Essays on Social Learning and Networks

Thesis by
Vadim Vadimovich Martynov

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2020
Defended April 27, 2020

Vadim Vadimovich Martynov
ORCID: 0000-0001-5357-4161

ACKNOWLEDGEMENTS

To the members of my dissertation committee, Omer Tamuz, Federico Echenique, Jakša Cvitanić, Luciano Pomato, I am deeply grateful for their guidance and advice through the whole process. To my adviser, Omer Tamuz, I am profoundly grateful for his unwavering support, humanity, mentoring and helpful suggestions regarding research throughout my graduate studies. To Federico Echenique, I am thankful for deepening my understanding of Economics and his challenging questions that prompted me to broaden my research perception. To Jakša Cvitanić, whose openness and approachability served as a vital component for my success, I am sincerely thankful for his great classes and motivating conversations with him. To Luciano Pomatto, my committee chair, I am indebted for his guidance and helpful suggestions for my presentations.

I would like to thank Antonio Rangel, Colin Camerer, Laura Doval, and Michael Gibilisco for their comments both in the classroom and outside of it.

I am deeply grateful to Kenneth Winston and Jakša Cvitanić who helped me find and pursue my professional interest through their classes, projects, mentorship, and extensive personal conversations.

I am thankful to the HSS department and the ISP that facilitated my academic learning and growth as a person, provided me with an opportunity to share my knowledge with others during the course of my studies here. I would especially like to thank Laurel Auchampaugh, Kapaūhi Stibbard, and Laura Flower Kim who were always there to help and resolve any complication, meanwhile making me feel that I belong to this place.

To my colleagues, I am thankful for their comments on my presentation and research, specifically to Alejandro Robinson-Cortes, Wade Hann-Caruthers, Saba Devdariani, Mali Zhang, Yimeng Li, and Shunto Kobayashi. Especially so to Hamed Hamze who would always listen to my ideas and hypotheses, providing helpful feedback.

To my friends, I am very thankful and honored to call you my friends with everything that comes with it.

To my family, I am eternally filled with gratitude for all you do for me and sincerely appreciate you always being there for me every step of the way. To Sasha, I am incredibly fortunate to have you by my side. I am thankful for all the help and support you give me as well as being the one who would always give me an outside perspective. I am deeply grateful to my parents, Vadim and Valentina, for their love and support. I candidly treasure the sacrifices they make so I can be the best version of myself. And to my grandmother, I always remember you, what you stand for, your sincere and unconditional love, and everything you do for the family.

ABSTRACT

This thesis offers a contribution to the study of Social Learning and Networks. It studies information aggregation and its effect on individual's actions (Chapter 2, 3) and social network (Chapter 4).

Chapter 2, co-authored with Omer Tamuz and Wade Hann-Caruthers, studies how quickly does the public belief converge to its true value when agents are able to observe actions of their predecessors. In the classical herding literature, agents receive a private signal regarding a binary state of nature, and sequentially choose an action, after observing the actions of their predecessors. When the informativeness of private signals is unbounded, it is known that agents converge to the correct action and correct belief. We study how quickly convergence occurs, and show that it happens more slowly than it does when agents observe signals. However, we also show that the speed of learning from actions can be arbitrarily close to the speed of learning from signals. In particular, the expected time until the agents stop taking the wrong action can be either finite or infinite, depending on the private signal distribution. In the canonical case of Gaussian private signals, we calculate the speed of convergence precisely, and show explicitly that, in this case, learning from actions is significantly slower than learning from signals.

In Chapter 3, I investigate how social planning can reduce the inefficiencies of social learning, stemming from herding and informational cascades. A social planner is introduced to the classical sequential social learning model. She can tax or subsidize players' actions in order to maximize social welfare, a discounted sum of agents' utilities. We solve or accurately approximate the expected utility of the social planner and the optimal pricing strategy for various signal distributions. In equilibrium, it is optimal to increase the price for the better action, causing a reduction in current agent's utility, but also a net gain, due to the information this action reveals. The addition of the social planner significantly improves social welfare and the asymptotic speed of learning.

Chapter 4 analyzes how different types of social connections between people shape their social networks. There are two possible types of ties between individuals, strong and weak, that differ in maintenance costs and reliability. A network formation game is played in which agents choose the number of ties of each type to maximize their chances of hearing about a new job opportunity. We find that in

equilibrium, people maintain both types of connections, which was not explained in previous theoretical models. Furthermore, in the socially optimal symmetric network, there are more strong ties than in the equilibrium one.

PUBLISHED CONTENT AND CONTRIBUTIONS

Hann-Caruthers, W., V. Martynov, and O. Tamuz (2018). “The speed of sequential asymptotic learning”. In: *Journal of Economic Theory* 173, pp. 383–409. DOI: https://doi.org/10.1016/j.jet.2017.11.009.

All authors conceived the study, worked on theorems and wrote the paper.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Chapter 1

INTRODUCTION

This thesis studies information aggregation and its effects on individuals' actions (Chapter 2, 3) and social network (Chapter 4). Chapters 2 and 3 analyze the speed of asymptotic social learning and how this process can be made more efficient from society's perspective. In Chapter 4, I analyze how information aggregation affects individuals' choices for different types of connections between people in their social network and rationalize why there is more than one type of links in the real-world networks.

In Chapter 2, with Omer Tamuz and Wade Hann-Caruthers, we consider a well-known model of sequential social learning introduced separately by Bihchandani, Hirshleifer, Welch (1992) and Banerjee (1992). In this study we calculate precisely the asymptotic speed of convergence of the log-likelihood belief to its true value ($+\infty$ or $-\infty$) when the strength of private signals is unbounded[1]. In doing so, we develop a novel technique that allows to approximate with a very high precision a highly non-linear recurrence equation. Furthermore, we solve the following open problem: is the expected time until people start taking the correct action finite or infinite? As it turns out, it depends on the distributions of individuals' private signals. We provide examples for both cases. Finally, we compare the learning speed to the one that we obtain when agents can observe not only actions of their predecessors, but also their actual signals. We show that in the latter case it is linear whereas in the former one it is always sublinear. Additionally, we can find a pair of signal distributions, one for each state of the world, such that the log-likelihood grows at a rate arbitrarily close to linear.

Chapter 3, which stems from the second one, studies how the aforementioned social learning mechanism can be made socially efficient. A problem arises when the log-likelihood ratio becomes very large in its absolute value. When this happens, we are unable to extract a lot of information about the private signal from the action that was taken, as we already anticipate the next person to follow the trend, given such a high prior belief. In order to solve this problem, I introduce a social planner who

[1] If their strength is bounded, then the log-likelihood is also bounded and cannot converge to its true value.

observes the sequence of actions taken by individuals. Based on its realization, she can set a price (positive or negative) for taking one of the actions[2]. This innovation does not change the ability to learn: if in the initial setup learning occurs (does not occur) with probability 1, then it also occurs (does not occur) under the new proposed mechanism. However, it does dramatically increase the social welfare and the accuracy of the final belief. In the case of bounded (binary) signals, I present a complete solution for the expected utility function and the pricing mechanism, which depends on the current sequence of actions. In the unbounded case, I prove that the general description of the social planner's strategy stays the same as in the binary case: she always sets a positive price for the better action (the one with higher likelihood belief). Furthermore, I present various numerical calculations of the optimal pricing strategy and the corresponding expected utility as a function of the public belief. Surprisingly, the new mechanism does not only increase the social welfare but also greatly improves (by more than 57%) the asymptotic learning speed as we get to extract more information from each individual action.

In Chapter 4, I investigate the underlying economic reason for why we observe different types of social connections between individuals in real life (strong and weak). In standard network models, if agents, who are facing a budget constraint, have a choice between different types of connections, they would choose the one type that maximizes their direct utility and completely disregard the rest of them. This raises the following question: is there an economic rationale for observing different types of ties between individuals in real life or is it just due to human psychological nature and there is no reason for economists to model them separately? To answer this question, I introduce a new model of network formation in which the underlying graph can have different clustering coefficients for various types of connections. This property of the model conforms with empirical data and provides a better intuition for the underlying process.

To elaborate more on this topic, imagine that there are two possible types of connections between individuals: strong – reliable and expensive, and weak – less reliable and cheap. This chapter makes two main contributions to the field. First, I give an economic explanation for why people do not choose only one type of ties by studying an equilibrium of this game. One of the largest shares of the agents' utility come from a group of people who are either weakly connected to some of her strong ties or who are strongly connected to some of her weak ties (so-called weak-strong

[2]This is equivalent to being able to set prices for both of them.

or strong-weak ties). To have a positive number of people in this group (weak-strong or strong-weak), there have to be both weak and strong types of connections present in the network.

Second, I compare this equilibrium network with a symmetric socially optimal one. I find that these two networks are very similar which means that agents are able to achieve an almost optimal outcome on their own. However, in the latter one agents have more strong ties. Intuitively, it happens because in the socially optimal case, people care not only about getting the signal, but also about sharing it with others, and the strong connections are more reliable sources of information.

Chapter 2

ASYMPTOTIC SPEED OF SOCIAL LEARNING

2.1 Introduction

When making decisions, we often rely on the decisions that others before us have made. Sequential learning models have been used to understand different phenomena that occur when many individuals make decisions based on the observed actions of others. These include herd behavior (cf. (Banerjee, 1992)), where many agents make the same choice, as well as informational cascades (e.g. (Bikhchandani, Hirshleifer, and Welch, 1992)), where the actions of the first few agents provide such compelling evidence that later agents no longer have incentive to consider their own private information.

Such results on how information aggregation can fail are complemented by results which demonstrate that when private signals are arbitrarily strong, learning is robust to this kind of collapse (L. Smith and Sørensen, 2000). In particular, in a process called asymptotic learning (see, e.g., (Acemoglu et al., 2011)), agents will eventually choose the correct action and their beliefs will converge to the truth. Two questions that has not been answered in the literature is: how quickly does this happen? And how does the speed of learning compare to a setting in which agents observe signals rather than actions?

We consider the classical setting of a binary state of nature and binary actions, where each of the two actions is optimal at one of the states. The agents receive private signals that are independent conditioned on the state. These signals are unbounded, in the sense that an agent's posterior belief regarding the state can be arbitrarily close to both 0 and 1. The agents are exogenously ordered, and, at each time period, a single agent takes an action, after observing the actions of her predecessors.

We measure the speed of learning by studying how the public belief evolves as more and more agents act. Consider an outside observer who observes the actions of the sequence of agents. The public belief is the posterior belief that such an outside observer assigns to the correct state of nature. It provides a measure of how well the population has learned the state. Since signals are unbounded, the public belief tends to 1 over time (L. Smith and Sørensen, 2000); equivalently, the corresponding log-likelihood ratio tends to infinity. As the outside observer may also be interested in learning the state, it is natural to ask how quickly she converges to the correct belief, and, in particular, to understand her asymptotic speed of learning when observing actions. Asymptotic rates of convergence are an important tool in the study of inference processes in statistical theory, and have also been studied in social learning models in the Economics literature (e.g., (Vives, 1993; Duffie and Manso, 2007; Duffie, Malamud, and Manso, 2009)).

When agents observe the *signals* (rather than actions) of all of their predecessors, this log-likelihood ratio is asymptotically linear. Thus, it cannot grow faster than linearly when the agents observe actions. Our first main finding is that when observing actions, the log-likelihood ratio always grows sub-linearly. Equivalently, the public belief converges sub-exponentially to 1. Our second main finding is that, depending on the choice of private signal distributions, the log-likelihood ratio can grow at a rate that is arbitrarily close to linear.

We next analyze the specific canonical case of Gaussian private signals. Here, we calculate precisely the asymptotic behavior of the log-likelihood ratio of the public belief. We show that learning from actions is significantly slower than learning from signals: the log-likelihood ratio behaves asymptotically as $\sqrt{\log t}$. To calculate this, we develop a technique that allows, much more generally, for the long-term evolution of the public belief to be calculated for a large class of signal distributions.

Since, in our setting of unbounded signals, agents eventually take the correct action, an additional, natural measure of the speed of learning is the expected time at which this happens: how long does it take until no more mistakes are made? We call this the *time to learn*.

We show that the expected time to learn depends crucially on the signal distributions. For distributions, such as the Gaussian, in which strong signals occur with very small probability, we show that the expected time to learn is infinite.[1] However, when strong signals are less rare, this expectation is finite.[2] Intuitively, when strong signals are rare, agents are more likely to emulate their predecessors, and so it may take a long time for a mistake to be corrected.

Finally, in the Gaussian case, we study another measure of the speed of learning. Namely, we consider directly how the probability of choosing the incorrect action varies as agents see more and more of the other agents' decisions before making

[1] In the benchmark case of observed signals this time is finite, for any signal distribution.

[2] This result disproves a conjecture of Sørensen (Sørensen, 1996, page 36).

their own. We find that this probability is asymptotically no less than $1/t^{1+\varepsilon}$ for any $\varepsilon > 0$. In contrast, when agents can observe the private signals of their predecessors, the probability of mistake decays exponentially, and so, also in this sense, learning from signals is much faster than learning from actions.

Related literature

Several previous studies have considered the same question. Chamley (Chamley, 2004) gives an estimate for the evolution of the public belief for a class of private signal distributions with fat tails. He also studies the speed of convergence in the Gaussian case using a computer simulation. Sørensen (Sørensen, 1996, Lemma 1.9) has published a claim related to our Theorem 1, with an unfinished proof. Also in (Sørensen, 1996), Sørensen shows that the expected time to learn is infinite for some signal distributions, and conjectures that it is always infinite, which we show to not be true. In (L. Smith and Sørensen, 1996), an early version of (L. Smith and Sørensen, 2000), the question of the time to learn is also addressed, and an example is given in which the time to learn is infinite, but is finite conditioned on one of the states. A concurrent paper by Rosenberg and Vieille (Rosenberg and Vieille, 2017) studies related questions. In particular, they study the time until the first correct action, as well as the number of incorrect actions—which are related to our time to learn—,and characterize when they have finite expectations.

A related model is studied by Lobel, Acemoglu, Dahleh and Ozdaglar (Lobel et al., 2009), who consider agents who also act sequentially, but do not observe all of their predecessors' actions. They study how the speed of learning varies with the network structure. Vives (Vives, 1993), in a paper with a very similar spirit to ours, studies the speed of sequential learning in a model with actions chosen from a continuum, and where agents observe a noisy signal about their predecessors' actions. He similarly shows that learning is significantly slower than in the benchmark case. An overview of this literature is given by Vives in his book (Vives, 2010, Chapter 6).

2.2 Model

Let $\theta \in \{-1, +1\}$ be the true state of the world, with each state a priori equally likely[3]. Each rational agent $t \in \{1, 2, \ldots\}$ receives a private signal s_t. The signals are i.i.d. conditioned on θ: if $\theta = +1$, they have cumulative distribution function (CDF)

[3]We make this simplification of a (1/2,1/2) prior to reduce the complexity of the presentation, but all results hold for general priors.

F_+ and if $\theta = -1$, they have CDF F_-.[4] We assume that F_+ and F_- are absolutely continuous with respect to each other, so that private signals never completely reveal the state.

Let

$$L_t = \log \frac{\mathbb{P}(\theta = +1|s_t)}{\mathbb{P}(\theta = -1|s_t)}$$

be the log-likelihood ratio of the belief induced by the agent's private signal. We assume that private signals are unbounded, in the sense that L_t is unbounded: for every $M \in \mathbb{R}$, the probability that $L_t > M$ is positive, as is the probability that $L_t < -M$. We denote by G_+ and G_- the conditional CDFs of L_t.

The agents act sequentially, with agent t acting after observing the actions of agents $\{1, \ldots, t-1\}$. The utility of the action $a_t \in \{-1, +1\}$ is 1 if $a_t = \theta$ and 0 otherwise.

Denote the public belief by

$$\mu_t = \mathbb{P}(\theta = +1|a_1, \ldots, a_{t-1}).$$

This is the posterior held by an outside observer after recording the actions of the first $t-1$ agents. We denote by ℓ_t the log-likelihood ratio of the public belief:

$$\ell_t = \log \frac{\mu_t}{1-\mu_t}.$$

In equilibrium, agent t chooses $a_t = +1$ iff[5]

$$\log \frac{\mathbb{P}(\theta = +1|a_1, \ldots, a_{t-1}, s_t)}{\mathbb{P}(\theta = -1|a_1, \ldots, a_{t-1}, s_t)} > 0.$$

A simple calculation shows that this occurs iff

$$\ell_t + L_t > 0.$$

Now, another straightforward calculation shows that when $a_t = +1$,

$$\ell_{t+1} = \ell_t + D_+(\ell_t), \tag{2.1}$$

where

$$D_+(x) = \log \frac{1 - G_+(-x)}{1 - G_-(-x)}.$$

[4] One could consider signals that take values in a general measurable space (rather than $\mathbb{R}$), but the choice of $\mathbb{R}$ is in fact without loss of generality, since all standard measurable spaces are isomorphic.

[5] For simplicity, we assume that agents choose action -1 when indifferent. This will have no impact on our results.

Likewise, when $a_t = -1$,

$$\ell_{t+1} = \ell_t + D_-(\ell_t),$$

where

$$D_-(x) = \log \frac{G_+(-x)}{G_-(-x)}.$$

We can interpret $D_+(\ell_t)$ and $D_-(\ell_t)$ as the contributions of agent t's action to the public belief.

2.3 The evolution of public belief

Consider a baseline model, in which each agent observes the private signals of all of her predecessors. In this case, the public log-likelihood ratio $\tilde{\ell}_t$ would equal the sum

$$\tilde{\ell}_t = \sum_{\tau=1}^{t} L_\tau.$$

Conditioned on the state, this is the sum of i.i.d. random variables, and so by the law of large numbers we have that the limit $\lim_t \tilde{\ell}_t/t$ would—conditioned on (say) $\theta = +1$—equal the conditional expectation of L_t, which is positive.[6]

Sub-linear public beliefs

Our first main result shows that when agents observe actions rather than signals, the public log-likelihood ratio grows sub-linearly, and so learning from actions is always slower than learning from signals.

Theorem 1. *It holds with probability 1 that* $\lim_t \ell_t/t = 0$.

Our second main result shows that, depending on the choice of private signal distributions, ℓ_t can grow at a rate that is arbitrarily close to linear: given any sub-linear function r_t, it is possible to find private signal distributions so that ℓ_t grows as fast as r_t.

Theorem 2. *For any* $r \colon \mathbb{N} \to \mathbb{R}_{>0}$, *such that* $\lim_t r_t/t = 0$ *there exists a choice of CDFs* F_- *and* F_+ *such that*

$$\liminf_{t\to\infty} \frac{|\ell_t|}{r_t} > 0$$

with probability 1.

[6] In fact, $\mathbb{E}(L_t|\theta = +1)$ is equal to the Kullback-Leibler divergence between F_+ and F_-, which is positive as long as the two distributions are different.

For example, for some choice of private signal distributions, ℓ_t grows asymptotically at least as fast as $t/\log t$, which is sub-linear but (perhaps) close to linear.

Long-term behavior of public beliefs

We next turn to estimating more precisely the long-term behavior of the public log-likelihood ratio ℓ_t. Since signals are unbounded, agents learn the state, so that ℓ_t tends to $+\infty$ if $\theta = +1$, and to $-\infty$ if $\theta = -1$. In particular ℓ_t stops changing sign from some t on with probability 1; all later agents choose the correct action.

We consider without loss of generality the case that $\theta = +1$, so that ℓ_t is positive from some t on. Thus, recalling (2.1), we have that from some t on,

$$\ell_{t+1} = \ell_t + D_+(\ell_t).$$

This is the recurrence relation that we need to solve in order to understand the long term evolution of ℓ_t. To this end, we consider the corresponding differential equation:

$$\frac{\mathrm{d}f}{\mathrm{d}t}(t) = D_+(f(t)).$$

Recall that G_- is the CDF of the private log-likelihood ratio L_t, conditioned on $\theta = -1$. We show (Lemma 27) that $D_+(x)$ is well approximated by $G_-(-x)$ for high x, in the sense that

$$\lim_{x\to\infty} \frac{D_+(x)}{G_-(-x)} = 1.$$

In some applications (including the Gaussian one, which we consider below), the expression for G_- is simpler than that for D_+, and so one can instead consider the differential equation

$$\frac{\mathrm{d}f}{\mathrm{d}t}(t) = G_-(-f(t)). \tag{2.2}$$

This equation can be solved analytically in many cases in which G_- has a simple form. For example, if $G_-(-x) = \mathrm{e}^{-x}$, then $f(t) = \log(t + c)$, and if $G_-(-x) = x^{-k}$ then $f(t) = ((k+1)\cdot t + c)^{1/(k+1)}$.

We show that solutions to this equation have the same long term behavior as ℓ_t, given that G_- satisfies some regularity conditions.

Theorem 3. *Suppose that G_- and G_+ are continuous, and that the left tail of G_- is convex and differentiable. Suppose also that $f\colon \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ satisfies*

$$\frac{\mathrm{d}f}{\mathrm{d}t}(t) = G_-(-f(t)) \tag{2.3}$$

for all sufficiently large t. Then conditional on $\theta = +1$,

$$\lim_{t\to\infty} \frac{\ell_t}{f(t)} = 1$$

with probability 1.

The condition[7] on G_- is satisfied when the random variables L_t (i.e., the log-likelihood ratios associated with the private signals), conditioned on $\theta = -1$, have a distribution with a probability density function that is monotone decreasing for all x less than some x_0. This is the case for the normal distribution and for practically every non-atomic distribution one may encounter in the standard probability and statistics literatures.

Gaussian signals

In the Gaussian case, F_+ is Normal with mean +1 and variance σ^2, and F_- is Normal with mean −1 and the same variance. A simple calculation shows that G_- is the Gaussian cumulative distribution function, and so we cannot solve the differential equation (2.2) analytically. However, we can bound $G_-(x)$ from above and from below by functions of the form $\mathrm{e}^{-c\cdot x^2}/x$. For these functions, the solution to (2.2) is of the form $f(t) = \sqrt{\log t}$, and so we can use Theorem 3 to deduce the following.

Theorem 4. *When private signals are Gaussian, then conditioned on $\theta = +1$,*

$$\lim_{t\to\infty} \frac{\ell_t}{(2\sqrt{2}/\sigma)\cdot\sqrt{\log t}} = 1$$

with probability 1.

Recall, that when private signals are observed, the public log-likelihood ratio ℓ_t is asymptotically *linear*. Thus, learning from actions is far slower than learning from signals in the Gaussian case.

The expected time to learn

When private signals are unbounded, then with probability 1 the agents eventually all choose the correct action $a_t = \theta$. A natural question is: how long does it take for that to happen? Formally, we define the *time to learn*

$$T_L = \min\{t \,:\, a_\tau = \theta \text{ for all } \tau \geq t\},$$

[7] By "the left tail of G_- is convex and differentiable", we mean that there is some x_0 such that, restricted to $(-\infty, x_0)$, G_- is convex and differentiable.

and study its expectation. Note that in the baseline case of observed signals T_L has finite expectation, since the probability of a mistake at time t decays exponentially with t.

We first study the expectation of T_L in the case of Gaussian signals. To this end, we define the *time of first mistake* by

$$T_1 = \min\{t \,:\, a_t \neq \theta\}$$

if $a_t \neq \theta$ for some t, and by $T_1 = 0$ otherwise. We calculate a lower bound for the distribution of T_1, showing that it decays at most as fast as $1/t$.

Theorem 5. *When private signals are Gaussian, then for every $\varepsilon > 0$ there exists a $k > 0$ such that for all t*

$$\mathbb{P}(T_1 = t) \geq \frac{k}{t^{1+\varepsilon}}.$$

Thus T_1 has a very thick tail, decaying far slower than the exponential decay of the baseline case. In particular, T_1 has infinite expectation, and so, since $T_L > T_1$, the expectation of the time to learn T_L is also infinite.

In contrast, we show that when private signals have thick tails—that is, when the probability of a strong signal vanishes slowly enough—, then the time to learn has finite expectation. In particular, we show this when the left tail of G_- and the right tail of G_+ are polynomial.[8]

Theorem 6. *Assume that $G_-(-x) = c \cdot x^{-k}$ and that $G_+(x) = 1 - c \cdot x^{-k}$ for some $c > 0$ and $k > 0$, and for all x greater than some x_0. Then $\mathbb{E}(T_L) < \infty$.*

An example of private signal distributions F_+ and F_- for which G_- and G_+ have this form is given by the probability density functions

$$f_-(x) = \begin{cases} c \cdot \mathrm{e}^{-x} x^{-k-1} & \text{when } 1 \leq x \\ 0 & \text{when } -1 < x < 1 \\ c \cdot (-x)^{-k-1} & \text{when } x \leq -1. \end{cases}$$

and $f_+(x) = f_-(-x)$, for an appropriate choice of normalizing constant $c > 0$. In this case, $G_-(-x) = 1 - G_+(x) = \frac{c}{k} x^{-k}$ for all $x > 1$.[9]

[8]Recall that G_- is the conditional cumulative distribution function of the private log-likelihood ratios L_t.

[9]Theorem 6 can be proved for other thick-tailed private signal distributions: for example, one could take different values of c and k for G_- and G_+, or one could replace their thick polynomial tails by even thicker logarithmic tails. For the sake of readability, we choose to focus on this case.

The proof of Theorem 6 is rather technically involved, and we provide here a rough sketch of the ideas behind it.

We say that there is an *upset* at time t if $a_{t-1} \neq a_t$. We denote by Ξ the random variable which assigns to each outcome the total number of upsets

$$\Xi = |\{t \,:\, a_{t-1} \neq a_t\}|.$$

We say that there is a *run* of length m from time t if $a_t = a_{t+1} = \cdots = a_{t+m-1}$. As we will condition on $\theta = +1$ in our analysis, we say that a run from time t is *good* if $a_t = 1$ and *bad* otherwise. A trivial but important observation is that the number of maximal finite runs is equal to the number of upsets, and so if $\Xi = n$, and if $T_L = t$, then there is at least one run of length t/n before time t. Qualitatively, this implies that if the number of upsets is small, and if the time to learn is large, then there is at least one long run before the time to learn.

We show that it is indeed unlikely that Ξ is large: the distribution of Ξ has an exponential tail. Incidentally, this holds for *any* private signal distribution:

Proposition 7. *For every private signal distribution, there exist $c > 0$ and $0 < \gamma < 1$ such that for all $n > 0$*

$$\mathbb{P}(\Xi \geq n) \leq c\gamma^n.$$

Intuitively, this holds because whenever an agent takes the correct action, there is a non-vanishing probability that all subsequent agents will also do so, and no more upsets will occur.

Thus, it is very unlikely that the number of upsets Ξ is large. As we observe above, when Ξ is small, then the time to learn T_L can only be large if at least one of the runs is long. When G_- has a thin tail, then this is possible; indeed, Theorem 5 shows that the first finite run has infinite expected length when private signals are Gaussian. However, when G_- has a thick, polynomial left tail of order x^{-k}, we show that it is very unlikely for any run to be long: the probability that there is a run of length n decays at least as fast as $\exp(-n^{k/(k+1)})$, and in particular runs have finite expected length. Intuitively, when strong signals are rare, then runs tend to be long as agents are likely to emulate their predecessor. Conversely, when strong signals are more likely, then agents are more likely to break a run, and so runs tend to be shorter.

Putting together these insights, we conclude that it is unlikely that there are many runs, and, in the polynomial signal case, it is unlikely that runs are long. Thus T_L has finite expectation.

Probability of taking the wrong action

Yet another natural metric of the speed of learning is the probability of mistake

$$p_t = \mathbb{P}(a_t \neq \theta).$$

Calculating the asymptotic behavior of p_t seems harder to tackle.

For the Gaussian case, while we cannot estimate p_t precisely, Theorem 5 immediately implies a lower bound: p_t is at least $k/t^{1+\varepsilon}$, for every $\varepsilon > 0$ and k that depends on ε. This is much larger than the exponentially vanishing probability of mistake in the revealed signal baseline case.

More generally, we can use Theorem 1 to show that p_t vanishes sub-exponentially for any signal distribution, in the sense that

$$\lim_{t\to\infty} \frac{1}{t} \log p_t = 0.$$

To see this, note that the probability of mistake at time $t-1$, conditioned on the observed actions, is exactly equal to

$$\min\{\mu_t, 1-\mu_t\};$$

where we recall that

$$\mu_t = \mathbb{P}(\theta = +1 | a_1, \ldots, a_{t-1}) = \frac{\mathrm{e}^{\ell_t}}{\mathrm{e}^{\ell_t}+1}$$

is the public belief. This is due to the fact that if the outside observer, who holds belief μ_t, had to choose an action, she would choose a_{t-1}, the action of the last player she observed, a player who has strictly more information than her. Thus

$$p_t = \mathbb{E}(\min\{\mu_t, 1-\mu_t\}) = \mathbb{E}\left(\frac{1}{\mathrm{e}^{|\ell_t|}+1}\right),$$

and since, by Theorem 1, $|\ell_t|$ is sub-linear, it follows that p_t is sub-exponential.

2.4 Conclusion

In this paper we consider a classical setting of sequential asymptotic learning from actions of others. We show that learning from actions is slow, as compared to the speed of learning when observing others' private signals, in the sense that the public log-likelihood ratio tends more slowly to infinity. However, it is possible to approach the linear rate of learning from signals and achieve any sub-linear rate.

We calculate the speed of learning precisely in the case of Normal private signals (among a large class of private signal distributions) and show that learning is very slow. We also show that in the Gaussian case the expected time to learn is infinite, as opposed to cases of more thick-tailed distributions, in which it is finite.

For the Gaussian case, we also provide a lower bound for the probability of mistake. Finding a matching upper bound seems beyond our reach at the moment, and provides a compelling open problem for further research.

Chapter 3

SUBSIDIZING LEMONS FOR EFFICIENT INFORMATION AGGREGATION

The aggregation of information in society is a complex and, at the same time, a very interesting process from an economic point of view. Especially it is interesting to see how it affects individuals' choices. People's decisions often rely on two types of information. The first one is their private knowledge about the choice they face. The second source is information received from society, in particular what other people did before. Social learning models have been used to analyze how people make decisions based on these types of information. They explain phenomena such as herding (Banerjee, 1992), informational cascades (Bikhchandani, Hirshleifer, and Welch, 1992), and asymptotic learning (L. Smith and Sørensen, 2000).

Usually, these models have huge inefficiencies. For example, people might end up in the wrong cascade. Furthermore, when herding occurs, people's actions convey much less information about their private signals than they do in the beginning. One of the reasons this happens is because people do not take into account how their actions affect future generations' utility. The main aftermath of this is a decrease in social welfare and in the asymptotic speed of learning (Hann-Caruthers, Martynov, and Tamuz, 2018). A natural question arises: how can we improve this social learning system?

In this paper, we start with the classical sequential learning model with binary states of nature and corresponding binary actions. Our leading example is the case in which there are two goods/technologies on the market: the old one, whose characteristics are well-known to everybody, and the new one, which is better than the former one in one state of the world and worse in the other. The actions here are represented by buying either the new or the old product. Agents have the same preferences: they want their actions to match the underlying state. Players get private signals that are i.i.d. conditioned on the realized state of the world. The players are exogenously ordered and, in every time period, one agent chooses which product to buy based on her private signal and the actions of her predecessors.

A social planner would like to maximize the discounted sum of players' utilities. She can choose relative prices, by taxing or subsidizing the goods in every period, that are publicly observed. In other words, she chooses a cost which can be positive, negative, or zero, of the new product, normalizing the price of the old one to be 0. The social planner's choice is based on the public information available at that time: prices in previous periods and corresponding players' actions.

We find the optimal pricing policy which, perhaps, is counterintuitive: tax the good that is more likely to be better. This results in an expected loss today, because agents whose private information is not strong enough are less likely to buy the better product. But at the same time, their choice will give the social planner more information about their private signals, which in its turn significantly increases the utility of future generations.

As a motivational example, imagine there are two drugs that treat the same disease but have unknown benefits/side effects, as they were tested only on a small sample of population. The only way the government can collect information about the real effects is by observing which one people bought given their personal and public knowledge. Here, the government plays the role of the social planner who can tax or subsidize one drug or the other. Without it, there is a significant risk that society will stick with the wrong medicine or that it will learn the truth very slowly. We show how these risks become lower when the social planner is involved in this process.

Our model overcomes a major criticism of the sequential social learning model—the assumption that players have to know the order in which the actions were taken by their predecessors, and then make a complicated calculation by Bayes rule to obtain the current public belief. In our setup, the optimal prices are functions of the public belief, which contains all the necessary information about the past. In equilibrium agents can recover the public belief from the price, and do not need to know the actions and specific order of their predecessors. Thus, agents only need to observe the current price and private signals to choose the optimal action.

