



Essays in Dynamic Games

Citation

Ishii, Yuhta. 2014. Essays in Dynamic Games. Doctoral dissertation, Harvard University.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:12274534>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Essays in Dynamic Games

A dissertation presented

by

Yuhta Ishii

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

May 2014

© 2014 Yuhta Ishii
All rights reserved.

Dissertation Advisors:
Professor Drew Fudenberg
Professor Attila Ambrus

Author:
Yuhta Ishii

Essays in Dynamic Games

Abstract

This dissertation presents three independent essays. Chapter 1, which is joint work with Mira Frick, studies a model of innovation adoption by a large population of long-lived consumers who face stochastic opportunities to adopt an innovation of uncertain quality. We study how the potential for social learning in an economy affects consumers' informational incentives and how these in turn shape the aggregate adoption dynamics of an innovation. For a class of Poisson learning processes, we establish the existence and uniqueness of equilibria. In line with empirical findings, equilibrium adoption patterns are either S-shaped or feature successions of concave bursts. In the former case, our analysis predicts a novel *saturation effect*: Due to informational free-riding, increased opportunities for social learning necessarily lead to temporary slow-downs in learning and do not produce welfare gains.

Chapter 2, which is joint work with Drew Fudenberg and Scott D. Kominers, extends the folk theorem of repeated games to settings in which players' information about others' play arrives with stochastic lags. To prove the folk theorem, we construct equilibria in "delayed-response strategies," which ensure that players wait long enough to respond to signals that with high probability all relevant signals are received before players respond. To do so, we extend past work on private monitoring to obtain folk theorems despite the small residual amount of private information.

Finally Chapter 3 demonstrates how uncertainty over patience can generate strong reputation effects that are weak when the long-run player's level of patience is common knowledge. With uncertainty over patience, these strong reputation effects are the result of a contagion effect initiated by very patient types: the most patient types have a strict

incentive to play the beneficial action in all equilibria which in turn incentivizes those with smaller levels of patience to also play this action. Our main result shows that even when very patient types are extremely small in probability, these contagion effects are very strong so that types with intermediate levels of patience obtain high payoffs in all equilibria.

Contents

Abstract	iii
Acknowledgments	x
Introduction	1
1 Innovation Adoption by Forward-Looking Social Learners	3
1.1 Introduction	3
1.1.1 Related Literature and Outline	9
1.2 Model	14
1.2.1 The Game	14
1.2.2 Learning	14
1.2.3 Equilibrium	17
1.3 Cooperative Benchmark	20
1.3.1 Cooperative Benchmark under Perfect Good News	21
1.3.2 Cooperative Benchmark under Perfect Bad News	23
1.4 Quasi-Single Crossing Property for Equilibrium Incentives	23
1.5 Perfect Bad News	26
1.5.1 Equilibrium Characterization	26
1.5.2 Shape of Adoption Curve	33
1.5.3 Welfare	36
1.5.4 The Effect of Increased Opportunities for Social Learning	39
1.6 Perfect Good News	46
1.6.1 Equilibrium Characterization	46
1.6.2 Shape of Adoption Curve	51
1.6.3 The Effect of Increased Opportunities for Social Learning	52
1.7 More Social Learning Can Hurt: An Example	58
1.8 Conclusion	61
2 Delayed Response Strategies in Repeated Games with Observation Lags	64
2.1 Introduction	64
2.1.1 Related Work	68

2.2	General Model	69
2.3	Bounded Lags	72
2.4	Lagged Almost-Perfect Monitoring with Two Players	73
2.4.1	Model	74
2.4.2	Auxiliary Repeated Game with Rare Observation Lags	75
2.4.3	The Repeated Game with Frequent Observation Lags	80
2.5	Lagged Public Monitoring	81
2.5.1	Model	81
2.5.2	Structure of the Observation Lags	83
2.5.3	The Folk Theorem	83
2.5.4	Private Monitoring Game with Communication	84
2.5.5	The Repeated Game with Observation Lags	93
2.6	Discussion and Conclusion	97
3	Contagious Commitment via Unknown Patience	99
3.1	Introduction	99
3.1.1	Literature Review	102
3.2	Model	104
3.3	Existence of Equilibrium	108
3.4	Reputation and Contagion	109
3.4.1	Proof Sketch	110
3.4.2	Details of the Proof	113
3.4.3	Concluding the Proof of Theorem 3.4.2	117
3.5	Negative Results	124
3.5.1	Necessity of Behavioral Commitment Types	124
3.5.2	Necessity of Arbitrarily Patient Rational Types	125
3.6	Numerical Simulations	127
3.6.1	Strictly Dominated Action: Product Choice Game	127
3.7	Conclusion and Discussion	131
	References	134
	Appendix A Appendix to Chapter 1	139
A.1	Preliminary Mathematical Tools for Equilibrium Analysis	139
A.1.1	General Properties of the Value to Waiting	139
A.1.2	Special Properties of the Value to Waiting under PBN	142
A.2	Quasi-Single Crossing Property for Equilibrium Incentives	143
A.2.1	Proof of Theorem 1.4.1 under Perfect Good News	143
A.2.2	Proof of Theorem 1.4.1 under Perfect Bad News	144

A.3	Equilibrium Uniqueness and Characterization	146
A.3.1	Equilibrium Characterization under Perfect Bad News	146
A.3.2	Equilibrium Characterization under Perfect Good News	151
A.4	Violation of Condition 1.5.2 under Perfect Bad News	152
A.5	Inefficiency of Equilibria	153
A.5.1	Inefficiency under PBN	153
A.6	Comparative Statics under PBN	154
A.6.1	Saturation Effect: Proof of Proposition 1.5.9	154
A.6.2	Learning Dynamics	155
A.6.3	Adoption Behavior	158
A.7	Comparative Statics under PGN	160
A.7.1	Adoption Behavior	160
A.8	Heterogeneous Discount Rate Example	161
A.9	Cooperative Benchmark	163
A.9.1	Perfect Bad News	163
A.9.2	Perfect Good News	164
Appendix B	Appendix to Chapter 2	165
B.1	Definitions of s_i^G and s_i^B	165
B.2	Details of the Proof of Theorem 2.4.3	166
B.2.1	Proof of Lemma 2.4.4	166
B.2.2	Proof of Lemma 2.4.5	169
B.3	Conditions Guaranteeing Small γ_i	170
Appendix C	Appendix to Chapter 3	171
C.1	Proof of Existence	171

List of Tables

3.1	Stage Game for Numerical Simulation	127
3.2	Stage Game for Example	131

List of Figures

1.1	Perfect Bad News	25
1.2	Perfect Good News	25
1.3	Partition of (p_t, Λ_t) when $\varepsilon < \rho$	32
1.4	Examples of S-shaped adoption curves (Source: Narayanan and O'Connor (2010), Figure 2.1.)	34
1.5	Adoption curve conditional on no breakdowns ($\varepsilon = 0$)	35
1.6	Ratio of equilibrium welfare to socially optimal welfare	38
1.7	Changes in adoption levels of a good product under perfect bad news ($\hat{\lambda} > \lambda$)	44
1.8	Adoption Curves under Perfect Good News (blue = breakthrough before t^* ; yellow = breakthrough after t^* ; pink = bad quality)	52
1.9	“Adoption” patterns for various blockbuster movies (Source: McLaren and DePaolo (2009))	53
3.1	Lower Bound on $\sigma_1[a_1^1]$	113
3.2	The payoff of type δ in the robust equilibrium.	132
3.3	The payoff of type δ in the robust equilibrium.	132
A.1	Partition of (p_t, Λ_t) when $\varepsilon \geq \rho$	152

Acknowledgments

I am deeply indebted to my advisors, Attila Ambrus, Drew Fudenberg, and Eric Maskin, for their constant guidance and encouragement. Always full of helpful advice and insights, learning from them has been an absolute pleasure. This dissertation would not have been possible without their generosity.

I am also very grateful to my coauthors from whom I learned much: Attila Ambrus, Mira Frick, Drew Fudenberg, Yuichiro Kamada, and Scott D. Kominers.

Personally, I want to also thank my wonderful parents and sister, and my always supportive partner Maisu. These acknowledgments would be incomplete without recognition of my classmates at Harvard University of which there are too many to name. Finally I am grateful for all of the long-lasting friendships I formed at Stanford University: Peter Ahn, Jose Benchimol, Can Baran, Momchil Filev, Carl Moore, Johnny Sun, and Rohan Vilms. I am very fortunate to be surrounded by such fantastic, supportive people.

Chapter 1, titled “Innovation Adoption by Forward-Looking Social Learners,” is joint work with Mira Frick. Chapter 2, titled “Delayed-Response Strategies in Repeated Games with Observation Lags,” was coauthored with Drew Fudenberg and Scott D. Kominers and published in the *Journal of Economic Theory* in 2014.

The Department of Economics at Harvard University provided generous financial support while I was working on this dissertation.

Dedicated to my parents, sister, and Maisu.

Introduction

This dissertation presents three independent essays. Chapter 1, which is joint work with Mira Frick, studies a model of innovation adoption by a large population of long-lived consumers that faces stochastic opportunities to adopt an innovation of uncertain quality. Consumers are *social learners*: Over time, news about the product's quality is generated endogenously, based on the experiences of past adopters. We analyze how the potential for social learning in an economy affects consumers' informational incentives and how these in turn shape the aggregate adoption dynamics of an innovation. Our main results highlight the importance of two features of the economy: The extent to which consumers are *forward-looking* and the nature of news events through which social learning occurs. When consumers are forward-looking social learners, the trade-off between the benefit of adopting the innovation at any given time and the option value of waiting for endogenous news can generate rich aggregate adoption dynamics, even in the absence of any consumer heterogeneity. The dynamics of this trade-off and the extent to which it is affected by increased opportunities for social learning interact in interesting ways with the news process of the economy. For a class of Poisson learning processes, we establish the existence and uniqueness of equilibria. In line with empirical findings, equilibrium adoption patterns are either S-shaped or feature successions of concave bursts. In the former case, our analysis predicts a novel *saturation effect*: Due to informational free-riding, increased opportunities for social learning necessarily lead to temporary slow-downs in learning and do not produce welfare gains.

Chapter 2, which is joint work with Drew Fudenberg and Scott D. Kominers, extends

the folk theorem of repeated games to two settings in which players' information about others' play arrives with stochastic lags. In our first model, signals are almost-perfect if and when they do arrive, that is, each player either observes an almost-perfect signal of period- t play with some lag or else never sees a signal of period- t play. In the second model, the information structure corresponds to a lagged form of imperfect public monitoring, and players are allowed to communicate via cheap-talk messages at the end of each period. In each case, we construct equilibria in "delayed-response strategies," which ensure that players wait long enough to respond to signals that with high probability all relevant signals are received before players respond. To do so, we extend past work on private monitoring to obtain folk theorems despite the small residual amount of private information.

Finally Chapter 3 demonstrates how uncertainty over patience can generate strong reputation effects that are unavailable when the long-run player's level of patience is common knowledge. With uncertainty over patience, these strong reputation effects are the result of a contagion effect initiated by very patient types: the most patient types have a strict incentive to play the beneficial action in all equilibria which in turn incentivizes those with smaller levels of patience to also play this action. Our main result shows that even when very patient types are extremely small in probability, these contagion effects are very strong so that types with intermediate levels of patience obtain high payoffs in all equilibria.

Chapter 1

Innovation Adoption by Forward-Looking Social Learners¹

1.1 Introduction

Suppose an innovation of uncertain quality, such as a novel medical treatment or a new piece of software, is released into the market. In recent years, the rise of internet-based review sites, retail platforms, search engines, video-sharing websites, and social networking sites (such as Yelp, Amazon, Google, YouTube, and Facebook) has greatly increased the *potential for social learning* about the innovation: An individual's treatment success story or discovery of a bug in the software is much more likely to find its way into the public domain; and there are more people than ever who have access to this common pool of consumer-generated information.

We analyze how the potential for social learning in an economy affects consumers' informational incentives and how these in turn shape the aggregate adoption dynamics of an innovation. Our main results highlight the importance of two features of the economy: The extent to which consumers are *forward-looking* and the nature of news events through which social learning occurs. In choosing whether to adopt an innovation, forward-looking

¹Co-authored with Mira Frick.

consumers recognize the option value of waiting for more information. With social learning, information is created *endogenously*, based on the consumption experiences of past adopters. In equilibrium, adoption levels must therefore strike a balance: If too many consumers adopt at any given time, then too much information is available in the future and all consumers would rather wait; conversely, if too few consumers adopt, it might not be worthwhile for anyone to wait. We show that the dynamics of this trade-off and the extent to which it is affected by increased opportunities for social learning depend crucially on the kind of information consumers expect to acquire by waiting. In line with numerous empirical findings, our analysis predicts adoption patterns that are either S-shaped or feature successions of concave bursts, suggesting novel micro-foundations for these observations. We also make new predictions regarding the impact of *increased* opportunities for social learning on consumer welfare, on equilibrium learning dynamics, and on observed adoption behavior.

In our model, an innovation of fixed, but uncertain quality (better or worse than the status quo) is introduced to a large population of forward-looking consumers. Consumers are (ex ante) identical, sharing the same prior about the quality of the innovation, the same discount rate, and the same tastes for good and bad quality. At each instant in continuous time, consumers receive stochastic opportunities to adopt the innovation. A consumer who receives an opportunity must choose whether to irreversibly adopt the innovation or to delay his decision until the next opportunity. In equilibrium, consumers optimally trade off the opportunity cost of delays against the benefit to learning more about the quality of the innovation.

Learning about the innovation is summarized by a public signal process, representing news that is obtained endogenously—based on the experiences of previous adopters; and possibly also from exogenous sources, such as professional critics or government watchdog agencies. Formally, we employ a variation of the Poisson learning models pioneered by Keller *et al.* (2005), Keller and Rady (2010), and Keller and Rady (2013). As in these models, our analysis distinguishes between *bad news* markets, in which signal arrivals (*breakdowns*)

indicate bad quality and the absence of signals makes consumers more optimistic about the innovation; and *good news* markets, in which signal arrivals (*breakthroughs*) suggest good quality and the absence of signals makes consumers more pessimistic. To capture social learning, we assume that the informativeness of signals is increasing in the number of previous adopters.

The automobile industry is an example of a market in which learning is predominantly via bad news events, as evidenced by the wide-spread social media coverage of a battery fire in a Tesla Model S electric car in October 2013 or of the 2009-2011 Toyota vehicle recalls. By contrast, in the market for (essentially side-effect free) herbal remedies or other alternative medical treatments, learning is mostly via good news: Occasional reports of success stories boost consumers' confidence in a treatment, while consumers grow more skeptical of its effectiveness in the absence of any such reports.²

The heart of our paper, Sections 1.5 and 1.6, analyzes and contrasts equilibrium adoption behavior in bad and good news markets. For tractability, we focus on *perfect* bad (respectively good) news environments, in which a *single* signal arrival *conclusively* indicates bad (respectively good) quality, so that equilibrium dynamics are non-trivial only in the absence of signals. A key insight facilitating our analysis is that consumers' equilibrium incentives across time must satisfy a quasi-single crossing property (Theorem 1.4.1): Absent signals, there can be at most one transition from strict preference for adoption to strict preference for waiting, or vice versa, with a possible period of indifference in between. This enables us to establish the existence of unique³ equilibria. Equilibrium adoption dynamics admit simple closed-form descriptions which are Markovian in current beliefs and in the mass of consumers who have not yet adopted.

Section 1.5 studies the perfect bad news case. In the absence of breakdowns, consumers grow increasingly *optimistic* about the innovation over time. As a result of the single-crossing

²Cf. Board and Meyer-ter-Vehn (2013) and MacLeod (2007) for additional examples of bad news and good news markets.

³Uniqueness is in terms of *aggregate* adoption behavior.

property, the unique equilibrium is then characterized by two times $0 \leq t_1^* \leq t_2^*$, which depend on the fundamentals (Theorem 1.5.1): Until time t_1^* , no adoption takes place and consumers acquire information only from exogenous sources; from time t_2^* on, all consumers adopt immediately when given a chance, unless a breakdown occurs, in which case adoption comes to a permanent standstill. If $t_1^* < t_2^*$, then throughout $[t_1^*, t_2^*)$ only *some* consumers adopt whenever given a chance, with the flow of new adopters uniquely determined by an ODE that guarantees consumers' indifference between adopting and delaying throughout this interval. Given that consumers are forward-looking, $t_1^* < t_2^*$ occurs in economies with a sufficiently large potential for social learning and not too optimistic consumers (by contrast, if consumers are myopic or if there are no possibilities for social learning, then necessarily $t_1^* = t_2^*$).

We highlight two key implications for aggregate adoption dynamics and consumer welfare:

First, provided $t_1^* < t_2^*$, the innovation's adoption curve (which plots the percentage of adopters in the population against time) has the characteristic S-shaped growth pattern that has been widely observed in empirical studies.⁴ Up to time t_1^* adoption is flat, on $[t_1^*, t_2^*)$ adoption levels increase convexly, and from time t_2^* there is a concave increase. Moreover, an increase in the potential for social learning prolongs the period of convex growth and leads to strictly lower expected adoption levels across time. The possibility of S-shaped adoption curves in our model is notable because we assume consumers to be (ex ante) identical, whereas most alternative explanations in the literature rely on specific distributions of consumer heterogeneity to generate a region of convex growth. In our model, convex growth is driven by informational incentives: As consumers grow increasingly optimistic, their opportunity cost to delaying goes up. To maintain indifference between adopting and delaying throughout $[t_1^*, t_2^*)$, this increase is offset by an increase in the flow of new adopters, which raises the odds that waiting will produce information allowing consumers to avoid a

⁴See, for example, Griliches (1957), Mansfield (1961), Mansfield (1968), Davies (1979), and Gort and Klepper (1982), among many others.

bad innovation.

Second, we predict a *saturation effect*: If the potential for social learning is great enough that $t_1^* < t_2^*$, then holding fixed other fundamentals, any additional increase in opportunities for social learning has no impact at all on (ex ante) equilibrium welfare levels. This is in stark contrast to the cooperative benchmark in which consumers coordinate on socially optimal adoption levels: Here increased opportunities for social learning are always strictly beneficial and can in fact be used to approximate first-best (complete information) payoffs in the limit. Relative to the cooperative benchmark, equilibrium adoption behavior displays two inefficiencies: First, adoption generally begins too late; second, once adoption begins it initially occurs at an inefficiently low rate, because during $[t_1^*, t_2^*)$ consumers who do not adopt when given a chance effectively free-ride on the information generated by consumers who do adopt. Increased opportunities for social learning exacerbate the second inefficiency by prolonging the period of free-riding. As a result, greater opportunities for social learning do not translate into uniformly faster learning about the quality of the innovation, but rather lead to strictly slower learning over some periods and faster learning over others. These two effects balance out to produce the saturation effect. In Section 1.7, we further build on this non-monotonicity in the speed of learning to construct an example involving consumers with heterogeneous discount rates, where increased opportunities for social learning are not only not beneficial, but in fact strictly *hurt* aggregate welfare.

