact “details”.

It turns out that results are reasonably robust with regard to how much network information is available. And experiments carried out in different countries and different cultures are consistent.

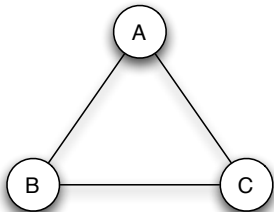
On the other hand, the 1-exchange rule is a definite factor impacting the results. In certain networks, substantially different findings would result if individuals could be negotiating two or more exchanges in a round. (In graph theory terms, this is called a b -matching when nodes can be adjacent to up to b edges in the matching.)

Anonymity is important. In particular, when knowing status, individuals of higher status tend to inflate their “options” and those of lower status tend to underplay their options.

Do all experiments converge in a consistent manner?

Another interesting observation is that for these simple networks, each round tends to come to consistent outcomes within the specified time limits.

However, there are networks where this is not the case. Consider the following triangle graph:



Note that any two of the nodes can wind up excluding the other. Hence we would expect that the final outcome in any round will be determined by the two nodes who get to settle just before the time deadline.

A mathematical perspective: The Nash Bargaining Solution

It would be good to have a model that will give us some insight into the kinds of results we see in the pure 1-exchange experiments (with anonymous participants).

John Nash (the same Nash who showed that all finite games have mixed equilibria) introduced the *Nash Bargaining Solution*. This will allow us to understand which outcomes will be *stable*. Note that without having a stable outcome, we cannot hope for participants to converge in any consistent way.

Conversely, we would expect that over enough rounds, participants would learn to converge to a stable outcome. Stable outcomes are equilibria and like most games, there can be many stable outcomes for a network exchange process.

Stable outcomes

We have already been implicitly discussing the idea of an outcome but for definiteness here is the definition for the case when every edge is worth 1\$. An outcome in a network exchange process on a graph $G = (V, E)$ is a pair (M, v) where $M \subseteq E$ is a matching and the value function $v : V \rightarrow [0, 1]$ satisfies:

- For every edge $e = (x, y) \in M$, $v_x + v_y = 1$.
- If a node $x \in V$ is not part of the matching M (i.e. does not appear as a vertex in any edge $(x, y) \in M$), then $v_x = 0$.

Intuitively, in a stable outcome, no agent (i.e. node) x can propose to an adjacent agent y , an offer that would improve both of their current outcomes. And we are now ready to define a stable outcome.

Stable Outcomes

An outcome (M, v) for a network exchange process is *stable* if for every edge $e = (x, y) \in E \setminus M$, $v_x + v_y \geq 1$.

Why are stable outcomes needed for an equilibrium?

Since we are assuming that each edge has exactly one \$ on each edge, clearly $v_x + v_y \leq 1$ for each edge $(x, y) \in M$, the matching.

Suppose $x + y < 1$ for an edge $(x, y) \notin M$. Then there is a surplus of $s = 1 - x - y$ that can be shared between x and y and there is no reason for them not to share this surplus so as to improve both their values.

Stable solutions are necessary but there can be many stable solutions and some are more natural (in the sense of corresponding to real behaviour) than others. What if x and y have other options other than to be in a given matching? Suppose that x (respectively, y) has an “outside option” of o_x (resp. o_y). Then $o_x + o_y \leq 1$ or else (x, y) could not be in a stable matching since they would eventually both take their outside options.

The Nash bargaining solution would be to put (x, y) in the matching and equally divide up any surplus from the outside options. That is, if

$s = 1 - o_x - o_y$, then set $v_x = o_x + \frac{s}{2} = \frac{o_x + 1 - o_y}{2}$ and

$v_y = o_y + \frac{s}{2} = \frac{o_y + 1 - o_x}{2}$. And hence we get:

$v_x + v_y = 1$ with (x, y) in the matching.

Why extreme outcomes are not real outcomes

As stated earlier in this chapter, in the three node path example, the theory thus far would predict that B will obtain the entire \$. But we are told that in experiments, more typically B gets a fraction $\frac{5}{6}$ and one other node gets a fraction $\frac{1}{6}$.

This can be explained once we understand that individuals (i.e., real people) are not driven solely by monetary payments. The “real value” to an individual may include some notion of fairness, pride, etc. Once we incorporate this into the framework, we can see why in these experiments, extreme solutions (which sometimes are the only stable solutions when viewed entirely in terms of monetary values) is not the outcome in these experiments.

In the following ultimatum game, we can perhaps better understand why participants tend to think beyond monetary rewards.

Another network exchange game: the so-called “Ultimatum Game”

We again are considering how two individuals divide a \$. But now we have the following experiment:

- One person (say A) is given one \$ and is told to propose a division of it to person B .
- Person B is then given the option of accepting the share offered or rejecting the offer.
- If B accepts, the game is over with the division as given by A . If B refuses then each person **gets nothing**.