We start with introducing the model, describing individuals' and the social planner's behaviors, and providing an intuitive illustration of how prices change the belief - action relationship. After this, we go to the binary signals case and calculate the expected utility function and its asymptotic characteristics when δ goes to 1. On the way to these results, we establish a few interesting and helpful properties of the binary case. We show that the public belief is a random walk that depends on the difference between the number of $High$ and Low signals up to this point. Another property is that when an agent takes into account their private signal, their expected utility is equal to the signals' precision.

In the next section, we study continuous distributions: bounded and unbounded. We find that there is still a difference in terms of asymptotic learning: when private signals have bounded strength, the underlying state is never revealed. However, the beliefs at which learning stops are closer to 0 and 1 than they would be without the social planner. For general signal distributions, this problem becomes very complex and we can not find the analytical solution. But we provide a good qualitative description for a specific unbounded signal distribution. This description looks similar to the one in the binary case: when one product seems to be better, it should be taxed in order to extract more information in the future periods and exploit the convexity of the expected utility function. Moreover, the prices are bounded away from 0 and 1 when the public belief converges to one of the extreme values. This implies that eventually people are going to buy the right product with probabilities close to 1, but at the same time, it is going to be positively priced to increase the social benefits. Finally, we calculate numerically the expected utility and the optimal pricing function for this signal distribution.

Literature review

Several previous studies consider a similar question: how prices might affect social learning in different scenarios.

Crapis, Ifrach, Maglaras, Scarsini (Crapis et al., 2016) consider a situation when people with heterogeneous preferences observe not only actions of their predecessors, but also their reviews of the product, the outcome, in a non-Baysian framework. Numerical experiments suggest that pricing policies that account for social learning may increase revenues considerably relative to policies that do not.

Papanastasiou, Savva (Papanastasiou and Savva, 2016) and Bhalla (Bhalla, 2012) allow both buyers and a monopolist to act strategicly over a finite number of time periods. The first one finds that the social learning increases the firm's expected profit and contrary to previous results in the literature, preannounced prices are not beneficial to the firm. The second one shows that prices are no longer submartingales, but that for some range of beliefs they can be super-martingales, too.

In a more classic setup, Bose, Orosel, Ottaviani, and Vesterlund (Bose et al., 2006) consider a binary model when a monopolists chooses pricing strategy in order to maximize its revenue and incurs some cost to produce the product. They find some qualitative results. For example the objective is convex, increases in number of periods the game is played, and that at some point herding occurs.

One more working paper (L. Smith, Sørensen, and Tian, 2020) studies optimal experimentation in herding. They have a similar setup where they try to maximize discounted sum of agents' utilities by introducing "an infinite-lived planner who devises individual choice rule". They find that planner's cutoff private belief increases with the increase in public belief (when signals satisfy log-concavity condition), which conforms with our finding that optimal prices rise in the public belief. They also suggest a mechanism which allows to make this socially optimal process decentralized, e.g. allows to get rid off the planner. However, they do not provide solutions for the optimal planner's choice rules or the value function as it has been technically challenging for a long time. We believe that these results, as well as other interesting comparative statics that are done in this paper, are an essential part if we want to apply this useful and interesting theory to real life applications or experiments.

There is also some literature on optimal pricing in networks. Candogan, Bimpikis and Ozdaglar (Candogan, Bimpikis, and Ozdaglar, 2012) study the optimal pricing strategy of a monopolist in a two-period game, where an agent's utility depends not only on her action, but also on decisions of her neighbors. They find that optimal price should depend on Bonacich centrality, on a markup term proportional to the influence that the network exerts and on a term that is independent of the network.

Another paper in this area is by Campbel (Campbell, 2013). He models a firm's ability to strategically influence the probability individual engages in WOM, Word of Mouth, through the price. The author derives the comparative static results of connectivity, mean-preserving spread of friendships, and clustering of friends on price.

3.1 Model

Let $\theta \in \{High, Low\}$ be the true state of the world, where both states a priori are equally likely.[1] At the beginning of the game, one state is realized and does not change.

There are countably many rational agents $t \in \{1, 2, \dots\}$, who receive private signals s_t. These signals are i.i.d. conditional on the state of the world: if $\theta = High$, they have cumulative distribution function (CDF) G_H and if $\theta = Low$ — G_L. The corresponding PDFs are g_H and g_L. We assume that signals never completely reveal

[1] We make this simplification of a (1/2,1/2) prior to reduce the complexity of the presentation, but all results hold for general priors

the true state, which is the same as saying that the conditional distributions are absolutely continuous with respect to each other.

Suppose there are two products on the market with different prices in each time period: the new one and the old one. When $\theta = High$, the new product is better than the old one and when $\theta = Low$ — vise versa. Each agent t has a decision to make: whether to buy the new product ($a_t = 1$) and pay the price c_t or to stick with the old one ($a_t = 0$). We normalize the price of the old product to be equal to 0. Utility of each agent from the action a_t is 1, if it matches the state ($a_t = 0$ and $\theta = Low$ or $a_t = 1$ and $\theta = High$)), minus the cost, if she buys the new product, -$\mathbb{1}\{a_t = 1\}c_t$. Players act sequentially and their decisions are based on two types of information: private information from their own private signal s_t and public information (history) h_t. The latter one includes actions that were taken before player t and the sequence of prices of the new good c_t, so $h_t = \{a_1, a_2, \ldots a_{t-1}, c_1, c_2, \ldots c_{t-1}\}$. Before we explain what is the real role of c_t, we need some notation.

Agent's decision process

Denote the posterior belief of the agent t that the new good is better by

$$\mu_t = \mathbb{P}(\theta = High | h_t, s_t).$$

We will also refer to it as the *total belief* as it combines public and private information. Also, let us call the corresponding likelihood ratio, $\mu_t/(1 - \mu_t)$, — the *total likelihood ratio*.

Note that μ_t also represents the expected utility of player t for taking action 1, not including the price. As h_t and s_t are independent of each other, the posterior belief has two components: the private belief $\mathbb{P}(\theta = 1 | s_t)$, which is known only to player t, and the public belief $p_t = \mathbb{P}(\theta = 1 | h_t)$, which is known to everyone who observed history up to time t. Also, denote by l_t the likelihood ratio of the public belief

$$l_t = \frac{p_t}{1 - p_t}.$$

We can see that there is a monotone bijection between the public belief and its likelihood ratio and we are going to use them interchangeably. As $p_t \in [0, 1]$ then $l_t \in [0, \infty]$. Also denote by $F_H(l_t)$, $F_L(l_t)$ the CDF's of l_t conditional on θ.

The posterior belief μ_t captures how confident the player is about buying the new product. We obtain p_t from $p_0 = 1/2$ and h_t using Bayes rule.

Now let us go back to c_t. Suppose that $c_t = 0$. As agents are expected utility maximizers, player t buys the new product if her posterior belief μ_t is greater than $1/2$. Now, if $c_t \neq 0$, then she buys it only if $\mu_t \geq 1/2 + c_t/2$: if she buys the new product then her expected utility is $\mu_t - c_t$ which has to be greater or equal to $1 - \mu_t$ — the utility from buying the old product. In other words, these prices reflect how confident you should be in the new product in comparison to the old one, in order for you to buy the former one.

We can summarize it in the following way

$$\begin{cases} a_t = 0, \text{ utility } 1 - \mu_t & \text{when } \mu_t < \frac{1}{2} + \frac{c_t}{2} \\ a_t = 1, \text{ utility } \mu_t - c_t & \text{when } \mu_t \geq \frac{1}{2} + \frac{c_t}{2} \end{cases}$$

The agent wants to guess the correct state of the world ($a_t = \theta$) in general. But there are some situations when it is more profitable to take the opposite action in order to avoid the cost c_t. Imagine that $\mu_t = 0.57$ and $c_t = 0.15$. Even though the total belief tells us to buy the new product ($\mu_t > 0.5$), we would get less utility by doing this ($0.57 - 0.15 = 0.42$) rather than buying the old one ($1 - 0.57 = 0.43$). Hence, the price forces some people with a not very strong belief to switch to a "non-optimal" action, while people with a strong belief are not affected by it.

We can also state this condition in terms of the likelihood ratio: $a_t = 1$ iff[2]

$$\frac{\mu_t}{1 - \mu_t} = \frac{p_t}{1 - p_t} \frac{g_H(s_t)}{g_L(s_t)} \geq k_t,$$

where

$$k_t = \frac{1 + c_t/2}{1 - c_t/2}.$$

If $c_t = 0$ ($k_t = 1$), we get the usual conditions for taking action 1 .We are going to use both c_t (for beliefs) and k_t (for likelihood ratios) as prices but in different settings. Notice that we can rewrite the condition above as follows

$$\frac{p_t}{(1 - p_t)k_t} \cdot \frac{g_H(s_t)}{g_L(s_t)} \geq 1.$$

We can interpret this as if there were no price, $k_t = 1$, but we had a lower public belief that corresponds to the likelihood ratio $p_t/((1 - p_t)k_t)$.

[2] For simplicity, we assume that agents choose action 1 when indifferent. This will have no impact on our results.

Definition 1. *Let us call $p_t/((1-p_t)k_t)$ — modified likelihood ratio. Similarly, we call the public belief that corresponds to the modified likelihood ratio — modified public belief.*

Last thing we mention here is how the public belief evolves after another player takes an action. If at time t the public belief is p_t and the price is k_t, then the public belief at time $t+1$ after taking action $a_t = 1$ satisfies the following formula

$$l_{t+1} = \frac{p_{t+1}}{1-p_{t+1}} = \frac{p_t}{1-p_t} \cdot \frac{1-F_H(y)}{1-F_L(y)},$$

where $y = (1-p_t)k_t/p_t$ and F_i - CDFs of the likelihood ratio conditioned on the true state. And if $a_t = 0$,

$$l_{t+1} = \frac{p_{t+1}}{1-p_{t+1}} = \frac{p_t}{1-p_t} \cdot \frac{F_H(y)}{F_L(y)}.$$

This is an application of the Bayes rule, given that player t buys (does not buy) the new product if her total likelihood ratio is above (below) k_t. This implies that her private likelihood ratio is above (below) $(1-p_t)k_t/p_t$.

Now let us look at this game from the social planner's perspective, who would like to maximize the discounted sum of the expected utilities of agents with a discount factor δ.

Social planner

We introduce a long run risk neutral social planner who chooses a price, by taxing or subsidizing the new product, in each period and then returns the collected money to players in the following way. If in time period $t-1$ c_{t-1} was collected from player $t-1$, then this money is put in the bank with the interest rate $1/\delta$ and returned to player t in the next period. This mechanism implies two things. First, redistribution does not affect players' decisions and they still act according to the previous subsection. This is true because player's t choice does not change how much money she gets back; this is already decided by the previous player. Second, the total amount of money that is taken from the players is 0, budget-balanced. Thus, the discounted sum of players' payoff is equal to the discounted sum of $u_t(\mu_t, c_t) = \mathbb{1}\{a_t(\mu_t, c_t) = 1\}\mu_t + \mathbb{1}\{a_t(\mu_t, c_t) = 0\}(1-\mu_t)$. Therefore, the utility function of the social planner who implements a pricing strategy $\{c_t\}_{t=1}^{\infty}$ is

$$u_{\{k_t\}}(p) = \sum_{t=1}^{\infty} u_t(\mu_t(p, h_{t-1}, s_t), k_t),$$

where $\mu_t(p, h_{t-1}, s_t)$ means that it is path dependent.

Due to the money redistribution defined above, if $\mu_t > 1/2$ and player t buys the new product, we treat her utility from society's prospective as just μ_t instead of subtracting the price, as it is returned later. But from the players perspective, they still take into account the price.

The optimal pricing strategy plays another crucial role in this model. Imagine that the social planner does not do anything ($\forall t\ k_t = 1$), then when the herd occurs, the difference between public beliefs in two consequitive periods converges to 0 extremely quickly. This is shown in (Hann-Caruthers, Martynov, and Tamuz, 2018). Without loss of generality, assume that they are buying the new product. But if we set a high price for the new product and observe that people are still buying it, then the posterior belief would grow faster than before. In other words, the asymptotic speed of learning would increase.

Let us denote by $u(p_t)$ the expected utility of the social planner with the public belief p_t when we use the optimal prices k_t^*. Due to the stationarity of the social planner's problem, the expected utility depends on the history only through the public belief. Then it should satisfy the following Bellman equation:

$$u(p_t) = \max_{k_t}\{(1-\delta)(\text{expected gain}(p_t) - \text{expected loss}(k_t, p_t)) + \delta\, \mathbb{E}\, u(p_{t+1})\}. \tag{3.1}$$

The expected gain calculates the expected utility in this period if there were no price, $\mathbb{E}\, u_t(\mu_t, 1)$. First, we calculate the expected μ_t, and then the expected payoff is equal to $\max(\mu_t, 1-\mu_t)$. If we receive a private signal with a likelihood ratio x, then the total likelihood ratio is equal to $y = p_t x/(1-p_t)$ and therefore, $\mu_t = y/(y+1)$. Depending on whether the corresponding likelihood ratio is above or below 1, the expected utility is either equal to μ_t or $1-\mu_t$, so we divide this into two cases: $0 \le x \le (1-p_t)/p_t$ and $x > (1-p_t)/p_t$. Therefore, the expected gain(p_t) is equal to

$$\int_0^{\frac{1-p_t}{p_t}} \left(1 - \frac{x\frac{p_t}{1-p_t}}{x\frac{p_t}{1-p_t}+1}\right)(p_t f_H(x) + (1-p_t) f_L(x))dx +$$
$$+ \int_{\frac{1-p_t}{p_t}}^{\infty} \left(\frac{x\frac{p_t}{1-p_t}}{x\frac{p_t}{1-p_t}+1}\right)(p_t f_H(x) + (1-p_t) f_L(x))dx.$$

The second term, the expected loss(k_t, p_t), calculates utility that agent loses due to a non-optimal action. Notice that if we apply price k_t, then the only loss that can

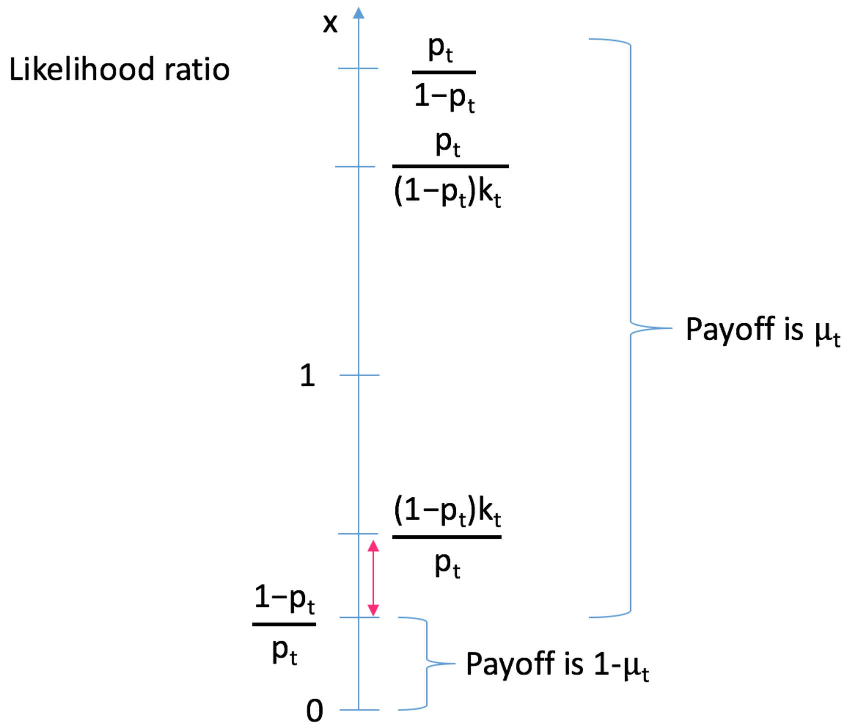

Red area corresponds to the private likelihood ratios when people take non-optimal actions. Here, $\mu_t = (x\frac{p_t}{1-p_t})/(x\frac{p_t}{1-p_t}+1)$.

Figure 3.1: Private likelihood ratios and today payoff.

occur is when the likelihood ratio of the private signal is strong enough to make the total likelihood ratio less than 1 for the modified likelihood ratio $(1-p_t)k_t/p_t$, but not for the initial one, $(1-p_t)/p_t$. When this happens, the likelihood ratio of the private signal x can be between $(1-p_t)/p_t$ and $(1-p_t)k_t/p_t$. If the total belief is μ_t, then the loss that player bears is $\mu_t-(1-\mu_t)=2\mu_t-1$. Therefore, the expected loss(k_t, p_t) is equal to

$$\int_a^b \left(\frac{2x\frac{p_t}{1-p_t}}{x\frac{p_t}{1-p_t}+1}-1\right)(p_t f_H(x)+(1-p_t)f_L(x))dx,$$

where $a=(1-p_t)/p_t$, $b=(1-p_t)k_t/p_t$.

The high complexity of this problem is due to the intricacy of the random walk of the public belief. Equation (3.1) shows the trade-off between losing some utility today due to the fact that some people (whose posterior likelihood is between 1 and k_t) take a non-optimal action (and paying $(2\mu_t-1)$ for this) and gaining utility through more of a disperse belief tomorrow. The latter one occurs due to the convexity of u which we prove a bit later.

In other words, strategic pricing can help to aggregate information more efficiently and increase the social welfare as well as speed up the asymptotic learning.

In the subsequent sections, we are going to consider both discrete and continuous, bounded and unbounded private signals and analyze how it affects properties of the optimal pricing policy and the corresponding solution. But before doing this, we are going to state a general property of the expected utility function.

Proposition 8. *The expected utility function $u(p)$ is convex, $u(1) = u(0) = 1$, and is symmetric around $1/2$, i.e $u(p) = u(1 - p)$.*

Most of the time, we are going to assume, without loss of generality, that $p_t \geq 0.5$, so the new product is a priory better, and will try to find the optimal price or at least a price that is better than no price at all. For $p_t < 0.5$, a symmetric analysis applies.

3.2 Binary signals

In this section, we are going to calculate the utility and the optimal strategy of the social planner when private signals are Bernoulli distributed. This means that each agent is going to be told whether the state is $High$ or Low as her private signal, and this information is going to be correct with probability $q > 1/2$.

Definition 2. *Define a* ***learning period****, LP_t to be all periods t' up to time t, conditioned on h_t, such that a player t' took into account her private signal when she chose an action $a_{t'}$.*

We abuse notation a bit and ignore the subscript t as it is going to be clear which period we have in mind. If people disregard their private signals when they take actions, then the public belief does not change, so in this sense, we *do not learn* in these periods. It is helpful to keep in mind this distinction between LP and not LP.

There are a few nice characteristics of the binary distribution. The first one is that if t is in the learning period, then agent t's action reveals her private signal. Indeed, if we know that people are going to take into account their private signals, then the only possibility is that they act according to them.

Lemma 9. *If $t \in LP$, then player t's action, a_t, reveals her private signal s_t.*

Thus, during the learning period actions not only convey but also reveal the information.

The second one is that there are only two nontrivial multipliers by which we update the likelihood ratio: either we observe the $High$ signal and multiply l_t by

$$\frac{\mathbb{P}(\theta = High|s_t = High)}{\mathbb{P}(\theta = Low|s_t = High)} = \frac{q}{1-q},$$

or we observe the Low signal and then multiply l_t by

$$\frac{\mathbb{P}(\theta = High|s_t = Low)}{\mathbb{P}(\theta = Low|s_t = Low)} = \frac{1-q}{q}.$$

When player t ignores s_t, we do not update the public belief, so l_t is multiplied by 1, and from now on all claims about the public belief assume that actions depend on the private signals. Thus, if player t's action is informative, then l_t is multiplied by either $q/(1-q)$ or $(1-q)/q$ depending on her action.

The third, and the best characteristic of the binary case is summarized in the following lemma.

Lemma 10. *In the binary case, the public belief is a random walk and its position depends on the difference in the number of times we observed $High$ and Low signals. Moreover, if we observed n $High$ signals and k Low ones between periods t and $t+n+k$, then the likelihood ratio at time $t+n+k$ is*

$$l_{n+k} = l_t \cdot \left(\frac{q}{1-q}\right)^{n-k}$$

Lemma 10 tells us that in order to find l_t, we just need to calculate the difference in number of times agents took actions 1 and 0 during the LP and take $q/(1-q)$ to the corresponding power.

Definition 3. *We say, the public belief (LR)* ***goes up*** *if we observed the $High$ signal during LP. Analogously, the public belief (LR)* ***goes down*** *if we observed the Low one.*

Notice that the public belief increases (decreases) when it goes up (down) as $q > 1/2$.

If at time t we have $p_t = 1/2$ and we go n times up, then by Bayes rule the public belief is

$$p_{t+n} = \frac{q^n}{q^n + (1-q)^n} \tag{3.2}$$

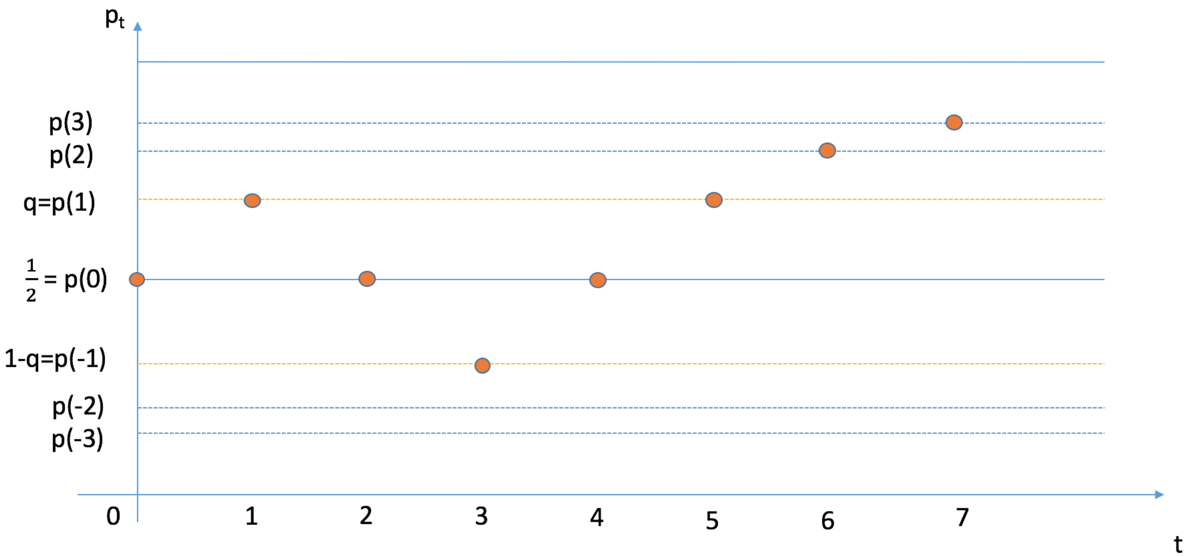

Where $p(k) = \frac{q^k}{q^k+(1-q)^k}$ and $p(-k) = \frac{(1-q)^k}{q^k+(1-q)^k}$ as in (3.4).

Figure 3.2: Random walk of the public belief with binary signals.

and if we go down n times from $p_t = 1/2$, then

$$p_{t+n} = \frac{(1-q)^n}{q^n + (1-q)^n}. \tag{3.3}$$

Definition 4. *For $k \geq 0$, define **level** k to be the public belief if it observed k more High signals than Low ones and denote it by $p(k)$. Similarly, define **level** $(-k)$ to be the public belief if we observed k more Low than High ones and denote it by $p(-k)$*

$$\begin{cases} p(k) = \dfrac{q^k}{q^k + (1-q)^k} \\ p(-k) = \dfrac{(1-q)^k}{q^k + (1-q)^k} \end{cases} \tag{3.4}$$

where $k > 0$.

For example, public belief $1/2$ corresponds to level 0.

The last property we mention is that if $k_t = 1$, then agents start ignoring their private signals (learning period stops) (a cascade[3] occurs) after one or two people take the

[3]This is an event in the classical model when people disregard their private signals and take the same action, thus public belief does not update after this. Once it has started, it does not stop. Although in our model we are able to stop it by introducing a price, which brings the public belief back to the region where agents act according to their private signals

same action. If we start with $p_t = 1/2$ and someone takes action 1, p_{t+1} will be equal to q. Even if the next player gets Low private signal, her posterior belief is going to be $1/2$, hence, she would take action 1 disregarding her private signal. This means we stop aggregating private information extremely quickly and therefore unable to get to a high public belief, which hurts social welfare.

To improve this situation, by maximizing the social welfare, we are going to implement the optimal pricing scheme. The following lemma will help in our analysis.

Lemma 11. *There exists an optimal strategy of the following form: the social planner* 1) *picks* $N \in \mathbb{N}$*, and* 2) *chooses prices such that actions reveal the private signals until the public belief reaches either level* N *or* $-N$*. After that there are no prices,* $k_t = 1$.

Notice that from the social planner's perspective this game is stationary. Thus, in this binary situation, there is only one way how she can affect the outcome: she can choose how long we are going to distinguish signals, so how long does the learning period continue. As we said above, due to the stationarity, there is no incentive to choose a non-zero price c_t once the learning period has ended. Notice that since we are not biased towards one or the other state, these two levels should be symmetric around $1/2$: upper bound — level N, lower bound — $-N$.

Thus, we pick $N \in \mathbb{N}$ at the beginning, and then in every period, until p_t hits either level N or $-N$, choose a price which allows us to separate $High$ and Low private signals. This is the pricing scheme.

How exactly do these prices look like? Suppose, at time t the likelihood ratio $l_t > 1$ (p_t is greater than $1/2$), which we assume is above $q/(1-q)$ (otherwise, do not need any price, we are still going to figure out the private signal from the action), then the social planner chooses a price $k_t > 1$ such that

$$\frac{l_t}{k_t} = \frac{q}{1-q} - \varepsilon > \frac{1}{2},$$

for a small, positive ε.

Now, we go back to the initial problem (3.1). Recall that one of our goals is to maximize the social welfare. To do this, first we need to calculate our expected utility from starting at the public belief $1/2$, $u(0.5)$, when private information is acquired until p_t hits one of the barriers that are at distance N from the initial belief $1/2$. And we are interested in N which maximizes $u(0.5)$.

Let us start with a few observations.

Lemma 12. *The expected utility that the social planner gets today in the learning period is q.*

There is a simple intuition behind this. As our action follows our signal in LP, we take the right action only with probability q. If the true state of the world is $High$, we receive the corresponding signal only with probability q. Hence, we are going to make the right action with the same probability, which equals to our expected utility today.

Now, let us go back to our Bellman equation and rewrite it for this case

$$u(p(n)) = (1-\delta)q + \delta\left(\mathbb{P}(\text{signal } High)u(p(n+1)) + \mathbb{P}(\text{signal } Low)u(p(n-1))\right), \tag{3.5}$$

where $\mathbb{P}(\text{signal } High) = (p_t q + (1-p_t)(1-q))$ and $\mathbb{P}(\text{signal } Low) = p_t(1-q) + (1-p_t)q$. We choose N in order to maximize utility at $1/2$, $p(0)$.

This is better than the general form (3.1) but still complicated, as the probability of going up or down depends on the current public belief.

Fortunately, in order to calculate $u(p(N))$, the analysis can be simplified. Recall that the only thing that the social planner controls is how far away are the absorbing boundaries from $1/2$, in other words she chooses N. After it is fixed, with probability $1/2$ we are going to be in the $High$ state and the probability of going up is just q instead of $p_t q + (1-p_t)(1-q)$, and the probability of going down is $(1-q)$ instead of $(1-p_t)q + p_t(1-q)$. Also, with probability $1/2$, we are in the Low state and again can simplify recurrence relation (3.5). Therefore, our utility at $1/2$ is going to be the average of $u_H(p(N))$ and $u_L(p(N))$, where $u_i(p(N))$ are defined by the following recurrence problems: in the $High$ state

$$\begin{cases} u_H(p(k)) = (1-\delta)q + \delta(q u_H(p(k+1)) + (1-q)u_H(p(k-1))) \\ u_H(p(2N)) = p(N) \ = \frac{q^N}{q^N+(1-q)^N} \\ u_H(p(0)) = p(-N) \ = \frac{(1-q)^N}{q^N+(1-q)^N} \end{cases} \tag{3.6}$$

and in the Low state

$$\begin{cases} u_L(p(k)) = (1-\delta)q + \delta((1-q)u_L(p(k+1)) + q u_L(p(k-1))) \\ u_L(p(N)) = 1 - p(N) \ = \frac{(1-q)^N}{q^N+(1-q)^N} \\ u_L(p(-N)) = 1 - p(-N) \ = \frac{q^N}{q^N+(1-q)^N} \end{cases} \tag{3.7}$$

Here, the boundary conditions come from the fact that when we reach levels N or $-N$, we stop learning and in every period just receive expected utility that is equal to the belief. To sum up the paragraph above, if we can solve two problems (3.6) and (3.7), we can get $u(p(0))$.

Theorem 13. *The expected utility of public belief* $1/2$ *when we stop learning upon arrival at levels* $2N$ *or* 0 *has the following form*

$$u\left(\frac{1}{2}\right) = \frac{\left(\frac{q^N}{q^N+(1-q)^N} - q\right) - \left(-\frac{(1-q)^N}{(1-q)^N+q^N} + q\right)\left(\frac{1-q}{q}\right)^N}{a_1^N + a_2^N} + q,$$

and

$$a_i = \frac{1}{2}\left(\frac{1}{\delta q} \pm \sqrt{\frac{1}{\delta^2 q^2} - \frac{4}{q} + 4}\right).$$

To prove this theorem, we use recurrence and linear equations, techniques. To get the utility function for some other initial level k (with belief $p(k)$), the same technique can be applied.

Even though this may not look very friendly, the numerator has a nice interpretation. The first term is a difference between the public belief at level N and our precision q, which is also our belief at level 1. Similarly, the first multiplier of the second term is a difference between our precision and the public belief at level $-N$. And the second multiplier is the likelihood ratio of observing N signals Low.

Given Theorem 13, we can solve for the optimal N and the utility at $1/2$ for any q and δ. For example, when $\delta = 0.9$, $q = 0.7$ the expected utility at $1/2$ as a function of stopping level N is depicted in Figure 3.3. The optimal $N^* = 4$ and $u(0.5) = 0.802$ in this case. If we do not have any prices, this expected utility is 0.7. Furthermore, in the optimal pricing case we end up with probability around $p_t = 0.97$ that we choose the right action (stopping public belief), comparing it to no price case where we end up with $p_t = 0.7$.

Moreover, it is going to be shown in the next subsection that for a fixed δ, the optimal N^* does not go to infinity no matter how we change q. This implies that it is impossible to learn the underlying state unless $\delta = 1$.

Furthermore, we not only increase the social welfare, but also end up with a much higher public belief about one of the states. In this case, instead of stopping at $q = 0.7$, we are going to stop at $p_t = 0.9674$ conditioned on the public belief being

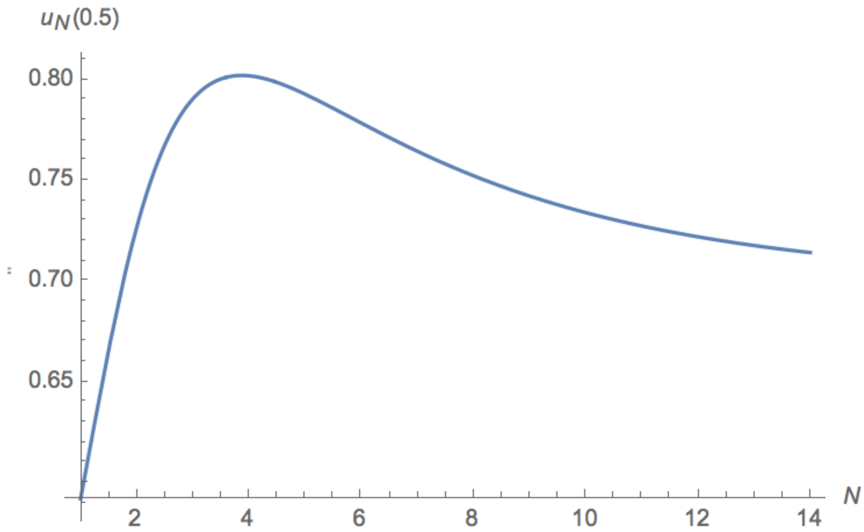

Figure 3.3: Expected utility at $p_0 = 1/2$ as a function of stopping time N.

absorbed at the upper bound. Recall, that if we do not have any price, then the random walk stops after one (two) steps because a cascade occurs. Given that we have a drift towards the underlying state, making boundaries further away from $1/2$ significantly decreases probability of ending up in the wrong cascade. We have this lemma to formalize it.

Lemma 14. *If it is optimal to stop the learning phase upon reaching level N or $-N$ then the probability of ending up in the wrong cascade is*

$$\frac{\left(\frac{1-q}{q}\right)^{2N} - \left(\frac{1-q}{q}\right)^{N}}{\left(\frac{1-q}{q}\right)^{2N} - 1}$$

Again, for $q = 0.7$, $\delta = 0.9$ the optimal N is 4 and so the probability of ending up in the wrong cascade is about 0.03, compared to $1 - q = 0.3$ in a situation without prices. It is easy to see that this probability goes to 0 as $((1-q)/q)^N$. Thus, acquiring information for a few steps dramatically decreases the probability of the wrong cascade.