In Section 1.6 we study learning via perfect good news. Here consumers grow increasingly *pessimistic* about the innovation in the absence of breakthroughs. Hence, the single-crossing property for equilibrium incentives implies adoption up to some time t^* (which depends on the fundamentals) and no adoption from t^* on, unless there is a breakthrough, after which all consumers adopt upon their first opportunity (Theorem 1.6.1). Interestingly, in contrast with the perfect bad news case, equilibrium adoption behavior is *all-or-nothing*: Regardless of the potential for social learning, there are no periods during which only some consumers adopt when given a chance. This highlights a fundamental way in which the nature of information transmission in an economy affects consumers'

adoption incentives. During a period of time when, absent signals, a consumer is prepared to adopt the innovation, he will be willing to delay his decision only if he expects to acquire *decision-relevant information* in the meantime: Since originally he is prepared to adopt the innovation, such information must make him strictly prefer *not* to adopt. When learning is via bad news, breakdowns have this effect, since they reveal the innovation to be bad. By contrast, breakthroughs in the perfect good news environment conclusively reveal the innovation to be good and hence cannot be decision-relevant to a consumer who is already willing to adopt.

The all-or-nothing nature of the good news equilibrium has the following implications for adoption dynamics and welfare:

First, adoption occurs in concave “bursts”: Up to time t^* adoption levels increase concavely, then adoption flattens out, possibly followed by another region of concave growth if a breakthrough occurs. While less commonly observed than S-shaped growth, this pattern is reminiscent of the “fast-break” product life cycles studied in the marketing literature⁵, with movies, music, and other “leisure-enhancing” products as canonical examples.⁶ We predict that increased opportunities for social learning bring forward t^* , compressing the initial period of concave growth, but do not affect the probability of adoption picking up again after coming to a temporary standstill.

Second, even in economies with rich opportunities for social learning, an increase in the potential for social learning is (essentially) always strictly beneficial and speeds up learning at all times. Nevertheless, equilibrium behavior is generally socially inefficient: Relative to the cooperative benchmark, adoption takes place at an optimal rate until time t^* , but consumers stop adopting too soon.

⁵Cf. Keillor (2007)

⁶For additional examples in the context of industrial process innovations, see Davies (1979).

1.1.1 Related Literature and Outline

Our paper proposes a model of innovation adoption by consumers who learn from each other's experiences and are forward-looking. Having a tractable model that can incorporate these two assumptions, examine the informational externalities they give rise to, and derive predictions for the effect of increased opportunities for social learning is desirable, as there is considerable empirical evidence for both assumptions. For example, a growing literature in development economics documents the effect of learning from others' experiences on the adoption of new agricultural technologies, as in Foster and Rosenzweig (1995) or Conley and Udry (2010). This literature also finds evidence for forward-looking behavior: Bandiera and Rasul (2006) analyze the decision of farmers in Mozambique to adopt a new crop, sunflower. They find that farmers whose network of friends and family contains *many* adopters of the new crop are *less* likely to initially adopt it themselves. Relatedly, Munshi (2004) compares farmers' willingness to experiment with new high-yield varieties (HYV) across rice and wheat growing areas in India. Farmers in rice growing regions, which compared with wheat growing regions display greater heterogeneity in growing conditions that make learning from others' experiences *less* feasible, are found to be *more* likely to experiment with HYV than farmers in wheat growing areas.

At a theoretical level, the key feature of our model is that social learning and forward-looking incentives jointly give rise to informational externalities that do not arise in the absence of either assumption. In relation to existing models of innovation adoption, this has at least two interesting implications.

First, many models of innovation adoption rely on consumer heterogeneity as a key ingredient in fitting observed adoption data. Our analysis suggests that in existing *learning-based* models⁷ heterogeneity is only crucial because of the common assumption that either consumers are forward-looking but news is generated purely exogenously, as in Jensen (1982), or that learning is social but consumers are myopic, as in Young (2009) or Ellison

⁷For comprehensive surveys of the literature, including also non-learning based explanations of innovation adoption, such as the epidemic model and the probit model of firm characteristics, see for example Geroski (2000) and Baptista (1999).

and Fudenberg (1993):⁸ In either case, a population of identical consumers would behave according to a simple cutoff rule, adopting the innovation at beliefs above a certain threshold and not adopting otherwise, and this rules out convex growth in adoption levels.⁹ By contrast, in our model consumers are assumed to be ex ante identical, but the combination of forward-looking behavior and social learning allows us to provide an alternative micro-foundation for convex growth in terms of purely informational incentives.

The literature also commonly appeals to variations in consumer heterogeneity in order to explain qualitative differences in adoption patterns across different products. For example, in his study of the diffusion of 22 post-war industrial process innovations in the UK, Davies (1979) uses symmetrical logistic distributions to fit the S-shaped adoption patterns characteristic of expensive and complex innovations, but lognormal distributions to fit the rapid, essentially concave growth in adoption levels he observes for less expensive and simpler innovations. Again, our analysis shows that when consumers are forward-looking social learners, these contrasting patterns can instead be explained through differences in the informational environment: S-shaped curves arise in bad news markets with a relatively large potential for social learning, while concave adoption patterns are characteristic of good news markets (or of bad news markets with little potential for social learning or with very optimistic consumers). Our focus on the role of the market learning process in shaping consumers' informational incentives and generating varied aggregate adoption dynamics is

⁸Two exceptions are Persons and Warther (1997) and Kapur (1995), who feature a form of forward-looking social learning, but differ substantially from our paper in terms of both setup and focus. Persons and Warther (1997) focuses on the combination of forward-looking incentives, endogenously generated news, and firm heterogeneity to provide rational foundations for seemingly irrational, fad-like patterns in the adoption of financial innovations. In Kapur (1995), a finite number of firms engage in a sequence of waiting contests to adopt a new technology, with each contest ending once a firm adopts. Restricting to MPE, he finds that if more information is revealed when more firms adopt during a given waiting contest, then the mean duration of waiting contests shrinks over time, suggesting a crude approximation of convex diffusion. Since both models are set in discrete time, they are less tractable and not suited to performing comparative statics analyses with respect to the potential for social learning in an economy. In addition, discrete time is less suited to highlighting the role of the market learning process in shaping aggregate adoption dynamics, because when the information process is sufficiently informative relative to the period length, adoption behavior is qualitatively similar across many news processes. By contrast, when the period length becomes short as in our continuous time model, differences become transparent.

⁹Adoption patterns can exhibit concave growth simply as a result of gradual depletion of the population of remaining consumers.

similar in spirit to Board and Meyer-ter-Vehn (2013), who in the context of a capital-theoretic model of quality and reputation, highlight the dependence of firms' reputational incentives on the news process and contrast reputational dynamics across different markets.

Second, in addition to providing an alternative explanation for observed data, the informational externalities that arise from the interaction between forward-looking behavior and endogenously generated information are important because they suggest caution in evaluating the effect of increased opportunities for social learning. In contrast to existing models, we predict that increased opportunities for social learning need not produce welfare gains and may lead to a temporary slowdown in learning and a strict decline in initial adoption levels. On the other hand, if learning is modeled as purely exogenous or consumers are assumed to be myopic, then increased opportunities for social learning necessarily speed up learning and are unambiguously welfare-improving.

The techniques and framework of this paper are closest to those employed in the strategic experimentation literature, e.g. Bolton and Harris (1999), Keller *et al.* (2005), Keller and Rady (2010), and Keller and Rady (2013). However, our paper differs in two key respects: First, in our model any individual consumer's influence on the information seen by others is negligible; second, adoption of the innovation is irreversible. The first assumption is natural in the context of the large market applications we have in mind, and for many new products (for example movies or books, for which consumption is usually a one-time event, or technologies that entail large switching costs) irreversibility is also more reasonable than the possibility of consumers continuously switching back and forth between the innovation and the status quo as in the strategic experimentation literature.¹⁰ In the strategic experimentation literature, consumers' direct influence on opponents' information and their ability to adjust their experimentation levels as a function of beliefs produces the so-called *encouragement effect*: There is an incentive to increase current experimentation in

¹⁰Moreover, if consumers could continuously switch back and forth between the two options, then under the large market assumption, consumers' equilibrium strategies would effectively reduce to myopic best response with respect to beliefs.

order to drive up beliefs and induce more future experimentation by others.¹¹ As a result of the encouragement effect, many comparative statics in those models differ substantially: For example, an increase in the rate of information transmission may cause consumers to begin to adopt earlier, whereas in our model, we observe that initially adoption rates always weakly decrease in response to such a change. Without the encouragement effect, we are more easily able to study comparative statics on adoption behavior, speed of learning, and welfare with respect to changes in the social learning environment. Moreover, we obtain equilibrium uniqueness (at the aggregate level) without any Markovian restriction on strategies.

A number of papers, including Rosenberg *et al.* (2007), Chamley and Gale (1994), and Murto and Välimäki (2011), also study the impact of informational externalities on adoption, investment, or exit behavior, but rely on the assumption that agents hold *private* information. Notably, Chamley and Gale (1994) obtain a result somewhat resembling our saturation effect, according to which in the limit, an increase in the number of players has no effect on the rate of investment or flow of information. In the context of a two-armed bandit problem in which the decision to switch to the *safe* arm is irreversible, Rosenberg *et al.* (2007) obtain a similar uniqueness result to ours in the limit as the number of players becomes large. However, the specifics of all these models differ substantially from ours, as agents obtain private information and make inferences about the quality of the product by observing others' *actions*, while in our model all relevant news is *public* and actions do not reveal additional information.

Finally, Bergemann and Välimäki (1997) and Bergemann and Välimäki (2000) study innovation adoption in the presence of pricing motives by sellers when learning is social. In these papers, prices that dynamically adjust through time act as an additional instrument through which the seller can affect the endogenous information generation process. Bergemann and Välimäki (1997) study a model in which one established firm (with known technology) and

¹¹There is no encouragement effect in the perfect good news environment of Keller *et al.* (2005), but consumers' ability to influence each other's beliefs as well as the reversibility of experimentation are once again crucial in generating asymmetric switching equilibria, in which consumers take turns in experimenting at different beliefs.

a new firm with a risky innovation compete through prices. They derive the Markov perfect equilibrium pricing strategies and adoption behavior and demonstrate that adoption is too fast (relative to the social optimum) when consumers are pessimistic and too slow when consumers are optimistic. The main difference with our paper is that consumers in their model best respond myopically at each point in time, so that adoption dynamics are driven purely by sellers' informational and pricing motives. By contrast, in our model consumers are more sophisticated and consider the option value to waiting, producing interesting adoption dynamics even in the absence of pricing motives.¹² Bergemann and Välimäki (2000) analyze a similar model in which consumers display forward-looking behavior. As in Bergemann and Välimäki (1997), they find that pricing motives cause experimentation to be excessive, which is in contrast to our finding that in the absence of pricing motives there is too little (and, under perfect bad news, too slow) adoption. They find additionally that when the innovation is launched in many markets simultaneously, adoption rates become socially optimal in the limit as the number of markets grows large. Much of the focus in our paper is on analyzing the effect of increased opportunities for social learning on consumers' informational incentives. In order to isolate the effect on the consumer side, our baseline model therefore abstracts away from pricing considerations.

The rest of the paper is organized as follows. Section 1.2 describes the model, defining formally the perfect bad news and perfect good news signal processes that we use throughout the paper as well as the equilibrium concept. Section 1.3 analyzes the cooperative (socially optimal) benchmark which selects an aggregate flow of adoption so as to maximize ex ante aggregate welfare. Section 1.4 establishes a quasi-single crossing property for equilibrium incentives that simplifies the equilibrium analysis in the following sections. Section 1.5 establishes existence of a unique equilibrium under perfect bad news and studies comparative statics with respect to changes in the potential for social learning. Section 1.6 performs the analogous exercise under perfect good news. Section 1.7 provides an example,

¹²The key distinction is again due to the assumption that adoption is irreversible in our model, so that potentially adopting a bad product incurs a cost on consumers. On the other hand, in Bergemann and Välimäki (1997), consumers adopt at every point in time and the adoption decision is freely reversible.

involving consumers with heterogeneous discount rates, where an increase in the potential for social learning strictly hurts ex ante welfare. Section 1.8 concludes. Appendix A.1 - A.9 contains proofs omitted from the main text.

1.2 Model

1.2.1 The Game

Time $t \in [0, +\infty)$ is continuous. At time $t = 0$, an innovation of unknown quality $\theta \in \{G = 1, B = -1\}$ and of unlimited supply is released to a continuum population of potential consumers of mass $\bar{N}_0 \in \mathbb{R}_+$. Consumers are ex ante identical: They have a common prior $p_0 \in (0, 1)$ that $\theta = G$; they are forward-looking with common discount rate $r > 0$; and they have the same actions and payoffs, as specified below.

At each time t , consumers receive stochastic opportunities to adopt the innovation. Adoption opportunities are generated independently across consumers and across histories according to a Poisson process with exogenous arrival rate $\rho > 0$.¹³ Upon an adoption opportunity, a consumer must choose whether to adopt the innovation ($a_t = 1$) or to wait ($a_t = 0$). If a consumer adopts, he receives an expected lump sum payoff of $\mathbb{E}_t[\theta]$, conditioned on information available up to time t , and drops out of the game. If the consumer chooses to wait or does not receive an adoption opportunity at t , he receives a flow payoff of 0 until his next adoption opportunity, where he faces the same decision again.

1.2.2 Learning

Over time, consumers observe public signals that convey information about the quality of the innovation. To capture the idea of social learning, the informativeness of the public signal at time t is increasing in the flow N_t of consumers newly adopting the innovation at t , which we define more precisely in Section 1.2.3.

¹³Stochasticity of adoption opportunities can be seen as capturing the natural assumption that consumers face cognitive and time constraints, making it impossible for them to ponder the decision whether or not to adopt the innovation at every instant in continuous time.

Formally, we employ a variation of the Poisson learning model pioneered by Keller *et al.* (2005), Keller and Rady (2010), and Keller and Rady (2013).¹⁴ Conditional on quality θ , public signals arrive according to an inhomogeneous Poisson process with arrival rate $(\varepsilon_\theta + \lambda_\theta N_t)dt$, where $\lambda_\theta > 0$ and $\varepsilon_\theta \geq 0$ are exogenous parameters that depend on the quality θ of the innovation. The signal process summarizes news events that are generated from two sources. First, the social learning term λN_t represents news generated endogenously, based on the experiences of other consumers: It captures the idea of a flow N_t of new adopters each generating signals at rate λdt .¹⁵ Thus, the greater the flow of consumers adopting the innovation at t , the more likely it is for a signal to arrive at t , and hence the absence of a signal at t is more informative the larger N_t . Second, we also allow for (but do not require) signals to arrive at a fixed exogenous rate εdt , which represents information generated independently of consumers' behavior, for example by professional critics or government watchdog agencies.

For tractability, we focus on learning via *perfect* Poisson processes, where a single signal provides *conclusive* evidence of the quality of the innovation. Learning is via *perfect bad news* if $\varepsilon_G = \lambda_G = 0$ and $\varepsilon_B = \varepsilon \geq 0$, $\lambda_B = \lambda > 0$, so that the arrival of a signal (called a *breakdown*) is conclusive evidence that the innovation is bad. Learning is via *perfect good news* if $\varepsilon_B = \lambda_B = 0$ and $\varepsilon_G = \varepsilon \geq 0$, $\lambda_G = \lambda > 0$, so that a signal arrival (called a *breakthrough*) is conclusive evidence for the innovation being good. As motivated in the Introduction, the distinction between bad news and good news can be seen to reflect the nature of news production in different markets. In addition, $\Lambda_0 := \lambda \bar{N}_0$ can be seen as a simple measure of

¹⁴Keller *et al.* (2005) have learning via perfect good news Poisson signals, Keller and Rady (2010) study imperfect good news learning, and Keller and Rady (2013) study perfect and imperfect bad news learning. For other recent work that prominently features learning via Poisson signals, see for example Che and Hörner (2013); Board and Meyer-ter-Vehn (2013); Halac *et al.* (2013).

¹⁵Note that by letting the social learning component of the signal arrival rate at time t , λN_t , depend only on the flow of adopters N_t at time t itself, we are effectively assuming that each each adopter can generate a signal only once, namely at the time of adoption. This assumption is natural for "innovations" such as new movies or medical procedures, for which "consumption" is a one-time event and quality is revealed upon consumption. For durable goods, such as cars or consumer electronics, it might be more natural to allow adopters to generate signals repeatedly over time, which can be captured by replacing λN_t with $\lambda \int_0^t N_s ds$. This would yield results that are qualitatively similar to those presented in the following sections.

the *potential for social learning* in an economy, summarizing both the likelihood λ with which individual adopters' experiences find their way into the public domain and the size \bar{N}_0 of the population which can contribute to and access the common pool of information.

We briefly summarize the evolution of consumers' beliefs under bad and good news:

Learning via Perfect Bad News

Under perfect bad news, consumers' posterior on $\theta = G$ permanently jumps to 0 at the first breakdown. Let p_t denote consumers' no-news posterior, i.e. the belief at t that $\theta = G$ conditional on no signals having arrived on $[0, t)$. Given a flow of adopters N , standard Bayesian updating implies that

$$p_t = \frac{p_0}{p_0 + (1 - p_0)e^{-\int_0^t (\varepsilon + \lambda N_s) ds}}.^{16} \quad (1.1)$$

In particular, if N_τ is continuous in an open interval $(s, s + \nu)$ for $\nu > 0$, then p_τ for $\tau \in (s, s + \nu)$ evolves according to the ODE:

$$\dot{p}_\tau = (\varepsilon + \lambda N_\tau) p_\tau (1 - p_\tau).$$

Note that the no-news posterior is continuous and *increasing*.