Aside: This is a little like the “I cut-you choose 2-person cake cutting algorithm” which insure “fairness”.

This is a one-shot experiment between people who do not know each other. **What do we expect to happen?**

Another network exchange game: the so-called “Ultimatum Game”

We again are considering how two individuals divide a \$. But now we have the following experiment:

- One person (say A) is given one \$ and is told to propose a division of it to person B .
- Person B is then given the option of accepting the share offered or rejecting the offer.
- If B accepts, the game is over with the division as given by A . If B refuses then each person **gets nothing**.

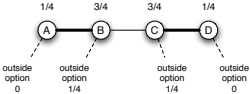
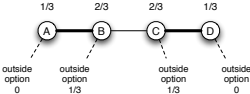
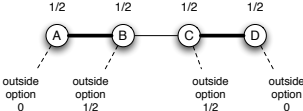
Aside: This is a little like the “I cut-you choose 2-person cake cutting algorithm” which insure “fairness”.

This is a one-shot experiment between people who do not know each other. **What do we expect to happen?**

Now in strictly monetary terms, person B should accept any offer (even a \$.01). But this is not what happens in experiments. In experiments, A tends to offer B about one third of the \$. **Why?**

Not all stable outcomes are “natural”

As we stated, there can be many stable outcomes for a given network. But some do not appear as natural as others and, in particular, stable outcomes can be “extreme solutions” that do not represent what we believe to be more realistic. Which of the following stable outcomes might be more expected “in practice”?



Balanced outcomes

It turns out that the $\frac{1}{3}, \frac{2}{3}$ split between A and B and also between C and D is what happens more in experiments and can be considered “more natural” in the following way.

The equal $\frac{1}{2}$ split amongst all parties does not reflect the relative much better bargaining position of B and C . In contrast, the $\frac{1}{4}, \frac{3}{4}$ split between A and B and also between C and D , seems to be giving B and C too much power given what we have been saying about how humans behave when taking say fairness, pride, etc into account.

Can we give a mathematical explanation for why the $\frac{1}{3}, \frac{2}{3}$ split should be a likely outcome?

It turns out that the $\frac{1}{3}, \frac{2}{3}$ split is the Nash Bargaining solution which we argued seemed like a fair way to divide up surpluses.

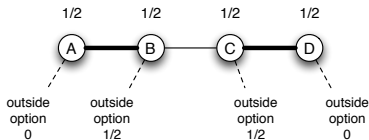
What is a balanced outcome?

Balanced outcomes

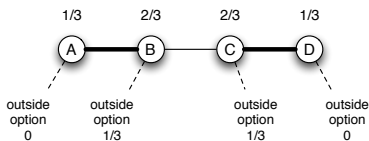
An outcome (M, v) is balanced if for every edge in the matching M , the split of money $\{v_x\}$ is the Nash bargaining solution for each node x , given the (best) outside options for each node.

Fact: For every exchange network, there is a balanced network.

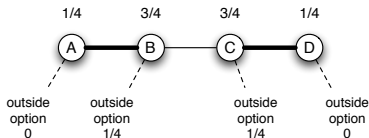
Balanced and unbalanced outcomes for the four node path



(a) *Not a balanced outcome*



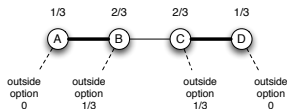
(b) *A balanced outcome*



(c) *Not a balanced outcome*

Checking that the balanced outcome is the Nash Bargaining solution

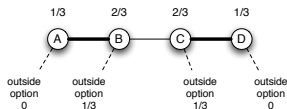
Let's check that the balanced outcome is indeed the Nash Bargaining solution.



Why is the best outside option for B (and similarly for C) equal to $\frac{1}{3}$?

Checking that the balanced outcome is the Nash Bargaining solution

Let's check that the balanced outcome is indeed the Nash Bargaining solution.



Why is the best outside option for B (and similarly for C) equal to $\frac{1}{3}$?

B has the option of offering $\frac{1}{3}$ (or maybe $\frac{1}{3} + \epsilon$ for some small $\epsilon > 0$) to entice C to leave its current match with D . Of course, A has no outside option so we can calculate that surplus for the matched edge (A, B) is $s = 1 - o_A - o_B = \frac{2}{3}$ and hence the Nash bargaining solution would be:

- $v_A = o_A + \frac{s}{2} = 0 + \frac{1}{3} = \frac{1}{3}$
- $v_B = o_B + \frac{s}{2} = \frac{1}{3} + \frac{1}{3} = \frac{2}{3}$

which is consistent with the balanced outcome.

Similarly we see that the values for nodes C and D are set in accordance with the Nash Bargaining solution.