One more thing should be mentioned before we start analyzing asymptotic properties of $u(p)$, when $\delta \to 1$. As we vary the stopping level N, the expected utilities at different levels increase (decrease) simultaneously. This again happens due to stationarity of the social planner's problem.

Proposition 15. *If $u(p(i))$ changes when we vary N, then it increases (decreases) iff $u(p(0))$ increases (decreases).*

In the proceeding section, we are going to analyze behavior of the optimal N and corresponding $u(0.5)$ when the social planner becomes patient, i.e. $\delta \to 1$.

Patient planner

It seems unlikely that there exists a closed form solution for the optimal N due to the complex form of u and because we are maximizing over natural numbers, which automatically prevents us from using the usual techniques. With numerical calculations, this optimal N can be easily found.

However, we can understand the behavior of N for a patient social planner. Let us unfold this statement. Consider the function $u(0.5)$ from Theorem 13 when q is close to 1. If δ is bounded away from 1, then $a_2 > 1$ and as $a_1 > 0$ and all terms in nominator are bounded by $1 - q$ $\forall N$, hence, N^* does not go to infinity (stays "finite"). This is a very natural result as if the precision of the private signals is high then we are more likely to get to a high public belief quicker and lose a significant utility for waiting extra rounds. If $q = 0.999$, then, after observing just one $High$ signal, public belief becomes $p_1 = 0.999$ and we know it can never be above 1, so there is no incentive to wait for more signals, the social planner should stop learning now.

On the other hand, if the social planner is extremely patient, i.e. δ goes to 1, then the denominator converges to 1 and the nominator is an increasing function of N, so N^* goes to ∞. There is also a good intuition for this fact. If he is very patient his utility today matters less, and also gaining some small amount, i.e. 0.01 extra, in the public belief can result in a significant increment of the total utility, even if we have to spend 10 extra periods to get it.

A natural question occurs: what happens to N^* if both q and δ go to 1? Does it go to infinity? It turns out that in this regime, N^* is of order $\ln(1-\delta)/\ln(1-q)$. We prove this by first showing that the optimal $N \in \mathbb{R}$ is of this order and then connect it to the optimal $N^* \in \mathbb{N}$.

Proposition 16. $\exists\, \varepsilon, \gamma > 0$ *such that for* $q > 1 - \varepsilon$ *and* $\delta > 1 - \gamma$*, the optimal* N^* *satisfies the following inequalities for some constants* r_1, r_2

$$r_1 \frac{\ln(1-\delta)}{\ln(1-q)} \geq N^* \geq r_2 \frac{\ln(1-\delta)}{\ln(1-q)}.$$

This result tells us that for N^* to remain constant, $(1-\delta)$ must go to 0 as fast, $(1-q)^c$ for some constant c. It is an interesting relationship between the precision of private signals and the patience of the social planner.

Another interesting case is when $\delta \to 1$, and q is fixed. We conjecture that as the planner becomes more patient, absorbing public beliefs become more extreme and $u(0.5)$ gets closer to 1, maximum possible utility. We calculate the rate at which it approaches 1 using a few facts from the proof of Proposition 16.

Proposition 17. *As δ goes to 1, utility at belief $1/2$ goes to 1 with the following rate*

$$u_N(0.5) = 1 - O((1-\delta)\ln(1-\delta)).$$

This result tells us that as we increase δ, and consequently increase N^*, $1 - u(0.5)$ goes almost exponentially quickly to 0.

3.3 Continuous signal distribution

We now consider a continuous distribution of private signals rather than binary. Continuity gives us more freedom in our actions as the social planner. For example, we have a finer trade-off between the expected loss and expected gain in period t. Also, it will allow us to make the expected loss that occurs due to a non-zero price as small as we want.

In the binary case, the expected loss is bounded away from 0 unless $k_t = 1$. Recall that in the binary case, the actuall loss, if it occured, was $p_t - q > 0$ if $p_t > 1/2$ $(1 - p_t - q > 0$ if it is less) and increasing with the public belief, when the latter gets closer to the boundary levels. Moreover, when we had non-zero prices, the loss occured with probability $1 - q$, making the expected loss bounded away from 0. This is an aspect of the discrete distributions because even if you have a very high public belief, there is still a probability of making the wrong action (conditioned on being in the learning period) which is bounded away from zero.

The likelihood ratio of the private signal $g_H(x)/g_L(x)$ has cumulative distribution function F_H in the *High* state and F_L in the *Low* state. Again, as in the classical model, we are going to consider two cases: when signals have bounded and unbounded strength. Let us start with the former one.

Bounded private signals

Assume that the private signals are bounded in a sense that[4]

$$1 > \frac{q}{1-q} \geq \frac{g_H(x)}{g_L(x)} \geq \frac{1-q}{q} > 0\cdot$$

So the agents can not get arbitrary strong information about either state. This implies that if $l_t > q/(1-q)$ and $k_t = 1$, then player t disregards her private signal as her total likelihood ratio is always above 1. Analogously, if $l_t < (1-q)/q$ and $k_t = 1$, then the total likelihood ratio is always below 1.

The main result in this case is similar to the one in the classical literature. Unless we have an extremely patient social planner, $\delta = 1$, it is impossible to learn the underlying state of the world. To put it differently, it is optimal to stop the LP before p_t reaches 1 or 0, as the expected loss is going to be bigger than the expected gain for high enough p_t. Notice that it also applies to situations when there are finitely many types of private signals, that are not completely revealing.

Theorem 18. *If distributions of the private belief are bounded and* $\delta < 1$*, then there exists* $\overline{p} < 1$ *and* $\underline{p} > 0$ *such that* $\underline{p} \leq p_t \leq \overline{p}$*.*

To prove this theorem, we first bound the expected gain in the public belief and then translate it to the expected gain in the utility. The latter one is compared to the expected loss.

This result can be explained by the choice of the form of utility function, discounted sum. It makes the social planner care more about the current generations rather than the ones that are far away in the future, even if δ is close to 1. Still, we can see from the proof that as δ increases, the public belief is able to get closer to the extreme beliefs. Which makes the society more certain about the realized state. We also conjecture that the expected utility significantly increases from the "no prices" case, as it did in the previous section.

In the next subsection, we are going to investigate how prices affect the outcome in the case of unbounded signals. In particular, we are going see what happens to the asymptotic speed of learning comparing to the classical model (Hann-Caruthers, Martynov, and Tamuz, 2018).

[4]We assume symmetry of private signals without loss of generality.

Unbounded private signals

Now we would like to see whether the prices in fact increase the asymptotic speed of learning at the same time as they increase the social welfare.

To remind what the asymptotic speed of learning is, recall that in the classical model with unbounded signals people eventually start choosing the correct action, 1 if the state is $High$ and 0 otherwise. Then we can see how quickly does l_t converge to the boundary value, i.e. find a function $f(t)$ such that $\lim_{t\to\infty} l_t/f(t) = 1$. This function $f(t)$ is called the *asymptotic speed of learning* or *ASL*.

To do this analysis in full generality, for any signal distribution, seems to be a very complex problem, so we choose a particular, well-known in the social learning literature, pair of distribution. Consider the following distributions of private signals for the $High$ and the Low states

$$G_H(x) = x^2$$
$$G_L(x) = 1 - (1-x)^2.$$

Then corresponding distributions of the likelihood ratios are

$$F_H(y) = \mathbb{P}\left(\frac{g_H(s)}{g_L(s)} \leq y \middle| \theta = High\right) = \frac{y^2}{(1+y)^2}$$
$$F_L(y) = \mathbb{P}\left(\frac{g_H(s)}{g_L(s)} \leq y \middle| \theta = Low\right) = \frac{y^2+2y}{(1+y)^2},$$

for $y \in [0, \infty]$. It is easy to see that this pair of distributions actually correspond to the likelihood beliefs, as $F'_H(y)/F'_L(y) = f_H(y)/f_L(y) = y$(L. Smith and Sørensen, 2000).

The following theorem tells us that it is indeed profitable for the social planner to choose prices that increase the asymptotic speed of learning. For example, if $p_t > 1/2$, then it is *better* to choose some constant price greater than 1 rather than to stick with no price. Furthermore, in this situation it is *not optimal* to choose the price that slows down the asymptotic speed, $k_t < 1$. Also, the optimal prices are bounded.

The fact that k_t is greater than 1 and $p_t > 1/2$ implies that the ASL increases in k_t times.

Theorem 19. *For high enough p_t there exists $k > 1$ such that $k_t = k$ gives a higher utility than no price at all. Furthermore, when it is better to choose $k_t > 1$ rather* *than* $k_t = 1$*, it is not optimal to choose* $k_t < 1$*. Moreover, prices are bounded:* $\exists$ $\underline{k} > 0$ *and* $\overline{k} < \infty$ *such that* $\underline{k} < k_t^* < \overline{k}$ *for the optimal* k_t^**.*

Under some mild conditions on $u'(p)$, we can relax the first statement "for high enough p_t" to "for $p_t > 1/2$".

This theorem does not only provide the desirable result about the increased ASL, as formally stated in Corollary 21, but also a surprising fact: the optimal prices k_t are bounded from 0 and ∞. This implies that as the public belief goes to 1, the optimal price c_t does not go to 1 (equivalent to k_t not going to ∞) in order to extract a lot of information from the action, as it is too costly.

It also implies that we can approximate the optimal outcome with a good precision by fixing some constant price when p_t is close to the boundary values, saving the cost of updating.

To understand why this theorem is very interesting, let us look at the big picture. The main problem with continuous signals is that p_t has a very complex behavior, and so the analytical solution of the utility function and hence, the optimal pricing policy, can not be obtained.

Given that we do not know the utility function, we are still able to provide a nice description of the optimal policy: the price should be against the belief, i.e. if $p_t > 1/2$, then $k_t > 1$, some constant pricing will already give us more utility then no prices at all, and these prices are bounded. As we will see in the next section, where we calculate $u(p)$ and the optimal k_t^* numerically, this is indeed a very good description of the optimal policy.

There are two main corollaries of Theorem 19. The first one is that the full learning occurs.

Corollary 20. *As* t *goes to infinity,* p_t *converges to* 1 *if* $\theta = High$ *and to* 0 *if* $\theta = Low$*.*

For the second one, notice that Theorem 19 tells us that from some point on, the optimal price, k_t, is above 1. If it stays above some $k > 1$, which we will see, in the next section is true, then the ASL is going to increase at least by factor k.

Corollary 21. *If the social planner chooses a constant price* $k_t = k > 1$ *and* $\theta = High$*, then the asymptotic speed of learning increases by a factor* k*.*

Therefore, prices indeed increase the social welfare and the ASL at the same time.

In the next section, we try a different approach to solve for the utility function u and the optimal pricing policy.

3.4 Numerical calculations

As we saw above, all theoretical calculations are already fairly complicated and it seems impossible to obtain the analytical solution for (3.1). Thus, we are going to provide a numerical solution of (3.1), which reaffirms claims and intuitions that we had before.

Again, suppose that conditional on the underlying state, the private likelihood ratio has CDF either $F_H(y)$ or $F_L(y)$ as in previous subsection and $\delta = 0.9$

$$F_H(y) = \frac{y^2}{(1+y)^2}$$
$$F_L(y) = \frac{y^2+2y}{(1+y)^2}.$$

We are going to look for a solution of the following form

$$\widetilde{u}(p) = \sum_{i=1}^{m} t_i T(i,p),$$

where t_i are constants and $T(i,p)$ is i's Chebyshev polynomial of the first kind. To remind readers, $T(0,p) = 1$, $T(1,p) = p$, $T(2,p) = 2p^2 - 1$, and $T(k,p) = 2pT(k-1,p) - T(k-2,p)$.

What we need to do is to find these constants, t_i's. This is done by a collocation method ("Collocation method" n.d.). The corresponding coefficients are

$$(t_i)_{i=0}^{9} = (136.6, -250.2, 195.8543, -129.0, 70.4, -30.9, 10.4, -2.4, 0.3, 0).$$

The graph of $\widetilde{u}$ is presented in Figure 3.4.

Let us check that this is a good approximation. To do so, we calculate the right hand side of (3.1) given $\widetilde{u}(p)$ and compare the maximal distance between these two functions. It is equal to $8.3 \cdot 10^{-4}$ which is extremely small.

Moreover, we can now see that as public belief approaches 1, optimal k_t increases and is significantly above 1. Here are a few examples

Here, $p(c_t)$ is the modified public belief.

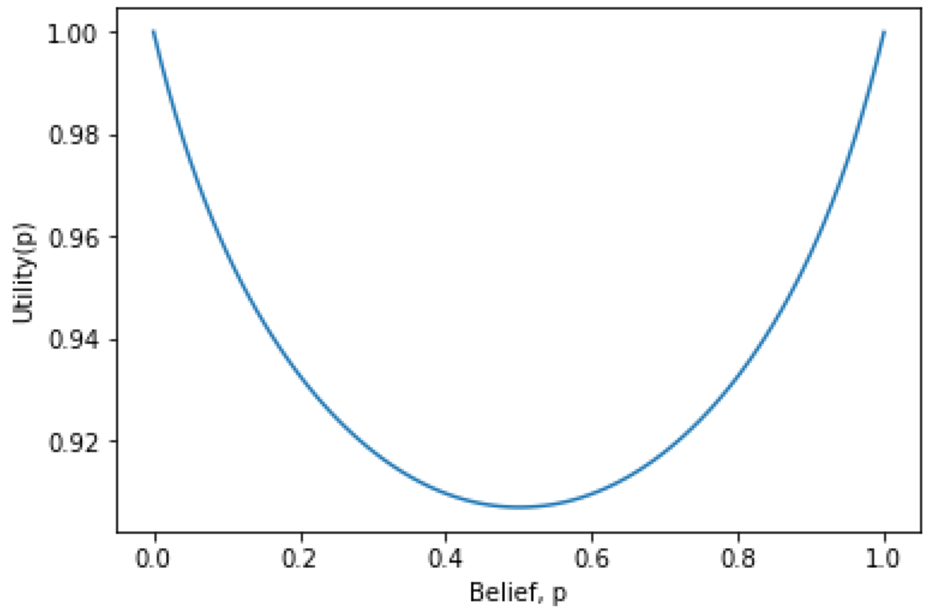

Figure 3.4: Approximation of the utility function with the optimal pricing, $u(p)$.

p_t	c_t	k_t	$p_t(c_t)$
0.7	0.12	1.26	0.64
0.8	0.2	1.49	0.73
0.9	0.32	1.94	0.82
0.95	0.4	2.37	0.89
0.99	0.5	3.04	0.97
0.999	0.53	3.28	0.996

Table 3.1: Optimal price and modified public belief given current public belief

We can also calculate the expected utility without any price in a similar way. Let us denote it by $u_{k=1}(p)$. Then, for $\delta = 0.9$, the difference in expected utilities as a percentage of $1 - u_{k=1}(p)$ is around 10%. The reason we are normalizing by $1 - u_{k=1}(p)$ and not by 1 or $u_{k=1}(p)$ is that this formula better captures how the expected utility functions flatten. In other words, it becomes more convex comparing to a 45 degree cone ($u_{k=1}(p)$) that we have when there is no price.

This difference may not be as dramatic as one would think. The reason for this is that we apply a significant non-zero price ($k_t \neq 1$) for high public beliefs, which results in utility gain. In order to bring back this utility growth to $p_t = 1/2$ from those high public beliefs, we need to take a large number of steps. Thus, the gain is significantly discounted. We can observe bigger improvements when we take higher values of δ.

This means that the optimal prices significantly increase the asymptotic speed of learning as well as the social welfare.

3.5 Conclusion

In this paper, we improve the main inefficiencies of the classical sequential learning model. We introduce a social planner whose objective is to maximize the social welfare of the agents by choosing the optimal prices for the new good in each period. We show that the optimal prices indeed increase the social welfare as well as have other positive effects on the public belief. In the case of bounded signals, society ends up with a much higher public belief than in the classical case, and for the unbounded case, it significantly increases the convergence speed.

We manage to provide a complete characterization of the utility function and the optimal strategy in the binary case. Even though calculations are very complex for general distributions of private signals, we provide the main properties and description of both the expected utility function and the optimal strategy for a specific continuous distribution. For example, the new product should cost more than the old one if the former one is believed to be better. Furthermore, when the public belief goes to its extreme value (1 or 0), the optimal price is bounded by a constant.

Chapter 4

WEAK AND STRONG TIES IN SOCIAL NETWORK

Many studies find (Jackson, 2007), (Jackson, 2011) that networks are important vehicles for passing information in various economic situations. As nowadays information is one of the most valuable resources, a deeper understanding of how it travels through the network can significantly benefit different aspects of our life. They range from career advancement to finding likeminded people and improving the quality of our leisure time.

As we know from our daily life experience, in social networks there are more than just two relationship states between people.[1] Different types of these ties between individuals have different effects. In this paper, we assume that there are three main types of connections. People might not know each other, however, if they do, there are two possibilities: they can be distant *acquaintances* or close *friends*. The question that we ask is: how do these various connections shape the network between individuals?

Social networks are known for delivering various types of information. One of the prime examples of this is learning about a job opening from someone in your social circle. According to a recent survey (Adler, 2016) published by LinkedIn, around 85% of all jobs are filled via networking. So a question arises: when searching for a job, should we rely more on our group of friends or acquaintances? What roles do these groups play in providing agents with informational opportunities?

In this paper, we study a two-stage model: first, agents strategically form a network; second, a piece of information, e.g. about a new job oppening, is planted in the network and is spread there for the next 2 time periods. Each player gets utility 1 if the information reaches her and 0 otherwise. Society is divided into villages: people within the same village have an opportunity to be friends with each other and players in different villages have an opportunity to be acquaintances. It is a real-life observation of the fact that it is easier to maintain a friendship with someone who is geographically, or in some other sense, close to you. Agents play a network formation game, where they choose how much effort, e.g. amount of time, they want to put into socializing with other people to connect to them. Both agents need

[1] Whether they know or do not know each other.

to exert a non-zero effort in order for them to have a chance to become connected. Once the network has been realized, people have to pay different costs to sustain two types of relationships. Friendship is more difficult to maintain but it is a more reliable source of information compared to acquaintanceship. After the agents have made their decisions, the network emerges and defines links between people. In the second stage, a random person in the network receives a piece of valuable information, and we allow this information to travel through the network for two periods.

In the case when information travels only for 1 period, the game becomes trivial. For the two types of connections, one will be more beneficial than the other in terms of the difference between the associated utility and the cost. Thus, agents will choose the more valuable one and completely disregard the other type of link. In other words, you should have only friends or only acquaintances. Also, the seminal paper by Granovetter (Granovetter, 1977) states that most job offers come from acquaintances (also called "weak ties") rather than from friends ("strong ties"). This phenomenon is called "strength of weak ties" and many papers (Contandriopoulos et al., 2016), (Weng et al., 2018) argue that they are significantly more helpful than the strong connections. We know from the sociological literature that social ties effect our well-being (Haythornthwaite. and B. Haythornthwaite, 1998), (C. Haythornthwaite and Wellman, 1990), (Sandstrom and Dunn, 2014), (Holt-Lunstad, T. Smith, and Layton, 2010). But why do we observe both types of connections in real life from economic stand point? Should we keep both strong and weak links or should we disregard one of them?

We believe the answer lies within the network's topology. In equilibrium, and in real-life, if B and C are both friends with A, then they are more likely to have a link between each other rather than if they were both A's acquaintances. It means that the friends' graph is denser compared to the acquaintances' one. This leads to the following intuition for observing both types of connections in real-life networks. An agent does not want to spend her whole budget on the strong ties as, at some point, a new friend is not going to bring many players into her network because of the high clustering[2] of the friends' graph. On the other hand, a new acquaintance will indirectly connect her to many people whom she does not know yet because their graph is sparser. At the same time, not having any friends is not optimal either, because in this case an agent can trade some weak ties for a few reliable

[2]Informally, clustering coefficient, or just clustering, says how likely two people who have a common friend/acquaintance know each other.

strong ones which, at this point, would not suffer from an aforementioned high clustering. However, the primary reason for observing both types of connections is the following. The main part of the agent's network (which brings the most utility) is constituted of weak-strong and strong-weak ties which combine the best of both worlds: 1) they are reliable as they include a strong tie, and 2) they bring information from diverse sources as they include a weak tie. Therefore, in order for agents to maximize their utility, which requires having weak-strong and strong-weak links present in equilibrium, individuals need to have both types of connections.

In this paper, we establish that a non-trivial symmetric Nash equilibrium can only be in a neighborhood of the ε-equilibrium of this game. In both these equilibria, it is the best response for people to have both friends and acquaintances at the same time. This explains why we observe both types of connections between individuals. We find that a bulk of agent's utility comes from weak-strong, strong-weak, and weak-weak connections, and as a result, agents always want to have these ties present in their network. We also perform multiple comparative statics of this equilibrium and check whether the friends' graph has a higher clustering coefficient than the acquaintances' one. Finally, we compare it to the optimal symmetric network and find that there are more strong ties in the latter one. People underinvest in their strong connections in the equilibrium. This suggests that if a social planner would "subsidize" the strong ties or "tax" the weak ones it would increase social welfare. The intuition for this is that it is socially beneficial not only to get the signal but also to share it with others through reliable strong connections.

The paper is organized in the following way. In the next section, we present our model and explain its different mechanics and properties. In the second one, we present the analytical findings that we described above. Next, we provide some computational results and graphs to confirm that our solution/equilibrium is a plausible approximation of the social networks that we observe. We conclude in the final chapter.

Literature review

In a seminal paper by Granovetter (Granovetter, 1977) the author describes a surprising finding of his field experiment: people find a new job through weak ties more often than from the strong ones (27% vs 16%). Many studies agree with Granovetter's opinion. Contandriopoulos et. al (Contandriopoulos et al., 2016) claims that it is significantly more valuable to play a "bridge" role in the network. This corresponds to having weak ties rather than "redundant" strong links: a researcher could gain up to one h-index point by making 3 weak connections. Furthermore, Weng et. al (Weng et al., 2018) find evidence that weak ties are used to collect "important" information.

There are empirical papers that address the ties' strength and their ability to help find a new job. Gee and et. al (Gee et al., 2017) construct a very large data set of deidentified individuals from 55 countries. They find that even though it is more likely to hear about a job opening from a week tie, the likelihood of working at a place where your friend/acquaintance works is increasing with the strength of your connection.

Furthermore, Kuzubas and Szabo (Kuzubas and Szabo, 2017) find that in Indonesia, people tend to use their strong ties more than the weak ones when the number of the latter ones is either on a small or on a large side of the range. At the same time, even with a medium-sized weak network, the probability of getting a job through a strong tie decreases by only 16%. However, positions found through weak links correspond to 10% higher wages.

In terms of theoretical work, there are fewer papers in this area. Golub and Livne (Golub and Livne, 2010) study how different levels of agent's neighborhoods affect the equilibrium network. They have a similar insight to ours, that player's utility does not only depend on her friends, but also, her friends of friends. The authors find, there are two equilibrium regimes in which the realized network can be either connected or fragmented when the costs are not too convex. However, they do not consider different types of connections that we observe in real life.

Boorman (Boorman, 1975) tries to capture the trade-off between strong and weak ties. One of his main results is that, depending on the probability of being unemployed, all-weak and all-strong networks are going to be equilibrium. The author assumes that there is no clustering not only in the acquaintances' network, but also in the friends one, which is a very unrealistic assumption. Furthermore, it seems that he does not take into account the agent's neighborhoods at a distance of more than 1, which we believe plays a crucial role in a deeper understanding of network effects.

Tümen (Tümen, 2017) claims that different types of links are used in different periods of our life. He connects weak links with an early career stable settlement while strong ties are associated with the amplified mobility that generates mismatch.

4.1 Model

In this paper, we study the following two-stage model: in the first stage, agents strategically form a network, and in the second stage, information spreads and defines the agents' utilities. We will start with the second stage.

There is a network of people with two different types of connections between each other. We represent it by an undirected graph (V, E), where V is a set of vertices and E is a set of edges. We write E as the disjoint union of E_s, the set of *strong ties*, and E_w, the set of *weak ties*. For $i, j \in V$, let $d(i, j)$ denote the length of the shorterst path between i and j.[3] Note that $d(i, j) = d(j, i)$, as the graph is undirected. We denote by $N_k(i)$ the neighborhood of agent i at distance k, which is a set of agents that are exactly at distance k from i:

$$N_k(i) = \{j \in V | \, d(i, j) = k\}.$$

Recall that a clustering coefficient (CC) of a graph is defined as

$$CC = \frac{3 \times \text{number of triangles}}{\text{number of all triplets}},$$

where the number of triangles is the number of unordered triplets (A, B, C), $A, B, C \in V$ such that there exists an edge between any pair from them.

In period 1, nature randomly selects a person i who receives a signal/piece of valuable information. In the first period, this player sends information to all of her friends and acquaintances. It travels through a strong tie with probability 1 and through a weak one with probability π_w[4]. In the second period, everyone who has received the signal from agent i sends it to their friends and acquaintances in the same way. If this information reaches a player in at least one of these two periods, then her utility is 1, and 0 otherwise.

We now define the first stage when agents form the underlying network, (V, E). Society is divided into $N + 1$ "villages", with K agents in each one. People in the same village can potentially be friends with each other and players in different ones can potentially be acquaintances.

For each player $j \neq i$, agent i chooses an effort level, $1 \geq p_{ij} \geq 0$, that she is willing to spend to become his friend or acquaintance, depending on whether they

[3] A path between two nodes, $i, j \in V$, is a sequence of nodes $n_1, n_2, \ldots n_t \in V$, such that $n_1 = i$, $n_t = j$, and $\forall i \in \{1, \ldots t - 1\}$ $e_{n_i n_{i+1}} \in E$.

[4] So with probability $1 - \pi_w$ information does not reach the other end of the week link.

are in the same village or not. Thus, a strategy for player i are effort levels for all other players, $\{p_{ij}\}\ \forall j \in V \setminus \{i\}$. We are going to focus on symmetric strategies: agent chooses one effort level for people in her own village and another level for people outside of it. Thus, a strategy for player i simplifies to $\{p_i, q_i\}$, where p_i is a socialization level within the village and q_i is a socialization level outside of it. After the strategies have been chosen, two agents i and j, who are in the same village, become friends with probability $p_i p_j$ and, if they are in different villages, they become acquaintances with probability $q_i q_j$.

Maintainting each strong and each weak link requires c_{strong} and c_{weak} units of time, respectively. People have a time budget, B, that they can spend on their social circle. Their utility vanishes if they exceed the budget and does not depend on c_{strong} and c_{weak} otherwise. To simplify our presentation in the following sections, we normalize c_{strong} to be 1 and allow c_{weak} and π_w to vary.

Here is a description of the game. Agent i, who is in a village $\mathcal{V}_i$, chooses her effort levels $1 \geq p_i, q_i \geq 0$ and becomes friends with player $j \in \mathcal{V}_i$ with probability $p_i p_j$ and she becomes acquainted with player $\ell \notin \mathcal{V}_i$ with probability $q_i q_\ell$. Agent i maximizes her expected utility U_N, where N stands for the number of villages other than hers, given by this expression:

$$U_N(p_i, q_i, p_{-i}, q_{-i}) = \begin{cases} \mathbb{E}_{(V,E)}\left(\mathbb{P}(i \text{ gets the signal})\right), \\ \text{if } c_{strong} p_i \sum\limits_{\substack{j \in \mathcal{V}_i \\ j \neq i}} p_j + c_{weak} q_i \sum\limits_{\ell \notin \mathcal{V}_i} q_\ell \leq B \\ \\ 0, \\ \text{otherwise,} \end{cases}$$

where $\{p_{-i}, q_{-i}\}$ stand for a strategy profile of every player except i and we are taking expectation over realizations of the random graph.

We are going to focus on symmetric equilibria, $p_i = p$ and $q_i = q\ \forall i$. There are two trivial equilibria: 1) $\forall i\ p_i = 0$; 2) $\forall i\ q_i = 0$. If everyone else is not exerting any effort for one type of connections, then it is strictly dominated for player i as well to choose a non-zero effort level for it. Whichever strategy player i chooses, this type of connection is not going to realize as the probability of it is the product of effort levels.

But as soon as p and q for other players are positive, agent i is going to choose p_i and q_i such that she uses all of the budget that is available to her, B. Otherwise, she can increase one of her effort levels and benefit from increasing the expected number of her connections. Therefore, we can treat this objective as if we have equality in the budget constraint.

From now on, we refer to nontrivial symmetric equilibria as equilibria unless stated otherwise. In the next section, we prove that an equilibrium is in a neighborhood of the unqiue ε-equilibrium of this game. We investigate their properties, including multiple comparative statics and welfare results.

4.2 Analytical results

We start our analysis by noting that the probability the signal reaches agent i can depend on the realized graph in a very complex way. As a simple example, consider the case when there are only two weak-strong paths between i and j and the latter one is given the signal by nature. Then, the probability that it reaches i is $1-(1-\pi_w)^2 = 2\pi_w - \pi_w^2$, instead of π_w if there was only one path. Fortunately, as we increase the number of villages, $N+1$, the chance that there are multiple connections between i and some other person j, that include at least one weak link, vanishes. This fact allows us to find a limit utility function U_∞ such that U_N uniformly converges to it as we increase $N+1$.

We are going to prove that when agents are maximizing U_∞ instead of U_N, there is a unique equilibrium, and as a result of this, there are equilibria of the initial function within the ε-neighborhood of it. Furthermore, there are no equilibria of the initial function that are not in a neighborhood of some equilibrium of U_∞. Moreover, it is clear to see that this equilibrium of the limit utility function is also an ε-equilibrium of the initial game. Let us first explain how we find this limit utility function, and then write it down explicitly.

Notice that when there are no overlaps in paths, that we described above (of length at most 2, which have at least one weak link in them), there are only 5 groups of people in our[5] social circle that affect our utility. We describe each group by the type of the shortest path between us: strong, weak, strong-strong, weak-strong and strong-weak (together in one group), weak-weak. Let us multiply the number of people in the weak, strong-weak and weak-strong groups by π_w, and the number of people in the weak-weak group by π_w^2[6]. If we now add the number of people in every group (after three of them have been normalized by π_w and π_w^2 in the previous sentence) and

[5]We use "our" and i's interchangeably.

[6]As these are the probabilities that a signal reaches player i through each of these paths.

add 1 (to include ourselves), then this sum, divided by the total number of people in society, is equal to the probability that the signal reaches us. Thus, by choosing p_i, q_i that maximize this number player i maximizes her utility. This happens because: 1) every individual is equally likely to be chosen by nature; 2) the signal travels for 2 time periods, and 3) there are no overlaps in paths of length 1 or 2, containing at least one weak tie, between us and some other player.

In Proposition 43, we evaluate how many people player i expects to have in each of these 5 groups. Denote the number of people in all those groups, some of which are normalized by π_w and π_w^2 correspondingly, by $\widetilde{U}_N$

$$\widetilde{U}_N(p_i, q_i, p, q) = \Big(1 + (K-1)p_i p + \pi_w K N q_i q + (K-1)(1 - p_i p) \times \\ \times (1 - (1 - p_i p^3)^{K-2}) + \pi_w K(K-1)N \left(p_i p q^2 + p^2 q_i q\right) + \\ + \pi_w^2 q_i q^3 K^2 N^2\Big)$$

The second and the third terms of $\widetilde{U}_N$ calculate how many strong and weak ties we have at distance 1. The next two terms are a bit more complicated. The fourth term stands for the number of friends of friends that are not our direct friends. The fifth one calculates the number of people that are at distance 2 such that there is either a *weak-strong* or *strong-weak* path between us. The last term calculates the number of people that are at distance 2 from us, connected through a *weak-weak* path[7]. The 1 stands for the agent i herself. We need to multiply the second and the last two terms by π_w and π_w^2, respectively, because information travels through a weak tie with probability π_w. So, for example, the expected utility from a group of weak links is proportional to the number of them, KNq_iq, multiplied by π_w.

Our $\widetilde{U}_N$ is a polynomial of two variables, p_i, q_i (keeping p and q fixed) that we need to maximize with respect to p_i and q_i subject to the budget constraint and symmetric equilibrium assumption ($p_i = p,\ q_i = q$). Notice that even though its order is fairly high, it does not depend on N. It is optimal for agent i to choose p_i and q_i such that her budget constraint is satisfied with equality. Thus, we can use this equation to write q_i (and q after applying a symmetric strategy assumption) in terms of p_i (p), and then substitute it back into $\widetilde{U}_N$ so it becomes a polynomial of one variable

[7]We multiply the number of weak-weak connections by $N/(N-1)$ which converges to 1 as we increase N.