Learning via Perfect Good News

Under perfect good news, consumers' posterior on $\theta = G$ permanently jumps to 1 at the first breakthrough. Given a flow of adopters N , Bayes' rule now implies that consumers' no-news posterior satisfies

$$p_t = \frac{p_0 e^{-\int_0^t (\varepsilon + \lambda N_s) ds}}{p_0 e^{-\int_0^t (\varepsilon + \lambda N_s) ds} + (1 - p_0)}. \quad (1.2)$$

¹⁶Definition 1.2.1 imposes measurability on N , so this expression is well-defined.

In particular, if N_τ is continuous in an open interval $(s, s + \nu)$ for $\nu > 0$, then p_τ for $\tau \in (s, s + \nu)$ must satisfy the ODE:

$$\dot{p}_\tau = -(\varepsilon + \lambda N_\tau) p_\tau (1 - p_\tau).$$

In contrast to the perfect bad news case, the no-news posterior is now continuous and *decreasing*.

1.2.3 Equilibrium

Since our main interest is in the aggregate adoption dynamics of the population, we take as the primitive of our equilibrium concept the aggregate flow $(N_t)_{t \geq 0}$ of consumers newly adopting the innovation over time and do not explicitly model individual consumers' behavior. Given our focus on *perfect* news processes, consumers' incentives are non-trivial only in the absence of signals: Under perfect bad news, no new consumers will adopt after a breakdown, while under perfect good news all remaining consumers will adopt when given a chance after there has been a breakthrough. Therefore, we henceforth let N_t denote the flow of new adopters at t *conditional on no signals up to time t* and define equilibrium in terms of this quantity. Reflecting the assumption that aggregate adoption behavior is predictable with respect to the news process of the economy, we require that N_t be a *deterministic* function of time. We consider all such functions which are feasible in the following sense:

Definition 1.2.1. A feasible flow of adopters is a right-continuous function $N: [0, +\infty) \rightarrow \mathbb{R}$ such that $N_t := N(t) \in [0, \rho \bar{N}_t]$ for all $t \in [0, +\infty)$, where $\bar{N}_t := \bar{N}_0 - \int_0^t N_s ds$.

Here \bar{N}_t denotes the mass of consumers remaining in the game at time t . We require that $N_t \leq \rho \bar{N}_t$ so that N_t is consistent with the remaining \bar{N}_t consumers independently receiving adoption opportunities at Poisson rate ρ .

Any feasible adoption process N defines an associated no-news posterior p_t^N as given by Equation 1.1 if learning is via perfect bad news and by Equation 1.2 if learning is via perfect good news.

In equilibrium, we require that at each time t , N_t is consistent with optimal behavior by the remaining \bar{N}_t forward-looking consumers: If a consumer receives an adoption opportunity at t , he optimally trades off his expected payoff to adopting against his value to waiting, given that he assigns probability p_t^N to $\theta = G$ and that he expects the population's adoption behavior to evolve according to the process N . For this we must first define the value to waiting at t .

Let Σ_t denote the set of all right-continuous functions $\sigma : [t, +\infty) \rightarrow \{0, 1\}$, each of which defines a potential set of future times at which, absent signals, a given consumer might adopt if given an opportunity. Under the Poisson process generating adoption opportunities, any $\sigma \in \Sigma_t$ defines a random time τ^σ at which, absent signals, the consumer will adopt the innovation and drop out of the game.¹⁷

Let $W_t^N(\sigma)$ denote the expected payoff to waiting at t and following σ in the future, given the aggregate adoption process N . Specifically, if learning is via perfect bad news, σ prescribes adoption at the random time τ^σ if and only if there have been no breakdowns prior to τ^σ , yielding

$$W_t^N(\sigma) := \mathbb{E} \left[e^{-r(\tau^\sigma - t)} \left(p_t^N - (1 - p_t^N) e^{-\int_t^{\tau^\sigma} (\varepsilon + \lambda N_s) ds} \right) \right],$$

where the expectation is with respect to the Poisson process generating adoption opportunities.

If learning is via perfect good news, then following σ means that at any adoption opportunity prior to τ^σ , adoption occurs only if there has been a breakthrough, and at τ^σ adoption occurs whether or not there has been a breakthrough. For any time $s \geq t$, denote by τ_s the random time at which the first adoption opportunity after s arrives. Then $W_t^N(\sigma)$

¹⁷Formally, we define τ^σ as follows. Let $(X_s)_{s \geq t}$ denote the stochastic process representing the number of arrivals generated on $[t, s]$ by a Poisson process with arrival rate ρ , and let $(X_{s^-})_{s > t}$ denote the number of arrivals on $[t, s)$. Then,

$$\tau^\sigma := \inf\{s \geq t : \sigma_s \times (X_s - X_{s^-}) > 0\},$$

where, as per convention, $\inf \emptyset := +\infty$. It is well known that the hitting time of a right-continuous process of an open set is an optional time. Therefore, the expectations in the definition of the value to waiting are well-defined.

is given by

$$\mathbb{E} \left[\left(p_t e^{-\int_t^{\tau^\sigma} (\varepsilon + \lambda N_s) ds} + (1 - p_t) \right) e^{-r(\tau^\sigma - t)} (2p_{\tau^\sigma} - 1) + p_t \int_t^{\tau^\sigma} (\varepsilon + \lambda N_s) e^{-\int_t^s (\varepsilon + \lambda N_k) dk} e^{-r(\tau_s - t)} ds \right],$$

where the expectation is again with respect to the Poisson process generating adoption opportunities.

The value to waiting at t is the payoff to waiting and behaving optimally in the future:

Definition 1.2.2. The value to waiting given a feasible adoption process N is the function $W_t^N : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ defined by $W_t^N := \sup_{\sigma \in \Sigma_t} W_t^N(\sigma)$ for all t .

We are now ready to define our equilibrium concept:

Definition 1.2.3. An equilibrium is a feasible adoption process $(N_t)_{t \geq 0}$ such that

1. $W_t^N \geq 2p_t^N - 1$ for all t such that $\rho \bar{N}_t > N_t$
2. $W_t^N \leq 2p_t^N - 1$ for all t such that $0 < N_t$.

Thus, Definition 1.2.3 requires that at any time t , the aggregate flow of new adopters N_t be consistent with the remaining \bar{N}_t consumers optimally trading off the expected payoff to immediate adoption, $2p_t^N - 1$, against the value to waiting, W_t^N .

Note that our definition of equilibrium is essentially Nash equilibrium, i.e. we do not require subgame perfection. The motivation for this is that in a continuum population any individual consumer's behavior has a negligible impact on the aggregate adoption levels so that any history not on the equilibrium path (in which a different number of consumers than expected previously adopted) is more than a unilateral deviation from the equilibrium path. Thus, off-path histories do not affect individual consumers' incentives on the equilibrium path and are unimportant for equilibrium analysis.

As usual, the equilibrium value to waiting W_t^N admits an alternative characterization as the solution to a functional equation, which we note here for use in future sections:

Lemma 1.2.4. *Suppose N is an equilibrium. If learning is via perfect bad news, W_t^N satisfies the functional equation*

$$V_t = \int_t^\infty \rho e^{-(r+\rho)(s-t)} \frac{p_t^N}{p_s^N} \max \left\{ \left(2p_s^N - 1 \right), V_s \right\} ds.$$

If learning is via perfect good news, W_t^N satisfies the functional equation

$$V_t = \int_t^\infty \rho e^{-(r+\rho)(s-t)} \left(p_t^N \left(1 - e^{-\int_t^s (\varepsilon + \lambda N_k) dk} \right) + \frac{p_t^N e^{-\int_t^s (\varepsilon + \lambda N_k) dk}}{p_s^N} \max \left\{ \left(2p_s^N - 1 \right), V_s \right\} \right) ds.$$

Proof. The proof is standard. □

1.3 Cooperative Benchmark

To establish a socially optimal benchmark, we first consider the cooperative problem: This selects an aggregate flow N of adopters that maximizes ex ante aggregate welfare, taking into account the effect of N on the public information process; we impose feasibility, but do not impose the incentive compatibility requirements of the equilibrium in Definition 1.2.3.¹⁸

Clearly, under perfect good news it is optimal to require adoption at the maximal possible rate once there has been a breakthrough. Similarly, under perfect bad news it is optimal to terminate adoption as soon as there has been a breakdown. Thus, the objective of the cooperative problem under perfect good news is:

$$\begin{aligned} \sup_N p_0 \int_0^\infty (\varepsilon_G + \lambda_G N_\tau) e^{-\int_0^\tau (\varepsilon_G + \lambda_G N_s) ds} & \left(\int_0^\tau e^{-rs} N_s ds + e^{-r\tau} \frac{\rho}{\rho + r} \left(\bar{N}_0 - \int_0^\tau N_s ds \right) \right) d\tau \\ & + p_0 e^{-\int_0^\infty (\varepsilon_G + \lambda_G N_s) ds} \int_0^\infty e^{-rs} N_s ds - (1 - p_0) \int_0^\infty e^{-rs} N_s ds, \end{aligned} \quad (19)$$

subject to the feasibility constraint that $N_t \in [0, \rho \bar{N}_t]$ for all t .

¹⁸We are not concerned with implementation here, but because beliefs are publicly observed, as long as we allow for transfers, the solution that we provide will be implementable while respecting individual rationality.

¹⁹We impose the convention that $e^{-\infty} = 0$. Thus whenever $\varepsilon_G > 0$, $e^{-\int_0^\infty (\varepsilon_G + \lambda_G N_s) ds} = 0$.

Under perfect bad news, the objective is:

$$\begin{aligned} \sup_N p_0 \int_0^\infty e^{-rs} N_s ds - (1 - p_0) \int_0^\infty (\varepsilon_B + \lambda_B N_\tau) e^{-\int_0^\tau (\varepsilon_B + \lambda_B N_s) ds} \int_0^\tau e^{-rs} N_s ds d\tau \\ - (1 - p_0) e^{-\int_0^\infty (\varepsilon_B + \lambda_B N_s) ds} \int_0^\infty e^{-rs} N_s ds,^{20} \end{aligned}$$

again subject to the feasibility constraint that $N_t \in [0, \rho \bar{N}_t]$ for all t .

Standard techniques show that the solution to both cooperative problems has an *all-or-nothing* form.²¹ In each problem, there is a cutoff time t^s (depending on the parameters) such that conditional on no signals, there is no (respectively maximal) adoption until time t^s under perfect bad (respectively good) news, and maximal (respectively no) adoption from t^s on:

Proposition 1.3.1. *In both problems, there exists an adoption flow N that attains the maximum. Furthermore, there exists an optimal adoption flow with the property that there exists t^s such that*

- $N_t = \rho \bar{N}_t$ for all t such that $(\lambda_G - \lambda_B)(t^s - t) > 0$;
- $N_t = 0$ for all t such that $(\lambda_G - \lambda_B)(t^s - t) < 0$.

Proof. See Appendix Section A.9. □

We now solve for the cutoff time, or equivalently the cutoff belief, under both signal structures.

1.3.1 Cooperative Benchmark under Perfect Good News

Under perfect good news, letting $\varepsilon := \varepsilon_G$ and $\Lambda_0 := \lambda_G \bar{N}_0$, the cutoff time t^s solves

$$\sup_{t^s \geq 0} \frac{\rho}{r + \rho} \left(1 - e^{-(r+\rho)t^s} \right) \bar{N}_0 (2p_0 - 1) + e^{-(r+\rho)t^s} \frac{\rho}{r + \rho} \bar{N}_0 \left(\pi^s + (1 - \pi^s) p^s \frac{\varepsilon}{\varepsilon + r} \right) \quad (1.3)$$

²⁰Again we assume that whenever $\varepsilon_B > 0$, $e^{-\int_0^\infty (\varepsilon_B + \lambda_B N_s) ds} = 0$.

²¹This is due to the linearity of the signal arrival rate in N_t .

where π^s and p^s denote, respectively, the probability of a breakthrough prior to time t^s and the no-news posterior at time t^s ; that is,

$$\begin{aligned}\pi^s &:= p_0 \left(1 - e^{-\varepsilon t^s - \Lambda_0(1 - e^{-\rho t^s})} \right), \\ p^s &:= \frac{p_0 e^{-\varepsilon t^s - \Lambda_0(1 - e^{-\rho t^s})}}{p_0 e^{-\varepsilon t^s - \Lambda_0(1 - e^{-\rho t^s})} + (1 - p_0)}.\end{aligned}$$

Taking the first order condition of the above, we obtain:

$$(r + \rho)(1 - \pi^s) \left(\left(2 - \frac{\varepsilon}{\varepsilon + r} \right) p^s - 1 \right) + p^s(1 - \pi^s) \left(\varepsilon + \Lambda_0 \rho e^{-\rho t^s} \right) \frac{r}{\varepsilon + r} = 0 \quad (1.4)$$

if an interior solution exists.

If the left-hand side of Equation 1.4 is non-positive at all times, then the cooperative cutoff satisfies $t^s = 0$, so that there is no adoption until a breakthrough. This happens if and only if

$$(r + \rho) \left((2p_0 - 1) - p_0 \frac{\varepsilon}{\varepsilon + r} \right) + p_0 (\varepsilon + \rho \Lambda_0) \frac{r}{\varepsilon + r} \leq 0. \quad (1.5)$$

On the other hand, if the left-hand side of Equation 1.4 is strictly positive at all times, then $t^s = +\infty$ and the cooperative solution calls for maximal adoption irrespective of whether or not there has been a breakthrough. This happens if and only if $\varepsilon = 0$ and $p_0 (1 + e^{-\Lambda_0}) \geq 1$.

We summarize this in the following proposition:

Proposition 1.3.2. *Under perfect good news, the cooperative cutoff time is as follows:*

- *If Inequality (1.5) holds, then $t^s = 0$.*
- *If $\varepsilon = 0$ and $p_0 (1 + e^{-\Lambda_0}) \geq 1$, then $t^s = +\infty$.*
- *Otherwise, t^s satisfies Equation (1.4).*

Note that the cutoff posterior p^s depends on the prior. This is in contrast to the strategic experimentation literature because of our assumption that the stock of remaining consumers is depleted as consumers drop out following adoption. In strategic experimentation, the cooperative solution only depends on the current belief and does not depend on the initial

conditions since experimenters remain in the game to potentially experiment further in the future.

1.3.2 Cooperative Benchmark under Perfect Bad News

Under perfect bad news, letting $\varepsilon := \varepsilon_B$ and $\Lambda_0 := \lambda_B \bar{N}_0$, the cutoff time t^s solves:

$$\sup_{t^s \geq 0} e^{-rt^s} \bar{N}_0 \left(p_0 \frac{\rho}{\rho + r} - (1 - p_0) e^{-\varepsilon t^s} \int_0^\infty \rho e^{-\varepsilon \tau - \Lambda_0 (1 - e^{-\rho \tau})} e^{-(r+\rho)\tau} d\tau \right).$$

Taking the first order condition, we obtain:

$$e^{-\varepsilon t^s} K(\Lambda_0) = \frac{r}{\varepsilon + r} \frac{\rho}{\rho + r} \frac{p_0}{1 - p_0}$$

where

$$K(\Lambda_0) := \int_0^\infty \rho e^{-\varepsilon \tau - \Lambda_0 (1 - e^{-\rho \tau})} e^{-(r+\rho)\tau} d\tau < \frac{\rho}{\varepsilon + \rho + r}.$$

Then an easy calculation yields the cutoff posterior:

$$p^s = \frac{K(\Lambda_0)}{\frac{r}{\varepsilon + r} \frac{\rho}{\rho + r} + K(\Lambda_0)} < \frac{\varepsilon + r}{\varepsilon + 2r}.$$

We summarize this in the following proposition:

Proposition 1.3.3. *Under perfect bad news, the cooperative solution is given by:*

$$N_t = \begin{cases} 0 & \text{if } p_t < p^s \\ \rho \bar{N}_t & \text{if } p_t \geq p^s, \end{cases}$$

where

$$p^s = \frac{K(\Lambda_0)}{\frac{r}{\varepsilon + r} \frac{\rho}{\rho + r} + K(\Lambda_0)}.$$

1.4 Quasi-Single Crossing Property for Equilibrium Incentives

We now proceed to equilibrium analysis. As a preliminary step, we first establish a useful property of equilibrium incentives under both perfect bad news and perfect good news.

Suppose that $N_{t \geq 0}$ is an *arbitrary* feasible flow of adopters, with associated no-news posterior $p_{t \geq 0}^N$ and value to waiting $W_{t \geq 0}^N$ as defined in Definition 1.2.2. In general, the dynamics of the trade-off between immediate adoption at time t (yielding expected payoff $2p_t^N - 1$) and delaying and behaving optimally in the future (yielding expected payoff W_t^N) can be quite difficult to characterize, with $(2p_t^N - 1) - W_t^N$ changing sign many times. However, when $N_{t \geq 0}$ is an *equilibrium* flow, then for any t ,

$$\begin{aligned} 2p_t^N - 1 < W_t^N &\implies N_t = 0; \text{ and} \\ 2p_t^N - 1 > W_t^N &\implies N_t = \rho \bar{N}_t; \end{aligned}$$

and this imposes considerable discipline on the dynamics of the trade-off. Indeed, the following theorem establishes that $2p_t^N - 1$ and W_t^N must satisfy a quasi-single crossing property:

Theorem 1.4.1. *Suppose that learning is either via perfect bad news ($\lambda_B > 0 = \lambda_G$) or via perfect good news ($\lambda_G > 0 = \lambda_B$). Let $N_{t \geq 0}$ be an equilibrium, with corresponding no-news posteriors $p_{t \geq 0}^N$ and value to waiting $W_{t \geq 0}^N$. Then $W_{t \geq 0}^N$ and $2p_{t \geq 0}^N - 1$ satisfy single-crossing, in the following sense:*

- *Whenever $(\lambda_B - \lambda_G)(W_t^N - (2p_t^N - 1)) < 0$, then $(\lambda_B - \lambda_G)(W_\tau^N - (2p_\tau^N - 1)) < 0$ for all $\tau > t$.*
- *Whenever $(\lambda_B - \lambda_G)(W_t^N - (2p_t^N - 1)) \leq 0$, then $(\lambda_B - \lambda_G)(W_\tau^N - (2p_\tau^N - 1)) \leq 0$ for all $\tau > t$.*

Proof. See Appendix Section A.2. □

The basic intuition is as follows. Consider first the case of learning via perfect bad news and suppose that immediate adoption is strictly better than waiting today (and hence also in the near future provided there are no breakdowns).²² Then all consumers adopt upon an opportunity in the near future, so the no-news posterior strictly increases, while the number of remaining consumers strictly decreases. Because information is generated endogenously,

²²The latter implication follows from the continuity of the equilibrium value to waiting, which is established in the Appendix.

this means that the flow of information must be decreasing over time. As a result, immediate adoption becomes even more attractive relative to waiting, and consequently immediate adoption continues to be strictly preferable in the future.