$\widetilde{U}_N(p_i, q_i(p_i), p, q(p))$. Keep in mind, that even after we do this, it is still tricky in general to get analytical results about a maximum of a high order polynomial, especially on $[0, 1]$. But before we proceed, we need to introduce some notations to make formulas more readable and intuitive.

Denote by M_{strong} and M_{weak} the following quantities

$$\begin{aligned} M_{strong} &= \frac{B}{(K-1)} \\ M_{weak} &= \frac{B}{KNc_{weak}}. \end{aligned} \tag{4.1}$$

These quantities come from $(K-1)p_i p + KNq_i q c_{weak} = B$, where we replaced c_{strong} with 1. Because all effort levels are non-negative, M_{strong} and M_{weak} are the upper bounds for $p_i p$ and $q_i q$, respectively. These bounds are achieved when the probability of having another type of connection is equal to 0. So, if player i decides not to have any acquaintances, $q_i = 0$, then $p_i p = M_{strong}$ and she has $M_{strong}(K-1)$ friends in expectation. Similarly, if she decides not to have any friends, $p_i = 0$, then $q_i q = M_{weak}$ and she is going to have $M_{weak} KN$ weak links in expectation. But the most helpful part of this notation is that we can express $q_i q$ in terms of $p_i p$ in a short form as $q_i q = M_{weak} - \frac{(K-1)}{KNc_{weak}} p_i p$.

Now, if we substitute $q_i q$ in $\widetilde{U}_N$ using the equation above, we get the limit utility function U_∞

$$\begin{aligned} U_\infty = {} & 1 + (K-1)p_i p + \frac{\pi_w}{c_{weak}}(K-1)(M_{strong} - p_i p) + (K-1)(1 - p_i p) \times \\ & \times (1 - (1 - p_i p^3)^{K-2}) + \frac{\pi_w}{c_{weak}}(K-1)^2 \Big(p_i p (M_{strong} - p^2) + p^2 \times \\ & \times (M_{strong} - p_i p) \Big) + \frac{\pi_w^2}{c_{weak}^2}(M_{strong} - p_i p)(M_{strong} - p^2)(K-1)^2. \end{aligned}$$

As we can see, U_∞ does not actually depend on N, which conforms with its name. Denote by p_∞ an equilibrium (trivial or non trivial) of U_∞ and by p_N an equilibrium of U_N. We show that U_N uniformly converges to U_∞. This, in its turn, implies that for any p_∞, there is a p_N within the ε-neighborhood of it, and vice versa.

Proposition 22. *For any $\varepsilon > 0 \; \exists \, \bar{N}$ such that $\forall \, N > \bar{N}$ and any fixed K*

$$|U_\infty(p_i, q_i, p, q) - U_N(p_i, q_i, p, q)| < \varepsilon.$$

Furthermore, $\forall\, \varepsilon > 0\ \exists\, \bar{N}$ *such that* $\forall N > \bar{N}$ *and for any equilibrium of* U_∞*,* p_∞*, there exists an equilibrium of* U_N*,* p_N*, within the* ε*-neighborhood of* p_∞*. Moreover, there are not any equilibria of* U_N *outside of those neighborhoods.*

From now on, we are going to assume that the agents' utilities are represented by the latter one, study the corresponding equilibria, and then connect them back to the ones of the initial function. Replacing U_N with U_∞ leads to the following objective function of player i:

$$\max_{p_i} \Bigg(1 + (K-1)p_i p + \frac{\pi_w}{c_{weak}}(K-1)(M - p_i p) + (K-1)(1 - p_i p)\times \\ \times (1 - (1 - p_i p^3)^{K-2}) + \frac{\pi_w}{c_{weak}}(K-1)^2 \times \\ \times \Big(p_i p(M - p^2) + p^2(M - p_i p)\Big) + \frac{\pi_w^2}{c_{weak}^2}(M - p_i p)(M - p^2)(K-1)^2\Bigg)$$

$$\text{subject to:}\quad (K-1)p_i p + KNq_i q\, c_{weak} = B, \\ 1 \geq p_i \geq 0, \\ 1 \geq q_i \geq 0.$$

Remember that in order to find symmetric equilibria, we need to find solutions to $(\partial U_\infty / \partial p_i)\big|_{p_i = p} = 0$. We already mentioned that there are two trivial equilibria. The first main result that we have is that there is only one non-trivial symmetric Nash equilibrium of U_∞, which we call p^*. It is also an ε-equilibrium of the initial game.[8] As a consequence of this and the previous proposition, there is an equilibrium of the initial utility function, U_N, in ε-neighborhood of p^*.

Theorem 23. *There is a unique equilibrium of* U_∞*,* p^**, if:*

- $\frac{\pi_w}{c_{weak}} < 1$
- $\frac{\pi_w}{c_{weak}} > e^{-\frac{B^2(K-3)}{(K-1)^2}} \left(1 - \frac{B}{K-1}\right)$

Otherwise, there are only trivial equilibria.

Because there are more potential acquaintances available (NK) than friends ($K-1$), we expect that the latter graph will be more clustered than the first one. This leads

[8] We leave this as an exercise.

to the following trade-off. When choosing between having another strong or weak tie, agents need to take two things into account: 1) how much utility does a new connection bring at distance 1 and 2) how much does it bring at distance 2. The first one asks about the relationship between c_{strong}, c_{weak}, and π_w to determine which type of link is more appealing if we were only getting utility from our *direct* ties (the ones that are at distance 1). The second one, more subtle, asks how many new connections does the new tie bring to our current neighborhood.

For example, if it is more beneficial at distance 1 to have a friend, we still might not want to spend all our time on the strong ties. Because, at some point, our friends' graph will be so clustered that a new strong connection will not significantly change our neighborhood[9], i.e. friends of my new friend are likely to be either my friends or friends of my friends already. Therefore, it is more beneficial to trade this new strong connection for a few weak ones, which will increase our utility by indirectly connecting us to people we do not know yet.[10] Alternatively, consider the case when we only have acquaintances and assume that we trade some of them for a few friends. The friends' graph is sparse at this point[11] so there is no clustering disadvantage that was described above. Furthermore, not only strong links are more beneficial at distance 1, but they also give us access to their own friends and acquaintances through a very *reliable* connection. So it is not optimal to have only one type of link in the equilibrium.

The uniqueness of a non-trivial equilibrium is a nice result as there is no ambiguity about which outcome is going to be implemented by society. This theorem tells us that unless the price for maintaining weak connections is extreme, we should observe both, weak and strong, connections between people.

Two constraints above are necessary for the existence of a nontrivial equilibrium. The first condition ensures that the equilibrium function, $\partial U_\infty / \partial p_i \Big|_{p_i=p}$ is nonnegative around $p = 0$. The second one, on the other hand, makes sure that it becomes negative at some point on $[0, \sqrt{M_{strong}}]$. Therefore, $\partial U_\infty / \partial p_i \Big|_{p_i=p}$ crosses 0, as it is continuous.

But there is more to these constraints than the technical explanations. If c_{weak} is smaller than π_w, then there are only trivial equilibrium solutions. The reason why the nontrivial one disappears is that weak connections are not only more appealing

[9] And, as a consequence, our expected utility.
[10] Because the acquaintances' graph is sparser than the friends' one.
[11] We do not have any friends yet, so from our point of view it is sparse/there is no clustering.

at distance 2 as their graph is more disperse, hence they can potentially bring in more people in our network, but they are also very beneficial at distance 1. A direct benefit from a friend is $1 - c_{strong} = 0$ and from an acquaintance is $\pi_w - c_{weak} > 0$. So in a case when $\pi_w / c_{weak} > 1$, it is rational to disregard strong connections and focus only on the weak ones.

Now let us look at the second constraint. Let us fix the left-hand side and change K first. When K is significantly big, the friends' graph becomes very sparse as the village gets larger, and the right-hand side of the constraint approaches 1. This requires c_{weak} to be close to π_w to make weak connections more attractive. Otherwise, we prefer to spend our budget only on friends. Which leads to the other trivial equilibrium. The same thing occurs if B is small. However, if we do have a lot of time, B is large, then the right-hand side decreases, which relaxes the bound on c_{weak} and allows it to be bigger. When we have a lot of time, keeping c_{weak} fixed, our friends' graph will become very clustered at some point and we would like to have a few acquaintances to maximize the effective size of our neighborhood.

The second part of Theorem 23 validates our choice of the approximation function. It shows that there is a nontrivial symmetric initial equilibrium (possibly multiple ones) within ε-neighborhood of the approximate equilibrium. Thus, we can learn about its behavior by analysing p^*.

While proving this proposition, we find that *weak-strong*, *strong-weak*, and *weak-weak* ties constitute the bulk of agent's utility. It does make sense if we think about it. Imagine that you are looking for a job at Google and you happen to know someone who works there. Why is this connection valuable to you? Is it because exactly this person is going to offer you a job? This is very unlikely. But she has many colleagues and one of them might give you an offer after she recommends you to him. So it makes a lot of sense to look at the number and the types of connections we have not only at distance 1, but also at distance 2.

Let us now analyze this approximate equilibrium p^* and see what happens to it as we change the parameters of the game. Does p^* increase when we increase the cost of maintaining weak ties, c_{weak}, or increase the number of people in the village, K? For this, we have the following comparative statics result.

Proposition 24. *As we do one of the following:*

- ***increase*** *the cost of weak ties,* c_{weak}*;*

- ***decrease*** *the probability of information traveling through a weak tie,* π_w*;*
- ***decrease*** *the number of people in the village,* K*,*

then the equilibrium effort level for making friends, p^**,* ***increases.***

The first two cases are very intuitive: as we make a weak tie option more attractive, we trade some of our friends for acquaintances as we decrease p^* and, hence, decrease the expected number of strong ties. But the last case, when we change K, seems more interesting in some sense. We have some fixed time budget that we can spend on our connections, so increase in K implies that the upper bound on p, M_{strong}, has to decrease. So it is not surprising that the equilibrium value p^* also decreases. However, the expected number of friends increases when we increase K as we will show in the next section.

Let us give an intuition for this result. Remember that Erdős - Rényi random graphs look like trees for small values of p and are significantly clustered for large values. If the latter one occurs, we have this trade-off between strong and weak ties: strong ties might give you higher utility at distance 1, but are more redundant at distance 2 because of the clustering. When we increase K, M_{strong} decreases which forces p to stay small enough so that the friends' graph does not get very clustered. Hence, the strong ties are more attractive than they were before because now they do not have (or have less) disadvantage at distance 2: they are also fairly sparse and bring more diverse information from distance 2. This is why the expected number of friends increases: ${p^*}^2$ does decrease but not as fast as K increases, so their product, $(K-1){p^*}^2$, increases.

There are two more questions that we want to answer in this section. First, in equilibrium, is friends' graph more clustered than the acquaintances' graph? Second, how does the socially optimal symmetric network differ from the equilibrium network? Should we have more friends or acquaintances to maximize the social welfare?

Let us start with the former one. We talked about the trade-off that we observe in this model: how much utility each type of links gives us at distance 1 vs how much it brings at distance 2 (how clustered is the graph for this connection type). We assumed that the friends' graph is more clustered than the acquaintances' one. But we did not, technically, force the CC for the former graph to be higher than the CC for the latter one. The following lemma makes sure that in the equilibrium, the friends' network is indeed more clustered.

Lemma 25. *If* $c_{weak} > \pi_w \frac{B}{B-1} + \frac{(K-1)(B+1)}{K(B-1)N}$*, then the clustering coefficient of the acquaintances' graph is smaller than the CC of the friends' graph.*

Thus, when c_{weak} is a bit bigger than its general bound, π_w (from Theorem 23), friends' graph is indeed more clustered. This tells us that our equilibrium graph mimics a very essential property of the real life networks.

Now that we have a good idea of how our equilibrium looks like and behaves, we can compare it with another network, the socially optimal one. To elaborate: what would be a socially efficient symmetric random network if we could tell people what strategy to play? In mathematical terms, what if we apply symmetry to $U_\infty(p_i, q_i(p_i), p, q)\big|_{p_i=p}$ and then maximize it with respect to p? Would we have more acquaintances or friends compared to the equilibrium graph? Would the new graph look very different from the equilibrium network? We have a very interesting result that the former one is not going to be dramatically different (solution probabilities and the equilibrium functions are going to be similar), but we are going to have more friends than in the latter one. So in a utopian world, where everyone can coordinate what to do, people should have more friends than in the equilibrium.

Theorem 26. *Let* $p_{optimal}$ *be a solution to* $U'_\infty(p, q(p), p, q(p)) = 0$*. If there exists an equilibrium* p^* *and*

- $\pi_w / c_{weak} \geq (1 - {p^*}^4)^{K-2}$,

then $p_{optimal} > p^*$*. In other words, there are more friends in the optimal network than in the equilibrium one.*

The intuition for this result is that if we want to maximize social welfare, we not only care about getting the signal, but also about sharing it with others, and the strong ties are more reliable sources of information.

This completes our analytical analysis and in the next section, we are going to present some graphs and quantitative results for the equilibrium.

4.3 Quantitative results

In this section, we show a few more interesting properties of the equilibrium and provide graphs for visualisation purpose.

We start with a graph of how many connections of different types people have in equilibrium, Figure 4.1. Recall that we fixed c_{strong} to be equal to 1 and vary π_w,

c_{weak}. As we saw in the last section, we care about their relative value towards each other. Let us fix $\pi_w = 0.4$ and vary c_{weak}.

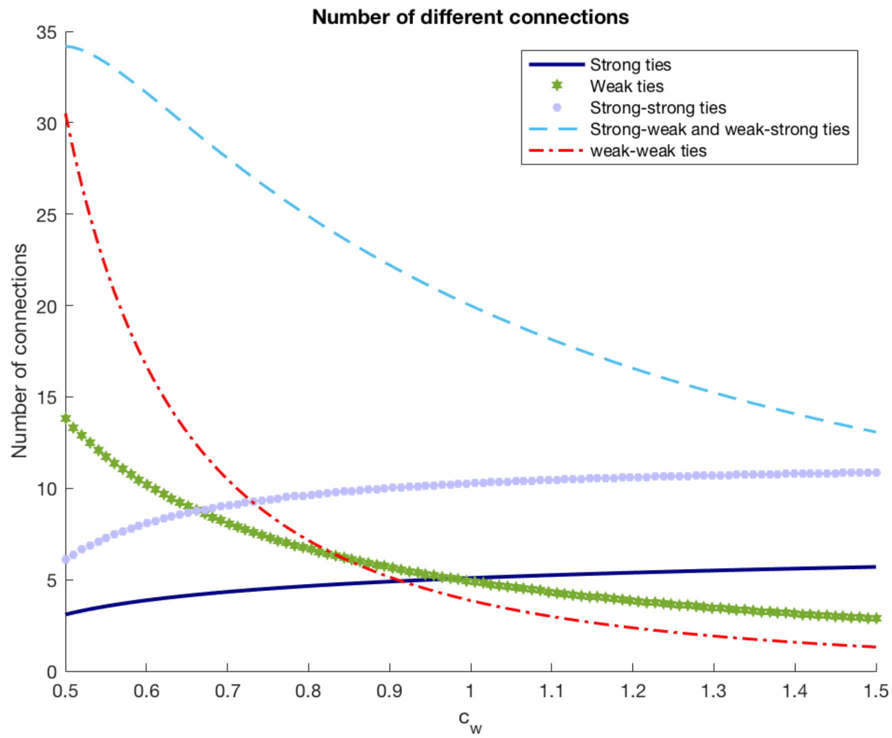

Figure 4.1: Number of different ties (normalized appropriately by π_w or π_w^2) people have in equilibrium when $K = 20$, $N = 100$, $c_{strong} = 1$, $\pi_w = 0.4$, $c_{weak} \in (0.45, 1.55)$.

Theorem 23 tells us that c_{weak} has to be bigger than π_w, so we choose c_{weak}-axis to be from 0.5 to 1.5.

As we can see, weak-strong and strong-weak ties (multiplied by π_w) correspond to the bulk of all agent's connections. Hence, they constitute the most utility out of all other types of connections. To have weak-strong and strong-weak connections present in the network, agents have to have a positive number of both weak and strong individual ties. This gives an explanation why we observe both types of connections in a real life, even though it might look unintiuitive at first.

The second biggest contributors to agent's utility are either weak-weak or strong-strong connections depending on the parameters of the model. As we see, most of agent's utility comes not from her direct links but from the indirect ones, which makes sense if we recall an example with Google's job offer after Theorem 23.

Notice how the number of friends vanishes when c_{weak} becomes closer to π_w and we approach a trivial equilibrium. It appears almost impossible to not satisfy the other constraint (when the weak connections disappear in equilibrium) of Theorem 23 as c_{weak} would have to be bigger than 209 (keeping other parameters the same).

The second property that we would like to show is mentioned after the Propositions 24. We proved that as we increase the number of people in each village, K, the equilibrium probability of making a new friend, p^{*^2}, decreases. However, the number of friends that an agent ends up having in equilibrium increases. Figure 4.2 illustrates these results.

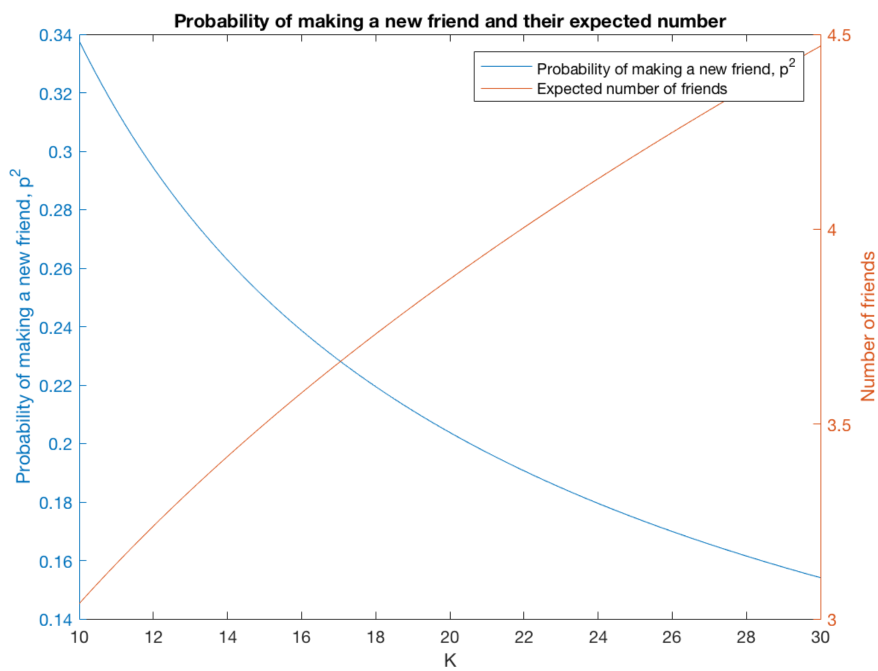

Figure 4.2: Value of p^{*^2} as we change the number of people in the village, K.

As we said before, p^2 intuitively should decrease because its upped bound, M_{strong}, decreases when K increases. But the former one does not do it quickly enough relatively to the increase in K and, as a result of this, their product, $(K-1)p^{*^2}$, increases.

When we increase the number of people in the village, keeping the budget fixed, the friends graph becomes sparser. Therefore, strong links have less disadvantage at distance 2. This motivates people to trade some of their acquaintances for new friends.

The last graph we would like to present compares the equilibrium and the optimal networks. In Theorem 26, we proved that it is socially optimal for all agents to increase the number of friends. In Figure 4.3, we can see that $p_{optimal}$ is indeed bigger than p^*. Interestingly, these values are very close to each other, so our equilibrium network does not differ a lot from the optimal one. This implies that society on its own can achieve a fairly efficient outcome without any interference from the social planner. This very positive result concludes our analysis.

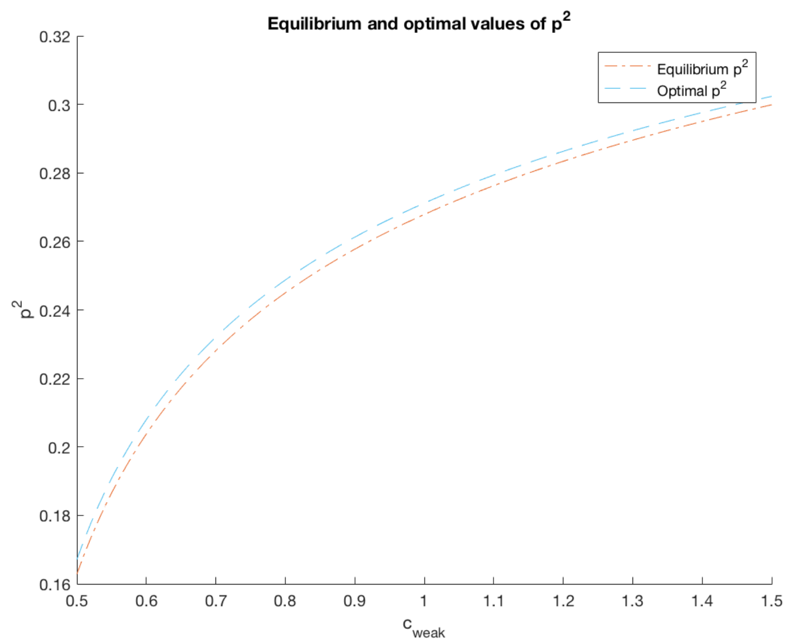

Figure 4.3: Value of p^{*^2} in the equilibrium and optimal networks as we change the cost of weak ties, c_{weak}.

4.4 Conclusion

In this study, we find necessary and sufficient conditions for the existence of a non-trivial equilibrium in which players choose both weak and strong ties. The reason why it is optimal for agents to choose both types of connections is because the bulk of their utility comes from weak-strong and strong-weak ties, which require positive amounts of both types of links to be present in the network. This provides an explanation for why we observe strong and weak ties in real life at the same time even though it might seem unintuitive at first. In the equilibrium (under a mild condition), the friends' graph is more clustered than the acquaintances' one, which complies with empirical evidence. We also provide comparative statics of the equilibrium.

Furthermore, we compare the equilibrium network with the socially optimal symmetric one. These networks are surprisingly similar, but in the latter, agents have more friends. Intuitively, when maximizing social welfare, agents care not only about receiving the information, but also sharing it with others. And the strong ties are more reliable in this case.

We would also like to note another aspect of this work. There are two main types of network models. The first one works with random graphs to represent society and ties between people. Whereas the second one uses game theory to make sure every link is consensual by both sides. Models of the latter type often require either a lot of symmetry from the network or simplifying assumptions due to very complex combinatorics issues. They also produce multiple equilibria, some of which do not look realistic form a network perspective. At the same time, the network does not appear completely randomly, but depends on agents' decisions and choices. In this paper, similar to Golub and Livne in (Golub and Livne, 2010), we are bringing these two approaches together as well as forming a bridge between sociological and economic literature.

BIBLIOGRAPHY

Acemoglu, D. et al. (2011). "Bayesian learning in social networks". In: *The Review of Economic Studies* 78.4, pp. 1201–1236.

Adler, L. (2016). *New Survey Reveals 85% of All Jobs are Filled Via Networking*. Available at https://www.linkedin.com/pulse/new-survey-reveals-85-all-jobs-filled-via-networking-lou-adler/.

Banerjee, A. (1992). "A simple model of herd behavior". In: *The Quarterly Journal of Economics*, pp. 797–817.

Bhalla, M. (2012). "Dynamic pricing under social learning with strategic consumers". In: *IIM Bangalore Research Paper No. 367*.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). "A theory of fads, fashion, custom, and cultural change as informational cascades". In: *Journal of political Economy*, pp. 992–1026.

Boorman, S. A. (1975). "A combinatorial optimization model for transmission of job information through contact networks." In: *Bell journal of Economics*, pp. 216–249.

Bose, S. et al. (2006). "Dynamic monopoly pricing and herding. The RAND Journal of Economics". In: *The RAND Journal of Economics*, pp. 910–928.

Campbell, A. (2013). "Word-of-mouth communication and percolation in social networks". In: *American Economic Review*, pp. 2466–98.

Candogan, O., K. Bimpikis, and A. Ozdaglar (2012). "Optimal pricing in networks with externalities". In: *Operations Research*, pp. 919–939.

Chamley, C. (2004). *Rational herds: Economic models of social learning*. Cambridge University Press.

"Collocation method" (n.d.). In: (). Available at http://hdl.handle.net/1721.1/64049.

Contandriopoulos, D. et al. (2016). "The impact of a researcher's structural position on scientific performance: An empirical analysis." In: *PloS one*.

Crapis, D. et al. (2016). "Monopoly pricing in the presence of social learning". In: *Management Science*, pp. 3586–3608.

Duffie, D., S. Malamud, and G. Manso (2009). "Information percolation with equilibrium search dynamics". In: *Econometrica* 77.5, pp. 1513–1574.

Duffie, D. and G. Manso (2007). "Information percolation in large markets". In: *The American Economic Review*, pp. 203–209.

Feller, W. (1968). *An introduction to probability theory and its applications: Volume 1*. Wiley.

Gee, L. et al. (2017). “The paradox of weak ties in 55 countries.” In: *Journal of Economic Behavior and Organization, 133*, pp. 362–372.

Golub, B. and Y. Livne (2010). “Strategic Random Networks and Tipping Points in Social Institutions.” In:

Granovetter, M. (1977). “The strength of weak ties”. In: *Social networks*, pp. 347–367.

Hann-Caruthers, W., V. Martynov, and O. Tamuz (2018). “The speed of sequential asymptotic learning”. In: *Journal of Economic Theory* 173, pp. 383–409.

Haythornthwaite, C. and B. Wellman (1990). “Different strokes from different folks: Community ties and social support.” In: *American Journal of Sociology, 96(3)*, pp. 558–588.

Haythornthwaite., C. and B. Haythornthwaite (1998). “Work, friendship, and media use for information exchange in a networked organization.” In: *Journal of the American Society for Information Science, 49(12)*, pp. 1101–1114.

Holt-Lunstad, J., T. Smith, and J. Layton (2010). “Social relationships and mortality risk: a meta-analytic review.” In: *PLoS medicine, 7(7)*.

Jackson, M. O. (2007). “The study of social networks in economics.” In:

– (2011). “An overview of social networks and economic applications.” In: *Handbook of social economics, Vol. 1*, pp. 511–585.

Kuzubas, T. U. and A. Szabo (2017). “Job search through weak and strong ties: Theory and evidence from Indonesia.” In: Available at `http://www.uh.edu/~aszabo2/paper_2017.pdf`.

Lobel, I. et al. (2009). “Rate of convergence of learning in social networks”. In: *Proceedings of the American Control Conference*.

Papanastasiou, Y. and N. Savva (2016). “Dynamic pricing in the presence of social learning and strategic consumers”. In: *Management Science*, pp. 919–939.

Rosenberg, D. and N. Vieille (2017). “On the efficiency of social learning”. In: *private communication*.

Sandstrom, G. and E. Dunn (2014). “Social Interactions and Well-Being: The Surprising Power of Weak Ties.” In: *Personality and Social Psychology Bulletin, 40(7)*, pp. 910–922.

Smith, L. and P. Sørensen (1996). *Pathological outcomes of observational learning*. Available at `http://hdl.handle.net/1721.1/64049`.

– (2000). “Pathological outcomes of observational learning”. In: *Econometrica* 68.2, pp. 371–398.

Smith, L., P. Sørensen, and J. Tian (2020). “Informational Herding, Optimal Experimentation, and Contrarianism”.

Sørensen, P. (1996). “Rational social learning”. PhD thesis. Massachusetts Institute of Technology.

Tümen, S. (2017). “Career choice and the strength of weak ties.” In: *Central Bank Review, 17(3)*, pp. 91–97.

Vives, X. (1993). “How fast do rational agents learn?” In: *The Review of Economic Studies* 60.2, pp. 329–347.

– (2010). *Information and learning in markets: the impact of market microstructure*. Princeton University Press.

Weng, L. et al. (2018). “Attention on weak ties in social and communication networks. In Complex Spreading Phenomena in Social Systems.” In: *Springer*, pp. 213–228.

Appendix A

APPENDIX TO CHAPTER 2

A.1 Sub-linear learning

Before proving our main theorems we make the observation (which has appeared before, e.g., (Chamley, 2004)) that *the log-likelihood ratio of the log-likelihood ratio is the log-likelihood ratio*. Formally, if ν_+ and ν_- are the conditional distributions of the private log-likelihood ratio L_t (i.e., have CDFs G_+ and G_-), then

$$\log \frac{\mathrm{d}\nu_+}{\mathrm{d}\nu_-}(x) = x.$$

It follows that

$$G_+(x) = \int_{-\infty}^{x} \mathrm{d}\nu_+(\zeta) = \int_{-\infty}^{x} \mathrm{e}^{\zeta}\,\mathrm{d}\nu_-(\zeta). \tag{A.1}$$

Our first lemma shows that asymptotically, D_+ behaves like the left tail of G_-, and D_- behaves like the right tail of G_+.

Lemma 27.

$$\lim_{x\to\infty} \frac{D_+(x)}{G_-(-x)} = 1 \textit{ and } \lim_{x\to-\infty} \frac{D_-(x)}{G_+(-x) - 1} = 1.$$

Proof. By definition,

$$D_+(x) = \log \frac{1 - G_+(-x)}{1 - G_-(-x)}.$$

Since $\log(1 - z) = -z + O(z^2)$, it holds for all x large enough that

$$D_+(x) > G_-(-x) - 2 \cdot G_+(-x).$$

Applying (A.1) yields

$$D_+(x) > \int_{-\infty}^{-x} (1 - 2\mathrm{e}^{\zeta})\,\mathrm{d}\nu_-(\zeta),$$

and so for any ε and all x large enough,

$$D_+(x) > (1 - \epsilon) \cdot \int_{-\infty}^{-x} \mathrm{d}\nu_-(\zeta) = (1 - \epsilon) G_-(-x).$$

Using the same approximation of the logarithm, we have that

$$D_+(x) < (1+\epsilon)G_-(-x) - G_+(-x) < (1+\epsilon)G_-(-x).$$

The statement for D_+ now follows by taking ε to zero. The corresponding bounds on D_- follow by identical arguments. □

Proof of Theorem 1. Condition on $\theta = +1$. Then ℓ_t is with probability 1 positive from some point on, and all agents take action +1 from this point on. Hence, for all t large enough,

$$\ell_{t+1} = \ell_t + D_+(\ell_t).$$

By Lemma 27, we know that $\lim_x D_+(x) = 0$. Hence for every $\epsilon > 0$ and all t large enough, $|\ell_{t+1} - \ell_t| < \epsilon$. It follows that the limit $\lim_t \ell_t/t = 0$. The analysis of the case $\theta = -1$ is identical. □

Proof of Theorem 2. Given r_t, we will construct private signal distributions such that $\liminf_t |\ell_t|/r_t > 0$ with probability one. These distributions will furthermore have the property that $D_+(x) = -D_-(-x)$. As a consequence we have that regardless of the action chosen by the agent, as long as the sign of the action is equal to that of ℓ_t (which happens from some point on w.p. 1),

$$|\ell_{t+1}| = |\ell_t| + D_+(|\ell_t|).$$

Intuitively, if we can choose private signal distributions that make $D_+(x)$ decay very slowly, then ℓ_t will be very close to being linear.

Formally, and by elementary considerations, the theorem will follow if, for every $Q\colon \mathbb{R} \to \mathbb{R}_{>0}$ with $\lim_{x\to\infty} Q(x) = 0$, we can find CDFs such that $D_+(x) = -D_-(-x)$ and $\liminf_{x\to\infty} D_+(x)/Q(x) > 0$.

Fix any Q such that $\lim_{x\to\infty} Q(x) = 0$, but assume without loss of generality that $Q(x)$ is monotone decreasing.[1] Define a finite measure ν on the integers by

$$\nu(n) = \frac{Q(n-1) - Q(n)}{\mathrm{e}^n}$$

and

$$\nu(-n) = Q(n-1) - Q(n)$$

[1] If Q is not monotone decreasing then consider instead $Q'(x) = \sup_{y \geq x} Q(y)$.

for all $n \geq 0$. Note that ν is indeed finite since

$$C := \sum_{n=-\infty}^{\infty} \nu(n) \leq 2Q(-1).$$

Note also that

$$\sum_{n=-\infty}^{\infty} \nu(n) \cdot \mathrm{e}^n$$

is likewise equal to C.