Similarly, suppose that learning is via perfect good news and that waiting is strictly more attractive than immediate adoption today (and hence also in the near future). Then in the near future, no consumers adopt and information is generated purely via the exogenous news source (or not at all if $\varepsilon = 0$). As a result, the no-news posterior decreases (weakly) while the number of remaining consumers does not change. This makes waiting even more attractive relative to adopting immediately, so that waiting continues to be strictly preferable in the future.




Figure 1.1: *Perfect Bad News*




Figure 1.2: *Perfect Good News*

Theorem 1.4.1 implies that any equilibrium features two threshold times $0 \leq t_1^* \leq t_2^* \leq +\infty$ given by²³

$$t_1^* := \inf\{t : (\lambda_B - \lambda_G) (2p_t^N - 1 - W_t^N) \geq 0\},$$

$$t_2^* := \inf\{t : (\lambda_B - \lambda_G) (2p_t^N - 1 - W_t^N) > 0\},$$

such that if there are no signal arrivals, then under perfect bad (respectively good) news,

²³With the usual convention that $\inf \emptyset = +\infty$.

waiting (respectively adoption) is strictly preferable before t_1^* , and adoption (respectively waiting) is strictly preferable after t_2^* , with indifference in between, as illustrated in Figures 1 and 2. In Sections 1.5 and 1.6 we will build on this observation to establish the existence of unique equilibria under both perfect bad news and good news. The threshold times, as well as the flow of adopters between t_1^* and t_2^* , are fully pinned down by the parameters.

Looking ahead to Section 1.6, we will see that under perfect good news, any equilibrium must in fact satisfy $t_1^* = t_2^*$.²⁴ Depending on parameters, the equilibrium takes three possible forms: (i) $0 = t_1^* = t_2^*$; (ii) $0 < t_1^* = t_2^* < +\infty$; or (iii) $0 < t_1^* = t_2^* = +\infty$.²⁵ By contrast, under perfect bad news in Section 1.5, the equilibrium takes one of six forms depending on parameters: (i) $0 = t_1^* = t_2^* < +\infty$; (ii) $0 = t_1^* < t_2^* < +\infty$; (iii) $0 < t_1^* = t_2^* < +\infty$; (iv) $0 < t_1^* < t_2^* < +\infty$; (v) $0 < t_1^* = t_2^* = +\infty$;²⁶ or (vi) $0 = t_1^* < t_2^* = +\infty$.²⁷ The possibility of a non-empty interval (t_1^*, t_2^*) of indifference will emerge as a key feature distinguishing bad news markets from good news markets. Maintaining indifference at times (t_1^*, t_2^*) requires a form of informational free-riding, which we term *partial adoption*, whereby only *some* consumers adopt when given the chance (i.e. $N_t \in (0, \rho \bar{N}_t)$ at each $t \in (t_1^*, t_2^*)$). We will see that partial adoption has important implications not just from an efficiency standpoint, but also for the shape of equilibrium adoption curves and for the impact of increased opportunities for social learning on welfare, learning, and adoption dynamics.

1.5 Perfect Bad News

1.5.1 Equilibrium Characterization

We now build on the analysis of the previous section to establish the existence of a unique equilibrium when learning is via perfect bad news. Fix parameters $r, \rho, \bar{N}_0 > 0, \varepsilon = \varepsilon_B$,

²⁴With the sole exception of $\varepsilon = 0$ and $p_0 = \frac{1}{2}$, in which case it is easy to see that $N \equiv 0$ and $t_1^* = 0 < t_2^* = \infty$.

²⁵This possibility will arise iff $\varepsilon = 0$ and $p_0 \left(1 + e^{-\lambda \bar{N}_0}\right) \geq 1$.

²⁶This possibility will arise iff $\varepsilon = 0$ and $p_0 < \frac{1}{2}$.

²⁷This possibility will arise iff $\varepsilon = 0$ and $p_0 = \frac{1}{2}$.

$\lambda = \lambda_B \geq 0$, and $p_0 \in (0, 1)$. Suppose $N_{t \geq 0}$ is an equilibrium flow of adopters. Let $p_{t \geq 0}$ and $W_{t \geq 0}$ be the corresponding no-news posterior and value to waiting, and let $\Lambda_{t \geq 0} := \lambda \bar{N}_{t \geq 0}$ describe the evolution of the economy's potential for social learning.²⁸ From Theorem 1.4.1, we know that there are times $0 \leq t_1^* \leq t_2^* \leq +\infty$ given by

$$\begin{aligned} t_1^* &:= \inf\{t : 2p_t - 1 \geq W_t\}, \\ t_2^* &:= \inf\{t : 2p_t - 1 > W_t\}, \end{aligned}$$

such that (appealing also to right-continuity) N must satisfy

$$\begin{cases} N_t = 0 & \text{if } t < t_1^*, \\ 2p_t - 1 = W_t & \text{if } t \in [t_1^*, t_2^*) \\ N_t = \rho \bar{N}_t & \text{if } t \geq t_2^*. \end{cases}$$

In the following we will show that t_1^* , t_2^* , and the evolution of N_t between t_1^* and t_2^* are uniquely pinned down by the parameters. We first introduce some notation. For any $p \in (0, 1)$ and $\Lambda \geq 0$, let

$$G(p, \Lambda) := \int_0^\infty \rho e^{-(r+\rho)\tau} \left(p - (1-p)e^{-(\varepsilon\tau + \Lambda(1-e^{-\rho\tau}))} \right) d\tau.$$

$G(p, \Lambda)$ represents the payoff to adopting at the next opportunity if there have been no breakdowns by then, *given* that the current belief is p , that the remaining potential for social learning is Λ , and that absent breakdowns the remaining $\frac{\Lambda}{\lambda}$ consumers adopt at their first opportunity in the future.

Define the posteriors \underline{p} , \bar{p} , and p^\sharp as follows. Let \underline{p} be the posterior given by $2\underline{p} - 1 = G(\underline{p}, 0)$; that is,

$$\underline{p} := \frac{(\varepsilon + r)(r + \rho)}{2(\varepsilon + r)(r + \rho) - \varepsilon\rho}.$$

Thus, \underline{p} is the lowest belief at which a consumer is willing to adopt given that he could also delay, obtain more information at rate ε and reevaluate his decision at his next adoption

²⁸Recall that $\bar{N}_t := \bar{N}_0 - \int_0^t N_s ds$ denotes the remaining population at time t .

opportunity which is generated at rate ρ .

Define $\bar{p} := \lim_{\rho \rightarrow \infty} \underline{p}$, that is,

$$\bar{p} = \frac{\varepsilon + r}{\varepsilon + 2r};$$

\bar{p} is the lowest belief at which a consumer would be willing to adopt given that he could also delay and obtain more information at rate ε and given that adoption opportunities arrive continuously in the future.

Define $p^\sharp := \lim_{\varepsilon \rightarrow \infty} \underline{p}$, that is,

$$p^\sharp = \frac{\rho + r}{\rho + 2r}.$$

p^\sharp is the lowest belief at which a consumer would be willing to adopt given that he could also delay until his next opportunity, which is generated at rate ρ , and given that all uncertainty is completely resolved by then.²⁹

Finally, define the function $\Lambda^* : (0, 1) \rightarrow \mathbb{R}_+ \cup \{+\infty\}$ as follows. Let $\Lambda^*(p) \equiv 0$ for all $p \leq \underline{p}$, $\Lambda^*(p) = +\infty$ for all $p \geq p^\sharp$, and for all $p \in (\underline{p}, p^\sharp)$, let $\Lambda^*(p) \in \mathbb{R}_+$ be the unique value such that

$$2p - 1 = G(p, \Lambda^*(p)).^{30}$$

Thus, if the current posterior is $p \in [\underline{p}, p^\sharp)$ and the current potential for social learning in the economy is $\Lambda^*(p)$, then consumers are indifferent between adopting now or at their next opportunity absent breakdowns, provided that all remaining $\frac{\Lambda^*(p)}{\lambda}$ consumers also adopt at their first opportunity in the future.

We are now ready to state the equilibrium characterization theorem:

Theorem 1.5.1. *Fix $r, \rho > 0$, $\varepsilon \geq 0$, and $p_0 \in (0, 1)$. Let $p^* := \min\{\bar{p}, p^\sharp\}$. For every $\lambda, \bar{N}_0 > 0$, there is a unique equilibrium. Furthermore, in the unique equilibrium, N_t is Markovian in (p_t, Λ_t)*

²⁹Note that for all $p > p^\sharp$, $\lim_{\Lambda \rightarrow \infty} G(p, \Lambda) < 2p - 1$ and for all $p < p^\sharp$, $\lim_{\Lambda \rightarrow \infty} G(p, \Lambda) > 2p - 1$.

³⁰Note that such a value must exist given that $p \in (\underline{p}, p^\sharp)$ and is unique because $\Lambda^*(p)$ is strictly increasing in p on this domain.

for all t and satisfies

$$N_t = \begin{cases} 0 & \text{if } p_t \leq p^* \text{ and } \Lambda_t > \Lambda^*(p_t), \\ \frac{r(2p_t-1)}{\lambda(1-p_t)} - \frac{\varepsilon}{\lambda} \in (0, \rho\bar{N}_t) & \text{if } p_t > p^* \text{ and } \Lambda_t > \Lambda^*(p_t) \\ \rho\bar{N}_t & \text{if } \Lambda_t \leq \Lambda^*(p_t). \end{cases} \quad (1.6)$$

A detailed proof of Theorem 1.5.1 is provided in Appendix Section A.3.1. Here we sketch the basic idea. Before we proceed, however, note the following two special cases of the theorem: First, if $\rho \leq \varepsilon$, so that $p^* := \min\{\bar{p}, p^\sharp\} = p^\sharp$, then by Equation (1.6) and because $\Lambda^*(p) = +\infty$ for all $p \geq p^\sharp$, Theorem 1.5.1 asserts that regardless of the other parameters, N_t takes an all-or-nothing form with cutoff belief p^\sharp : $N_t = 0$ whenever $p_t < p^\sharp$ and $N_t = \rho\bar{N}_t$ whenever $p_t \geq p^\sharp$. Second, if $\varepsilon = 0$ and $p_0 \leq \frac{1}{2}$, then it is easy to see that Theorem 1.5.1 asserts that regardless of the other parameters, the unique equilibrium is given by $N_t = 0$ for all t .

Throughout Section 1.5, we will be particularly interested in the implications of N featuring a *partial adoption* region, in which $N_t \in (0, \rho\bar{N}_t)$ is as described by the second line of Equation (1.6). Since the two special cases above preclude the existence of such a region regardless of other parameters, we rule out these cases for the remainder of Section 1.5 by imposing the following two conditions:³¹

Condition 1.5.2. The rate at which exogenous information arrives is small relative to the rate at which consumers obtain adoption opportunities: $\varepsilon < \rho$. Thus, $p^* = \bar{p} < p^\sharp$.

Condition 1.5.3. Either $\varepsilon > 0$ or $p_0 \in (\frac{1}{2}, 1)$.

Given these two conditions, we now sketch the derivation of Theorem 1.5.1. In order to obtain the Markovian description of N_t in Equation (1.6), we note the following lemma, which we prove in the Appendix. This provides an alternative characterization of the threshold times t_1^* and t_2^* , relating these times to the evolution of (p_t, Λ_t) :

³¹In Section A.4 in the Appendix, we discuss in more detail the case where $\rho \leq \varepsilon$.

Lemma 1.5.4. Fix $r, \rho > 0$, $\varepsilon \geq 0$ and $p_0 \in (0, 1)$ satisfying Conditions 1.5.2 and 1.5.3. Let $N_{t \geq 0}$ be an equilibrium with corresponding no-news posterior $p_{t \geq 0}$ and threshold times t_1^* and t_2^* , and let $\Lambda_{t \geq 0} := \lambda \bar{N}_{t \geq 0}$ describe the evolution of the economy's potential for social learning. Then

1. $t_2^* = \inf\{t : \Lambda_t < \Lambda^*(p_t)\}$; and
2. $t_1^* = \min\{t_2^*, \sup\{t : p_t < \bar{p}\}\}$.³²

Proof. See Appendix Section A.3.1. □

By Lemma 1.5.4 the first line of Equation (1.6) corresponds to times $t \leq t_1^*$, the second line to $t \in (t_1^*, t_2^*)$, and the third line to $t \geq t_2^*$. Thus, the first and third lines are immediate from the definition of these threshold times. We now give a heuristic argument outlining the derivation of the second line, i.e. the equilibrium flow of adoption at times $t \in (t_1^*, t_2^*)$, where adoption is partial. At all these times, consumers must be exactly indifferent between adopting today and waiting for more information. Maintaining consumer indifference at these times requires that the cost and benefit of delaying be equal:

$$\underbrace{(\varepsilon + \lambda N_t)(1 - p_t)dt}_{\text{Probability of breakdown}} \underbrace{(0 - (-1))}_{\text{Benefit: Avoid Bad Product}} = \underbrace{(1 - (\varepsilon + \lambda N_t)(1 - p_t)dt)}_{\text{Probability of no breakdown}} \underbrace{(2p_{t+dt} - 1)rdt}_{\text{Cost: Discounting}}. \quad (1.7)$$

Delaying one's decision by an instant is beneficial if a breakdown occurs at that instant, allowing a consumer to permanently avoid the bad product. The gain in this case is $(0 - (-1))$, and this possibility arises with an instantaneous probability of $(\varepsilon + \lambda N_t)(1 - p_t)dt$. On the other hand, if no breakdown occurs, which happens with instantaneous probability $1 - (\varepsilon + \lambda N_t)(1 - p_t)dt$, then consumers incur an opportunity cost of $(2p_{t+dt} - 1)rdt$, reflecting the time cost of delayed adoption.³³ Ignoring terms of order dt^2 and

³²With the convention that if $\{t \geq 0 : p_t < \bar{p}\} = \emptyset$, then $\sup\{t : p_t < \bar{p}\} = 0$.

³³Note that ρ does not enter into this expression, because in the indifference region consumers obtain the same continuation payoff regardless of whether or not they obtain an adoption opportunity in the time interval $(t, t + dt)$ and hence are indifferent between receiving an opportunity to adopt or not.

rearranging yields $N_t = \frac{r(2p_t-1)}{\lambda(1-p_t)} - \frac{\varepsilon}{\lambda}$, as in Equation (1.6).³⁴

Finally, Figure 1.3 illustrates how from Equation 1.6, we obtain a unique equilibrium as a function of the parameters. Regions (2) and (3) represent values of (p_t, Λ_t) corresponding to the first line of Equation (1.6), so that no adoption takes place in these regions. Region (4) corresponds to partial adoption as given by the second line of Equation (1.6). Finally, region (1) corresponds to the third line of Equation (1.6) and thus to immediate adoption.

If (p_0, Λ_0) is in region (2), then initially no adoption takes place and the no-news posterior drifts upward according to the law of motion $\dot{p}_t = p_t(1-p_t)\varepsilon$, while Λ_t remains unchanged at Λ_0 . This yields a unique time $0 < t_1^* = t_2^*$ at which (p_t, Λ_t) hits the boundary separating regions (2) and (1); from then on consumers adopt immediately upon an opportunity so that $N_t = \rho e^{-\rho(t-t_2^*)} \bar{N}_{t_2^*}$ uniquely pins down the evolution of (p_t, Λ_t) . If (p_0, Λ_0) is in region (3), then again no initial adoption occurs and the no-news posterior drifts upward according to the law of motion $\dot{p}_t = p_t(1-p_t)\varepsilon$, while Λ_t remains unchanged at Λ_0 . However, now this yields a unique time $0 < t_1^*$ at which (p_t, Λ_t) hits the boundary separating regions (3) and (4), and at this time $\Lambda_{t_1^*} = \Lambda_0 > \Lambda(p_{t_1^*}) = \Lambda(\bar{p})$, so that we must have $t_1^* < t_2^*$. From t_1^* on the evolution of (p_t, Λ_t) is uniquely pinned down by the second line of Equation (1.6).³⁵ Thus, t_2^* is uniquely given by the first time t at which $\Lambda_t = \Lambda^*(p_t)$, at which point (p_t, Λ_t) enters region (1). Similar arguments show that when (p_0, Λ_0) starts in region (4), we have $t_1^* = 0$ and $t_2^* > t_1^*$ is the first time at which (p_t, Λ_t) , evolving according to the second line of Equation (1.6), enters region (1). Finally, if (p_0, Λ_0) is in region (1), then $0 = t_1^* = t_2^*$ and absent breakdowns all consumers adopt upon their

³⁴A bit more precisely, ignoring terms of order dt^2 , the right hand side of Equation 1.7 is given by $(1 - (\varepsilon + \lambda N_t)(1 - p_t)dt)(2(p_t + \dot{p}_t dt) - 1)rdt = r(2p_t - 1)dt$. Further rearrangement yields the desired expression.

³⁵Specifically, combining the second line of Equation (1.6) with Equation (1.1) yields the ODE:

$$\dot{p}_t = rp_t(2p_t - 1),$$

which pins down p_t uniquely given the initial value $p_{t_1^*} = \bar{p}$:

$$p_t = \frac{p_{t_1^*}}{2p_{t_1^*} - e^{r(t-t_1^*)}(2p_{t_1^*} - 1)}.$$

Plugging this back into $N_t = \frac{r(2p_t-1)}{\lambda(1-p_t)} - \frac{\varepsilon}{\lambda}$ uniquely pins down $\Lambda_t = \lambda \bar{N}_t$. Note that since $p_{t_1^*} > \frac{1}{2}$, p_t given above is strictly increasing and reaches $p^\#$ in finite time. Thus $t_2^* = \inf\{t : \Lambda_t < \Lambda^*(p_t)\} < +\infty$.




Figure 1.3: Partition of (p_t, Λ_t) when $\epsilon < \rho$

first opportunity from the beginning. This completes the description of the equilibrium.

As seen above, whether or not the equilibrium features a period of partial adoption depends on the fundamentals. More specifically, we can show that if consumers are forward-looking and not too optimistic, then $t_1^* < t_2^*$ arises whenever the potential for social learning in the economy is sufficiently large. To state this precisely, first note that from the Markovian description of equilibrium dynamics, it is easy to see that $\Lambda_0 = \lambda \bar{N}_0$ is a sufficient statistic for equilibrium when other fundamentals are fixed:

Lemma 1.5.5. Fix $r, \rho > 0$, $p_0 \in (0, 1)$, and $\epsilon \geq 0$. Suppose that $\hat{\lambda} \hat{N}_0 = \lambda_0 \bar{N}_0$. Let \hat{N}_t and N_t denote the unique equilibrium adoption flows under $(\hat{\lambda}, \hat{N}_0)$ and (λ, \bar{N}_0) , respectively, and let $\hat{p}_t, \hat{t}_1^*, \hat{t}_2^*$ and p_t, t_1^*, t_2^* denote the corresponding equilibrium beliefs and cutoff times. Then

1. $\hat{t}_i^* = t_i^*$ for $i = 1, 2$;
2. $\hat{p}_t = p_t$ for all t

3. and $\hat{\lambda}\hat{N}_t = \lambda N_t$ for all t .

Proof. Immediate from the proof of Theorem 1.5.1. \square

With this, the condition for partial adoption to arise in equilibrium can be stated as follows:

Lemma 1.5.6. *Fix ρ , ε and p_0 satisfying Conditions 1.5.2 and 1.5.3. Assume $p_0 < p^\#$. Then for all $r > 0$, there exists $\bar{\Lambda}_0(r) > 0$ such that $t_1^*(\Lambda_0) < t_2^*(\Lambda_0)$ if and only if $\Lambda_0 > \bar{\Lambda}_0(r)$.*

Proof. Set $\bar{\Lambda}_0(r) := \max\{\Lambda^*(p_0), \Lambda^*(\bar{p})\}$ and see Section A.6.1 in the Appendix. \square

On the other hand, if learning is *purely exogenous* ($\lambda = 0$ and $\varepsilon > 0$) or if consumers are *myopic* (" $r = +\infty$ "), then there is *never* any partial adoption, regardless of other parameters. In the former case, $0 = \Lambda_t < \Lambda^*(p)$ for all $p > \underline{p}$, so by Theorem 1.5.1 no consumers adopt until the no-news posterior hits \underline{p} (at $t_1^* = t_2^*$) and from then on all consumers adopt immediately when given a chance. The latter case corresponds to $\underline{p} = \bar{p} = \frac{1}{2}$ and $\Lambda^*(p) = +\infty$ for all $p > \frac{1}{2}$, so $t_1^* = t_2^* = \inf\{t : p_t > \frac{1}{2}\}$. Thus, the possibility of partial adoption in equilibrium hinges crucially both on consumers being forward-looking and on there being opportunities for social learning.

1.5.2 Shape of Adoption Curve

With the equilibrium characterization in place, we can explore implications for the shape of an innovation's *adoption curve*, which plots the percentage of adopters in the population against time. Conditional on no breakdowns up to time t , this is given by

$$A_t := \int_0^t \frac{N_s}{N_0} ds.$$

Conditional on the innovation being good, observed adoption levels at t will be exactly A_t . If the innovation is bad, then observed adoption levels follow A_t until the first breakdown (which occurs at a stochastic time), and remain constant from then on. As a result of the




Figure 1.4: Examples of S-shaped adoption curves (Source: Narayanan and O'Connor (2010), Figure 2.1.)

equilibrium characterization in Theorem 1.5.1, we obtain the following prediction for the shape of the adoption curve:

Corollary 1.5.7. *In the unique equilibrium of Theorem 1.5.1, the adoption curve A_t conditional on no breakdowns up to time t has the following shape:*

- for $0 \leq t < t_1^*$, $A_t = 0$
- for $t_1^* \leq t < t_2^*$, A_t is strictly increasing and convex in t
- for $t \geq t_2^*$, A_t is strictly increasing and concave in t .

In particular, if $t_1^* < t_2^*$, then Corollary 1.5.7 predicts that, possibly after an initial period of no adoption, the adoption curve conditional on no breakdowns exhibits an *S-shaped* (i.e. convex-concave) growth pattern. In the empirical literature on innovation adoption,³⁶ S-shaped adoption patterns have been widely documented for many different innovations over the past century, including new agricultural seed varieties, such as hybrid corn;

³⁶See, for example, Griliches (1957), Mansfield (1961), Mansfield (1968), Davies (1979), and Gort and Klepper (1982), among many others.

household electronics, such as refrigerators and color television; and industrial and medical innovations, such as the diesel locomotive and electrocardiographs. Figure 1.4 illustrates this for a selection of household technologies. Figure 1.5 represents a typical adoption curve generated in our model when $\varepsilon = 0$.




Figure 1.5: Adoption curve conditional on no breakdowns ($\varepsilon = 0$)

The intuition for S-shaped adoption curves in our model is as follows: There is no adoption before time t_1^* , because initially consumers are pessimistic about the quality of the innovation and strictly prefer to wait for information from the exogenous news source rather than risk adopting a bad product. The adoption curve is concave from time t_2^* on, because now consumers are sufficiently optimistic to strictly prefer adopting the innovation when given the chance, so that the flow of new adopters is depleted at the rate ρ at which adoption opportunities are generated.

More interestingly, *the period of convex growth coincides precisely with the period of informational free-riding (in the form of partial adoption)*. The reason for this is the fundamental trade-off between adopting now and waiting for more information that arises when consumers are forward-looking social learners. During the period (t_1^*, t_2^*) of partial adoption, consumers are indifferent between adopting immediately and delaying. Conditional on no breakdowns during this period, consumers grow increasingly optimistic about the quality

of the innovation, which increases their opportunity cost of delaying adoption. In order to maintain indifference as captured by equation (1.7), the benefit to delaying adoption must also increase over time: Since consumers are forward-looking, this can be achieved by increasing the arrival rate of future breakdowns, which improves the odds that waiting will allow consumers to avoid the bad product. Since consumers are social learners, the arrival rate of information is increasing in the flow N_t of new adopters. Thus, whenever there is informational free-riding, N_t is strictly increasing over time. Since N_t represents the rate of change of the proportion A_t of adopters in the population, this is equivalent to A_t being convex.

Once again, this result relies crucially on our two modeling assumptions that learning is social and that consumers are forward-looking. As we pointed out following Lemma 1.5.6, if learning is purely exogenous or if consumers are myopic, then $t_1^* = t_2^*$, in which case the adoption curve does not feature a region of convex growth. In order to generate S-shaped adoption patterns in the absence of either of our assumptions, alternative models appeal to specific distributions of consumer heterogeneity, for example Jensen (1982) (in a model of exogenous learning with forward looking consumers) or Young (2009) (in a model of myopic social learning). The interplay of social learning and forward-looking consumers allows us to explain convex growth in terms of purely *informational* incentives, thus suggesting a novel micro-foundation for S-shaped curves that remains valid even when consumers are fully homogeneous.

1.5.3 Welfare

We now examine ex ante consumer welfare, as captured by the time 0 equilibrium value to waiting, W_0 . Fix $r, \rho > 0$, $\varepsilon \geq 0$, and $p_0 \in (0, 1)$ satisfying Conditions 1.5.2 and 1.5.3. Then Lemma 1.5.5 and Lemma 1.2.4 imply that $W_0 = W_0(\Lambda_0)$ depends only on the potential for social learning in the economy. The key finding is the possibility of a saturation effect: For sufficiently large Λ_0 , additional increases in opportunities for social learning are welfare-neutral.

Nature of Inefficiency

We first note that, as is to be expected, the equilibrium is in general inefficient relative to the socially optimal cooperative benchmark:

Proposition 1.5.8. *Fix $r, \rho > 0$, $\varepsilon \geq 0$, and $p_0 \in (0, 1)$ satisfying Conditions 1.5.2 and 1.5.3. The unique equilibrium in Theorem 1.5.1 is socially optimal if and only if $\Lambda_0 < \Lambda^*(p_0)$.*

Proof. See Appendix Section A.5.1. □

Note that if $\Lambda_0 < \Lambda^*(p_0)$, then in equilibrium all consumers adopt immediately upon first opportunity, which is exactly as prescribed by the cooperative benchmark in Proposition 1.3.3. Whenever $\Lambda_0 > \Lambda^*(p_0)$, then the proof of Proposition 1.5.8 demonstrates two sources of inefficiency relative to the cooperative benchmark. First, provided we also have $p_0 \leq \bar{p}$ (so that $t_1^* > 0$), then adoption begins *too late* in equilibrium. Second, provided we also have $\Lambda_0 > \Lambda^*(\bar{p})$ (so that $t_1^* < t_2^*$), then even once consumers begin to adopt, the initial rate of adoption is *too slow* due to partial adoption on (t_1^*, t_2^*) . Note that both types of inefficiency rely on a sufficiently large potential for social learning. Moreover, for any $p_0 > 1/2$, the first type arises only if ε is sufficiently large or r is sufficiently small (in particular, if learning is purely social or if consumers are myopic, then $t_1^* = 0$). On the other hand, the second inefficiency relies on consumers being forward-looking, but can arise even if $\varepsilon = 0$.

Saturation Effect

The fact that the equilibrium can feature inefficiencies relative to the cooperative benchmark is to be expected. However, the second type of inefficiency discussed above, which arises when there is free-riding in the form of partial adoption, has the following more surprising implication:

Proposition 1.5.9. *Fix $r, \rho > 0$, $\varepsilon \geq 0$, and $p_0 \in (0, 1)$ satisfying Conditions 1.5.2 and 1.5.3. Let $\bar{\Lambda}_0 := \max\{\Lambda^*(p_0), \Lambda^*(\bar{p})\}$. Then in the unique equilibrium of Theorem 1.5.1, $W_0(\Lambda_0)$ satisfies the following:*

1. $W_0(\Lambda_0)$ is strictly increasing in Λ_0 whenever $\Lambda_0 < \bar{\Lambda}_0$;
2. $W_0(\Lambda_0) = W_0(\bar{\Lambda}_0)$ is constant in Λ_0 for all $\Lambda_0 \geq \bar{\Lambda}_0$.

Proof. See Appendix Section A.6.1. □

When $p_0 < p^\sharp$ so that $\bar{\Lambda}_0$ is finite, Proposition 1.5.9 states that an economy's ability to harness its potential for social learning is subject to a *saturation effect*: If Λ_0 is small, increases in Λ_0 are strictly beneficial; however, once Λ_0 is sufficiently large, any additional increase in Λ_0 is completely welfare-neutral. This is in stark contrast to the cooperative benchmark: There increases in Λ_0 are always strictly beneficial and for any $p_0 > \frac{1}{2}$ the first-best (complete information) payoff of $\frac{\rho}{r+\rho}p_0$ can be approximated in the limit as $\Lambda_0 \rightarrow \infty$. We illustrate this in Figure 1.6,³⁷ which for varying levels of Λ_0 plots the ratio of equilibrium and socially optimal welfare levels.




Figure 1.6: Ratio of equilibrium welfare to socially optimal welfare

To see the intuition for Proposition 1.5.9, suppose that $p_0 > \bar{p}$, so that $t_1^* = 0$. Then as long as $\Lambda_0 < \bar{\Lambda}_0$, all consumers adopt immediately upon their first opportunity until there is a breakdown. In this case, a slight increase in Λ_0 does not change consumers' behavior in the absence of a breakdown; however, conditional on the innovation being *bad*, it does

³⁷Parameters used to generate the figure are: $\varepsilon = 0$, $r = 1$, $p_0 = 0.6$, and $\rho = 1$.

increase the probability of a breakdown occurring prior to any given time—this is clearly welfare-improving, as more consumers are able to avoid the bad product. On the other hand, whenever $\Lambda_0 > \bar{\Lambda}_0$, then $t_1^* < t_2^*$, so that consumers must initially be *indifferent* between adopting and delaying. Then irrespective of the value of $\Lambda_0 \geq \bar{\Lambda}_0$, equilibrium incentives immediately imply that $W_0(\Lambda_0) = 2p_0 - 1$. In the next section, we provide some more intuition for the source of the saturation effect by studying the impact of increases in Λ_0 on equilibrium learning and adoption dynamics.

1.5.4 The Effect of Increased Opportunities for Social Learning

To further elucidate the saturation effect, this section examines the impact of an increase in Λ_0 on equilibrium learning dynamics and adoption levels. We find that the saturation effect corresponds to the following two surprising implications of partial adoption: Increased opportunities for social learning lead to strictly *less* learning over some periods of time and to a strict reduction in adoption of both good and bad innovations at all times.

Throughout this section we fix $r, \rho > 0$, $\varepsilon \geq 0$ and p_0 satisfying Conditions 1.5.2 and 1.5.3. To isolate the role of partial adoption, which is the inefficiency driving the saturation effect, we assume that $p^\sharp > p_0 > \bar{p}$, so that $t_1^* = 0$ and $\Lambda^*(p_0) < +\infty$. With these parameters fixed, Lemma 1.5.5 implies that Λ_0 is a sufficient statistic for all quantities we consider in this section. The following preliminary observation is central to the main results of this section:

Lemma 1.5.10. *Suppose that $\hat{\Lambda}_0 = \hat{\lambda}\hat{N}_0 > \Lambda_0 = \lambda\bar{N}_0 > \Lambda^*(p_0)$, with corresponding equilibrium flows of adoption \hat{N} and N . Then*

1. $0 < t_2^*(\Lambda_0) < t_2^*(\hat{\Lambda}_0)$.
2. For all $t < t_2^*(\Lambda_0)$, $\lambda N_t = \hat{\lambda}\hat{N}_t$.

Proof. See Appendix Section A.6.2. □

Point (ii) states that at all times during which there is partial adoption under both $\hat{\Lambda}_0$ and Λ_0 , the rate of social learning is the same. Intuitively, this is because in order to

maintain indifference between immediate adoption and an instantaneous delay, Equation (1.7) uniquely pins down the instantaneous arrival rate of breakdowns in the partial adoption region. The first bullet point states that increased opportunities for social learning prolong the initial period of partial adoption. To see the intuition, consider Figure 1.3: For any posterior p , $\Lambda^*(p)$, which represents the amount of future social information required to make consumers indifferent between adopting immediately at p and delaying, is the same under both $\hat{\Lambda}_0$ and Λ_0 . However, since $\hat{\Lambda}_0 > \Lambda_0$ and since by (ii) the evolution of beliefs in region (4) is the same under $\hat{\Lambda}_0$ and Λ_0 , it takes longer to reach the Λ^* -curve from the initial point $(p_0, \hat{\Lambda}_0)$.

Non-Monotonicity of Learning

In this section, we consider the effect of increased opportunities for social learning on the evolution of equilibrium beliefs. The following proposition states a non-monotonicity result: Increases in Λ_0 do not necessarily translate into increases in p_t at all times t . Specifically, if $\Lambda_0 \geq \Lambda^*(p_0)$, corresponding to the cutoff for the saturation effect, then upon an increase in Λ_0 there is a period of times at which p_t is strictly lower:

Proposition 1.5.11. *Fix r, ρ, ε , and p_0 satisfying Conditions 1.5.2 and 1.5.3 and such that $p_0 \in (\bar{p}, p^\sharp)$. Consider $\hat{\Lambda}_0 = \hat{\lambda}\hat{N}_0$ and $\Lambda_0 = \lambda\bar{N}_0$ such that $\hat{\Lambda}_0 > \Lambda_0 \geq \Lambda^*(p_0)$. Then there exists some $\bar{t} \in (t_2^*(\Lambda_0), +\infty)$ such that*

- $p_t^{\Lambda_0} = p_t^{\hat{\Lambda}_0}$ for all $t \leq t_2^*(\Lambda_0)$,
- $p_t^{\Lambda_0} > p_t^{\hat{\Lambda}_0}$ for all $t \in (t_2^*(\Lambda_0), \bar{t})$,
- $p_t^{\Lambda_0} < p_t^{\hat{\Lambda}_0}$ for all $t > \bar{t}$.

However, when $\Lambda_0 < \Lambda^*(p_0)$, then $p_t^{\Lambda_0}$ is strictly increasing in Λ_0 for all t .

Proof. See Appendix Section A.6.2. □

Note that by Equation (1.1), the probability of a breakdown occurring prior to time t conditional on the innovation being bad is given by

$$1 - e^{-\int_0^t (\varepsilon + \lambda N_s) ds} = 1 - \frac{p_0(1 - p_t)}{p_t(1 - p_0)}, \quad (1.8)$$

which is increasing in p_t . Thus, Proposition 1.5.11 has the surprising implication that whenever Λ_0 is large enough, any additional increase in opportunities for social learning will result in consumers being strictly *less* likely to find out about a bad product over a period of times.

The intuition for Proposition 1.5.11 is closely related to whether or not there is free-riding in the form of partial adoption (and hence relies on consumers being forward-looking social learners). Whenever $\Lambda_0 < \Lambda^*(p_0)$, then $0 = t_1^* = t_2^*$, so that absent breakdowns all consumers adopt immediately upon their first opportunity. In this case, it is easy to see from Theorem 1.5.1 that the rate λN_t at which social learning occurs is strictly increasing in Λ_0 : We have $\lambda N_t = \rho e^{-\rho t} \Lambda_0$ for all t . Thus, by Equation (1.8) increasing Λ_0 necessarily speeds up learning at all times.

On the other hand, if $\Lambda_0 > \Lambda^*(p_0)$, then $0 = t_1^* < t_2^*(\Lambda_0)$ and the equilibrium features an initial region of partial adoption. In this case, an increase to $\hat{\Lambda}_0 > \Lambda_0$ has the following effect. By Lemma 1.5.10, free-riding occurs over a longer period of time: $t_2^*(\hat{\Lambda}_0) > t_2^*(\Lambda_0)$; moreover, at all times $t \leq t_2^*(\Lambda_0)$ where there is free-riding under both Λ_0 and $\hat{\Lambda}_0$, the rate of social learning is the *same*: $\lambda N_t = \hat{\lambda} \hat{N}_t$. This explains the first bullet point in Proposition 1.5.11. The strict slowdown in learning at times just after $t_2^*(\Lambda_0)$ is due to the following: The proof of Theorem 1.5.1 shows that whenever $t_1^* < t_2^*$, then the flow of adopters N_t is continuous at all times *except* at exactly t_2^* , where there is a discontinuous increase. This is evident from the adoption curve in Figure 1.5 where a visible non-differentiability exists at the point of transition from partial adoption to immediate adoption. Since $t_2^*(\Lambda_0) < t_2^*(\hat{\Lambda}_0)$, this means that at $t_2^*(\Lambda_0)$ the difference between $\hat{\lambda} \hat{N}_t$ and λN_t jumps from 0 to a strictly negative value, resulting in the temporary slowdown in learning stated in the second bullet point.