Let the private signal distributions be given by

$$\mathbb{P}(s_t = n|\theta = +1) = C^{-1}\nu(n)\mathrm{e}^n$$

and

$$\mathbb{P}(s_t = n|\theta = -1) = C^{-1}\nu(n).$$

Then

$$L_t = \log \frac{\mathbb{P}(s_t|\theta = +1)}{\mathbb{P}(s_t|\theta = -1)} = s_t,$$

the distribution of L_t is identical to that of s_t, and so $G_+ = F_+$ and $G_- = F_-$. By our definition of F_-, we have that for $x > 0$

$$G_-(-x) = C^{-1} \cdot Q(\lceil x \rceil - 1). \tag{A.2}$$

Now, by Lemma 27, we know that

$$(1 - \epsilon) \cdot G_-(-x) < D_+(x) < (1 + \epsilon) \cdot G_-(-x),$$

for any $\epsilon > 0$ and all x large enough. It follows that

$$\liminf_{x\to\infty} \frac{D_+(x)}{Q(x)} = \liminf_{x\to\infty} \frac{G_-(-x)}{Q(x)},$$

which, by (A.2) equals

$$\liminf_{x\to\infty} \frac{C^{-1}Q(\lceil x \rceil - 1)}{Q(x)} \geq C^{-1}.$$

□

A.2 Long-term behavior of public belief

The primary goal of this section is to prove Theorem 3, which states that public belief is asymptotically given by the solution to the differential equation (2.3). The proof of this theorem uses two general lemmas regarding recurrence relations. We state these lemmas now and prove them later. The first lemma states that two similar recurrence relations yield similar solutions. The second shows that the solution to a recurrence relation (of the type we are interested in) is well approximated by the solution to the corresponding differential equation.

Lemma 28. *Let $A, B\colon \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ be continuous, eventually monotone decreasing, and tending to zero.*

Let (a_t) and (b_t) be sequences satisfying the recurrence relations

$$a_{t+1} = a_t + A(a_t)$$
$$b_{t+1} = b_t + B(b_t).$$

Suppose

$$\lim_{x\to\infty} \frac{A(x)}{B(x)} = 1.$$

Then

$$\lim_{t\to\infty} \frac{a_t}{b_t} = 1.$$

Lemma 29. *Assume that $A\colon \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ is a continuous function with a convex differentiable tail, and that $A(x)$ goes to 0 as x goes to ∞. Let (a_t) be any sequence satisfying the recurrence equation $a_{t+1} = a_t + A(a_t)$, and suppose there is a function $f : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ with $f'(t) = A(f(t))$ for all sufficiently large t. Then*

$$\lim_{t\to\infty} \frac{f(t)}{a_t} = 1.$$

Given these lemmas, we are ready to prove our theorem.

Proof of Theorem 3. Let (a_t) be any sequence in $\mathbb{R}_{>0}$ satisfying:

$$a_{t+1} = a_t + G_-(-a_t).$$

Then by Lemma 29, the sequence (a_t) is well approximated by $f(t)$, the solution to the corresponding differential equation:

$$\lim_{t\to\infty} \frac{a_t}{f(t)} = 1.$$

Now, conditional on $\theta = +1$, all agents take action $+1$ from some point on with probability 1. Thus, with probability 1,

$$\ell_{t+1} = \ell_t + D_+(\ell_t)$$

for all sufficiently large t. Further, by Lemma 27,

$$\lim_{x\to\infty} \frac{D_+(x)}{G_-(-x)} = 1.$$

So by Lemma 28,

$$\lim_{t\to\infty} \frac{\ell_t}{a_t} = 1$$

with probability 1. Thus, we have

$$\lim_{t\to\infty} \frac{\ell_t}{f(t)} = \lim_{t\to\infty} \frac{\ell_t}{a_t} \cdot \frac{a_t}{f(t)} = 1$$

with probability 1. □

Proofs of Lemmas 28 and 29

Proof of Lemma 28. We prove the claim in two steps. First, we show that for every $\varepsilon > 0$ there are infinitely many times t such that

$$(1-\varepsilon)a_t \le b_t \le (1+\varepsilon)a_t. \tag{A.3}$$

Second, we show that if (A.3) holds for some t large enough, then it holds for all $t' > t$, proving the claim.

We start with step 1. Assume without loss of generality that $a_t \le b_t$ for infinitely many values of t. Fix $\varepsilon > 0$. To show that $(1-\varepsilon)a_t \le b_t \le (1+\varepsilon)a_t$ holds for infinitely many values of t, let $x_0 > 1$ be such that for all $x > x_0$ it holds that A and B are monotone decreasing,

$$A(x), B(x) < \varepsilon < 1$$

and

$$(1-\varepsilon/2)A(x) < B(x) < (1+\varepsilon/2)A(x). \tag{A.4}$$

Assume that $a_t, b_t > x_0$; this will indeed be the case for t large enough, since A and B are positive and continuous, and so both a_t and b_t are monotone increasing and tend to infinity. So

$$B(b_t) < (1 + \varepsilon/2)A(b_t) \leq (1 + \varepsilon/2)A(a_t),$$

where the first inequality follows from (A.4), and the second follows from the fact that A is monotone decreasing and $a_t < b_t$. Since $B(b(t)) = b_{t+1} - b(t)$ and $A(a_t) = a_{t+1} - a(t)$ we have shown that

$$b_{t+1} - b_t < (1 + \varepsilon/2)(a_{t+1} - a_t),$$

and so eventually $b_t \leq (1 + \varepsilon)a_t$. Also, notice that the first time this obtains, we also have that the left inequality in (A.3) holds at the same moment:

$$b_t > b_{t-1} > a_{t-1} = a_t - (a_t - a_{t-1}) > a_t - \varepsilon > a_t - \varepsilon a_t = (1 - \varepsilon)a_t.$$

This completes the first step. Now we go to step 2. Here we show that if (A.3) holds for large enough t then it holds for all $t' > t$.

Fix $\varepsilon > 0$, and let x_0 be defined as above. Suppose that $(1 - \varepsilon)a_t < b_t < (1 + \varepsilon)a_t$, with $a_t, b_t > x_0$. Assume without loss of generality that $b_t \geq a_t$. Then our assumptions and (A.4) imply

$$\begin{aligned} b_{t+1} &= b_t + B(b_t) \\ &< (1 + \varepsilon)a_t + (1 + \varepsilon)A(b_t). \end{aligned}$$

Because $a_t \leq b_t$ and A is decreasing we have

$$\begin{aligned} b_{t+1} &< (1 + \varepsilon)a_t + (1 + \varepsilon)A(a_t) \\ &= (1 + \varepsilon)a_{t+1}. \end{aligned}$$

For the other direction, note first that

$$b_{t+1} > b_t \geq a_t,$$

by assumption. We can write $a_t = (1 - \varepsilon)a_t + \varepsilon a_t$, and since $a_t > x_0 > 1$, $\varepsilon a_t > (1 - \varepsilon)\varepsilon$, and so

$$b_{t+1} > (1 - \varepsilon)a_t + (1 - \varepsilon)\varepsilon.$$

Now, $\varepsilon > A(a_t)$ since $a_t > x_0$, and so

$$\begin{aligned} b_{t+1} &> (1-\varepsilon)a_t + (1-\varepsilon)A(a_t) \\ &= (1-\varepsilon)a_{t+1}. \end{aligned}$$

Thus

$$(1-\varepsilon)a_{t+1} < b_{t+1} < (1+\varepsilon)a_{t+1}, \tag{A.5}$$

as required.

□

Proof of Lemma 29. We restrict the domain of f to the interval (t_0, ∞) such that for $t > t_0$ it already holds that $f'(t) = A(f(t))$. Since A is continuous, $\lim_{t\to\infty} f(t) = \infty$, and so we can also assume that in the interval $(f(t_0), \infty)$ it holds that A is convex and differentiable.

Since f is strictly increasing in (t_0, ∞), it has an inverse f^{-1}. For x large enough define $B(x) = f(f^{-1}(x) + 1) - x$.

Now, let (b_t) be any sequence satisfying the recurrence relation

$$b_{t+1} = b_t + B(b_t).$$

In order to apply Lemma 28, we will first show that

$$\lim_{x\to\infty} \frac{B(x)}{A(x)} = 1.$$

Let $t = f^{-1}(x)$. Such a t exists and is unique for all sufficiently large x, because f is monotone. Notice that by the definitions of $B(x)$ and $f'(x)$

$$\begin{aligned} B(x) &= f(f^{-1}(x) + 1) - x \\ &= f(f^{-1}(x) + 1) - x - f'(f^{-1}(x)) + f'(f^{-1}(x)) \\ &= f(t+1) - f(t) - f'(t) + A(f(t)), \end{aligned}$$

where in the last equality we substitute $t = f^{-1}(x)$. Because f' is positive and decreasing (f is concave) then $f(t+1) - f(t) \geq f'(t+1)$, and so

$$B(x) \geq f'(t+1) - f'(t) + A(f(t)).$$

By the definition of f, $f'(t) = A(f(t))$, and so

$$B(x) \geq A(f(t+1)) - A(f(t)) + A(f(t)) = A(f(t+1)).$$

Again, due to concavity of f we have $f(t+1) \leq f(t) + f'(t)$ and as A is decreasing and convex we get

$$\begin{aligned} B(x) &\geq A(f(t) + f'(t)) \\ &\geq A'(f(t))f'(t) + A(f(t)) \\ &= A'(f(t))A(f(t)) + A(f(t)). \end{aligned}$$

We now substitute back $x = f(t)$:

$$\begin{aligned} B(x) &\geq A'(x)A(x) + A(x) \\ &= A(x)(A'(x) + 1) \end{aligned}$$

so in particular, since $A'(x) \to 0$ as $x \to \infty$,

$$\liminf_{x \to \infty} \frac{B(x)}{A(x)} \geq 1.$$

Now we are going to show that $\limsup_{x \to \infty} \frac{B(x)}{A(x)} \leq 1$ which will conclude the proof. By the definitions of $f^{-1}(x)$ and $B(x)$

$$B(x) = B(f(t)) = f(t+1) - f(t) = \int_t^{t+1} f'(\zeta)\, \mathrm{d}\zeta.$$

As f' is decreasing it follows that

$$B(x) \leq \int_t^{t+1} f'(t)\, \mathrm{d}\zeta = f'(t) = A(f(t)) = A(x).$$

Therefore,

$$\limsup_{x \to \infty} \frac{B(x)}{A(x)} \leq 1.$$

Hence, from these two inequalities we get that

$$\lim_{x \to \infty} \frac{B(x)}{A(x)} = 1.$$

Now notice that, by construction, $f(t+1) = f(t) + B(f(t))$. Thus, by Lemma 28,

$$\lim_{n \to \infty} \frac{f(t)}{a_t} = 1.$$

□

Monotonicity of solutions to a differential equation

We now prove a general lemma regarding differential equations of the form $a'(t) = A(a(t))$. It shows that the solutions to this equation are monotone in A. This is useful for calculating approximate analytic solutions whenever it is impossible to find analytic exact solutions, as is the case of Gaussian signals, in which we use this lemma.

Lemma 30. *Let $A, B\colon \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ be continuous, and let $a, b\colon \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ satisfy $a'(t) = A(a(t))$ and $b'(t) = B(b(t))$ for all sufficiently large t.*

Suppose that

$$\liminf_{x\to\infty} \frac{A(x)}{B(x)} > 1.$$

Then $a(t) > b(t)$ for all sufficiently large t.

Proof. Notice that $a(t)$ and $b(t)$ are eventually monotone increasing and tend to infinity as t tends to infinity. Thus for all x greater than some $x_0 > 0$ large enough, a and b have inverses that satisfy the following differential equations:

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}x} a^{-1}(x) &= \frac{1}{A(x)} \\ \frac{\mathrm{d}}{\mathrm{d}x} b^{-1}(x) &= \frac{1}{B(x)}. \end{aligned}$$

Since $\liminf_x A(x)/B(x) > 1$, we can furthermore choose x_0 so that for all $x \geq x_0$, $A(x) > (1+\varepsilon)B(x)$ for some $\varepsilon > 0$. Thus, for $x > x_0$

$$\begin{aligned} a^{-1}(x) &= a^{-1}(x_0) + \int_{x_0}^{x} \frac{1}{A(x)}\, \mathrm{d}x \\ b^{-1}(x) &= b^{-1}(x_0) + \int_{x_0}^{x} \frac{1}{B(x)}\, \mathrm{d}x \end{aligned}$$

and so

$$\begin{aligned} a^{-1}(x) &< a^{-1}(x_0) + \frac{1}{1+\varepsilon} \int_{x_0}^{x} \frac{1}{B(x)}\, \mathrm{d}x \\ &= a^{-1}(x_0) + \frac{1}{1+\varepsilon}(b^{-1}(x) - b^{-1}(x_0)) \end{aligned}$$

and thus

$$a^{-1}(x) - b^{-1}(x) < -\frac{\varepsilon}{1+\varepsilon} b^{-1}(x) + \left[a^{-1}(x_0) - \frac{1}{1+\varepsilon} b^{-1}(x_0) \right].$$

Since $b^{-1}(x)$ tends to infinity as x tends to infinity, it follows that for all sufficiently large x, $a^{-1}(x) < b^{-1}(x)$. Thus, for all sufficiently large t

$$t = a^{-1}(a(t)) < b^{-1}(a(t)),$$

and so, since $b(t)$ is monotone increasing,

$$b(t) < a(t).$$

□

Eventual monotonicity of public belief update

We end this section with a lemma that shows that under some technical conditions on the left tail of G_-, the function $u_+(x) = x + D_+(x)$ (i.e., the function that determines how the public log-likelihood ratio is updated when the action +1 is taken) is eventually monotone increasing.

Lemma 31. *Suppose G_- has a convex and differentiable left tail. Then the map $u_+(x) = x + D_+(x)$ is monotone increasing for all sufficiently large x.*

Proof. Recall that

$$D_+(x) = \log \frac{1 - G_+(-x)}{1 - G_-(-x)}.$$

Since G_- has a differentiable left tail, it has a derivative $g_-(-x)$ for all x large enough. It then follows from (A.1) that G_+ also has a derivative in this domain, and

$$\begin{aligned} u_+'(x) &= 1 + \frac{g_+(-x)}{1 - G_+(-x)} - \frac{g_-(-x)}{1 - G_-(-x)} \\ &= 1 + \frac{\mathrm{e}^{-x} g_-(-x)}{1 - G_+(-x)} - \frac{g_-(-x)}{1 - G_-(-x)}. \end{aligned}$$

Since $1 - G_-(-x)$ and $1 - G_+(-x)$ tend to 1 as x tends to infinity,

$$\lim_{x \to \infty} u_+'(x) = \lim_{x \to \infty} 1 + \mathrm{e}^{-x} g_-(-x) - g_-(-x).$$

Since G_- is eventually convex, $g_-(-x)$ tends to zero, and therefore

$$\lim_{x \to \infty} u_+'(x) = 1.$$

In particular, $u_+'(x)$ is positive for x large enough, and hence $u_+(x)$ is eventually monotone increasing. □

A.3 Gaussian private signals

Preliminaries

We say that private signals are Gaussian when F_- is the normal distribution with mean -1 and variance σ^2, and F_+ is the normal distribution with mean $+1$ and variance σ^2. To calculate the evolution of ℓ_t, we need to calculate G_+ and G_-, the conditional distributions of the private log-likelihood ratio L_t. Notice that in this case

$$L_t = \log \frac{e^{-(s_t-1)^2/2\sigma^2}}{e^{-(s_t-(-1))^2/2\sigma^2}} = 2s_t/\sigma^2,$$

so that L_t is simply proportional to the signal s_t. It follows that L_t is also normally distributed, conditioned on the state θ, and that G_+ and G_- are cumulative distribution functions of Gaussians, with variance $4/\sigma^2$.

Notation

In this section and those that follow, we denote by ℓ_t^* the public log-likelihood ratio when all agents before agent t take the correct action. Formally,

$$\ell_t^* = \log \frac{\mathbb{P}(\theta = +1 \mid a_1 = \cdots = a_{t-1} = +1)}{\mathbb{P}(\theta = -1 \mid a_1 = \cdots = a_{t-1} = +1)}.$$

For convenience, we will also use the notation $\mathbb{P}_+(\cdot)$ as shorthand for $\mathbb{P}(\cdot \mid \theta = +1)$.

The evolution of public belief

Proof of Theorem 4. Let $f : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ be any function such that $f'(t) = G_-(-f(t))$ for all sufficiently large t. Then by Theorem 3,

$$\lim_{t\to\infty} \frac{\ell_t}{f(t)} = 1$$

with probability 1.

Recall from above that L_t is distributed normally, and $G_-(-x)$ is the CDF of a normal distribution with variance $\tau^2 = 4/\sigma^2$.

For $1 > \eta \geq 0$, define

$$F_\eta(x) = \frac{e^{-\frac{1-\eta}{2\tau^2}x^2}}{x}$$

$$f_\eta(t) = \frac{\sqrt{2}\tau}{\sqrt{1-\eta}}\sqrt{\log(t) + \log\frac{(1-\eta)^2}{2\tau^2}}.$$

By a routine application of L'Hospital's rule, F_0 and F_η are lower and upper bounds for G_-, in the sense that

$$\lim_{x\to\infty} \frac{G_-(-x)}{F_0(x)} = \infty$$
$$\lim_{x\to\infty} \frac{F_\eta(x)}{G_-(-x)} = \infty, \eta > 0.$$

Since $f_\eta'(t) = F_\eta(f_\eta(t))$ for all sufficiently large t, we have by Lemma 30 that for any $\eta > 0$,

$$f_0(t) < f(t) < f_\eta(t)$$

for all sufficiently large t. So

$$\liminf_{t\to\infty} \frac{f(t)}{\sqrt{2}\tau\sqrt{\log t}} = \liminf_{t\to\infty} \frac{f(t)}{f_0(t)} \geq 1$$

and for any $\eta > 0$,

$$\limsup_{t\to\infty} \frac{f(t)}{\sqrt{2}\tau\sqrt{\log t}} = \frac{1}{\sqrt{1-\eta}} \cdot \limsup_{t\to\infty} \frac{f(t)}{f_\eta(t)} \leq \frac{1}{\sqrt{1-\eta}}.$$

Thus,

$$\lim_{t\to\infty} \frac{f(t)}{\sqrt{2}\tau\sqrt{\log t}} = \lim_{t\to\infty} \frac{f(t)}{(2\sqrt{2}/\sigma)\sqrt{\log t}} = 1$$

so with probability 1,

$$\lim_{t\to\infty} \frac{\ell_t}{(2\sqrt{2}/\sigma)\sqrt{\log t}} = \lim_{t\to\infty} \frac{\ell_t}{f(t)} \cdot \frac{f(t)}{(2\sqrt{2}/\sigma)\sqrt{\log t}} = 1.$$

□

To prove Theorem 5, we will need two lemmas. The first is general, and will be used several times in the sequel, while the second deals exclusively with the Gaussian case.

Denote by E_t the event that $a_\tau = +1$ for all $\tau \geq t$; that is, that there are no more mistakes after time t. The next lemma provides a uniform bound for the probability of E_t, conditioned on the public belief. It implies, in particular, that the probability of E_1 is positive, which we will use in the proof of Theorem 5.

Lemma 32. *Suppose G_- and G_+ are continuous, and G_- has a convex and differentiable left tail. Then for every $L \in \mathbb{R}$, there is some $m_L > 0$ such that for any t, $x \geq L$ implies $\mathbb{P}_+(E_t \mid \ell_t = x) \geq m_L$.*

Proof. Recall the definition of the public belief $\mu_t = \mathbb{P}(\theta = +1 | a_1, \ldots, a_{t-1})$. The process $(\mu_1, \mu_2, \ldots)$ is a bounded martingale, and therefore, by a standard argument on bounded martingales, if we condition on $\mu_t = q$, then the probability that $\mu_\tau \leq 1/2$ for some $\tau > t$ is at most $2(1-q)$.[2] This event is precisely the complement of E_t, and therefore we have that $\mathbb{P}(E_t \mid \mu_t = q)$ is at least $2q - 1$. Hence, conditioning on $\theta = +1$, we have that $\mathbb{P}_+(E_t \mid \mu_t = 1 - q) \geq (2q-1)/q$, which is positive for all $q > 1/2$.

Since $\mu_t = q$ is equivalent to $\ell_t = \log q/(1-q)$, what we have shown implies that there is an $\varepsilon > 0$ such that for all $x \geq 1$ (here the choice of 1 is arbitrary and can be replaced with any positive number)

$$\mathbb{P}_+(E_t \mid \ell_t = x) > \varepsilon.$$

Now, for any $L < 1$, the compactness of the interval $[L, 1]$, together with the continuity of G_- and G_+, implies that there is an n_L such that if $\ell_t \geq L$, and if agents t through $t + n_L - 1$ take action $+1$, then $\ell_{t+n_L} > 1$. Further, since the probability of agents t through $t + n_L - 1$ all taking action $+1$ conditional on $\ell_t = x$ is continuous in x, there is a $p_L > 0$ such that

$$\mathbb{P}_+(E_t \mid \ell_t = x) \geq p_L \cdot \varepsilon$$

since with probability at least p_L there are no mistakes up to time $t + n_L$, and thence there are no mistakes with probability at least ε.

□

Lemma 33. *Assume private signals are Gaussian. For every $\varepsilon > 0$ there exists some $k > 0$ such that for all t,*

$$\mathbb{P}_+(a_t = -1 \mid a_\tau = +1 \textit{ for all } \tau < t) > \frac{k}{t^{1+\varepsilon}}.$$

Proof. By the definitions of ℓ_t^* and G_+,

$$\begin{aligned}\mathbb{P}_+(a_t = -1 \mid a_\tau = +1 \text{ for all } \tau < t) &= \mathbb{P}_+(a_t = -1 \mid \ell_t = \ell_t^*) \\ &= G_+(-\ell_t^*).\end{aligned}$$

[2] Intuitively, if I assign high belief now to the event $\theta = +1$, then the probability that I assign this event low belief in the future must be small.

Now, by Theorem 4, for every $\beta > 0$, $\ell_t^* < (1+\beta)\frac{2\sqrt{2}}{\sigma}\sqrt{\log t}$ for all sufficiently large t. Further, it follows from a routine application of L'Hopital's rule (or from the standard asymptotic expansion for the CDF of a normal distribution) that for all sufficiently large x,

$$G_+(-x) > \frac{\mathrm{e}^{-(\sigma^2/8)x^2}}{x}.$$

Let $\varepsilon > 0$, and take $\beta < \sqrt{1+\varepsilon} - 1$. Then by monotonicity of $G_+(-x)$ and a straightforward calculation,

$$\begin{aligned} G_+(-\ell_t^*) &> G_+(-(1+\beta)\frac{2\sqrt{2}}{\sigma}\sqrt{\log t}) \\ &> \left[\frac{1}{(1+\beta)\frac{2\sqrt{2}}{\sigma}}\right] \cdot \frac{t^{(1+\varepsilon)-(1+\beta)^2}}{\sqrt{\log t}} \cdot \frac{1}{t^{1+\varepsilon}} \\ &> \frac{1}{t^{1+\varepsilon}} \end{aligned}$$

for all sufficiently large t. From this, the claim follows immediately.

□

Proof of Theorem 5. Denote by C_t be the event that $a_\tau = +1$ for all $\tau < t$, and note that the event $T_1 = t$ is simply the intersection of C_t with the event that $a_t = -1$.

Let $\varepsilon > 0$. By Lemma 33 there is some $k' > 0$ such that for all t,

$$\mathbb{P}_+(a_t = -1 \mid C_t) > \frac{k'}{t^{1+\varepsilon}}.$$

Now, put $\gamma = \mathbb{P}_+(a_\tau = +1 \text{ for all } \tau \geq 1)$, the probability that all agents take the correct action. By Lemma 32, $\gamma > 0$, so this provides a lower bound on the probability of the first $t-1$ agents taking the correct action. Formally,

$$\mathbb{P}_+(C_t) \geq \mathbb{P}_+(a_\tau = +1 \;\text{ for all } \tau \geq 1) = \gamma.$$

Thus,

$$\begin{aligned} \mathbb{P}_+(T_1 = t) &= \mathbb{P}_+(a_t = -1, C_t) \\ &= \mathbb{P}_+(a_t = -1 \mid C_t) \cdot \mathbb{P}_+(C_t) \\ &\geq \frac{\gamma k'}{t^{1+\varepsilon}} \end{aligned}$$

for all t.

□

A.4 Upsets and runs

We recall a few definitions from Section 2.3. We say that there is an *upset* at time t if $a_{t-1} \neq a_t$. We denote by Ξ the random variable which assigns to each outcome the total number of upsets, and by Ξ_t the total number of upsets at times up to and including t. We say that there is a *run* of length m from t if $a_t = a_{t+1} = \cdots = a_{t+m-1}$. Note that this definition does not preclude a run from being part of a longer run; we will refer to a run of finite length which is not strictly contained in any other run as *maximal*. We say that a run from t is *good* if $a_t = +1$ and *bad* otherwise.

Notice that the number of maximal runs is exactly equal to the number of upsets. We use this observation now to show that the probability of having many maximal runs is very small, so that most of the probability is concentrated in the outcomes with few maximal runs.

Proof of Proposition 7. Denote by Υ the random variable which assigns to each outcome the number of finite maximal good runs it contains; note that with probability 1, Υ is finite.

By Lemma 32, there is a $\beta > 0$ such that for any $x \geq 0$, if $\ell_t = x$, then the probability that all agents from t on take the correct action is at least β. Formally,[3]

$$\mathbb{P}_+(a_\tau = +1 \text{ for all } \tau \geq t \,|\, \ell_t = x) \geq \beta.$$

Thus, whenever $a_{t-1} = -1$ and $a_t = +1$ (or $t = 1$), the probability that there is exactly one more maximal good run is at most $1 - \beta$. It follows that for $n \geq 0$,

$$\mathbb{P}_+(\Upsilon = n + 1) \leq (1 - \beta)\mathbb{P}_+(\Upsilon = n)$$

and thus, for any $n \geq 0$,

$$\mathbb{P}_+(\Upsilon = n) \leq (1 - \beta)^n \mathbb{P}_+(\Upsilon = 0)$$

and so

$$\mathbb{P}_+(\Upsilon \geq n) \leq \frac{\mathbb{P}_+(\Upsilon = 0)}{\beta} \cdot (1 - \beta)^n.$$

Finally, since $\Upsilon = \lfloor \Xi/2 \rfloor$, we have for any n:

$$\mathbb{P}_+(\Xi \geq n) \leq \mathbb{P}_+(\Upsilon \geq \lfloor n/2 \rfloor) \leq c \cdot \gamma^n$$

where $c = \mathbb{P}_+(\Upsilon = 0)/\beta$ and $\gamma = (1 - \beta)^{\frac{1}{3}}$.

□

[3]We remind the reader that $\mathbb{P}_+(\cdot)$ is shorthand for $\mathbb{P}(\cdot \,|\, \theta = +1)$.

Whenever asymptotic learning occurs (that is, whenever the probability that all agents take the correct action from some point on is equal to 1), the total number of upsets is almost surely finite. In particular, the probability that Ξ_t is logarithmic in t tends to zero as t tends to infinity. Using Proposition 7, we can show that in fact this probability tends to 0 quickly:

Corollary 34. *Let c, γ be as in Proposition 7. Then*

$$\mathbb{P}(\Xi_t \geq -\frac{2.1}{\log \gamma} \log t) \leq c \cdot \frac{1}{t^{2.1}}.$$

Proof.

$$\begin{aligned}\mathbb{P}(\Xi_t \geq -\frac{2.1}{\log \gamma} \log t) &\leq \mathbb{P}(\Xi \geq -\frac{2.1}{\log \gamma} \log t) \\ &\leq c \cdot \gamma^{-\frac{2.1}{\log \gamma} \log t} \\ &= c \cdot \frac{1}{t^{2.1}}.\end{aligned}$$

□

In fact, it is equally easy to show the same statement for exponents larger than 2.1, but this will suffice for our purposes.

One important consequence of Corollary 34 is that with high probability, there is at least one maximal run before time t which is long relative to t. Thus, much of the dynamics is controlled by what happens during long runs.

We previously analyzed only long runs that start at time 1, when the public log-likelihood ratio is equal to 0. If a long run starts at some public belief $\ell_t \neq 0$ then its evolution is different from the former case. However, if the run is long enough then the analysis above can still be applied. The following lemma states that if a run starts at some $\ell_t > 0$ then we can bound the future public belief from below using ℓ^*.

Lemma 35. *Suppose that G_- has a convex and differentiable left tail. Then there exists a $z > 0$ such that, if there is a good run of length s from t, then $\ell_{t+s} \geq \ell^*_{s-z}$.*

Proof. Let $u_+(x) = x + D_+(x)$. Then by (2.1), whenever agent t takes action +1, $\ell_{t+1} = u_+(\ell_t)$.

Since G_- is eventually convex and differentiable, $u_+(x)$ is monotone increasing for sufficiently large x, by Lemma 31. Take

$$z = \min\{t \in \mathbb{N} \colon u_+(x) \text{ is monotone on } (\ell^*_t - 1, \infty)\}.$$

Now, let $\mu = \inf_{x \in [0,\ell_z^*]} D_+(x)$. By continuity of $D_+(x)$ and compactness of $[0, \ell_z^*]$, $\mu > 0$, since $D_+(x) > 0$ for all x. Put $N = \lceil \frac{\ell_z^*}{\mu} \rceil$. Then for all $x \in [0, \ell_z^*]$, $u_+^N(x) \geq \mu \cdot N \geq \ell_z^*$. Further, since $u_+(x) > x$ for all x, it follows that whenever there is a run of length N from t, $\ell_{t+N} > \ell_z^*$.

This implies that if there is a good run from t of length $s \geq N$, then $\ell_{t+s} \geq \ell_{s-z}^*$.

□

A.5 Distributions with polynomial tails

In this appendix we prove Theorem 6, showing that for private log-likelihood distributions with polynomial tails, the expected time to learn is finite.

As in the setting of Theorem 6, assume that the conditional distributions of the private log-likelihood ratio satisfy

$$G_+(x) = 1 - \frac{c}{x^k} \text{ for all } x > x_0 \tag{A.6}$$

$$G_-(x) = \frac{c}{(-x)^k} \text{ for all } x < -x_0 \tag{A.7}$$

for some $x_0 > 0$.

We remind the reader that we denote by ℓ_t^* the log-likelihood ratio of the public belief that results when the first $t-1$ agents take action $+1$. It follows from Theorem 3 that in this setting, ℓ_t^* behaves asymptotically as $t^{1/(k+1)}$. Notice also that, by the symmetry of the model, the log-likelihood ratio of the public belief that results when the first $t-1$ agents take action -1 is $-\ell_t^*$.

We begin with the simple observation that a strong enough bound on the probability of mistake is sufficient to show that the expected time to learn is finite. Formally, we have the following lemma. We remind the reader that $\mathbb{P}_+(\cdot)$ is shorthand for $\mathbb{P}(\cdot \mid \theta = +1)$.

Lemma 36. *Suppose there exist k, $\varepsilon > 0$ such that for all $t \geq 1$, $\mathbb{P}_+(a_t = -1) < k \cdot \frac{1}{t^{2+\varepsilon}}$. Then $\mathbb{E}_+(T_L)$ is finite.*

Proof. Since $T_L = t$ only if $a_{t-1} = -1$, $\mathbb{P}_+(T_L = t) \le \mathbb{P}_+(a_{t-1} = -1)$. Thus

$$\begin{aligned}\mathbb{E}_+(T_L) &= \sum_{t=1}^{\infty} t \cdot \mathbb{P}_+(T_L = t) \\ &\le \mathbb{P}_+(T_L = 1) + \sum_{t=2}^{\infty} t \cdot \mathbb{P}_+(a_{t-1} = -1) \\ &\le 1 + k \sum_{i=2}^{\infty} \frac{t}{(t-1)^{2+\varepsilon}} \\ &< \infty.\end{aligned}$$

□

Accordingly, this section will be primarily devoted to studying the rate of decay of the probability of mistake, $\mathbb{P}_+(a_t = -1)$. In order to bound this probability, we will need to make use of the following lemmas, which give some control over how the public belief is updated following an upset.

Lemma 37. *For* G_+ *and* G_- *as in* (A.6) *and* (A.7), $|\ell_{t+1}| \le |\ell_t|$ *whenever* $|\ell_t|$ *is sufficiently large and* $a_t \ne a_{t+1}$.

Proof. Assume without loss of generality that $a_t = +1$ and $a_{t+1} = -1$, so that

$$\ell_{t+1} = \ell_t + D_-(\ell_t).$$

Thus, to prove the claim we compute a bound for D_-. To do so we first obtain a bound for the left tail of G_+. By assumption, for $x > x_0$ (with x_0 as in (A.6) and (A.7)),

$$g_-(-x) = G'_-(-x) = \frac{ck}{x^{k+1}}$$

and so by (A.1),

$$g_+(-x) = \mathrm{e}^{-x} g_-(-x) = ck \frac{\mathrm{e}^{-x}}{x^{k+1}}.$$

Hence,

$$G_+(-x) = \int_{-\infty}^{-x} g_+(\zeta)\,\mathrm{d}\zeta = \int_{-\infty}^{-x} ck \frac{\mathrm{e}^{\zeta}}{(-\zeta)^{k+1}}\,\mathrm{d}\zeta = ck \int_x^{\infty} \zeta^{-k-1}\mathrm{e}^{-\zeta}\,\mathrm{d}\zeta.$$

For ζ sufficiently large, ζ^{-k-1} is at least, say, $\mathrm{e}^{-.1\zeta}$. Thus, for x sufficiently large,

$$G_+(-x) \ge ck \int_x^{\infty} \mathrm{e}^{-1.1\zeta}\,\mathrm{d}\zeta = \frac{ck}{1.1}\mathrm{e}^{-1.1x}.$$

It follows that for x sufficiently large,

$$D_-(x) = \log \frac{G_+(-x)}{G_-(-x)} \geq \log \frac{ck}{1.1} - 1.1x + k \log x \geq -1.2x.$$

Thus, for ℓ_t sufficiently large,

$$\ell_{t+1} = \ell_t + D_-(\ell_t) = \ell_t + \log \frac{G_+(-\ell_t)}{G_-(-\ell_t)} \geq \ell_t + 1.2(-\ell_t) = -.2\ell_t$$

so in particular, $|\ell_{t+1}| < |\ell_t|$.