Finally, learning under $\hat{\Lambda}_0$ must eventually overtake learning under Λ_0 , because at time 0 the payoff to immediate adoption is the same under both Λ_0 and $\hat{\Lambda}_0$ and in both cases consumers are indifferent between adopting immediately and delaying. This relates back to

the saturation effect for welfare observed in Proposition 1.5.9 as follows. By Lemma 1.2.4, ex-ante welfare W_0 under $\Lambda_0 > \Lambda^*(p_0)$ can be written as

$$W_0(\Lambda_0) = \int_0^\infty \rho e^{-(r+\rho)\tau} \frac{p_0}{p_\tau^{\Lambda_0}} (2p_\tau^{\Lambda_0} - 1) d\tau = p_0 \int_0^\infty \rho e^{-(r+\rho)\tau} \left(2 - (p_\tau^{\Lambda_0})^{-1}\right) d\tau,$$

and similarly for $\hat{\Lambda}_0 > \Lambda_0$. The non-monotonicity result for beliefs then has the following implication. If a consumer obtains his first adoption opportunity prior to $t_2^*(\Lambda_0)$, his expected payoff is the same under Λ_0 and $\hat{\Lambda}_0$; if his first adoption opportunity is during $(t_2^*(\Lambda_0), \bar{t})$, he is strictly worse off under $\hat{\Lambda}_0$, because in case the innovation is bad he is less likely to have found out by then; finally, if his first opportunity is after \bar{t} , he is strictly better off under $\hat{\Lambda}_0$. Depending on $\hat{\Lambda}_0$, \bar{t} adjusts endogenously to balance out the benefits, which arrive at times after \bar{t} , with the costs incurred at times $(t_2^*(\Lambda_0), \bar{t})$. This produces the saturation effect in Proposition 1.5.9.

Even more strongly, in Section 1.7, we exploit the non-monotonicity result for beliefs to construct an example involving consumers with heterogeneous discount rates in which an increase in Λ_0 is not only not beneficial, but in fact strictly *hurts* aggregate welfare.

Slowdown in Adoption

We now consider the effect of increased opportunities for social learning on observed adoption levels, analyzing separately the case of a good innovation and of a bad innovation.

Adoption Conditional on a Good Product: Recall that A_t denotes the percentage of consumers in the population who adopt the innovation by time t conditional on no breakdowns before t , which is the same as the percentage of adopters at t conditional on the innovation being good:

$$A_t(G) = A_t := \int_0^t \frac{N_s}{N_0} ds.$$

Note also that by Lemma 1.5.5, Λ_0 is a sufficient statistic for the equilibrium levels of A_t holding fixed r, ρ, p_0 and ε , because $\frac{N_s}{N_0} = \frac{\lambda N_s}{\Lambda_0}$ and Λ_0 is a sufficient statistic for λN_s .

For a good innovation, we show that when the potential for social learning Λ_0 is small, additional small increases in opportunities for social learning have no effect on adoption levels, but when Λ_0 is sufficiently large, increases strictly drive down adoption levels at all times. Once again, the cutoff is given by the level $\Lambda^*(p_0)$ above which partial adoption occurs.

Proposition 1.5.12. *Fix r, ρ, ε , and p_0 satisfying Conditions 1.5.2 and 1.5.3 and such that $p_0 \geq \bar{p}$. Then for all t , $A_t(\Lambda_0, G)$ is constant in Λ_0 for all $\Lambda_0 \leq \Lambda^*(p_0)$ and strictly decreasing in Λ_0 for all $\Lambda_0 > \Lambda^*(p_0)$.*

Proof. See Appendix Section A.6.3. □

The reason why $A_t(\Lambda_0, G)$ is constant for all $\Lambda_0 \leq \Lambda^*(p_0)$ is familiar: For all such Λ_0 , consumers adopt upon their first opportunity and $A_t = 1 - e^{-\rho t}$. If $\Lambda_0 > \Lambda^*(p_0)$, then the strict slowdown in adoption is due to increased free-riding in the form of partial adoption. More precisely, an increase from $\Lambda_0 > \Lambda^*(p_0)$ to $\hat{\Lambda}_0$ has two effects, as summarized in Lemma 1.5.10: First, on the *extensive margin*, increased opportunities for social learning push out t_2^* and lead to a longer period of free-riding under $\hat{\Lambda}_0$. Second, on the *intensive margin*, the increase strictly drives down the growth rate of A_t at all times prior to $t_2^*(\Lambda_0)$:

$$\dot{A}_t = \frac{N_t}{N_0} = \frac{\lambda N_t}{\Lambda_0} = \frac{\hat{\lambda} \hat{N}_t}{\Lambda_0} > \frac{\hat{\lambda} \hat{N}_t}{\hat{\Lambda}_0} = \frac{\hat{N}_t}{\hat{N}_0} = \hat{A}_t.$$

Figure 1.7 illustrates these two effects and their implications for a strict slowdown in adoption. Finally, from $t_2^*(\Lambda_0)$ adoption occurs at a maximal rate under Λ_0 , so that from then on \hat{A}_t must remain below A_t by feasibility.

Two remarks are in order. First, our prediction of a strict slowdown of adoption of the good product in response to increased opportunities for social learning once again relies crucially on consumers being forward-looking. If consumers are myopic, then by the first part of Proposition 1.5.12 adoption levels at all times remain unchanged following the increase. More interestingly, if consumers are myopic, it is not possible to generate this prediction under perfect bad news even if we allow for an *arbitrary* distribution of




Figure 1.7: Changes in adoption levels of a good product under perfect bad news ($\hat{\lambda} > \lambda$)

heterogeneity in tastes. Thus, while models of innovation adoption by myopic social learners, such as Young (2009), can generate S-shaped adoption curves by imposing suitable distributions of consumer heterogeneity, the prediction in our model of a strict reduction in initial adoption of a good innovation is novel.

Second, Proposition 1.5.12 implies that conditional on a good product, increased opportunities for social learning are welfare-neutral at best (if $\Lambda_0 < \Lambda^*(p_0)$) and potentially strictly harmful (if $\Lambda_0 \geq \Lambda^*(p_0)$), since adoption levels are unchanged in the former case and in the latter case adoption is strictly delayed. Therefore any potential welfare gains due to increased opportunities for social learning must result from more consumers being able to avoid the *bad* product. We now study this point by analyzing the effect of increases in Λ_0 on adoption levels of a bad product.

Adoption Conditional on a Bad Product: Conditional on a bad innovation, adoption is stochastic, following A_t until the first breakdown, which occurs at a random time, and remaining constant from then on. We therefore study the effect of increased opportunities for social learning on the *expected* percentage of adopters at time t conditional on a bad

product, which is given by:

$$\begin{aligned} A_t(B) &:= \int_0^t (\varepsilon + \lambda N_\tau) e^{-\int_0^\tau (\varepsilon + \lambda N_s) ds} \left(\int_0^\tau \frac{N_s}{\bar{N}_0} ds \right) d\tau + e^{-\int_0^t (\varepsilon + \lambda N_s) ds} \int_0^t \frac{N_s}{\bar{N}_0} ds \\ &= \int_0^t \frac{N_\tau}{\bar{N}_0} e^{-\int_0^\tau (\varepsilon + \lambda N_s) ds} d\tau, \end{aligned}$$

where the second line is obtained by integrating the first expression by parts. Again, Λ_0 is a sufficient statistic for $A_t(B)$ when r , ρ , p_0 , and ε are fixed. For bad innovations, increased opportunities for social learning always produce strict decreases in the expected level of adoption at all times, irrespective of the original level of Λ_0 :

Proposition 1.5.13. *Fix r , ρ , ε , and p_0 satisfying Conditions 1.5.2 and 1.5.3 and such that $p_0 \geq \bar{p}$. Then $A_t(\Lambda_0, B)$ is strictly decreasing in Λ_0 for all $t > 0$.*

Proof. See Appendix Section A.6.3. □

If $\Lambda_0 < \Lambda^*(p_0)$, this is immediate since by Proposition 1.5.11 and Proposition 1.5.12 adoption levels conditional on no breakdowns are the same, but breakdowns prior to any time are more likely for higher values of $\Lambda_0 < \Lambda^*(p_0)$. If $\Lambda_0 \geq \Lambda^*(p_0)$, then there is a tension: On the one hand, Proposition 1.5.12 implies that an increase in Λ_0 leads at all times to strictly lower adoption levels conditional on no breakdowns, but on the other hand, the non-monotonicity result for learning implies that there are times before which a breakdown is strictly more likely under lower Λ_0 . We show that the former effect always strictly dominates.

Proposition 1.5.12 and Proposition 1.5.13 relate to the saturation effect observed in Proposition 1.5.9 as follows: If $\Lambda_0 < \Lambda^*(p_0)$, then small increases in opportunities for social learning do not affect adoption conditional on the good product, but strictly decrease the number of consumers adopting the bad product by any time, leading to an overall welfare gain. On the other hand, if $\Lambda_0 \geq \Lambda^*(p_0)$, then increased opportunities for social learning strictly decrease adoption both for good products (which is harmful) and for bad products (which is beneficial), making welfare predictions a priori ambiguous. However, the

saturation effect illustrates that in welfare terms these two implications balance out exactly.

1.6 Perfect Good News

1.6.1 Equilibrium Characterization

We now turn to study equilibrium behavior when learning is via perfect good news. As under perfect bad news, the unique equilibrium is Markovian in the state variables (p_t, Λ_t) . Surprisingly, however, regardless of the potential for social learning in the economy, the unique equilibrium under perfect good news does not exhibit any region of partial adoption and adoption at each time is *all-or-nothing*:

Theorem 1.6.1. *Let $r, \rho, \bar{N}_0 > 0$, $p_0 \in (0, 1)$, and $\lambda, \varepsilon \geq 0$. There exists a unique equilibrium. In the unique equilibrium, N_t is Markovian in (p_t, Λ_t) (or equivalently (p_t, \bar{N}_t)) for all t and satisfies:*

$$N_t = \begin{cases} \rho \bar{N}_t & \text{if } p_t > p^* \\ 0 & \text{if } p_t \leq p^*. \end{cases} \quad (1.9)$$

where

$$p^* = \frac{(\varepsilon + r)(\rho + r)}{2(\varepsilon + \rho)(\varepsilon + r) - \varepsilon\rho}.$$

To prove Theorem 1.6.1 we again invoke the quasi-single crossing property for equilibrium incentives established in Theorem 1.4.1. Suppose $N_{t \geq 0}$ is an equilibrium flow of adopters. Let $p_{t \geq 0}$ and $W_{t \geq 0}$ be the corresponding no-news posterior and value to waiting, and let $\Lambda_{t \geq 0} := \lambda \bar{N}_{t \geq 0}$ describe the evolution of the economy's potential for social learning. By Theorem 1.4.1, there are times³⁸

$$t_1^* := \inf\{t : 2p_t - 1 \leq W_t\},$$

$$t_2^* := \inf\{t : 2p_t - 1 < W_t\},$$

³⁸With the usual convention that $\inf \emptyset = +\infty$.

such that (appealing also to right-continuity) N must satisfy

$$\begin{cases} N_t = \rho \bar{N}_t & \text{if } t < t_1^*, \\ 2p_t - 1 = W_t & \text{if } t \in [t_1^*, t_2^*) \\ N_t = 0 & \text{if } t \geq t_2^*. \end{cases}$$

In the following, we build on this fact to establish the existence of a unique equilibrium as a function of the parameter values. The following lemma establishes the all-or-nothing nature of the perfect good news equilibrium:

Lemma 1.6.2. *Suppose either $\varepsilon > 0$ or $p_0 \neq \frac{1}{2}$.³⁹ Let $N_{t \geq 0}$ be an equilibrium with associated threshold times t_1^* and t_2^* . Then $t_1^* = t_2^* =: t^*$.*

Proof. See Appendix Section A.3.2. □

Thus, absent breakthroughs, *all* consumers adopt immediately if given the chance prior to t^* , and after t^* , consumers stop adopting altogether and rely solely on information generated by exogenous sources (if $\varepsilon = 0$, both adoption and learning come to a permanent standstill at this point). If a breakthrough occurs at any time (prior to or after t^*), then from then on all consumers adopt the innovation whenever given a chance.

To see the intuition for the all-or-nothing nature of the equilibrium, suppose we had $t_1^* < t_2^*$. Then consumers would be indifferent between adopting and delaying at each time $t \in (t_1^*, t_2^*)$. As with perfect bad news, we can compare a consumer's payoff to adopting at t with the payoff to delaying his decision by an instant and decompose the difference into two terms:

$$r(2p_t - 1)dt + p_t(\lambda N_t + \varepsilon)dt \left(1 - \frac{\rho}{r + \rho}\right).$$

The first term represents the gain to immediate adoption if no breakthrough occurs between t and $t + dt$, which happens with instantaneous probability $1 - p_t(\lambda N_t + \varepsilon)dt$. Just as with perfect bad news, the gain to adopting immediately in this case is $r(2p_{t+dt} - 1)dt$,

³⁹If $\varepsilon = 0$ and $p_0 = \frac{1}{2}$, then it is easy to see that the unique equilibrium must be $N \equiv 0$, so that $t_1^* = 0 < t_2^* = +\infty$.

representing time discounting at rate r and the fact that at $t + dt$ the consumer remains indifferent between adopting if given the chance and delaying. The second term represents the gain to immediate adoption if there is a breakthrough between t and $t + dt$, which happens with instantaneous probability $p_t(\lambda N_t + \varepsilon)dt$. Now the situation is very different from the perfect bad news setting: A breakthrough conclusively signals *good* quality, so a consumer who delays his decision by an instant will adopt immediately at his next opportunity. This results in a discounted payoff of $\frac{\rho}{r+\rho}$, reflecting the stochasticity of adoption opportunities. On the other hand, by adopting at t , the consumer receives a payoff of $1 > \frac{\rho}{r+\rho}$ immediately. Thus, regardless of whether or not there is a breakthrough between t and $t + dt$, there is a strictly positive gain to adopting immediately at t , which contradicts indifference at t .

The above argument illustrates a fundamental difference between the bad news and good news setting. In order to maintain indifference over a period of time between immediate adoption and waiting, it must be possible to acquire *decision-relevant information* by waiting: Consumers who are prepared to adopt at t will be willing to delay their decision by an instant only if there is a possibility that at the next instant they will no longer be willing to adopt. In the bad news setting, this is indeed possible: If there is a breakdown between t and $t + dt$, then the innovation is revealed to be bad and no one is willing to adopt from $t + dt$ on. On the other hand, if learning is via good news, this cannot happen: A breakthrough between t and $t + dt$ reveals the innovation to be good, so consumers strictly prefer to adopt from $t + dt$ on; if there is no breakthrough, then consumers remain indifferent at $t + dt$, so in either case the information obtained is not decision-relevant.⁴⁰

With Lemma 1.6.2, the derivation of Theorem 1.6.1 is straightforward. To this end, we show that any equilibrium can be characterized in terms of a cutoff posterior that only depends on ε , ρ , and r . Given any equilibrium $N_{t \geq 0}$ with associated no-news posteriors

⁴⁰Note that breakthroughs do of course convey decision-relevant information at beliefs where consumers strictly prefer to delay. But during a region of indifference, this cannot be the case.

$p_{t \geq 0}$, value to waiting $W_{t \geq 0}$, and cutoff time t^* , define

$$H_t := p_t \int_0^\infty (\varepsilon + \lambda N_{t+\tau}) e^{-(\varepsilon\tau + \int_t^{t+\tau} \lambda N_s ds)} \frac{\rho}{r + \rho} e^{-(r+\rho)\tau} d\tau.$$

Thus, H_t represents a consumer's expected value to waiting at time t given that from t on he adopts only if there has been a breakthrough and given that the population's flow of adoption follows N . By optimality of W_t , we must have $H_t \leq W_t$ for all t . We can define a lower bound for H_t : For any posterior $p \in (0, 1)$, let

$$H(p, 0) := p \int_0^\infty \varepsilon e^{-\varepsilon\tau} \frac{\rho}{r + \rho} e^{-(r+\rho)\tau} d\tau = p \frac{\rho\varepsilon}{(r + \rho)(\varepsilon + r + \rho)}.$$

$H(p, 0)$ represents a consumer's expected value to waiting at posterior p , given that he adopts only once there has been a breakthrough and given that breakthroughs are only generated exogenously. Note that for all t , we have $H(p_t, 0) \leq H_t$: When breakthroughs are generated both exogenously and at rate λN_t , then the probability that a breakthrough is generated by any given time is (weakly) greater than if learning is purely exogenous; this benefits a consumer who only adopts once there has been a breakthrough. Moreover, for all $t \geq t^*$, we have $W_t = H(p_t, 0)$.

Recall the definition of p^* in Theorem 1.6.1, $p^* := \frac{(\varepsilon+r)(\rho+r)}{2(\varepsilon+\rho)(\varepsilon+r)-\varepsilon\rho}$, and note that p^* is the unique solution to $2p^* - 1 = H(p^*, 0)$. By definition of p^* , if $p_t \leq p^*$ at any time t , then

$$2p_t - 1 \leq H(p_t, 0) \leq H_t \leq W_t,$$

so for all $t < t^*$, we must have $p_t > p^*$. Conversely, if $t^* < +\infty$ and $t \geq t^*$, then

$$2p_t - 1 \leq W_t = H(p_t, 0),$$

so $p_t \leq p^*$ for all $t \geq t^*$. We summarize this in the following lemma:

Lemma 1.6.3. *Let $N_{t \geq 0}$ be an equilibrium with corresponding cutoff time t^* and no-news posterior $p_{t \geq 0}$. Then*

$$p_t > p^* \Leftrightarrow t < t^*.$$

Given Lemma 1.6.2 and Lemma 1.6.3, the equilibrium characterization under perfect

good news follows readily. Equation (1.9) is immediate from Lemma 1.6.3. For fixed parameters, we then obtain the unique equilibrium as follows: If $p_0 \leq p^*$, then $t^* = 0$ and $N_t = 0$ for all t . If $p_0 > p^*$, then we must have $t^* > 0$ and $N_t = \rho e^{-\rho t} \bar{N}_0$ for all $t < t^*$; if in addition $\varepsilon > 0$ or $p_0 \left(1 + e^{-\lambda \bar{N}_0}\right) < 1$, then $t^* < +\infty$ is uniquely determined as the solution to

$$p_t = \frac{p_0}{p_0 + (1 - p_0) e^{t + (1 - e^{-\rho t}) \bar{N}_0}} = p^*. \quad (1.10)$$

If instead $p_0 > p^*$ and $\varepsilon = 0$ and $p_0 \left(1 + e^{-\lambda \bar{N}_0}\right) \geq 1$, then Equation (1.10) does not admit a solution, and we must have $t^* = +\infty$: In this case, the potential for social learning in the economy is so small that even a bad innovation is eventually adopted by all consumers, despite the fact that no breakthroughs are ever generated.

As highlighted at the beginning of the section, the equilibrium under perfect good news is Markovian in (p_t, Λ_t) . However, in marked contrast to the bad news case, if $\varepsilon = 0$, then adoption behavior is independent of the discount rate r : Even very patient consumers will behave entirely myopically, adopting the innovation at all posteriors above $\frac{1}{2}$ and not adopting otherwise. If $\varepsilon > 0$, then consumers' forward-looking nature is reflected by the fact that the cutoff posterior p^* below which consumers are unwilling to adopt is $\frac{(r+\rho)(r+\varepsilon)}{2(r+\rho)(r+\varepsilon)-\rho\varepsilon} > \frac{1}{2}$. In both cases, the cutoff posterior does not depend on λ or \bar{N}_0 : Social learning only affects the *time* t^* at which adoption ceases conditional on no breakthroughs. Moreover, as under perfect bad news, it is easy to see that holding fixed other parameters, $\Lambda_0 = \lambda \bar{N}_0$ is a sufficient statistic for equilibrium behavior:

Lemma 1.6.4. *Fix $r, \rho > 0$, $p_0 \in (0, 1)$, and $\varepsilon \geq 0$. Suppose that $\hat{\lambda} \hat{N}_0 = \lambda_0 \bar{N}_0$. Let \hat{N}_t and N_t denote the unique equilibrium adoption flows under $(\hat{\lambda}, \hat{N}_0)$ and (λ, \bar{N}_0) , respectively, and let \hat{p}_t , \hat{t}^* and p_t , t^* denote the corresponding equilibrium beliefs and cutoff times. Then*

1. $\hat{t}^* = t^*$;
2. $\hat{p}_t = p_t$ for all t
3. and $\hat{\lambda} \hat{N}_t = \lambda N_t$ for all t .

Proof. Immediate from the proof of Theorem 1.6.1. □

1.6.2 Shape of Adoption Curve

Theorem 1.6.1 has the following implication for the shape of adoption curves in good news markets:

Corollary 1.6.5. *In the unique equilibrium of Theorem 1.6.1, the proportion of adopters in the population is strictly increasing and concave for all $t < t^*$ and given by*

$$A_t := \int_0^t \frac{N_s}{\bar{N}_0} ds = 1 - e^{-\rho t}.$$

If there is a breakthrough prior to t^ , then the proportion of adopters is given by $1 - e^{-\rho t}$ for all t ; if the first breakthrough occurs at $s > t^*$,⁴¹ then adoption comes to a temporary standstill between t^* and s , and for all $t \geq s$, the proportion of adopters is strictly increasing and concave and given by $1 - e^{-\rho(t^*+t-s)}$.*

Thus, as illustrated in Figure 1.8,⁴² adoption proceeds in *concave “bursts”*: Up to time t^* , all consumers adopt the innovation upon their first opportunity, with the flow of new adopters declining at the rate ρ at which adoption opportunities arrive. Conditional on no breakthroughs, adoption comes to a standstill at time t^* , because by that point consumers are pessimistic enough about the product to prefer to delay adoption. If $\varepsilon > 0$, then exogenous news sources might generate a breakthrough after t^* , in which case a second concave burst in adoption occurs.

While less common than the S-shaped curves we predicted under bad news,⁴³ this type of adoption pattern also corresponds to recurrent empirical findings. For instance, the

⁴¹This occurs only if $\varepsilon > 0$.

⁴²The parameters used to generate the figure are: $\varepsilon = 1/2$, $r = 1$, $\rho = 1$, $\lambda = 0.5$, and $p_0 = 0.7$.

⁴³Note that in our model purely concave adoption curves can also arise under bad news if the economy’s potential for social learning is relatively limited or consumers are very optimistic (so that $t_1^* = t_2^*$). The key difference is that under perfect good news adoption curves are *necessarily* concave, even in economies with a large potential for social learning or with fairly pessimistic and forward-looking consumers.

marketing literature⁴⁴ has coined the term “fast-break” product life cycle (PLC) to describe goods with large initial sales volumes accompanied by a gradual decline in new purchases (implying a concave adoption pattern), in contrast to S-shaped PLCs that initially feature low sales volumes accompanied by a gradual increase in the number of new purchases. The textbook example for fast-break PLCs is the movie industry,⁴⁵ as illustrated in Figure 1.9. Given that the movie industry is also sometimes cited as a typical example of a good news market⁴⁶ with learning occurring predominantly via positive events such as awards and recommendations in social media, this finding appears to be in line with our model.




Figure 1.8: Adoption Curves under Perfect Good News (blue = breakthrough before t^* ; yellow = breakthrough after t^* ; pink = bad quality)

1.6.3 The Effect of Increased Opportunities for Social Learning

To further illustrate the distinction between good news and bad news markets, we now study the effect of increased opportunities for social learning under good news. In contrast

⁴⁴Cf. Keillor (2007) pp. 51-61.

⁴⁵Additional evidence can be found in Davies (1979)’s study of the diffusion of 22 post-war process innovations among industries in the UK. In the context of his probit model of innovation diffusion, he finds that while S-shaped (logistic) diffusion paths are characteristic of complex and expensive innovations, they are less suited to fitting the diffusion paths of simpler and less expensive innovations, which typically feature rapid, essentially concave growth from the beginning and are better approximated by a lognormal model.

⁴⁶Cf. Board and Meyer-ter-Vehn (2013)




Figure 1.9: “Adoption” patterns for various blockbuster movies (Source: McLaren and DePaolo (2009))

to our results under perfect bad news, we find that increased opportunities for social learning (essentially) always speed up learning, leave initial adoption levels unaffected, and are strictly welfare-improving—all three results are due to the absence of partial adoption regions under good news. Throughout this section, we fix $r, \rho > 0$, $p_0 \in (0, 1)$, and $\varepsilon \geq 0$, and let p^* denote the equilibrium cutoff posterior:

$$p^* = \frac{(r + \rho)(r + \varepsilon)}{2(r + \rho)(r + \varepsilon) - \rho\varepsilon'}$$

which is independent of the potential for social learning. Given that all other parameters are fixed, Lemma 1.6.4 implies that $\Lambda_0 = \lambda \bar{N}_0$ is a sufficient statistic for all the quantities we consider in this section.

Learning Speeds Up

We first turn to the effect of increased opportunities for social learning on equilibrium beliefs. As a result of the all-or-nothing nature of the perfect good news equilibrium, we can see that learning necessarily speeds up—this is in contrast to the possibility of nonmonotonicities due to partial adoption under perfect bad news. More precisely:

Proposition 1.6.6. Fix $\hat{\Lambda}_0 > \Lambda_0$ ⁴⁷ and let $t^*(\hat{\Lambda}_0)$, $p_t^{\hat{\Lambda}_0}$ and $t^*(\Lambda_0)$, $p_t^{\Lambda_0}$ denote the corresponding equilibrium cutoff times and posteriors conditional on no breakthrough.

1. If $p_0 > p^*$, then

- $0 < t^*(\hat{\Lambda}_0) < t^*(\Lambda_0)$
- $p_t^{\hat{\Lambda}_0} < p_t^{\Lambda_0}$ for all $t > 0$
- $p_{t^*(\hat{\Lambda}_0)+k}^{\hat{\Lambda}_0} = p_{t^*(\Lambda_0)+k}^{\Lambda_0}$ for all $k \geq 0$.

2. If $p_0 \leq p^*$, then

- $t^*(\hat{\Lambda}_0) = t^*(\Lambda_0) = 0$
- $p_t^{\hat{\Lambda}_0} = p_t^{\Lambda_0}$ for all t .

If $p_0 > p^*$, then conditional on no breakthroughs, all consumers adopt immediately upon an opportunity until the time t^* at which the cutoff posterior p^* is reached. By Theorem 1.6.1, there is never any partial adoption, so that an increase from Λ_0 to $\hat{\Lambda}_0$ directly translates into a faster rate of social learning at all times t prior to $\min\{t^*(\hat{\Lambda}_0), t^*(\Lambda_0)\}$: $\lambda N_t = \rho e^{-\rho t} \Lambda_0 < \rho e^{-\rho t} \hat{\Lambda}_0 = \hat{\lambda} \hat{N}_t$. Since the cutoff posterior p^* is independent of social learning, this implies that $t^*(\hat{\Lambda}_0) < t^*(\Lambda_0)$ and that learning is strictly faster under $\hat{\Lambda}_0$ at all times. However, once the cutoff posterior is reached, information is generated at the constant exogenous rate ε , which means that conditional on $t > t^*$, beliefs depend only on $t - t^*$, as summarized in the third bullet point under (i).

On the other hand, if $p_0 \leq p^*$, then all consumers rely entirely on the exogenous news source from the beginning, so the potential for social learning is irrelevant.

No Initial Slowdown of Adoption

The all-or-nothing nature of the perfect good news equilibrium also implies that increased opportunities for social learning do not affect initial adoption levels—this is again in contrast

⁴⁷If $\varepsilon = 0$ we assume that $p_0(1 + e^{-\Lambda_0}) < 1$ so that $t^*(\Lambda_0) < \infty$.

to the possibility of initial slowdowns due to partial adoption under perfect bad news. More precisely:

Proposition 1.6.7. *Suppose $\hat{\Lambda}_0 > \Lambda_0$.*

1. *If $p_0 > p^*$, then:*

- *For all $t \leq t^*(\hat{\Lambda}_0)$, $A_t(\hat{\Lambda}_0; \theta) = A_t(\Lambda_0; \theta) = 1 - e^{-\rho t}$ for $\theta = B, G$.*
- *For all $t > t^*(\hat{\Lambda}_0)$, $A_t(\hat{\Lambda}_0; \theta) < A_t(\Lambda_0; \theta)$ for $\theta = B, G$.*

2. *If $p_0 \leq p^*$, then for all t :*

- $A_t(\Lambda_0; B) = A_t(\hat{\Lambda}_0; B) = 0$;
- $A_t(\Lambda_0; G) = A_t(\hat{\Lambda}_0; G) = \left(1 - \frac{\rho}{\rho - \varepsilon} e^{-\varepsilon t}\right) + \frac{\varepsilon}{\rho - \varepsilon} e^{-\rho t}$.

Until $t^*(\hat{\Lambda}_0)$ all consumers adopt immediately upon an opportunity under both Λ_0 and $\hat{\Lambda}_0$ regardless of the quality of the innovation. However, from $t^*(\hat{\Lambda}_0)$ on, expected adoption levels are strictly lower under $\hat{\Lambda}_0$ than under Λ_0 : If the innovation is bad, this is because adoption comes to a permanent standstill under $\hat{\Lambda}_0$ (until a further breakthrough generated by the exogenous information ε), but continues until $t^*(\Lambda_0)$ under Λ_0 . If the innovation is good, the result is again immediate for all $t \leq t^*(\Lambda_0)$ since adoption occurs at the maximal rate under Λ_0 . For $t > t^*(\Lambda_0)$, there are two opposing effects: On the one hand, the guaranteed lower bound on adoption is higher under Λ_0 , but on the other hand the probability of a breakthrough occurring prior to time t is always higher under $\hat{\Lambda}_0$. We show in the Appendix Section A.7.1 that the former effect dominates.

On the other hand, if $p_0 \leq p^*$, then increased opportunities for social learning once again have no effect at all on adoption levels, because no consumers adopt until the exogenous news source generates a breakthrough.

No Saturation Effect

Proposition 1.6.7 showed that from time $t^*(\hat{\Lambda}_0)$ on, adoption levels for both good and bad quality products are strictly lower under $\hat{\Lambda}_0 > \Lambda_0$ than under Λ_0 . In welfare terms, the

former effect is harmful while the latter is beneficial. This raises the question whether welfare under perfect good news might be subject to a similar saturation effect as under bad news. Provided $p_0 > p^*$ and $\varepsilon > 0$, the answer is negative:

Proposition 1.6.8. *Suppose $\hat{\Lambda}_0 > \Lambda_0$.*

- *If $p_0 > p^*$ and $\varepsilon > 0$, then $W_0(\hat{\Lambda}_0) > W_0(\Lambda_0)$.*
- *If $p_0 \leq p^*$ or $\varepsilon = 0$, then $W_0(\hat{\Lambda}_0) = W_0(\Lambda_0)$.*

Thus, in contrast to the perfect bad news case, increased opportunities for social learning are always strictly beneficial, except in two cases: If consumers rely entirely on exogenous information ($p_0 \leq p^*$), or if there is no exogenous information ($\varepsilon = 0$). Welfare-neutrality in these two exceptional cases is clear: Increased opportunities for social learning can have an effect on welfare only if there are histories at which a consumer's decision whether to adopt or delay is affected by information generated as a result of social learning. If $p_0 \leq p^*$, then consumers' behavior depends only on information obtained exogenously (and no adoption ever takes place if $\varepsilon = 0$). If $\varepsilon = 0$ and $p_0 > p^* = \frac{1}{2}$, then consumers are willing to adopt at all histories, since no matter how large Λ_0 , the equilibrium posterior always remains weakly above $\frac{1}{2}$.

On the other hand, if $p_0 > p^*$ and $\varepsilon > 0$, then under both Λ_0 and $\hat{\Lambda}_0$ consumers adopt immediately upon first opportunity until p^* is reached and from then on delay adoption until there has been a breakthrough. Moreover, the probability π^* of a breakthrough occurring prior to p^* being reached is the same under both Λ_0 and $\hat{\Lambda}_0$: $\pi^* = \frac{1-p_0}{1-p^*}$. And because learning occurs at the same rate once p^* is reached, the continuation value W^* conditional on p^* being reached is also the same: $W^* = p^* \int_0^\infty \varepsilon e^{-(\varepsilon+r)t} \frac{\rho}{r+\rho} dt = 2p^* - 1$. So the only difference is that conditional on no breakthroughs, the time t^* at which p^* is reached occurs earlier under $\hat{\Lambda}_0$. To see that this is strictly beneficial, note that W_0 is composed of the following two terms:

$$W_0(\Lambda_0) = \left(1 - e^{-(r+\rho)t^*(\Lambda_0)}\right) \frac{\rho}{r+\rho} (2p_0 - 1) + e^{-(r+\rho)t^*(\Lambda_0)} \left(\pi^* \frac{\rho}{r+\rho} + (1 - \pi^*) W^*\right),$$

and similarly for $\hat{\Lambda}_0$. The first term represents the case when a consumer receives an adoption opportunity prior to time t^* , and the second represents the case when a consumer's first adoption opportunity occurs after t^* . Conditional on either of these cases occurring, the expected payoff is the same under both Λ_0 and $\hat{\Lambda}_0$, but the time-discounted probability $e^{-(r+\rho)t^*}$ with which the second case occurs is strictly greater under $\hat{\Lambda}_0$. This is strictly beneficial, because the expected payoff in the second case is strictly greater:

$$\left(\pi^* \frac{\rho}{r+\rho} + (1-\pi^*)(2p^*-1) \right) - \frac{\rho}{r+\rho} (2p_0-1) = \frac{r}{r+\rho} (1-\pi^*)(2p^*-1) > 0.$$

Intuitively, in the second case the consumer adopts the innovation only once it has been revealed to be good while in the first case he adopts it regardless of its quality, and the resulting benefit from avoiding a bad innovation outweighs the cost of possibly having to delay adoption of a good innovation.

Nature of Inefficiency: Even though there is no saturation effect and consumers are able to always benefit from increased opportunities for social learning, equilibrium adoption behavior is not in general socially optimal. Let p^s denote the cutoff posterior for the cooperative benchmark derived in Proposition 1.3.2.

Proposition 1.6.9. *If $\varepsilon = 0$, equilibrium adoption behavior is socially optimal if and only if either $p_0(1 + e^{-\Lambda_0}) \geq 1$ or Inequality 1.5 holds. If $\varepsilon > 0$, then equilibrium adoption behavior is socially optimal if and only if $p^s \geq p_0$.*

Consider first the case where $\varepsilon = 0$. Then if $p_0(1 + e^{-\Lambda_0}) \geq 1$, we have that $t^* = t^s = +\infty$; and if Inequality 1.5 holds, then $t^* = t^s = 0$. For the converse and to deal with the case when $\varepsilon > 0$, it then suffices to show that $p^s < p^*$: This implies that whenever $p_0 > p^s$, then conditional on no breakthroughs *adoption ends too soon* in equilibrium (or doesn't take place at all if $p_0 \leq p^*$ even though the cooperative benchmark prescribes some initial adoption). On the other hand, if $p_0 \leq p^s$, then both the cooperative benchmark and the equilibrium prescribe no adoption until there has been an exogenously generated breakdown. Note that adoption ending too soon under the perfect good news equilibrium is the analog of adoption beginning inefficiently late under perfect bad news. However, since the perfect

good news equilibrium does not feature regions of partial adoption, there is no analog of the second type of inefficiency that arose under perfect bad news: Whenever adoption does occur under perfect good news, it takes place at an optimal rate.

To see that $p^s < p^*$, note that

$$\left(2 - \frac{\varepsilon}{\varepsilon + r} \frac{\rho}{\rho + r}\right) p^* - 1 = 0.$$

Using the above equality and evaluating the derivative of the objective function of the cooperative problem in Equation 1.3 at p^* , we obtain:

$$(1 - \pi^*) p^* \Lambda_0 \rho e^{-\rho t^*} \frac{r}{\varepsilon + r} \left(e^{-(r+\rho)t^*} \frac{\rho}{\rho + r} \bar{N}_0 \right) > 0.$$

This shows that $t^s > t^*$ and so $p^s < p^*$ as the objective function of the cooperative problem is single-peaked.

1.7 More Social Learning Can Hurt: An Example

In Proposition 1.5.9 we established the saturation effect, whereby increased opportunities for social learning under perfect bad news are welfare-neutral when Λ_0 is sufficiently large relative to the other fundamentals. Nevertheless, under the assumption of completely homogeneous consumers in the previous sections, increases in Λ_0 never produced ex ante welfare losses. In this section, we establish the surprising result that when consumers are heterogeneous, increased opportunities for social learning can *strictly hurt* some consumers and bring about Pareto-decreases in ex ante welfare. To illustrate this, we introduce some heterogeneity in consumers' patience levels.