□

We will make use of the following lemma, which bounds the range of possible values that ℓ_t can take.

Lemma 38. *For G_+ and G_- as in* (A.6) *and* (A.7)*, there exists an $M > 0$ such that for all $t \geq 0$, $|\ell_s| \leq M \cdot \ell_t^*$ for all $s \leq t$.*

Proof. For each $\tau \geq 0$, define

$$M_\tau = \max \frac{|\ell_\tau|}{\ell_\tau^*}$$

where the maximum is taken over all outcomes. Note that there are at most 2^τ possible values for this expression, so M_τ is well-defined and finite. Put

$$M = \sup_{\tau \geq 0} M_\tau.$$

To establish the claim, we must show that M is finite. To do this, it suffices to show that for τ sufficiently large, $M_{\tau+1} \leq M_\tau$.

Now, let $u_+(x) = x + D_+(x)$ and $u_-(x) = x + D_-(x)$. Then as shown in the section about the model, whenever agent τ takes action $+1$, $\ell_{\tau+1} = u_+(\ell_\tau)$, and whenever agent τ takes action -1, $\ell_{\tau+1} = u_-(\ell_\tau)$.

By Lemma 31, u_+ and u_- are eventually monotonic. Thus, there exists $x_0 > 0$ such that u_+ is monotone increasing on (x_0, ∞) and u_- is monotone decreasing on $(-\infty, -x_0)$.

For τ sufficiently large, $\ell_\tau^* > x_0$. Further, it follows from Lemma 37 that for τ sufficiently large, $|\ell_{\tau+1}| < |\ell_\tau|$ whenever $a_\tau \neq a_{\tau+1}$ and $|\ell_\tau| > |\ell_\tau^*|$. Let (a_τ) be any sequence of actions with $\frac{|\ell_{\tau+1}|}{\ell_{\tau+1}^*} = M_{\tau+1}$. If $a_\tau \neq a_{\tau+1}$

$$M_{\tau+1} = \frac{|\ell_{\tau+1}|}{\ell_{\tau+1}^*} \leq \frac{|\ell_\tau|}{\ell_\tau^*} \leq M_\tau.$$

If $a_\tau = a_{\tau+1}$, then either $M_{\tau+1} = 1$, in which case $M_{\tau+1} \le M_\tau$, or $M_{\tau+1} > 1$. If $M_{\tau+1} > 1$, then since $|D_+|$ and $|D_-|$ are decreasing on (x_0, ∞) and $(-\infty, -x_0)$ respectively, $|\ell_{\tau+1} - \ell_\tau|/|\ell_\tau| \le |\ell^*_{\tau+1} - \ell^*_\tau|/|\ell^*_\tau|$. So

$$M_{\tau+1} = \frac{|\ell_{\tau+1}|}{\ell^*_{\tau+1}} = \frac{|\ell_\tau| + |\ell_{\tau+1} - \ell_\tau|}{\ell^*_\tau + |\ell^*_{\tau+1} - \ell^*_\tau|}$$

where the second equality follows from the fact that ℓ_τ and $\ell_{\tau+1}$ have the same sign. Finally,

$$M_{\tau+1} = \frac{|\ell_\tau|}{\ell^*_\tau} \cdot \frac{1 + |\ell_{\tau+1} - \ell_\tau|/|\ell_\tau|}{1 + |\ell^*_{\tau+1} - \ell^*_\tau|/\ell^*_\tau} \le \frac{|\ell_\tau|}{\ell^*_\tau} \le M_\tau.$$

Thus, for all sufficiently large τ, $M_{\tau+1} \le M_\tau$.

□

Proposition 39. *There exists $\kappa > 0$ such that $\mathbb{P}_+(a_t = -1) < \kappa t^{-2.1}$ for all $t > 0$.*

Proof. Let $\beta = -2.1/\log\gamma$, where γ is as in Proposition 7. To carry out our analysis, we will divide the event that $a_t = -1$ into three disjoint events and bound each of them separately:

$$A = (a_t = -1) \text{ and } (\Xi_t > \beta \log t)$$
$$B_1 = (a_t = -1) \text{ and } (\Xi_t \le \beta \log t) \text{ and } (|\{s\colon s < t, a_s = +1\}| \ge \frac{1}{2}t)$$
$$B_2 = (a_t = -1) \text{ and } (\Xi_t \le \beta \log t) \text{ and } (|\{s\colon s < t, a_s = +1\}| < \frac{1}{2}t).$$

First, by Corollary 34 we have a bound for $\mathbb{P}_+(A)$

$$\mathbb{P}_+(A) \le c \cdot \frac{1}{t^{2.1}}.$$

Next, we bound $\mathbb{P}_+(B_1)$. This is the event that the number of upsets so far is small and the majority of agents so far have taken the correct action.

Since there are at most $\beta \log t$ upsets, there are at most $\frac{1}{2}\beta \log t$ maximal good runs. Since, furthermore, there are at least $\frac{1}{2}t$ agents who take action $+1$, there is at least one maximal good run of length at least $t/(\beta \log t)$.

Thus, $\mathbb{P}_+(B_1)$ is bounded from above by the probability that there are some $s_1 < s_2 < t$ such that there is a good run of length $s_2 - s_1 \ge t/(\beta \log t)$ from s_1 and $a_{s_2} = -1$.

For fixed s_1, s_2, denote by E_{s_1,s_2} the event that there is a good run of length $s_2 - s_1$ from s_1. Denote by Γ_{s_1,s_2} the event $(E_{s_1,s_2}, a_{s_2} = -1)$. Then

$$\begin{aligned}\mathbb{P}_+(\Gamma_{s_1,s_2}) &= \mathbb{P}_+(a_{s_2} = -1|E_{s_1,s_2}) \cdot \mathbb{P}_+(E_{s_1,s_2}) \\ &\leq \mathbb{P}_+(a_{s_2} = -1|E_{s_1,s_2}).\end{aligned}$$

By Lemma 35, there exists a $z > 0$ such that E_{s_1,s_2} implies that $\ell_{s_2} \geq \ell^*_{s_2-s_1-z}$. Therefore,

$$\mathbb{P}_+(\Gamma_{s_1,s_2}) \leq G_+(-\ell^*_{s_2-s_1-z}).$$

Since for t sufficiently large $\ell^*_t > t^{\frac{1}{k+2}}$ and since $G_+(-x) \leq \mathrm{e}^{-x}$ by (A.1),

$$\mathbb{P}_+(\Gamma_{s_1,s_2}) \leq \mathrm{e}^{-\alpha(s_2-s_1-z)^{\frac{1}{k+2}}} \leq \mathrm{e}^{-\alpha(t/(\beta\log t)-z)^{\frac{1}{k+2}}}.$$

To simplify, we further bound this last expression to arrive at, for some $c > 0$,

$$\mathbb{P}_+(\Gamma_{s_1,s_2}) \leq c\mathrm{e}^{-t^{\frac{1}{k+3}}}$$

for all t. Since B_1 is covered by fewer than t^2 events of the form Γ_{s_1,s_2} (as s_1 and s_2 are less than t), it follows that

$$\mathbb{P}_+(B_1) < ct^2\mathrm{e}^{-t^{\frac{1}{k+3}}} < \frac{1}{t^{2.1}}$$

for all t large enough.

Finally we bound $\mathbb{P}_+(B_2)$. This is the event that the number of upsets so far is small and the majority of agents so far have taken the wrong action. As in B_1, there is a maximal bad run of length at least $t/(\beta\log(t))$.

Denote by R the event that there is at least one bad run of length $t/(\beta\log(t))$ before time t and by R_s the event that agents s through $s + t/(\beta\log t) - 1$ take action -1. Since B_2 is contained in R, and since R is contained in the union $\cup_{s=1}^t R_s$, we have that

$$\mathbb{P}_+(B_2) \leq \mathbb{P}_+(R) \leq \sum_{s=1}^t \mathbb{P}_+(R_s).$$

Taking the maximum of all the addends in the right hand side, we can further bound the probability of B_2:

$$\mathbb{P}_+(B_2) \leq t \cdot \max_{1\leq s\leq t} \mathbb{P}_+(R_s).$$

Conditioned on ℓ_s, the probability of R_s is

$$\mathbb{P}_+(R_s \mid \ell_s) = \prod_{r=s}^{s+t/(\beta\log t)-1} G_+(-\ell_r).$$

By Lemma 38, there exists $M > 0$ such that $|\ell_r| \leq M\ell_t^*$, for all $r \leq t$. Therefore, since G_+ is monotone,

$$\mathbb{P}_+(R_s) \leq G_+(M\ell_t^*)^{t/(\beta\log t)}.$$

It follows that

$$\mathbb{P}_+(B_2) \leq t \cdot G_+(M\ell_t^*)^{t/(\beta\log t)}.$$

Since $G_+(x) = 1 - c \cdot x^{-k}$ for x large enough, and since ℓ_t^* is asymptotically at most $t^{1/(k+0.5)}$, we have that

$$\log G_+(M\ell_t^*) \leq -cM^{-k} \cdot t^{-k/(k+0.5)}.$$

Thus

$$\mathbb{P}_+(B_2) \leq t \cdot \exp\left(-cM^{-k} \cdot t^{1/(2k+1)}/(\beta\log t)\right) \leq t^{-2.1},$$

for all t large enough. This concludes the proof, because $\mathbb{P}_+(a_t = -1) = \mathbb{P}_+(A) + \mathbb{P}_+(B_1) + \mathbb{P}_+(B_2) \leq \kappa\frac{1}{t^{2.1}}$ for some constant κ.

□

Given this bound on the probability of mistakes, the proof of the main theorem of this section follows easily from Lemma 36.

Proof of Theorem 6. By Proposition 39, there exists $\kappa > 0$ such that $\mathbb{P}(a_t = -1 \mid \theta = +1) < \kappa\frac{1}{t^{2.1}}$ for all $t \geq 1$. Hence, by Lemma 36 $\mathbb{E}(T_L \mid \theta = +1) < \infty$. By a symmetric argument the same holds conditioned on $\theta = -1$. Thus, the expected time to learn is finite. □

Appendix B

APPENDIX TO CHAPTER 3

B.1 Binary distributions

The primary goal of this chapter is to prove Theorem 13, which gives the expected utility when the initial public belief equal to $1/2$ and the social planner chooses to stop learning upon arrival at levels N or $-N$. The proof uses some properties of the public belief, when signals are binary. For example, we show that it is a random walk in Lemma 10 and that the utility in the current period, that is in LP, is constant and equals to signal's precision, q. The latter one is showed in Lemma 12. Furthermore, we prove some general facts about the expected utility function and how it changes for different beliefs by varying stopping N.

Proof of Proposition 8. This is a standard fact as each strategy $\{k_t\}_{t=1}^{\infty}$ gives a linear function in belief p and maximum of convex functions is convex.

Moreover, if $p \in \{0, 1\}$ then if we just go with our belief, always choose 1 when $p = 1$ and always choose 0 otherwise, then our discounted utility is already going to be 1 which is the maximal possible utility we can get.

Also, the game is completely symmetric around $1/2$ and so our utility function $u(p)$ is also symmetric around $1/2$.

□

Proof of Lemma 9. If t is in the learning period then player t has to act according to her private signal: if $s_t = High$ choose 1, if $s_t = Low$ choose 0. Otherwise, she either goes against her signal and chooses a non optimal action or chooses the same action regardless of the s_t. Both situations are not allowed either by the statement of the lemma or by the assumption that agents are the expected utility maximizers. □

Proof of Lemma 10. Suppose there were n actions 1 and k actions 0 between times t_0 and t_{n+k}. Moreover, $\forall t' \in \{0, \dots n+k\}\ t' \in LP$, i.e. players' actions non trivially depended on their private signals. By lemma 9 we know that there were n $High$ signals and k Low ones.

Let us think in terms of the likelihood ratio. As we said before, l_t is going to be multiplied by $q/(1-q)$ n times and $(1-q)/q$ k times. These factors do not depend on l_t, therefore it also does not matter in which order we observe these actions, the final likelihood is going to be equal to

$$l_t \cdot \left(\frac{q}{1-q}\right)^n \left(\frac{1-q}{q}\right)^k = l_t \cdot \left(\frac{q}{1-q}\right)^{n-k}.$$

Therefore, the public belief is a random walk on a fixed lattice and its position is defined by the difference in the number of times we observed signals *High* and *Low* and $l_0 = 1$.

□

Proof of Lemma 11. First we show that once a we stop learning (no prices, $k_t = 1$) there is no reason to reenact LP latter. Notice that from the social planner's perspective this game is stationary. Therefore, if it is optimal to stop at time t_0 at levels N and $-N$ then at any time $t' > t_0$ there is no incentives to choose k_t different from 1 as we face the same problem as at time t_0 when it is optimal to stop.

Second, suppose we choose to stop upon arrival at either level $h > 0$ or $l < 0$. Then the expected utility at $1/2$ satisfies the following relation

$$u(0.5) = a_l + \mathbb{E}(\delta^t)u(l) + a_h + \mathbb{E}(\delta^t)u(h),$$

where a_h and a_l correspond to the utility that we get on the way to the corresponding boundary and does not depend on u. Let us sum up the first two terms and the last two and choose the biggest one. If they are equal - pick one. Let us say it is the second one, that corresponds to level h. Then due to the symmetry of u it is profitable for the social planner to choose the lower level to be $-h$ instead of l. This concludes the proof.

□

Proof of Lemma 12. Suppose at time t the public belief $p_t > q$ and we are still acquiring information/in the learning period. Depending on the private signal we get, our posterior will be equal to

$$\mu_t = \begin{cases} \dfrac{p_t q}{p_t q + (1-q)(1-p_t)} & \text{if we get the } High \text{ signal} \\ \dfrac{p_t(1-q)}{p_t(1-q) + q(1-p_t)} & \text{if we get the } Low \text{ signal.} \end{cases}$$

Moreover, the former signal happens with probability $(p_t q + (1 - p_t)(1 - q))$ as with probability p_t we are in the *High* state and with $(1 - p_t)$ - in the *Low* state. The analogous calculation gives that the latter signal occurs with probability $p_t(1-q) + q(1-p_t)$. Also, remember when we get the *Low* private signal we take action 0, so our expected utility is $1 - \mu_t$. Therefore, the expected utility today is equal to

$$\frac{p_t q}{(p_t q + (1-q)(1-p_t))}(p_t q + (1-q)(1-p_t)) +$$
$$+\left(1 - \frac{p_t(1-q)}{(p_t(1-q) + q(1-p_t))}\right)(p_t(1-q) + q(1-p_t)) = q.$$

And similarly we get the same expected utility if $p_t < 1 - q$.

□

Now we are ready to prove the main result.

Proof of Theorem 13. For notation convenience we write u_n instead of $u_H(p(n))$. We are going to start with solving (3.6) and then explain why $u_L(0.5)$ from (3.7) is the same. We know that the solution for the recurrence equation

$$u_n = (1-\delta)q + \delta(q u_{n+1} + (1-q)u_{n-1})$$

is

$$u_n = c_1\left(\frac{1}{2}\left(\frac{1}{\delta q} - \frac{\sqrt{1 - 4q\delta^2 + 4q^2\delta^2}}{\delta q}\right)\right)^n + c_2\left(\frac{1}{2}\left(\frac{1}{\delta q} + \frac{\sqrt{1 - 4q\delta^2 + 4q^2\delta^2}}{\delta q}\right)\right)^n + q,$$

for some constants c_1 and c_2. It is also straightforward to verify that these u_n's indeed satisfy our recurrence equation. Now we need to solve for c_1 and c_2 using boundary conditions $u_{-N} = (1-q)^N/((1-q)^N + q^N)$ and $u_N = q^N/((1-q)^N + q^N)$. This results in two equations

$$\begin{cases} \dfrac{(1-q)^N}{q^N + (1-q)^N} = c_1 a_1^{-N} + c_2 a_2^{-N} + q \\ \dfrac{q^N}{q^N + (1-q)^N} = c_1 a_1^N + c_2 a_2^N + q \end{cases}$$

where

$$a_i = \left(\frac{1}{2}\left(\frac{1}{\delta q} \pm \frac{\sqrt{1 - 4q\delta^2 + 4q^2\delta^2}}{\delta q}\right)\right).$$

Thus,

$$\begin{cases} c_1 = \left(\frac{(1-q)^N}{q^N + (1-q)^N} - q\right) a_1^N - c_2 a_2^{-N} a_1^N \\ c_2 \left(a_2^N - a_1^{2N} a_2^{-N}\right) = \frac{q^N}{q^N + (1-q)^N} - q + \left(q - \frac{(1-q)^N}{q^N + (1-q)^N}\right) a_1^{2N}. \end{cases}$$

It follows that

$$\begin{cases} c_2 = \frac{\left(\frac{q^N}{q^N+(1-q)^N} - q\right) + \left(q - \frac{(1-q)^N}{q^N+(1-q)^N}\right) a_1^{2N}}{a_2^N - a_1^{2N} a_2^{-N}} \\ c_1 = \frac{-\left(\frac{q^N}{q^N+(1-q)^N} - q\right) a_1^N a_2^N - \left(q - \frac{(1-q)^N}{q^N+(1-q)^N}\right) a_2^{-N} a_1^N}{a_2^N - a_1^{2N} a_2^{-N}} \end{cases}$$

Now we can plug this into our solution

$$\begin{aligned} u_0 &= c_1 a_1^0 + c_2 a_2^0 + q \\ &= \frac{-\left(\frac{q^N}{q^N+(1-q)^N} - q\right) a_1^N a_2^N - \left(q - \frac{(1-q)^N}{q^N+(1-q)^N}\right) a_2^N a_1^N}{a_2^N - a_1^{2N} a_2^{-N}} + \\ &+ \frac{\left(\frac{q^N}{q^N+(1-q)^N} - q\right) + \left(q - \frac{(1-q)^N}{q^N+(1-q)^N}\right) a_1^{2N}}{a_2^N - a_1^{2N} a_2^{-N}} \\ &= \frac{(a_2^N - a_1^N)\left(\left(\frac{q^N}{q^N+(1-q)^N} - q\right) - \left(q - \frac{(1-q)^N}{q^N+(1-q)^N}\right) a_1^N a_2^N\right)}{a_2^{2N} - a_1^{2N}} \\ &= \frac{\left(\frac{q^N}{q^N+(1-q)^N} - q\right) - \left(q - \frac{(1-q)^N}{q^N+(1-q)^N}\right)\left(\frac{1-q}{q}\right)^N}{a_2^N + a_1^N}, \end{aligned}$$

where in the third equation we multiplied everything by $a_2 \neq 0$ and in the last one we used the fact that

$$a_1 a_2 = \frac{1}{4}\left(\frac{1}{\delta^2 q^2} - \frac{1 - 4q\delta^2 + 4q^2\delta^2}{q^2\delta^2}\right) = \frac{1-q}{q}.$$

What is left to explain is why $u_L(0.5)$ would be the same as $u_H(0.5)$ which would save us time solving the second analogous system. Before, in state $High$, we had a random walk with a drift q towards the boundary with higher utility, level N. When we condition on state being Low we have the same drift but in the opposite direction, towards level $-N$. But notice that utilities of level $-N$ in state Low and level N in state $High$ are equal to each other. The same is true for the other two boundary utilities. Moreover, in both states we have the same underlying lattice for the random walk. Therefore, problem (3.7) is the same problem as (3.6) up to renaming levels. This concludes the proof.

□

Proof of Lemma 14. This is a classic Gambler's ruin problem (Feller, 1968)

□

Proof of Proposition 15 . For notation simplicity we will write $u(n)$ instead of $u(p(n))$.

We know, that the social planner adopts a symmetric strategy of stopping at levels N or $-N$. Suppose that if we increase the stopping boundaries from N to $N+1$ and from $-N$ to $-N-1$ correspondingly such that the utility at level N, $u(N)$, increases. Notice that the utility at level 0, given this strategy, is equal to the utility at the boundary levels multiplied by the expected discount factor plus some utility that we collect on the way to it. The randomness comes from the hitting the boundary levels time. We divide this expression in two parts: for the level N and $-N$.

$$u(0.5) = a_{-N} + u(-N)\,\mathbb{E}_{-N}(\delta^t) + a_N + u(N)\,\mathbb{E}_N(\delta^t),$$

where a_{-N}, a_N are constants $\mathbb{E}_{\pm N}(\delta^t)$ is the expected discounted factor until we hit the corresponding boundary.

Recall that u is symmetric and so $u(-N) = u(N)$, therefore if we increase $u(N)$ then we increase $u(0.5)$ also and vice versa. This concludes the proof.

□

Patient planner

Now we would like to see what happens to the optimal N^* and the expected utility, when the social planner becomes more patient. In other words, how does the optimal N^* and $u(0.5)$ behave, when $\delta \to 1$. We first establish a condition for when N^* does not go to ∞ when $\delta \to 1$ in Proposition 16: precision also has to go to 1 at a certain speed in this case. Secondly, we look at our expected utility as δ increases, but q stays fixed in Proposition 17. There, we stumble upon $(1-q)/q$ factor for the third time.

Proof of Proposition 16. Recall that the expected utility when we start with prior 1/2 (at level N) and stop the random walk upon arrival at levels $2N$ or 0 is

$$u(N) = \frac{-\left(-\frac{(1-q)^N}{(1-q)^N+q^N} + q\right)\left(\frac{1-q}{q}\right)^N + \left(\frac{q^N}{q^N+(1-q)^N} - q\right)}{a_1^N + a_2^N} + q,$$

where $a_1 < a_2$ and

$$a_i = \frac{1}{2}\left(\frac{1}{\delta q} \pm \sqrt{\frac{1}{(\delta q)^2} - \frac{4}{q} + 4}\right).$$

In order to find bounds on optimal $N \in \mathbb{R}$ we are going to focus on the first term of $u(N)$ as the second one is just a constant that does not affect optimal N, and let $q = 1 - \varepsilon$, $\delta = 1 - \gamma$ where $\varepsilon, \gamma \to 0$.

Notice, that $u(N)$ is single-peaked in N, therefore, if N maximizes it over $\mathbb{R}$ and N^* maximizes over $\mathbb{N}$ then N^* is either $\lceil N \rceil$ or $\lfloor N \rfloor$. Thus, bounds on N are going to give very tight bounds on N^*. First, consider a_i:

$$\begin{aligned}
a_i &= \frac{1}{2}\frac{1 \pm \sqrt{1 - 4\delta^2 q + 4q^2\delta^2 - 4q\delta + 4q\delta}}{\delta q} \\
&= \frac{1}{2}\frac{1 \pm \sqrt{(2q\delta - 1)^2 + 4q\delta(1-\delta)}}{\delta q} \\
&= \frac{1}{2}\frac{1 \pm (2q\delta - 1 + 2q\delta\gamma c_1) + o(\gamma)}{\delta q},
\end{aligned}$$

where the thirds equality comes from Taylor series expansion of $\sqrt{a^2 + x}$ around 0. Therefore,

$$\begin{aligned}
a_1 &= \frac{1}{\delta q} - 1 - O(\gamma) \\
a_2 &= 1 + O(\gamma)
\end{aligned}$$

Now we can take derivative of g.

$$g'(N) = \frac{\left(\left(\frac{(1-q)^N \ln(1-q)((1-q)^N+q^N) - (1-q)^N((1-q)^N \ln(1-q) + q^N \ln q)}{((1-q)^N+q^N)^2}\frac{(1-q)^N}{q^N} + m_1\left(\frac{1-q}{q}\right)^N \ln\frac{q}{1-q}\right)}{(a_1^N + a_2^N)^2} + \\
+ \frac{\left(\frac{q^N \ln q((1-q)^N+q^N) - q^N((1-q)^N \ln(1-q) + q^N \ln q)}{((1-q)^N+q^N)^2}\right)\Big)(a_1^N + a_2^N) - (a_1^N \ln a_1 + a_2^N \ln a_2)m_2}{(a_1^N + a_2^N)^2},$$

where

$$q \geq m_1 = \left(-\frac{(1-q)^N}{(1-q)^N + q^N} + q\right) \geq 2q - 1$$

$$1 - q \geq m_2 = -\left(-\frac{(1-q)^N}{(1-q)^N + q^N} + q\right)\left(\frac{1-q}{q}\right)^N + \left(\frac{q^N}{q^N + (1-q)^N} - q\right) \geq 0^1$$

In order to satisfy F.O.C. nominator should be equal to 0

$$\left(-\frac{(1-q)^N q^N \ln \frac{q}{1-q}}{((1-q)^N + q^N)^2}\frac{(1-q)^N}{q^N} + m_1\left(\frac{1-q}{q}\right)^N \ln\frac{q}{1-q} + \frac{q^N(1-q)^N \ln\frac{q}{1-q}}{((1-q)^N + q^N)^2}\right)(a_1^N + a_2^N) -$$
$$- (a_1^N \ln a_1 + a_2^N \ln a_2) m_2 = 0$$

Notice that the left-hand side is smaller than

$$\leq \left(m_1\left(\frac{1-q}{q}\right)^N \ln\frac{q}{1-q}\right)(c_1 a_2^N) - (a_2^N \ln(1 + O(\gamma)))m_2$$
$$\leq \left(m_1 c_1\left(\frac{1-q}{q}\right)^N \ln\frac{q}{1-q}\right) a_2^N - a_2^N \gamma m_2,$$

where $c_i > 1$.

In order for this to be non negative we need N to satisfy the following constraint

$$N \geq \frac{\ln\dfrac{m_2\gamma}{c_1 m_1 \ln(q/(1-q))}}{\ln\dfrac{1-q}{q}} \geq \frac{r_2 \ln\gamma}{\ln\varepsilon},$$

where r_2 goes to 1 as γ and δ go to 0. At the same time, notice that LHS is bigger than

$$\geq m_1\left(\frac{1-q}{q}\right)^N \ln\left(\frac{q}{1-q}\right) a_2^N - a_2^N \gamma c_3 m_2,$$

where $c_3 > 1$. In order for this to be non positive N should satisfy the following constraint

$$N \leq \frac{\ln\dfrac{\gamma c_3 m_2}{m_1 \ln(q/(1-q))}}{\ln\dfrac{1-q}{q}}.$$

[1] As N increases m_1 gets very close to q and m_2 to $1-q$.

Notice, that if $(\ln c_3\gamma)$ is smaller or proportional to $\ln(-\ln((1-q)/q))$ then N is always finite and the lower bound is satisfied when $q, \delta \to 1$, as

$$\frac{\ln\left(-\ln\frac{1-q}{q}\right)}{\ln\frac{1-q}{q}} \to 0$$

Otherwise, we get that

$$N \leq \frac{\ln\gamma}{r_1 \ln\varepsilon},$$

for some constant r_1. Moreover, if N satisfies these constraints then N^* satisfies

$$\left\lfloor \frac{\ln\gamma}{r_1\ln\varepsilon} r \right\rfloor \geq N^* \geq \left\lceil \frac{r_2\ln\gamma}{\ln\varepsilon} \right\rceil.$$

This concludes the proof.

□

Proof of Proposition 17. Recall that

$$u(N) = \frac{-\left(-\frac{(1-q)^N}{(1-q)^N+q^N} + q\right)\left(\frac{1-q}{q}\right)^N + \left(\frac{q^N}{q^N+(1-q)^N} - q\right)}{a_1^N + a_2^N} + q.$$

$$a_1 = \frac{1}{(1-\gamma)q} - 1 - O(\gamma)$$

$$a_2 = 1 + O(\gamma)$$

After plugging expressions for a_1 and a_2 in $u(N)$ we get

$$\begin{aligned} u(N) - q &= \frac{-\left(-\frac{(1-q)^N}{(1-q)^N+q^N} + q\right)\left(\frac{1-q}{q}\right)^N + \left(\frac{q^N}{q^N+(1-q)^N} - q\right)}{\left(\frac{1-q}{q} - O(\gamma)\right)^N + (1+O(\gamma))^N} \\ &= \frac{-\left(-\frac{(1-q)^N}{(1-q)^N+q^N} + q\right)\left(\frac{1-q}{q}\right)^N + \left(\frac{q^N}{q^N+(1-q)^N} - q\right)}{\left(\frac{1-q}{q}\right)^N - o(\gamma) + 1 + O(\gamma N)} \\ &= \frac{-\left(-\frac{(1-q)^N}{(1-q)^N+q^N} + q\right)\left(\frac{1-q}{q}\right)^N + \left(\frac{q^N}{q^N+(1-q)^N} - q\right)}{\left(\frac{1-q}{q}\right)^N + 1} - O(\gamma N). \end{aligned}$$

From the proof of Proposition 16 we know that when q is fixed N behaves as $O(\ln\gamma)$. Now, let us transform the first term into a more clear form

$$u(N) - q = \left(\frac{q^N}{q^N + (1-q)^N} - q\right) - \frac{\left(\frac{q^N-(1-q)^N}{q^N+(1-q)^N}\right)\left(\frac{1-q}{q}\right)^N}{\left(\frac{1-q}{q}\right)^N + 1} - O(\gamma N)$$

As $N \to \infty$ and $q > \frac{1}{2}$, $((1-q)/q)^N \to 0$. Suppose $((1-q)/q)^N = \nu$ then

$$\frac{q^N}{q^N + (1-q)^N} \geq 1 - \nu$$

and

$$\frac{\left(\frac{q^N-(1-q)^N}{q^N+(1-q)^N}\right)\left(\frac{1-q}{q}\right)^N}{\left(\frac{1-q}{q}\right)^N + 1} \leq \frac{\nu}{\nu + 1}$$
$$\leq \nu.$$

From the proof of Theorem 16 $\nu = O(\gamma)$. Hence,

$$u(N) - q \geq 1 - q - 2\nu - O(\gamma N) = 1 - q - O(\gamma \ln\gamma).$$

It means that as $N \to \infty$, $u(N)$ goes to its maximal value as $C((1-q)/q)^N \cdot N \to 0$, which is quicker than $((1-q)/q)^{Nk}$, for any $k < 1$. Also, as optimal N^* increases then absorbing beliefs are further away from $1/2$.

□

B.2 Continuous distributions

In this section we study continuous distributions. For our purposes they differ in a few aspects from the binary case. First of all, now we have to deal with more complicated expressions for expected loss and gain. Second of all, the set of prices that we can choose and that result in different outcomes is now continuous rather than binary. This also complicates the analysis, especially if we do not have an explicit expression for the expected utility. And the last, and probably the main difference, is that the public belief is not a random walk anymore, but has an intricate behavior.

We start with bounded signals and establish an analogous result, to the one in the classical sequential model (L. Smith and Sørensen, 2000): the underlying state of the world is never fully revealed if signals have bounded strength. To prove this, we look at what are the expected loss and gain, when the social planner chooses a price $k_t \neq 1$. As we do not know the exact formula for the expected utility function, we first bound the expected gain in the public belief and then translate it to the expected utility gain. The latter one uses the fact that $u(p)$ is convex and the absolute value of its derivative is less than 1.

Bounded private signals

Suppose without loss of generality that the modified LR satisfies $q/(1-q) \geq \frac{p_t}{(1-p_t)k_t} \geq 1/2$, so the $High$ state is more likely and we are still in the learning period (less than $q/(1-q)$). Given this, the total modified belief of agent t will be in the interval

$$\left[\frac{p_t}{(1-p_t)k_t} \cdot \frac{1-q}{q}; \frac{p_t}{(1-p_t)k_t} \cdot \frac{q}{1-q}\right].$$

In order to get the final modified likelihood-ratio equal to y in the interval above we need

$$\frac{g_H(x)}{g_L(x)} = y \cdot \frac{(1-p_t)k_t}{p_t}.$$

Before we calculate the expected loss and the expected gain from the price k_t we need a few more facts.

Note 40. *The public belief and the corresponding likelihood ratio satisfy the following relation*

$$p_t = \frac{l_t}{1+l_t}.$$

Furthermore, if the total belief $\mu_t > 1/2$ and agent t takes the action 0*, then the expected loss is $\mu_t - (1-\mu_t) = 2\mu_t - 1$.*

The following lemma helps us understand how the change in the LR translates to the change in the public belief.

Lemma 41. *If $l_{t+1} = l_t(1+\delta)$ then*

$$p_{t+1} - p_t = \delta p_t(1-p_t) + o(\delta(1-p_t)).$$

Proof of Lemma 41. By the definition of the likelihood ratio

$$\frac{p_t}{1-p_t} = x$$
$$\frac{p_{t+1}}{1-p_{t+1}} = x(1+\delta),$$

then

$$\begin{aligned} p_{t+1} - p_t &= \frac{x + x\delta}{1 + x + x\delta} - \frac{x}{1+x} \\ &= \frac{x\delta}{(1+x)(1+x+x\delta)} \\ &= \frac{\delta p_t(1-p_t)}{(1-p_t+p_t)(1-p_t+p_t(1+\delta))} \\ &= \delta p_t(1-p_t) + o(\delta(1-p_t)). \end{aligned}$$

□

Furthermore, the following relation is satisfied between the PDFs of the likelihood ratio (L. Smith and Sørensen, 2000)

$$f_H(x) = x f_L(x).$$

Therefore,

$$\begin{aligned} F_H(x) &= \int_{\frac{1-q}{q}}^{x} f_H(y)dy \\ &= \int_{\frac{1-q}{q}}^{x} y f_L(y)dy \\ &\leq x F_L(x) \end{aligned} \tag{B.1}$$

These facts help us prove one of the main results, that if private signals are bounded then the full revelation of the underlying state is not possible, unless $\delta = 1$.

Proof of Theorem 18. Let us start with the expected loss which is equal to

$$\begin{aligned} &\int_a^b \left(2\frac{x\frac{p_t}{1-p_t}}{1+x\frac{p_t}{1-p_t}} - 1\right)(p_t f_H(x) + (1-p_t) f_L(x))dx \\ &\geq c_1 p_t \left(F_H\left(\frac{1-p_t}{p_t}k_t\right) - F_H\left(\frac{1-q}{q}\right)\right) + (1-p_t)\left(F_L\left(\frac{1-p_t}{p_t}k_t\right) - F_L\left(\frac{1-q}{q}\right)\right) \\ &= c_1\left(p_t F_H\left(\frac{1-p_t}{p_t}k_t\right) + (1-p_t)F_L\left(\frac{1-p_t}{p_t}k_t\right)\right), \end{aligned}$$

where $a = (1-q)/q$, $b = (1-p_t)k_t/p_t$ and c_1 gets arbitrary close to 1 as p_t goes to 1. Notice that we can make $(1-p_t)k_t/p_t$ as close to $(1-q)/q$ as we want, which is the lowest value for the likelihood ratio.

Now let us find the expected gain. Remember that the modified public belief is $p_t/((1-p_t)k_t)$. Hence we get the following expressions for $p_{t+1}/(1-p_{t+1})$

$$\overline{\left(\frac{p_{t+1}}{1-p_{t+1}}\right)} = \frac{p_t}{1-p_t} \cdot \frac{1-F_H\left(\frac{(1-p_t)k_t}{p_t}\right)}{1-F_L\left(\frac{(1-p_t)k_t}{p_t}\right)}$$

if player t buys the new product (random walk goes up) and

$$\underline{\left(\frac{p_{t+1}}{1-p_{t+1}}\right)} = \frac{p_t}{1-p_t} \cdot \frac{F_H\left(\frac{(1-p_t)k_t}{p_t}\right)}{F_L\left(\frac{(1-p_t)k_t}{p_t}\right)}$$

if she does not (random walk goes down). Here, underline and overline represent possible LRs that can result from the current public belief p_t and price k_t depending on the action of player t.

As utility function on $[0.5, 1]$ is convex and symmetric around 0.5 we know that expected gain is less than $u(\overline{p_{t+1}}) - u(p_t)$ (as if we go up with probability 1). Which in its turn is less than $\overline{p_{t+1}} - p_t$ as u is also above 45 degree line and at 1 is equal to 1. In order to calculate the latter we should find $\overline{(p_{t+1}/(1-p_{t+1}))}$. We are thinking about $(1-p_t)k_t/p_t$ as a point close to the left end of the domain of F_H and F_L, so in the following calculation we are using notation $\varepsilon = (1-p_t)k_t/p_t$.

$$\begin{aligned}
\overline{\left(\frac{p_{t+1}}{1-p_{t+1}}\right)} &= \frac{p_t}{1-p_t}\left(\frac{1-F_H\left(\frac{(1-p_t)k_t}{p_t}\right)}{1-F_L\left(\frac{(1-p_t)k_t}{p_t}\right)}\right) \\
&\leq \frac{p_t}{1-p_t}\Big(1-c_2F_L(\varepsilon))(1+F_L(\varepsilon)+O(F_-^2(\varepsilon))\Big) \\
&= \frac{p_t}{1-p_t}(1+F_L(\varepsilon)(1-c_2)+O(F_L^2(\varepsilon)))
\end{aligned}$$

where $c_2 > 0$ is a constant comes from (B.1). And hence,

$$\begin{aligned}
\delta(\overline{p_{t+1}} - p_t) \leq \delta\frac{F_L(\varepsilon)(1-c_2)(1-p_t)}{p_t} + O(F_L^2(\varepsilon)) \\
= \frac{F_L(\varepsilon)(\delta(1-p_t+p_tc_2-c_2))}{p_t} + O(F_L^2(\varepsilon)).
\end{aligned}$$

Comparing it to $F_L(\varepsilon)((1-\delta)(1-p_t+p_tc_2))c_1$ tells us that for $p_t \to 1$ coefficient $\delta(1-p_t+p_tc_2-c_2)/(p_t)$ goes to 0 whereas the other one goes to $(1-\delta)c_2 > 0$. □

Unbounded private signals

Now we are going to consider an example of unbounded private signals. This is the only situation, when it is appropriate to talk about the asymptotic speed of learning, as in other cases the public belief does not converge to 0 or 1. The second main result is Theorem 19, which says that the optimal prices, k_t, are bounded away from 0 and ∞ and that it is optimal to choose a positive price, $k_t > 1$, when the public belief is high enough. We need to assume that p_t is high due to the lack of information about u and its derivative. As we mentioned above, if we use the results from the numerical calculation, it can be generalized for $p_t \geq 1/2$.

One of the corollaries is that we indeed have the asymptotic learning and the asymptotic speed of learning is also significantly increased, due to taxation of the good that is more likely to be better. The situation when $p_t < 1/2$ is symmetric.

Recall that the CDFs of the likelihood ratios are

$$F_H(y) = \mathbb{P}\left(\frac{g_H(s)}{g_L(s)} \leq y \middle| \theta = High\right) = \frac{y^2}{(1+y)^2}$$
$$F_L(y) = \mathbb{P}\left(\frac{g_H(s)}{g_L(s)} \leq y \middle| \theta = Low\right) = \frac{y^2+2y}{(1+y)^2},$$

for $y \in [0, \infty]$.

Proof of Theorem 19. Recall that if we apply non zero price k_t then we have the expected loss due to non optimal actions and the expected gain from bigger expected increase in public belief p_{t+1}. We start with the former one.

Suppose we have a public belief p_t and a price at this period is k_t. As was stated in the Section 3.1, agent t is going to buy the new product **iff**

$$\frac{p_t}{(1-p_t)k_t} \cdot \frac{g_H(x)}{g_L(x)} \geq 1.$$

Notice, that a player t will take a non optimal action only if her likelihood belief with price k_t is less than 1 and her likelihood belief without the price is above 1. More formally private likelihood ratio can take the following values

$$\frac{(1-p_t)k_t}{p_t} \geq \frac{g_H(x)}{g_L(x)} \geq \frac{1-p_t}{p_t}$$

If a person t gets a private likelihood belief in this interval and her total likelihood (without accounting for the price) is equal to y then the loss, due to taking a non optimal action, is equal to

$$\frac{2y}{1+y} - 1.$$

Therefore, expected loss $l(p_t, k_t)$ has the following form

$$\begin{aligned} l(p_t,k_t) &= \int_a^b \left(\frac{2x\frac{(1-p_t)}{p_t}}{x\frac{(1-p_t)}{p_t}+1} - 1 \right) (p_t f_H(x) + (1-p_t) f_L(x))\, dx \\ &\leq \left(\frac{2k_t}{k_t+1} - 1 \right) \int_a^b (p_t f_H(x) + (1-p_t) f_L(x))\, dx \\ &= \left(\frac{k_t-1}{k_t+1} \right) \left(p_t \left(F_H\left(\frac{(1-p_t)k_t}{p_t} \right) - F_H\left(\frac{1-p_t}{p_t} \right) \right) + \right. \\ &\left. + (1-p_t) \left(F_L\left(\frac{(1-p_t)k_t}{p_t} \right) - F_L\left(\frac{1-p_t}{p_t} \right) \right) \right), \end{aligned}$$

where $a = \frac{(1-p_t)}{p_t}$ and $b = \frac{(1-p_t)k_t}{p_t}$. After pluging in the expressions for F_i we get

$$l(p_t,k_t) = \frac{((-1+k_t)^2 p_t^2 (1-p_t)^2)}{(k_t+p_t-k_t p_t)^2}.$$

Now let us go to the expected gain term. Suppose that in period t the public belief $p_t > 1/2$ is high enough. If there is no price, $k_t = 1$, then the public belief in the next period goes to either a or b depending on the action of player t (buying the new and the old products correspondingly). And if we apply a price k_t then the public belief in the next period goes to a' and b' correspondingly

$$a = \frac{2-p_t}{3-2p_t},\ b = \frac{p_t}{(1+2p_t)},$$
$$a' = \frac{(2k_t(1-p_t)+p_t)}{(2k_t(1-p)+1)},\ b' = \frac{k_t p_t}{(k_t+2p_t)}$$

As the public belief is a martingale then in expectation it does not move. This means, that the expected gain is equal to the distance between two chords that connect $u(a), u(b)$ and $u(a'), u(b')$ at point p_t. This is depicted in Figure B.1. In other words, the expected gain is equal to the distance between d and d', where d belongs to the chord $u(a), u(b)$ and d' to the chord $u(a')u(b')$ and their p-coordinate is p_t. So the distance between d' and d is alongside the y-coordinate.

Let us calculate this distance. Denote by y_a and y_b - y-coordinates of a and b. Notice, that the y-coordinate of a' and b' are equal to $y_a + a_1(a'-a)$ and $y_b - b_1(b'-b)$ where a_1, b_1 are some constants that are in $[u'(a), u'(a')]$ and $[-u'(b'), -u'(b)]$.

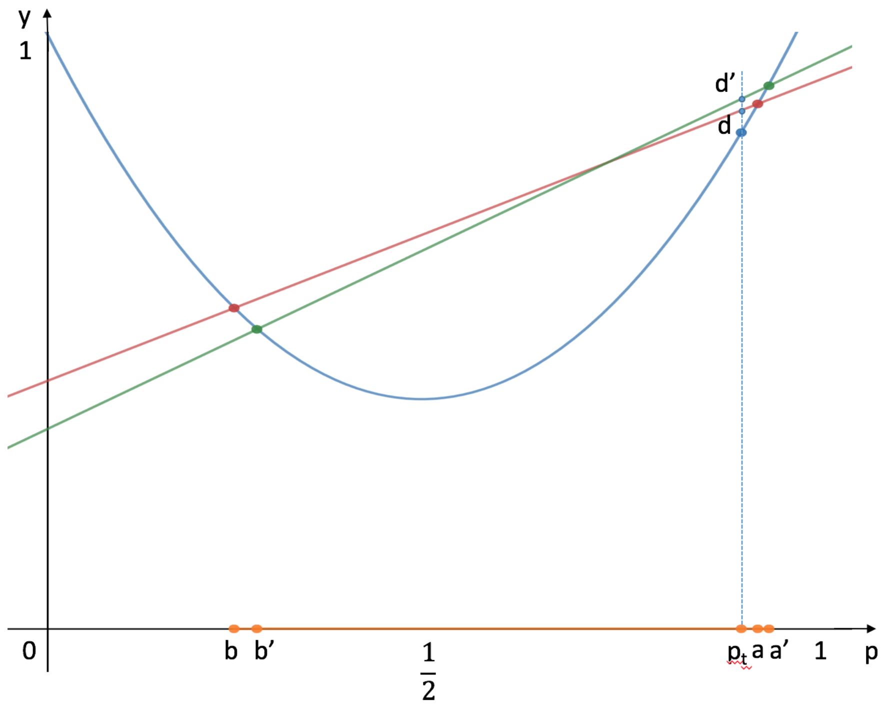

Where d and d' are the y-coordinates of point p_t on the cords $u(a)u(b)$ and $u(a')u(b')$ correspondingly.

Figure B.1: The expected gain induced by the cost.

$$d = y_b + \frac{y_a - y_b}{a - b}(p_t - b)$$
$$d' = y_b - b_1(b' - b) + \frac{y_a - y_b + a_1(a' - a) + b_1(b' - b)}{a' - b'}(p - b').$$

Pluging in the expressions for a, b, a', b' gives us the following formula

$$\text{expected gain}(p_t, k_t) = d' - d = \frac{((2a_1(-1 + k_t)(-1 + p_t)^2 p_t^2)}{((3 - 2p_t)(k_t + p_t - k_t p_t)^2))} -$$
$$- \frac{(2b_1(-1 + k_t)k_t(-1 + p_t)^2 p_t^2)}{((1 + 2p_t)(k_t + p_t - k_t p_t)^2)} +$$
$$+ \frac{((-1 + k_t)(-1 + p_t)^2 p_t^2(-1 - 3k_t + 2(-1 + k_t)p_t)(y_a - y_b))}{(k_t + p_t - k_t p_t)^2} =$$
$$= \frac{(-1 + k_t)(1 - p_t)^2 p_t^2}{(k_t + p_t - k_t p_t)^2} \left(\frac{2a_1}{3 - 2p_t} - \frac{2b_1 k_t}{1 + 2p_t} - (1 + 3k_t - 2(-1 + k_t)p_t)(y_a - y_b) \right).$$

Denote by P the terms in the parenthesis

$$P(p_t, k_t) = \left(\frac{2a_1}{3 - 2p_t} - \frac{2b_1 k_t}{1 + 2p_t} - (1 + 3k_t - 2(-1 + k_t)p_t)(y_a - y_b) \right).$$

Notice, that $y_a \geq y_b$ as a is closer to 1 than b is to 0 for $p_t \geq 1/2$

$$\frac{2-p_t}{3-2p_t} - \left(1 - \frac{p_t}{1+2p_t}\right) = \frac{2p_t - 1}{(3-2p_t)(1+2p_t)} \geq 0.$$

Let us now compare the expected gain and the expected loss due to the price k_t. We do this by subtracting the former one by the latter and taking the discount factor into account

$$\frac{\delta}{1-\delta}\text{expected gain} - \text{expected loss} = \frac{(-1+k_t)(1-p_t)^2 p_t^2}{((k_t+p_t-k_t p_t)^2}\cdot$$
$$\left(\frac{\delta}{1-\delta}\left(\frac{2a_1}{3-2p_t} - \frac{2b_1 k_t}{1+2p_t} - (1+3k_t-2(-1+k_t)p_t)(y_a-y_b)\right) - (k_t-1)\right) =$$
$$= \left(\frac{\delta}{1-\delta}P(p_t,k_t) - k_t + 1\right)\cdot\frac{(-1+k_t)(1-p_t)^2 p_t^2}{((k_t+p_t-k_t p_t)^2}$$

We can see that the sign of this expression is defined by the parenthesis and $k_t - 1$ terms.

Suppose that the expression in the parenthesis is positive for some $k_t > 1$ (and $p_t \notin \{0,1\}$)

$$\frac{\delta}{1-\delta}\left(\frac{2a_1}{3-2p_t} - \frac{2b_1 k_t}{1+2p_t} - (1+3k_t-2(-1+k_t)p_t)(y_a-y_b)\right) - (k_t-1) > 0 \tag{B.2}$$

Then $d' - d$ is positive as the sign here is defined by this parenthesis and $(k_t - 1)$. If we now make k less then 1, the expression above will increase and $k_t - 1$ will become negative, which will make the expected gain also negative. Therefore, if (B.2) holds for some $k_t > 1$ then the optimal k_t^* is not less than 1.

Now let see whether there exists $k_t > 1$ which gives higher utility than $k_t = 1$ for high enough p_t. As we said before, it is enough to show that (B.2) holds for some $k_t > 1$. The main challenge here is that we do not the utility function u and how convex it is.

Still, we can say the following: $y_a - y_b < a_1(a-b) < a_1/3$ and $b_1 < a_1$. This means that for p_t close enough to 1

$$P(p_t,k_t) > \left(\frac{2a_1}{3-2p_t} - \frac{2a_1 k_t}{1+2p_t} - (1+3k_t-2(-1+k_t)p_t)\frac{a_1}{3})\right).$$

For $k_t = 1$ this function is increasing in p_t and at $p_t = 1$ it is 0. The expected loss is also 0 for $k_t = 1$. But as the bounds for $y_a - y_b$ and b_1 are not tight, $P > 0$ at 1. As P and $-(k_t - 1)$ are also continuous functions of k_t there exists a threshold p such that for any $p_t > p$ the expected gain is positive for some $k_t = 1 + \varepsilon$. In other words, there exists $k_t > 1$ such that it is better to choose than $k_t = 1$ for p_t high enough. Moreover, if we use tighter constraints that we get from section 3.4 we can see that it holds for $p_t > 1/2$.

Furthermore, there exists neighborhoods of 0 and ∞ such that the optimal k_t^* does not lie in them for any p_t.

Let us start with 0

$$\begin{aligned} P(p_t, 0) =& \frac{2a_1}{3 - 2p_t} - (1 + 2p_t)(y_a - y_b) \\ \geq& \frac{2a_1}{3 - 2p_t} - \frac{(1 + 2p_t)a_1}{3}, \end{aligned}$$

where the inequality comes from the fact that $(y_a - y_b) \leq a_1(a - b) \leq a_1/3$. Furthermore,

$$\frac{2}{3 - 2p_t} > \frac{1 + 2p_t}{3},$$

as $(3 - 2p_t)(1 + 2p_t) \leq 4$. Thus, for k in some neighborhood of 0 $P(p_t, k) > 0$ and $k - 1 < 0$. Hence, there exists $\underline{k} > 0$ such that $k^* > \underline{k}$.

Also, it is obvious that for high enough k_t, $P(p_t, k_t)$ is negative. There is another case when b' jumps to the right side of $1/2$ such that y_b' becomes higher than y_b. Then we have a bit different expression for d'

$$d' = y_b + b_1(b' - (1 - b)) + \frac{y_a - y_b + a_1(a' - a) - b_1(b' - (1 - b))}{a' - b'}(p - b').$$

And hence,

$$\begin{aligned} d' - d - \text{expected loss}(p_t, k_t) = \frac{(-1 + p_t)^2}{(k_t + p_t - k_t p_t)^2} \Bigg(\frac{2a_1(-1 + k_t)p_t^2}{(3 - 2p_t)} + \\ + \frac{b1 k_t(-2p_t(1 + p_t) + k_t(-1 + 2p_t^2))}{1 + 2p_t} + (-1 + k_t)p_t^2(-1 - 3k_t + 2(-1 + k_t)p_t) \times \\ \times (y_a - y_b)) - (k_t - 1)^2 p_t^2 \Bigg). \end{aligned}$$

Define the term in the big parenthesis by $P'(p_t, k_t)$ and notice that

$$P'(p_t, k_t) \leq \left(2a_1(k_t - 1)p_t^2 + \frac{b_1 k_t \cdot k_t(2p_t^2 - 1)}{2} - (k_t - 1)^2 p_t^2 \right).$$

If we look at this as a function of k_t then it is a downward looking parabola, which implies that for k_t high enough P' is negative and, as a consequence, the expected gain is smaller than the expected loss. Thus, k_t does not go to ∞ as p_t increases.

Therefore there exist $\overline{k} > 1$ such that $k^* < \overline{k}$. This concludes the proof.

□

Proof of Corollary 20. This is an implication of two facts. The first one is that prices are bounded, so it is impossible to stop aggregating information. And second, public belief is a martingale and hence, converges. For more details see (L. Smith and Sørensen, 2000).

□

Proof of Corollary 21. Suppose at time t the likelihood ratio is l_t and $\theta = High$. Let us calculate l_{t+1} when agent t buys the new product with price k

$$l_{t+1} = l_t \cdot \frac{1 - F_H\left(\frac{k}{l_t}\right)}{1 - F_L\left(\frac{k}{l_t}\right)} = l_t \left(2\frac{k}{l_t} + 1)\right) = 2k + l_t.$$

This means that when everybody start taking the same correct action instead of adding 2 to the LR we going to add $2k$. Therefore, the convergence speed increases in k times.

□

Appendix C

APPENDIX TO CHAPTER 4

C.1 Agents' objective and its approximation

Before we prove Proposition 22 we need a few auxiliary results. As we said above, when the number of villages increases, the probability that there are multiple connections between i and some other person j of length at most 2, that include at least one weak link, vanishes. Furthermore, conditioning on the complement of this event does not change our expected utility at the limit as we show in the following proposition.

Proposition 42. *For a fixed K denote by A^c an event that there are multiple pathes, that include weak tie(s), between agent i and some other player of length* 1 *or* 2 *or there is a weak-weak connection that ends in village $\mathcal{V}_i$ in the realized graph (V, E). Then*

$$\lim_{N\to\infty} \mathbb{P}(A^c) = 0$$
$$\lim_{N\to\infty} \left| U_N(p_i, q_i, p, q)|A - U_N(p_i, q_i, p, q) \right| = 0$$

Proof of Proposition 42. Before we get to the event A, let us consider an auxiliary event A_0 that player i has no more than S_{weak} acquaintances. We know that the expected number of strong and weak connections are bounded by B and B/c_{weak} respectively because of the budget constraint. Let us prove that the probability that there are more than $S_{weak} = N^{\frac{1}{6}}$ weak connections goes to 0 as N increases. For agent i there are NK potential acquaintances and each link realizes independently with probability $q_i q$. It means that a random variable Z that equals the number of realized weak connections is distributed binomially with parameters $(NK, q_i q)$. There is the following well-known bound for the upper tail of the binomial distribution:

$$\mathbb{P}(Z \geq r) \leq exp\left(-NKD\left(\frac{r}{NK}\middle\| q_i q\right)\right),$$

where $D(a\|p) = a \log \frac{a}{p} + (1-a) \log \frac{1-a}{1-p}$ is a relative entropy between *Binomial*(a) and *Binomial*(p) and $NKq_iq \leq r \leq NK$. Notice, that $q_iq \leq B/(NKc_{weak})$. Hence,

$$\begin{aligned}
\mathbb{P}(Z \geq N^{\frac{1}{6}}) &\leq exp\left(-NK\left(\frac{N^{\frac{1}{6}}}{NK}\log\left(\frac{\frac{N^{\frac{1}{6}}}{NK}}{q_i q}\right)+\left(1-\frac{N^{\frac{1}{6}}}{NK}\right)\log\left(1-\frac{N^{\frac{1}{6}}}{NK}\right)\right)\right) \\
&\leq exp\left(-\left(N^{\frac{1}{6}}\log\left(\frac{\frac{N^{\frac{1}{6}}}{NK}}{q_i q}\right)+\left(NK-N^{\frac{1}{6}}\right)\left(-\frac{N^{\frac{1}{6}}}{NK}+O\left(\frac{1}{N^{\frac{5}{3}}}\right)\right)\right)\right) \\
&\leq exp\left(-\left(N^{\frac{1}{6}}\log\left(\frac{N^{\frac{1}{6}}}{\frac{B}{c_{weak}}}\right)+NK\left(-\frac{N^{\frac{1}{6}}}{NK}+O\left(\frac{1}{N^{\frac{5}{3}}}\right)\right)\right)\right) \\
&\leq exp\left(-N^{\frac{1}{6}}\right).
\end{aligned}$$

In the first inequality we used the fact that $\log(1-a)/(1-p)$ is negative, so decreasing p to 0 only increases the right-hand side. In the second line we used Taylor series for $\log(1-x)$ for small x. To get the third inequality recall that $q_i q \leq B/(NKc_{weak})$. Thus,

$$\mathbb{P}(Z < N^{\frac{1}{6}}) \geq 1 - exp^{-N^{\frac{1}{6}}},$$

which converges to 1 as N increases. Let us condition on the event, A_0, that there are at most $S_{weak} = N^{\frac{1}{6}}$ ties. Then the number of weak, weak-strong, strong-weak and weak-weak ties is bounded by $S_{weak} + 2S_{weak}B + S_{weak}^2 < 3N^{\frac{1}{3}}$. Denote this bound by $\bar{S}$. We are now going to show that the probability that these pathes do not overlap with each other goes to 1 as we increase the number of villages N.

For each path we are going to pick 1 out of N villages where it ends[1]. Note that if all these $\bar{S}$ villages that we chose are distinct then none of the pathes can overlap with each other. Denote by X a number of times that we pick some village that was already chosen before. Let us calculate the expected value of this random variable. We are going to choose villages sequentially with replacement. For the n-th pick denote by $\mathbb{1}_n$ an indicator function which equals to 1 if the n-th village we choose has already been chosen before. Then the expectation of X equals to the expected sum of these indicator functions from 1 to $\bar{S}$.

The first path can not be assigned to the village that was already chosen, as it is the first one. The second one will have the same village as the first path with probability $1/N$. For the third one, the probability is less than $2/N$. Let us eleborate this part. The first two pathes can be either in the same or in different villages with some probabilities, p_1 and $1-p_1$. Then, the probability the third one is in the same village

[1] For each weak-weak path we choose two villages: one for the first weak link and one for the second one.

as one of the first two is $p_1 1/N + (1 - p_1)2/N < 2/N$. The analogous argument is applied to the subsequent ones as well. Therefore,

$$\mathbb{E}(X) \leq 0 + \frac{1}{N} + \frac{2}{N} + \cdots + \frac{\bar{S}-1}{N} = \frac{\bar{S}(\bar{S}-1)}{2N}$$

Hence,

$$\frac{S(S-1)}{2N} \geq \mathbb{E}(X) = \sum_{i=1}^{k} \mathbb{P}(X=i)i \geq \sum_{i=1}^{k} \mathbb{P}(X=i) = 1 - \mathbb{P}(X=0)$$

Thus,

$$\mathbb{P}(X=0) \geq 1 - \frac{\bar{S}(\bar{S}-1)}{2N} > 1 - \frac{4N^{\frac{2}{3}}}{2N} = 1 - \frac{2}{N^{\frac{1}{3}}}$$

which goes to 1 as N increases. Let us call the event when $X = 0$, A_1 and its complement A_1^c.

Now we will calculate the probability of the event that there is at least one weak-weak path that comes back to village $\mathcal{V}_i$. Denote by $n_w^{\mathcal{V}_i}$ the number of weak-weak pathes that end up in $\mathcal{V}_i$ and by A_2 an event that $n_w^{\mathcal{V}_i} = 0$. This means that none of the acquaintance of i has an acquaintance in i's village. If everyone has at most $\bar{S}$ acquaintances and each of them has at most $\bar{S}$ weak links then the probability of A_2^c, that at least one weak-weak path comes back to $\mathcal{V}_i$, is equal to

$$\mathbb{P}(n_w^{\mathcal{V}_i} > 0) = 1 - \mathbb{P}(n_w^{\mathcal{V}_i} = 0) = 1 - \left(1 - \frac{1}{N}\right)^{\bar{S}^2},$$

which converges to 0 as $N \to \infty$. Then the probability of either events A_1^c or A_2^c happening is less than or equal to the sum of their corresponding probabilities.

$$\mathbb{P}(A_1^c \cup A_2^c) \leq \mathbb{P}(A_1^c) + (A_2^c) \leq 1 - \left(1 - \frac{1}{N}\right)^{\bar{S}^2} + \frac{\bar{S}(\bar{S}-1)}{2N},$$

which goes to 0 as N goes to ∞. Therefore, $\mathbb{P}(A_1 \cap A_2)$ converges to 1.

Notice, that we calculated probabilities of $A_1 \cap A_2$ conditioned on A_0. But because all their corresponding probabilities converge to 1 the limit of the probability of the unconditional event $A_1 \cap A_2$ is also 1.

$$\lim_{N\to\infty}\mathbb{P}(A_1\cap A_2)=\lim_{N\to\infty}\Big(\mathbb{P}(A_1\cap A_2|A_0)\mathbb{P}(A_0)+\mathbb{P}(A_1\cap A_2|A_0^c)\mathbb{P}(A_0^c)\Big)=1.$$

Now we will calculate the absolute difference in expected utility, U_N, when we condition on $A = A_1 \cap A_2$ and when we do not.

$$\begin{aligned}|\mathbb{E}(U|A)-\mathbb{E}(U)|=&\left|\frac{\mathbb{E}(U\mathbb{1}_A)}{1-\mathbb{P}(A^c)}-\mathbb{E}(U)\right|\\ =&\left|\frac{\mathbb{E}(U)\mathbb{P}(A^c)-\mathbb{E}(U\mathbb{1}_{A^c})}{1-\mathbb{P}(A^c)}\right|\\ \leq&4||U||_\infty\mathbb{P}(A^c),\end{aligned}$$

as the denominator is bigger than $1/2$ and the numerator is less than $2||U||_\infty\mathbb{P}(A^c)$. The right-hand side converges to 0, because U is bounded and $\mathbb{P}(A^c)$ goes to 0 as N increases. Therefore, we can use our approximation of the utility function.

□

Now, if we calculate $U_N|A$ we will almost have U_∞. Recall, that when we condition on the event A, there are only 5 groups of people that affect i's utility: $N_1^s(i)$, $N_1^w(i)$, $N_1^{ss}(i)$, $N_2^{ws}(i)$, N_2^{ww}. The following proposition calculates the expected number of people in these different types of neighborhoods and $U_N|A$ itself.

Proposition 43.

$$\begin{aligned}U(p_i,q_i,p,q)|A=&\Big(1+(K-1)p_ip+HKNq_iq+(K-1)(1-p_ip)\times\\ &\times(1-(1-p_ip^3)^{K-2})+\\ &+H\left(p_ipq^2K(K-1)N+p^2q_iqK(K-1)N\right)+\\ &+H^2q_iq^3K^2N(N-1)\Big)\end{aligned}$$

Proof of Proposition 43. Before we calculate the expected number of people in different groups we need some notation. Denote by

- $N_1^s(i)$ – a set of strong connections of i;
- $N_1^w(i)$ – a set of weak connections of i;
- $N_2^{ss}(i)$ – a set of strong-strong connections of i, that are not in $N_1^s(i)$;
- $N_2^{ws}(i)$ – a set of weak-strong or strong-weak connections of i;

- $N_2^{ww}(i)$ – a set of weak-weak connections of i.

Let us proceed. Expectation is a linear operator, hence we can calculate the expected number of people for each of those five groups seperately. Let us start with $N_i^s(i)$. There are $K-1$ people in the village besides player i and she is going to connect to each of them with probability $p_i p$. Therefore, $\mathbb{E}_{(V,E)}\,|N_i^s(i)| = (K-1)p_i p$. Analogously, $\mathbb{E}_{(V,E)}\,|N_i^w(i)| = KNq_i q$.

In order for an agent to be in $N_2^{ss}(i)$ she can not be connected directly to i but there has to be exactly one person between them. The probability that player j is not connected to i through person k (i, j, k are in the same village $\mathcal{V}_i$) is $(1-p_i p_k p_k p_j)$. As all edges appear (or not) independently of each other, the probability that i is connected to j at distance 2 is

$$1 - \prod_{\substack{k\in\mathcal{V}_i\\ k\neq i,j}} (1 - p_i p_k p_k p_j).$$

The subtrahend is the probability that i is not connected to j through any player k in the same village. After we apply symmetry (every player, except i, plays strategy (p,q)) and remember that j can not be at distance 1 from i, i.e. there can not be a strong edge between i and j, we get the desired formula:

$$\mathbb{E}_{\{p_t,q_t\}}\,|N_2^{ss}(i)| = (K-1)(1-p_i p)(1-(1-p_i p^3)^{K-2})$$

Now let us calculate the expected number of strong-weak connections. For each individual j who is a firend of player i, $i, j \in \mathcal{V}_i$, we need to calculate how many acquaintances j has, n_j^w, and add them up.

$$\mathbb{E}\left(\sum_{j\in\mathcal{V}_i} \mathbb{1}_{e_{ij}\in E_s}\, n_j^w\right) = \sum_{j\in\mathcal{V}_i} \mathbb{E}\left(\mathbb{1}_{e_{ij}\in E_s}\, n_j^w\right) = \sum_{j\in V_i} \mathbb{E}\left(\mathbb{1}_{e_{ij}\in E_s}\right) \mathbb{E}\, n_j^w =$$
$$= p_i p(K-1)q^2 KN.$$

In the equations above we used linearity of expectation and the fact that $\mathbb{1}_{e_{ij}\in E_s}$ and n_j^w are independent, hence, expectation of their product is the product of their expectations.

For the weak-strong and weak-weak connections we can do the analogous calculations. Denote by n_k^s the number of strong friends that player k,$k \notin \mathcal{V}_i$, has.

$$N_2^{ws}(i) = \mathbb{E}\left(\sum_{k\notin\mathcal{V}_i} \mathbb{1}_{e_{ik}\in E_w}\, n_k^s\right) = \sum_{k\notin\mathcal{V}_i} \mathbb{E}\left(\mathbb{1}_{e_{ik}\in E_w}\, n_k^s\right) = \sum_{k\notin\mathcal{V}_i} \mathbb{E}\left(\mathbb{1}_{e_{ik}\in E_w}\right) \mathbb{E}\, n_k^s =$$

$$= KNq_iq(K-1)p^2,$$

where n_k^s and $\mathbb{1}_{e_{ik}\in E_w}$ are independent.

For the weak-weak links denote by n_k^w the number of acquaintances player $k \notin \mathcal{V}_i$ has.

$$N_2^{ww}(i) = \mathbb{E}\left(\sum_{k\notin\mathcal{V}_i} \mathbb{1}_{e_ik\in E_w}\, n_k^w\right) = \sum_{k\notin\mathcal{V}_i} \mathbb{E}\left(\mathbb{1}_{e_ik\in E_w}\right) \mathbb{E}\, n_k^w = NKq_iq(N-1)Kq^2.$$

In the last equation we have $N(N-1)$ instead of N^2 because in the event A weak-weak links do not come back to the same village that player i is. This means that i's acquaintances can only pick from $N-1$ other villages to create weak-weak connections.

Recall that because we are conditioning on event A then none of these pathes, that we calculated above, overlap with each other. Now we just subsitute these calculations into the objective function to complete the proof.

□

Now we are ready to prove Proposition 22.

Proof of Proposition 22. From Proposition 42 we know that U_N uniformly converges to $U|A$. Define L to be $M_{weak}(K-1)/B$. Then we have $q_iq = L(M_{strong} -$

$p_i p$). Notice that

$$\begin{aligned}
U|A =& \Big(1 + (K-1)p_i p + \pi_w KNL(M_{strong} - p_i p) + (K-1)(1 - p_i p)\times \\
& \times (1 - (1 - p_i p^3)^{K-2}) + \\
& + \pi_w \left(p_i p L(M_{strong} - p^2)K(K-1)N + p^2 L(M_{strong} - p_i p)K(K-1)N\right) + \\
& + \pi_w^2 L(M_{strong} - p_i p)L(M_{strong} - p^2)K^2 N(N-1)\Big) \\
=& 1 + (K-1)p_i p + \frac{\pi_w}{c_{weak}}(K-1)(M_{strong} - p_i p) + (K-1)(1 - p_i p)\times \\
& \times (1 - (1 - p_i p^3)^{K-2}) + \\
& + \frac{\pi_w}{c_{weak}}(K-1)^2\Big(p_i p(M_{strong} - p^2) + \\
& + p^2(M_{strong} - p_i p)\Big) + \frac{\pi_w^2 (N-1)}{c_{weak}^2 N}(M_{strong} - p_i p)\times \\
& \times (M_{strong} - p^2)(K-1)^2 \\
=& U_\infty - \frac{\pi_w^2}{c_{weak}^2}\frac{1}{N}(M_{strong} - p_i p)(M_{strong} - p^2)(K-1)^2 \\
=& U_\infty + O\left(\frac{1}{N}\right),
\end{aligned}$$

where in the first equation we used notation from (4.1).

This means that $U|A$ uniformly converges to U_∞. Therefore, U_N also uniformly converges to U_∞.

Let us show that equilibrium of U_N, p_N, can only be within ε-neighborhood of the equilibrium of U_∞. First, notice that U_N and U_∞ both have two trivial equilibria: 1) all players invest only in friends and 2) all players invest only in acquaintances.

Now let us deal with the non trivial one. For each p, $U_\infty(p_i, q_i(p_i), p, q(p))$ has a unique maximum with respect to p_i, as it is a concave function. Denote by $f(p)$ the value of p_i at which the corresponding maximum is attained. This is a continuous function as U_∞ is a polynomial of a fixed degree of p_i and p, hence, a small change in p will require a small change in p_i to maintain $\partial U_\infty / \partial p_i = 0$. Then $h(p) = U_\infty(p, q(p), p, q(p)) - U_\infty(f(p), q(f(p)), p, q(p))$ is also a continuous function due to a triangle inequality. Therefore, h attains its minimum, δ, on $[\varepsilon, p^* - \varepsilon] \cup [p^* + \varepsilon, \sqrt{M_{strong}} - \varepsilon]$ which is positive. Hence, for $\bar{N}$ big enough such that $U_\infty - U_N < \delta/3$ for $N > \bar{N}$ there are no equilibria of the initial utility function ouside of those neighborhoods of the equilibria of U_∞. This concludes one direction.

Now let us show the other direction. There are three equilibria of U_∞: two trivial and one non trivial. Also, U_N has the same two trivial equilibria. Let us prove that there is a non trivial equilibrium of U_N in some negihborhood of p^*. Let $p^1 = p^* - \varepsilon$ and $p^2 = p^* + \varepsilon$. For both p^1 and p^2 let us find N_1 and N_2 such that the maximums of $U_\infty(p_i, q_i(p_i), p^1, q(p^1))$ and $U_\infty(p_i, q_i(p_i), p^2, q(p^2))$ with respect to p_i are within the δ-neighborhood of the maximums of $U_N(p_i, q_i(p_i), p^1, q(p^1))$ and $U_N(p_i, q_i(p_i), p^2, q(p^2))$ correspondingly. Notice, that $U_N(p_i, q_i(p), p, q(p))$ can not have maximums outside of these neighborhoods as $U_\infty(p_i, q_i(p), p, q(p))$ is single-peaked for a fixed p with respect to p_i. Pick δ to be equal to $\min(h(p^1), h(p^2))$. We can do this because U_N uniformly converges U_∞. Let us pick $N_0 = \max(N_1, N_2)$. Then we know that at $p = p^1$ the maximum of $U_N(p_i, q_i(p_i), p^1, q(p^1))$ is to the right of $p_i = p^1$[2] and at $p = p^2$ the maximum of $U_N(p_i, q_i(p_i), p^2, q(p^2))$ is to the left of $p_i = p^2$ for all $N > N_0$. Hence, at some point $p^1 \leq p_N \leq p^2$ the maximum of $U_N(p_i, q_i(p_i), p, q(p))$ crosses the line of $p_i = p$. This is a symmetric non trivial equilibrium of U_N in p^*'s the neighborhood.

Therefore, there are equilibria of U_N within the ε-neighborhood of (trivial and non trivial) equilibria of U_∞. Moreover, there are no equilibria of the initial utility function outside of those neighborhoods.

□

C.2 Equilibrium and comparative statics

Proof of Theorem 23. In order to prove this theorem we need to show that $(\partial U_\infty / \partial p_i)\Big|_{p_i = p} = 0$ has a unique non trivial solution.

[2] $\partial U_\infty / \partial p_i\big|_{p_i = p}$ is positive at $p = p^1$ and negative at $p = p^2$.

$$
\begin{aligned}
\left.\frac{\partial U_\infty}{\partial p_i}\right|_{p_i=p} =& p{\pi_w}^2K^2N^2L^2\left(p^2-(M_{strong}-p^2)\right)+\\
&+p\pi_w(K-1)KNL\left((M_{strong}-p^2)-p^2\right)-\\
&-\pi_w(K-1)KNp^3L-p\pi_wKNL+p(K-1)\left(1-p^4\right)^{K-2}+\\
&+(K-2)(K-1)p^3\left(1-p^2\right)\left(1-p^4\right)^{K-3}\\
=&(K-2)(K-1)p^3\left(1-p^2\right)\left(1-p^4\right)^{K-3}+p(K-1)\left(1-p^4\right)^{K-2}+\\
&-\pi_wKNLp\Big(1-(K-1)(M_{strong}-2p^2)+\pi_wKNL(M_{strong}-p^2)\Big)\\
=&(K-2)(K-1)p^3\left(1-p^2\right)\left(1-p^4\right)^{K-3}+p(K-1)\left(1-p^4\right)^{K-2}+\\
&-\frac{\pi_w}{c_{weak}}(K-1)p\Bigg(1-(K-1)(M_{strong}-2p^2)+\\
&+\frac{\pi_w}{c_{weak}}(K-1)(M_{strong}-p^2)\Bigg).
\end{aligned}
$$

Notice, that we can do this if $q_iq > 0$, otherwise, all terms that include a weak tie become 0 and we do not differentiate them. When $q_iq = 0$, U_∞ it is a trivial equilibrium of U_∞, as well as of U_N, to choose $p_i = \sqrt{M_{strong}}\,\forall\, i$. Denote $\left.(\partial U_\infty/\partial p_i)\right|_{p_i=p}$ by F.

$$
\begin{aligned}
F =&(K-2)(K-1)p^3\left(1-p^2\right)\left(1-p^4\right)^{K-3}+p(K-1)\left(1-p^4\right)^{K-2}+\\
&-\frac{\pi_w}{c_{weak}}(K-1)p\left(1-(K-1)(M_{strong}-2p^2)+\frac{\pi_w}{c_{weak}}(K-1)(M_{strong}-p^2)\right)\\
=&T_1(p)+T_2(p).
\end{aligned}
$$

where T_1 is the term on the 1st line and T_2 – on the 2nd one. We will prove that F has a unique non trivial root.

Under the assumptions of the theorem, F has the following form: it is 0 at $p = 0$, has a positive derivative there (when $c_{weak} > \pi_w$), it is negative at $\sqrt{M_{strong}}$ (when $e^{\frac{-B^2}{K-1}}(1 - M_{strong}) < \pi_w/c_{weak}$) and crosses 0 only once on $(0, \sqrt{M_{strong}})$. Let us show this. It is clear that $F(0) = 0$, so we will skip it.

$$
\begin{aligned}
F'(0) = (K-1)\Bigg(&\frac{\pi_w\left(\pi_w(K-1)\left(3p^2 - M_{strong}\right)\right.}{c_{weak}^2}+ \\
&+ \frac{-c_{weak}\left(6(K-1)p^2 - KM_{strong} + M_{strong}\right)\left.\right)}{c_{weak}^2} \\
&+ \frac{1}{c_{weak}^2} - \left(1-p^2\right)^{K-3}\left(p^2+1\right)^{K-4}\Big((K-1)(4K-7)p^6 + (6K-11)p^4+ \\
&+ (5-3K)p^2 - 1\Big)\Bigg)\Bigg|_{p=0} \\
&= (K-1)\left(\frac{\pi_w\left((K-1)M_{strong}(c_{weak}-\pi_w) - c_{weak}\right)}{c_{weak}^2} + 1\right) \\
&= (K-1)\left(\frac{\pi_w\left((K-1)M(c_{weak}-\pi_w) + c_{weak}\left(\frac{c_{weak}}{\pi_w} - 1\right)\right)}{c_{weak}^2}\right) \\
&= (K-1)\left(\frac{\pi_w\left((c_{weak}-\pi_w)\left((K-1)M_{strong} + \frac{c_{weak}}{\pi_w}\right)\right)}{c_{weak}^2}\right) \\
&> 0,
\end{aligned}
$$

as $c_{weak} > \pi_w$.

This implies that within some ε-neighborhood of $p = 0$, F is positive. Now we are going to show that at the other end, at $p = \sqrt{M_{strong}}$, it is negative.

$$
\begin{aligned}
F(\sqrt{M_{strong}}) =& (K-1)\Bigg(M_{strong}^{\frac{1}{2}}(1-M_{strong})^{K-2}(M_{strong}+1)^{K-3}\times \\
&\times((K-1)M_{strong}+1) - \frac{\pi_w M_{strong}^{\frac{1}{2}}}{c_{weak}}(1+(K-1)M_{strong})\Bigg) \\
=& -(K-1)(1+(K-1)M_{strong})M_{strong}^{\frac{1}{2}}\Bigg(\frac{\pi_w}{c_{weak}} - (1-M_{strong}^2)^{K-3}\times \\
&\times(1-M_{strong})\Bigg) \\
=& -(K-1)(1+(K-1)M_{strong})M_{strong}^{\frac{1}{2}}\Bigg(\frac{\pi_w}{c_{weak}} - \\
&-\left(1-\frac{B^2}{(K-1)^2}\right)^{K-3}\left(1-\frac{B}{K-1}\right)\Bigg) \\
\le& -(K-1)(1+(K-1)M_{strong})M_{strong}^{\frac{1}{2}}\Bigg(\frac{\pi_w}{c_{weak}} - e^{-\frac{B^2(K-3)}{(K-1)^2}}\times \\
&\times\left(1-\frac{B}{K-1}\right)\Bigg) \\
<& 0,
\end{aligned}
$$

as $\frac{\pi_w}{c_{weak}} > e^{-\frac{B^2(K-3)}{(K-1)^2}}\left(1-\frac{B}{K-1}\right)$.

The fact that $F(\sqrt{M_{strong}}) < 0$, $F(\varepsilon) > 0$ for some small $\varepsilon > 0$ and it is a continuous function implies that $\exists\, p^* \in (0, \sqrt{M_{strong}})$ such that $F(p^*) = 0$. Now we just need to make sure that such non trivial equilibrium p^* is unique.

Because our function is positive and increasing near the 0 and is negative near $\sqrt{M_{strong}}$ then, at some point, $p^\star$, its derivative has to become negative, $F'(p^\star) < 0$, before $F(p)$ becomes negative.

$$
\begin{aligned}
F'(p) =& (K-1)(1-p^4)^{K-4}(1-p^2)\Bigg(1+(3K-5)p^2-(6K-11)p^4-(K-1)\times \\
&\times(4K-7)p^6 - \frac{3(K-1)p^2\left(2-\frac{\pi_w}{c_{weak}}\right)\frac{\pi_w}{c_{weak}}}{(1-p^4)^{K-4}(1-p^2)}\Bigg)+ \\
&+(K-1)\frac{\pi_w}{c_{weak}}\left(M_{strong}(K-1)\left(1-\frac{\pi_w}{c_{weak}}\right)-1\right)
\end{aligned}
$$

We would like $F'(p)$ to stay negative after the first time it becomes less than 0. Then it can not cross 0 more than once. Assume that the constant term above is positive[3]. It is a part of the derivative of the second summand of $F(p)$, $T_2(p)$. Notice, that $T_2(p)$ has to be negative at the non trivial equilibrium, because $T_1(p)$ is always positive and their sum is 0. Also, $T_2(p)$ is concave, $T_2(0) = 0$ and $T_2'(0) \geq 0$, therefore, before T_2 becomes negative, its derivative has to become less than 0. Denote by p_1 a point at which $T_2'(p_1) = 0$. Then we can rewrite $F'(p)$ as

$$F'(p) = (K-1)(1-p^4)^{K-4}(1-p^2)A(p),$$

where

$$A(p) = ap^2 - 2p^2 - bp^4 - cp^6 - \frac{ad(p^2 - p_1^2)}{(1-p^2)\left(1-p^4\right)^{K-4}} + 1,$$

and $a = 3(K-1)$, $b = (6K-11)$, $c = (K-1)(4K-7)$, $d = (2-\pi_w/c_{weak})\pi_w/c_{weak}$.

It will be sufficient to show that once A becomes negative it stays negative as $(1-p^4)^{K-4}(1-p^2)$ is always positive for $p \in [0, \sqrt{M_{strong}}]$. Notice that $A(0) > 0$, so before it gets negative its derivative has to become negative.

$$\begin{aligned} A'(p) =& 2p\Big(- ad\left(p^2+1\right)^{3-K}\left(1-p^2\right)^{2-K}\Big(2(K-4)p^4 + p^2\left((-2K+7){p_1}^2+1\right) - \\ & - {p_1}^2 + 1\Big) + a - 2bp^2 - 3cp^4 - 2\Big) \end{aligned}$$

Now, call the term in parentheses $B(p)$, then B' is negative for $p > 0$.

$$\begin{aligned} B'(p) =& -4p\left(1-p^4\right)^{-K}\Big(\left(1-p^4\right)^K\left(b+3cp^2\right) + ad\left(p^2-1\right)\left(p^2+1\right)^2 \times \\ & \times \Big(-(K-4)(2K-7)p^6 + (2K-7)p^4\left((K-3){p_1}^2 - 1\right) + \\ & + p^2\Big(2(K-3){p_1}^2 - 3K + 10\Big) + (K-3){p_1}^2 - 1\Big)\Big) \\ =& -4p\left(1-p^4\right)^{-K}\Big(\left(1-p^4\right)^K\left(b+3cp^2\right) + ad\left(p^2-1\right)\left(p^2+1\right)^2 \times \\ & \times \Big(-1 - (2K-7)p^4\Big((K-4)p^2 + 1 - (K-3){p_1}^2\Big) - \\ & - p^2\Big(-2(K-3){p_1}^2 + 2K + 7\Big) - (K-3)(p^2 - {p_1}^2)\Big)\Big) \\ & < 0. \end{aligned}$$

---

[3] If it is negative then instead of having $p^2 - p_1^2$ below we are going to have $p^2 + p_1^2$ which will only help with the analogous proof.

Thus, $B'(p)$ is negative. This implies the following. When A' becomes negative it stays negative. As $A(0) > 0$, before A becomes negative its derivative, A', has to become negative. Therefore, once A starts decreasing it continues decreasing from that point on. So when A becomes negative it stays negative and this is what we wanted. Hence, $F(p)$ crosses 0 only once on $(0, \sqrt{M_{strong}})$. Let us call this point p^*. This is a unique non trivial equilibrium of U_∞.

Now, assume that $\pi_w \geq c_{weak}$ then $\pi_w / c_{weak} \geq 1 > (1 - M_{strong}^2)^{K-3}(1 - M_{strong})$. Furthermore, it means that $F(0) = 0$ and F' is negative on $(0, \sqrt{M_{strong}})$. Thus, there is only one equilibrium of U_∞: $\forall\, i\ p_i = 0$ and $q_i = M_{weak}$. When $\pi_w > c_{weak}$ weak links are more beneficial at both distance 1 and 2 and so no one wants to invest in friends.

On the other hand, if $\pi_w / c_{weak} < (1 - M_{strong}^2)^{K-3}(1 - M_{strong})$ then $\pi_w < c_{weak}$. It means that $F(\sqrt{M_{strong}}) > 0$. Hence, $F(p) = 0$ does not have a solution and is positive on $(0, \sqrt{M_{strong}}]$. So without the two conditions of the theorem we have only trivial symmetric equilibria.

To finish the proof we will show that we found the maximum and not the minimum.

$$\left.\frac{\partial^2 U_\infty}{\partial p_i^2}\right|_{p_i = p} = -\,(K-2)(K-1)p^4\left(1 - p^2\right)^{K-3}\left(p^2 + 1\right)^{K-4}\left((K-1)p^2 + 2\right)$$

$$<0.$$

Thus, the extremum we found is indeed the maximum. This concludes the proof.

□

Proof of Proposition 24. To prove this proposition, let us calculate derivatives of $F(p)$ with respect to c_{weak}, π_w and K. We know from Theorem 23 the derivative of F at p^* is negative. Hence, if we know the derivatives of F and their signs with respect to those parameters we will be able to calculate the corresponding comparative statics. Notice, that only the second summand of F depends on these 3 parameters, which simplifies calculations.

$$1)\ \ \frac{\partial F(p^*)}{\partial c_{weak}} = \partial\Bigg(-\frac{\pi_w(K-1)}{c_{weak}} p^* \left(1-(K-1)(M_{strong}-2p^{*^2})\right) +$$

$$+\frac{\pi_w}{c_{weak}}(K-1)(M_{strong}-p^{*^2})\Bigg)/\partial c_{weak}$$

$$= \frac{\pi_w p^* c_{weak}\left(1-(K-1)(M_{strong}-2p^{*^2})\right)}{c_{weak}^3} +$$

$$+\pi_w p^* \frac{2\pi_w(K-1)\left(M_{strong}-p^{*^2}\right)}{c_{weak}^3}.$$

Remember, that the second summand of F is negative at the equilibrium, $T_2(p^*) < 0$ from Theorem 23, therefore

$$\frac{\partial F(p^*)}{\partial c_{weak}} = \frac{-T_2(p^*) + \frac{\pi_w^2 p^*(K-1)(M_{strong}-p^{*^2})}{c_{weak}^2}}{c_{weak}}$$
$$> 0.$$

Therefore, as we increase c_{weak}, p^* also increases.

$$2)\ \ \frac{\partial F(p^*)}{\partial H} = -\frac{p^*\left(1-(K-1)(M_{strong}-2p^{*^2}) + 2\frac{\pi_w}{c_{weak}}(K-1)\left(M_{strong}-p^{*^2}\right)\right)}{c_{weak}}$$
$$< 0.$$

We get the last inequality in the same way we argued in 1) that $T_2(p^*)/\pi_w$ minus something negative is negative at p^*.

$$3)\ \ \frac{\partial F'(p)}{\partial K} = p\Bigg(\left(1-p^2\right)\left(1-p^4\right)^{K-3}\left(\left((K-1)p^2+1\right)\log\left(1-p^4\right)+p^2\right) -$$

$$-p^2\left(2-\frac{\pi_w}{c_{weak}}\right)\frac{\pi_w}{c_{weak}}\Bigg)$$

Notice, that $\log(1-p^4) < 0$. Furthermore, at the approximate equilibrium p^* we have

$$F(p^*) = 0$$

$$0 =(K-1)p^* \left(1-p^{*^2}\right)^{K-2} \left(p^{*^2}+1\right)^{K-3} \left((K-1)p^{*^2}+1\right) -$$
$$- \frac{\pi_w(K-1)}{c_{weak}} p^* \left(1-(K-1)(M_{strong}-2p^{*^2}) + \frac{\pi_w}{c_{weak}}(K-1)(M_{strong}-p^{*^2})\right).$$

Hence,

$$\left(1-p^{*^2}\right)\left(1-p^{*^4}\right)^{K-3} = \frac{\pi_w \left(1-(K-1)(M_{strong}-2p^{*^2})\right)}{c_{weak}\left((K-1)p^{*^2}+1\right)} + \\ + \frac{\frac{\pi_w^2}{c_{weak}}(K-1)(M-p^{*^2})}{c_{weak}\left((K-1)p^{*^2}+1\right)}$$

Let us substitute this into $\partial F'(p)/\partial K$, but skip the term with the logarithm and the p term outside the parentheses.

$$\frac{\pi_w \left(1-(K-1)(M_{strong}-2p^{*^2}) + \frac{\pi_w}{c_{weak}}(K-1)(M_{strong}-p^{*^2})\right)}{c_{weak}\left((K-1)p^{*^2}+1\right)} p^{*^2} -$$
$$- p^{*^2}\left(2-\frac{\pi_w}{c_{weak}}\right)\frac{\pi_w}{c_{weak}} =$$
$$\frac{\pi_w p^{*^2}}{c_{weak}((K-1)p^{*^2}+1)}\Bigg(1-(K-1)M_{strong}\left(1-\frac{\pi_w}{c_{weak}}\right) + p^{*^2}(K-1)\left(2-\frac{\pi_w}{c_{weak}}\right) -$$
$$-\left(2-\frac{\pi_w}{c_{weak}}\right)((K-1)p^{*^2}+1)\Bigg) =$$
$$\frac{\pi_w p^2}{c_{weak}((K-1)p^2+1)}\left(-1+\frac{\pi_w}{c_{weak}} - (K-1)M_{strong}\left(1-\frac{\pi_w}{c_{weak}}\right)\right) < 0.$$

Thus, the whole $\partial F'(p)/\partial K$ is also negative at $p = p^*$. So the probability of having a strong friend is decreasing as we increase K. This concludes our proof.

□

Proof of Lemma 25. First we are going to show that $p^{*^2} > q^{*^2}$ in equilibrium. Remember that from the proof of Theorem 23

$$p^{*^2} > M\frac{c_{weak}-\pi_w}{2c_{weak}-\pi_w} - \frac{c_{weak}}{(K-1)(2c_{weak}-\pi_w)} = \frac{(B-1)c_{weak}-B\pi_w}{(K-1)(2c_{weak}-\pi_w)}$$

as the $T_2(p)$ term has to be negative. Hence,

$$q^{*^2} = L(M - p^{*^2}) < \frac{(K-1)}{KNc_{weak}} \left(\frac{B(2c_{weak} - \pi_w) - (B-1)c_{weak} + B\pi_w}{(K-1)(2c_{weak} - \pi_w)} - \right) = \frac{(B+1)}{KN(2c_{weak} - \pi_w)}.$$

Let us look at their difference

$$p^{*^2} - q^{*^2} > \frac{(B-1)c_{weak} - B\pi_w}{(K-1)(2c_{weak} - \pi_w)} - \frac{(B+1)}{KN(2c_{weak} - \pi_w)} = \frac{(B-1)c_{weak}KN - B\pi_w KN - (B+1)(K-1)}{(K-1)KN(2c_{weak} - \pi_w)}.$$

For this difference to be greater than 0 we need to have

$$c_{weak} > \frac{B}{(B-1)}\pi_w + \frac{(B+1)(K-1)}{(B-1)KN}.$$

Now let us remember a well-known fact about Erdős - Rényi random graphs: the golbal clustering coefficient for friends' network is equal to $p^{*^2} + O\left((KN)^{-0.5}\right)$ and the global CC for acquanintances' network is $q^{*^2} + O\left((KN)^{-0.5}\right)$. Using this fact and the inequality that we got above we can conclude that when c_{weak} is not too close to π_w or when N is big enough (so $p^* > 0$ and $q^* < p^*$ for big enough N) the CC of the friends' network is bigger than the CC of the acquaintances'. This concludes the proof. □

C.3 Socially optimal network

Proof of Theorem 26. To prove this theorem we are going to differentiate $U_\infty(p, q(p), p, q(p)) \equiv U_{optimal}(p)$ and find its non-trivial root $p_{optimal}$, which is also unique. Then we will show that $p_{optimal}$ has to be bigger than p^*.

$$U'_{optimal}(p) = 2p(K-1)\Bigg(2(K-2)p^2\left(1-p^2\right)^{K-2}\left(p^2+1\right)^{K-3} + \left(1-p^4\right)^{K-2} - \frac{\pi_w}{c_{weak}}\left(1 - 2(K-1)(M_{strong} - 2p^2)\right) + \frac{2\pi_w}{c_{weak}}(K-1)\left(M_{strong} - p^2\right)\Bigg).$$

As we can see it greatly resembles $\left(\partial U_\infty/\partial p_i\right)\Big|_{p_i=p}$. Before we proceed, let us show that $U'_{optimal}$ has the same shape as $\left(\partial U_\infty/\partial p_i\right)\Big|_{p_i=p}$: at 0 it equals 0 and has a positive derivative. Moreover, once its derivative becomes negative it stays negative.

$$
\begin{aligned}
U'_{optimal}(0) =& 0 \\
U''_{optimal}(0) =& \left(\frac{\pi_w \left(2(K-1)M_{strong}\left(1-\frac{\pi_w}{c_{weak}}\right) - 6(K-1)p^2\left(2-\frac{\pi_w}{c_{weak}}\right)-1\right)}{c_{weak}} - \right. \\
& - \left(1-p^2\right)^{K-3}\left(p^2+1\right)^{K-4}\left((2K-3)(4K-7)p^6+(8K-15)p^4+\right. \\
& \left.\left.\left. + (11-6K)p^2-1\right)\right)\right|_{p=0} \\
=& \frac{\pi_w\left(2(K-1)M_{strong}\left(1-\frac{\pi_w}{c_{weak}}\right)-1\right)}{c_{weak}} + 1 \\
>& 0,
\end{aligned}
$$

as $c_{weak} > \pi_w$.

$$
\begin{aligned}
U'_{optimal}(\sqrt{M_{strong}}) =& 2\sqrt{M_{strong}}(K-1)\left(\left(1-p^2\right)^{K-2}\left(p^2+1\right)^{K-3}\times\right. \\
& \times\left((2K-3)p^2+1\right) - \\
& - \frac{\pi_w}{c_{weak}}\left(1-2(K-1)(M_{strong}-2M_{strong})\right) + \\
& \left. + \frac{2\pi_w}{c_{weak}}(K-1)\left(M_{strong}-M_{strong}\right)\right) \\
=& -\left(\frac{\pi_w}{c_{weak}} - (1-M_{strong})^{-2+K}(1+M_{strong})^{-3+K}\right)\times \\
& \times(1+(-2+2K)M_{strong}) - (1-M_{strong})^{-2+K}\times \\
& \times(1+M_{strong})^{-3+K}M_{strong} \\
<& 0.
\end{aligned}
$$

We got the last inequality from the same condition we had in Theorem 23, which is necessary for an existance of the equilibrum, p^*.

$$
\begin{aligned}
U''_{optimal}(p) = & 2(K-1)\Bigg(\left(1-p^2\right)^{K-3}\left(p^2+1\right)^{K-4}\Big(-(2K-3)(4K-7)p^6- \\
& -(8K-15)p^4+(6K-11)p^2+1\Big)+ \\
& +\frac{\pi_w}{c_{weak}}\Bigg(2\left((K-1)M_{strong}\left(1-\frac{\pi_w}{c_{weak}}\right)-3(K-1)p^2\left(2-\frac{\pi_w}{c_{weak}}\right)\right)- \\
& -1\Bigg)\Bigg)
\end{aligned}
$$

As we can see $U''_{optimal}$ is positive at $p = \varepsilon$ for ε small enough, so its derivative has to become negative before $U'_{optimal}$ becomes negative. To prove that once $U'_{optimal}$ becomes negative it stays negative we can do an analogous exercise to the one we did in Theorem 23.

Now let see if the solution to $\partial U_{optimal}/\partial p = 0$, $p_{optimal}$, is bigger than p^*, the non trivial solution to $\left(\partial U_\infty/\partial p_i\right)\Big|_{p_i=p} = 0$. To do this we are going to look at the sign of $U'_{optimal}(p^*)$.

$$
\begin{aligned}
U'_{optimal}(p^*) =& 2p^*(K-1)\Bigg(2(K-2){p^*}^2\left(1-{p^*}^2\right)^{K-2}\left({p^*}^2+1\right)^{K-3}+\left(1-{p^*}^4\right)^{K-2}-\\
&-\frac{\pi_w}{c_{weak}}\bigg(1-2(K-1)(M_{strong}-2{p^*}^2)+2\frac{\pi_w}{c_{weak}}(K-1)\times\\
&\times(M_{strong}-{p^*}^2)\bigg)\Bigg)\\
=& p^*(K-1)\Bigg(4(K-2){p^*}^2\left(1-{p^*}^2\right)^{K-2}\left({p^*}^2+1\right)^{K-3}+2\left(1-{p^*}^4\right)^{K-2}-\\
&-\frac{\pi_w}{c_{weak}}\bigg(2-4(K-1)(M_{strong}-2{p^*}^2)+4\frac{\pi_w}{c_{weak}}(K-1)\times\\
&\times(M_{strong}-{p^*}^2)\bigg)\Bigg)\\
=& p^*(K-1)\Bigg(2(K-2){p^*}^2\left(1-{p^*}^2\right)^{K-2}\left({p^*}^2+1\right)^{K-3}-\\
&-\frac{\pi_w}{c_{weak}}\bigg(-2(K-1)(M_{strong}-2{p^*}^2)+2\frac{\pi_w}{c_{weak}}(K-1)\times\\
&\times(M_{strong}-{p^*}^2)\bigg)\Bigg)\\
=& 2p^*(K-1)\Bigg((K-2){p^*}^2\left(1-{p^*}^2\right)^{K-2}\left({p^*}^2+1\right)^{K-3}+\\
&-\frac{\pi_w}{c_{weak}}\bigg(-(K-1)(M_{strong}-2{p^*}^2)+\frac{\pi_w}{c_{weak}}(K-1)\times\\
&\times(M_{strong}-{p^*}^2)\bigg)\Bigg),
\end{aligned}
$$

where in the third equation we used the fact that $\left(\partial U_\infty/\partial p_i\right)\Big|_{p_i=p}(p^*)=0$. We are going to use it one more time now to get

$$
U_{optimal}(p^*) = 2p^*(K-1)\left(-\left(1-{p^*}^4\right)^{K-2}+\frac{\pi_w}{c_{weak}}\right)
$$

Therefore, if

$$
\frac{\pi_w}{c_{weak}} > \left(1-{p^*}^4\right)^{K-2}
$$

then $U'_{optimal}(p^*) > 0$, hence, $p_{optimal} > p^*$.

To conclude the proof, let us show that $p_{optimal}$ is indeed a maximum not a minimum. Notice, that $U_{optimal}(p)$ is a polynomial of one variable that has one extreme point on $(0, \sqrt{M_{strong}}]$. As we showed before, $U'_{optimal}$ is non negative in $\varepsilon-$neighborhood of 0 and positive outside of $p = 0$. Furthermore, $U'_{optimal}(\sqrt{M_{strong}}) < 0$. This means that the maximum of our function is not attained at the endpoints. Thus, the extreme point that we found is the maximum. This concludes our proof.

□