Consider a population consisting of two types of consumers: There is a mass M_0^p of patient types with discount rate $r_p > 0$ and a mass M_0^i of impatient types with discount rate $r_i > r_p$. To simplify the analysis we assume that $\varepsilon = 0$ and $p_0 > 1/2$, although our arguments easily extend to the case where $\varepsilon > 0$. Because our purpose is simply to construct an example illustrating the possibility of welfare loss, we restrict attention to a perfect bad news setting.

To construct our example, we begin by examining equilibria in which only types with discount rate r_p exist in the economy. Recall from Section 1.5.1 that for any discount rate $r > 0$, we can define the function Λ_r^* implicitly for every $p \in (\frac{1}{2}, \frac{\rho+r}{\rho+2r})$ by

$$2p - 1 = G_r(p, \Lambda_r^*(p)) := \int_0^{\infty} \rho e^{-(r+\rho)\tau} \left(p - (1-p)e^{-\Lambda_r^*(p)(1-e^{-\rho t})} \right) d\tau.$$

Then by Theorem 1.5.1, whenever $p_0 < \frac{\rho+r_p}{\rho+2r_p}$ and $\hat{\lambda}M_0^p > \lambda M_0^p > \Lambda_{r_p}^*(p_0)$, then in the game consisting solely of consumers of type r_p , the two equilibria corresponding to information structures λ and $\hat{\lambda}$ both feature initial regions of partial adoption.

The main argument in the construction of our example is to consider heterogeneous economies where the mass M_0^i of impatient types is very small, holding fixed the mass of patient types at M_0^p . More specifically, we show that when the mass of impatient types is sufficiently small, the equilibrium behavior of the patient types in both equilibria (under information process $\hat{\lambda}$ and λ) approximates the behavior in the corresponding equilibria when only patient types are present. Then using arguments about the properties of equilibria in the game with only patient types, in particular the non-monotonicity result for learning established in Proposition 1.5.11, we can obtain the following result:

Theorem 1.7.1. *Suppose $0 < r_p < r_i < +\infty$. Fix $M_0^p > 0$ and $\hat{\lambda} > \lambda > 0$ such that $\hat{\lambda}M_0^p > \lambda M_0^p > \Lambda_{r_p}^*(p_0)$. Then there exists $\eta > 0$ such that whenever $M_0^i < \eta$, $W_0^i(\hat{\lambda}) < W_0^i(\lambda)$ and $W_0^p(\hat{\lambda}) = W_0^p(\lambda)$. Thus, whenever $M_0^i < \eta$, the ex ante payoff profile in the λ -equilibrium Pareto-dominates the ex ante payoff profile in the $\hat{\lambda}$ -equilibrium and*

$$M_0^i W_0^i(\hat{\lambda}) + M_0^p W_0^p(\hat{\lambda}) < M_0^i W_0^i(\lambda) + M_0^p W_0^p(\lambda).$$

Here we sketch the main arguments of the theorem. Consider first an economy consisting only of types with discount rate r_p : $M_0^i = 0$ and $M_0^p > 0$. If $\hat{\lambda}M_0^p > \lambda M_0^p > \Lambda_{r_p}^*(p_0)$, then the two equilibria corresponding to information structures λ and $\hat{\lambda}$ both feature initial regions of partial adoption. Thus $W_0^p(\hat{\lambda}) = W_0^p(\lambda) = 2p_0 - 1$.

Now consider the payoffs that a hypothetical type r_i (even though such a type does not

exist in this economy) would obtain if he were to behave optimally when faced with the flow of information generated in each of these equilibria. Because an optimal strategy (there will be a continuum of optimal strategies) of a consumer of type r_p is to adopt upon first opportunity absent breakdowns, it is straightforward to show that an optimal strategy of such a hypothetical type r_i would also be to adopt upon first opportunity.

Given this, the payoffs of the hypothetical type r_i in the two equilibria are given by the following two expressions:

$$W_0^i(\hat{\lambda}) = \int_0^{\infty} \rho e^{-(r_i+\rho)\tau} \frac{p_0}{p_\tau^{\hat{\lambda}}} (2p_\tau^{\hat{\lambda}} - 1) d\tau$$

$$W_0^i(\lambda) = \int_0^{\infty} \rho e^{-(r_i+\rho)\tau} \frac{p_0}{p_\tau^\lambda} (2p_\tau^\lambda - 1) d\tau.$$

Furthermore, patient types begin in a partial adoption phase in both equilibria:

$$2p_0 - 1 = W_0^p(\hat{\lambda}) = \int_0^{\infty} \rho e^{-(r_p+\rho)\tau} \frac{p_0}{p_\tau^{\hat{\lambda}}} (2p_\tau^{\hat{\lambda}} - 1) d\tau$$

$$2p_0 - 1 = W_0^p(\lambda) = \int_0^{\infty} \rho e^{-(r_p+\rho)\tau} \frac{p_0}{p_\tau^\lambda} (2p_\tau^\lambda - 1) d\tau.$$

Recall from Proposition 1.5.11 that there exists $\bar{t} > t^* := t_2^*(\lambda)$ such that $p_\tau^{\hat{\lambda}} = p_\tau^\lambda$ for all $\tau \leq t^*$, $p_\tau^{\hat{\lambda}} < p_\tau^\lambda$ for all $\tau \in (t^*, \bar{t})$ and $p_\tau^{\hat{\lambda}} > p_\tau^\lambda$ for all $\tau > \bar{t}$. We now exploit the expressions for the value to waiting of the two types together with the deceleration of learning at times just after t^* to obtain the result. Intuitively, because $W_0^p(\hat{\lambda}) = W_0^p(\lambda) = 2p_0 - 1$, the deceleration in learning followed by a later acceleration must balance out exactly so that the patient type r_p obtains the same ex ante payoff under λ and $\hat{\lambda}$. But then these adjustments must strictly hurt the less patient hypothetical type r_i , because relative to type r_p , type r_i weights the losses due to the slow down of learning more heavily than the benefits that arrive at later times.⁴⁸

To complete the proof, we can show that even when $M_0^i > 0$, as long as M_0^i is sufficiently

⁴⁸A formal argument is provided in Appendix Section A.8.

small, we must still have $W_0^i(\hat{\lambda}) < W_0^i(\lambda)$ and $W_0^p(\hat{\lambda}) = W_0^p(\lambda)$. The first inequality is the result of a simple continuity argument. The second equality comes from the fact that even upon perturbing M_0^i slightly, the patient type must continue to partially adopt initially in both equilibria.

Note that a crucial assumption underlying the above argument is that adoption opportunities are *stochastic and limited*. When ρ is finite, because of a natural delay in adoption, the impatient types may not receive any adoption opportunities for a long time. As a result, if an impatient type obtains his first adoption opportunity late in the game, then the information available at that point in time would be strictly lower under the equilibrium with information process $\hat{\lambda}$ than λ . This decrease in information (due to increased free-riding of the patient types) when impatient types receive adoption opportunities late in the game is precisely the cause of the impatient type's welfare loss. If on the other hand consumers were able to adopt freely at any time, then the impatient types would incur no losses as they would adopt at exactly time 0 in both the λ and $\hat{\lambda}$ -equilibria. Thus the example here illustrates an interesting interaction between heterogeneity and delays due to limited opportunities for adoption.

1.8 Conclusion

This paper develops a model of innovation adoption when consumers are forward-looking and learning is social. Our analysis isolates the effect of purely informational incentives on aggregate adoption dynamics, learning, and welfare, and highlights the way in which these incentives vary across different informational environments. The possibility of free-riding in the form of partial adoption is found to be particularly important, because it casts doubt on the received wisdom that the recent internet-driven surge in opportunities for social learning should speed up learning and benefit consumers. Owing to the advantages of continuous time and Poisson learning, the model is very tractable, yielding closed-form expressions for key quantities and allowing us to compute numerous comparative statics.

We briefly discuss some questions for ongoing and future research that could build on

the modeling framework and techniques developed in this paper. Current work in progress relaxes the assumption of perfect Poisson learning to allow for signals that while indicative of bad (respectively good) quality are *not* conclusive. Serving as a robustness check for our results obtained under perfect Poisson learning, preliminary results suggest that many key qualitative features are preserved, for example the possibility of partial adoption regions in bad news markets (which once again coincide with convex growth in adoption levels) as well as the absence of such regions under good news learning. In addition, the extension to imperfect Poisson learning introduces interesting new questions that cannot be studied when signals are conclusive. For example, when $\varepsilon = 0$, then under imperfect (but not under perfect) bad news it is possible for good innovations to fail, because even good products can generate strings of breakdowns that might permanently halt adoption. This suggests investigating the “fragility” of the adoption process as a function of parameters such as the initial market belief and the relative rates at which bad and good products generate breakdowns.

Further work in progress relaxes the assumption that signals are public. To see the idea, suppose that learning is social, but that signals derived from past adopters’ experiences are observed privately and independently (at rate λN_t) by each remaining consumer, instead of publicly and simultaneously as in the model in this paper. This captures the intuition of *decentralized* social learning, for example when consumers frequent different blogs and social media platforms. Assuming that at any time consumers make inferences based only on their own private signals and on the expected number of adopters in the population, another interesting difference between bad news and good news markets emerges: Under bad news, a consumer who privately observes a breakdown will never adopt the innovation in the future and hence will never generate any signals himself; this has a *dampening* effect on the production of information in the economy and reduces free-riding incentives. By contrast, consumers who privately observe breakthroughs under good news will adopt the innovation at their next opportunity, thus *amplifying* information production in the future and increasing free-riding incentives. This difference has important implications

for aggregate adoption dynamics and for the impact of increased opportunities for social learning.

Finally, moving beyond our focus in this paper on the purely informational aspects of the problem, one could explore the implications of incorporating consumer heterogeneity and pricing motives into the model. As we saw in Section 1.7, heterogeneity can interact in interesting ways with informational free-riding incentives, sometimes rendering increases in the potential for social learning strictly harmful. A more general characterization of this interaction under more complex distributions of consumer heterogeneity appears challenging but desirable. As for pricing, assume that the innovation is sold by a forward-looking monopolist who does not have any influence on the quality of the innovation and has access to exactly the same public information as consumers, but can influence the endogenous production of information via the price. As a simple first step, we could restrict the monopolist to setting a single fixed price and compute comparative statics on this price and on welfare under increased opportunities for social learning. More challengingly, we could allow the monopolist to commit to a time path of prices, examining for instance how the fact that information is generated endogenously by consumer purchases affects the monopolist's incentives for intertemporal price-discrimination relative to the well-known complete information results of Stokey (1979). We leave these two topics as interesting avenues for future research.

Chapter 2

Delayed Response Strategies in Repeated Games with Observation Lags¹

2.1 Introduction

Understanding when and why individuals cooperate in social dilemmas is a key issue not just for economics but for all of the social sciences,² and the theory of repeated games is the workhorse model of how and when concern for the future can lead to cooperation even if all agents care only about their own payoffs. The clearest expression of this idea comes as players become arbitrarily patient; here various folk theorems provide conditions under which approximately efficient payoffs can be supported by equilibrium strategies. Because of the influence of these results, it is important to understand which of their assumptions are critical and which are merely convenient simplifications; a large literature (discussed below) has extended the folk theorems under successively weaker assumptions about the

¹Co-authored with Drew Fudenberg and Scott D. Kominers and published in the *Journal of Economic Theory* (2014).

²See e.g., Ahn *et al.* (2003); Gächter *et al.* (2004).

“monitoring structures” that govern the signals players receive about one another’s actions.

Here we relax an assumption which is maintained throughout most of the prior repeated games literature: the assumption that signals of the actions taken in each period (simultaneously) arrive immediately after players’ actions in that period. Instead, we consider repeated games in which the players’ signals about other player’s actions arrive with stochastic and privately observed lags. Our folk theorems for settings with lagged signals show that the assumption that signals are observed immediately is not necessary for repeated play to support cooperation.

To prove these folk theorems, we use the idea of “delayed-response” strategies, under which players wait to respond to signals of a given period’s play for long enough that it is likely (although not certain) that every player has observed the relevant signals by the time players respond to signal information. Although the observation lags generate a form of imperfect private monitoring, the private information here has a special form that allows delayed-response strategies to construct the same set of limit equilibrium payoffs as if the lags were not present.

More specifically, we suppose that players act simultaneously each period, and that players’ actions jointly determine a probability distribution over signals, but that players

- do not observe signals immediately and
- might observe signals asynchronously.

The times at which observation occurs are private information and may be infinite, that is, a particular signal may never arrive. Some sort of observation lags seem plausible in many cases; for example there may be a small probability that a player is momentarily inattentive and temporarily does not see their partner’s actions; more strongly, in some cases a player may never learn just what happened during moments of inattention. Moreover, information lags of multiple periods seem especially appropriate in settings for which the time period under consideration is extremely short (Fudenberg and Levine (2007a, 2009); Sannikov and Skrzypacz (2010)), and in continuous-time models, where the “period length” is effectively

0 (Bergin and MacLeod (1993); Sannikov (2007); Sannikov and Skrzypacz (2007); Faingold and Sannikov (2011)).³

To prove our folk theorems, we construct *delayed-response strategies*, in which the repeated game is divided into a finite number of “threads,” with play in each thread independent of play in the other threads. Section 2.3 examines the simplest application of this idea, which is to the case of bounded lags, where there is a K such that every signal arrives within K periods of play. Then, using strategies that have $K + 1$ threads, we can ensure that each thread is equivalent to an instance of the original game (with the original game’s underlying monitoring structure), a smaller discount factor, and *no lag*. Hence if the folk theorem holds in a given repeated game (with any sort of contemporaneous monitoring), the associated strategies can be used to establish a folk theorem—in delayed-response strategies—in the corresponding game with bounded observation lags.⁴

The rest of the paper allows the lag distribution to have unbounded support, and also allows for a small probability that some signals never arrive at all (corresponding to an infinite observation lag). In these cases the use of delay strategies reduces but does not eliminate the impact of lags, and the game played in each thread has some additional decision-relevant private information. Section 2.4 considers the case where signals are almost-perfect if and when they do arrive—that is, each player either observes an almost-perfect signal of period- t play with some lag, or else never sees a signal of period- t play.⁵ In our second model, presented in Section 2.5, players are allowed to communicate (via cheap talk) each period, and the underlying information structure is one of imperfect public monitoring.⁶ In each case, players do not know whether and when other players observe

³Indeed, physics suggests that the speed of light is a constraint on the speed with which signals can travel.

⁴The Ellison (1994) study of contagion equilibria uses threads for a rather different purpose: to substitute for public randomization as a way to weaken the effect of a grim-trigger punishment as the discount factor tends to 1. In Sections 2.3 and 2.4, we use threads only as a way for the players to wait for lagged signals to arrive; in Section 2.5, we also use threading in order to weaken the effect of grim-trigger punishments.

⁵In the case of lagged almost-perfect monitoring, we consider only games with two players, so that we may invoke results of Hörner and Olszewski (2006). We do not know whether the folk theorem extends to the analogous setting with n players.

⁶Our analysis assumes that messages are received the instant they are sent, but the results extend to cases in

the signals associated with each period's play, so there is a special but natural form of private information.

For both of our main results, we use a similar proof technique: First, we consider an auxiliary game with "rare" lags in which each player sees a private signal immediately with probability close to (but not equal to) 1. After proving a folk theorem for the auxiliary game with rare lags, we relate the perturbed game with rare lags to the game with possibly long lags by identifying the event in which the signal does not arrive immediately with the event that the signal arrives after some large time T . We then construct equilibria in the game with lags by using delayed-response strategies as described above. For the first main result we prove the folk theorem for the auxiliary game by extending the bloc-strategy construction of Hörner and Olszewski (2006) (henceforth HO2006) to treat as "erroneous" any history of the auxiliary game in which some player observes another's action with a strictly positive lag; this corresponds to a "real lag" that is longer than the number of threads. The HO2006 construction does not directly apply here, as signals about past play may arrive outside of the relevant block, but we construct equilibria that are belief free for the past T periods provided that the probability of lagged observation is sufficiently small.

To prove the second main result, we first consider a game with private monitoring, communication, and no observation lag. In this game, each player either observes the true action profile or a null signal. We relate this to a game with a public signal that is observed by all players, but where the game ends each period with a fixed small probability, corresponding to strategies in the original game that will use reversion to static Nash equilibrium whenever the reported signals disagree. We prove a sort of folk theorem here using the techniques of Fudenberg *et al.* (1994) (henceforth FLM) and then again use threads and delayed responses to extend this to a proof for the original game.

which messages are received much more quickly than observations. We attribute the difference in speeds to the fact that messages are crafted to be easily processed, while processing and interpreting signals can take longer.

2.1.1 Related Work

The repeated games literature has explored successively weaker assumptions on players' monitoring structures, while maintaining the assumption that signals arrive immediately after play. The first wave of repeated-games models established folk theorems under the assumption that players observe each others' actions without error at the end of each round of play (Aumann and Shapley (1976), Friedman (1971), Rubinstein (1994), and Fudenberg and Maskin (1986)). Subsequent work extended the folk theorem to cases where agents receive imperfect signals of other agents' actions, where these signals can either be public (FLM) or private but accompanied by cheap-talk public messages (Compte (1998), Kandori and Matsushima (1998), and Obara (2009)),⁷ or private and without communication (e.g., Sekiguchi (1997), Mailath and Morris (2002), Hörner and Olszewski (2006), and Hörner and Olszewski (2009)). As one step in our argument for the case of lagged almost-perfect monitoring (Section 4), we extend the Hörner and Olszewski (2006) construction to almost-perfect monitoring with rare lags.⁸ With each type of signal structure, the key assumptions relate to the qualitative nature of the information that signals provide: Roughly speaking, in order for the folk theorem to obtain, signals must be informative enough to "identify deviations" in a statistical sense.⁹

The papers of Fudenberg and Olszewski (2011) and Bhaskar and Obara (2011) are the closest to the present work, as in each, the time at which signals arrive is private information. Fudenberg and Olszewski (2011) studied the effect of short privately-known lags in observing the position of a state variable that evolves in continuous time, so that a player observing the state variable at slightly different times would get different readings.

⁷We allow public messages in Section 2.5. The role of such messages has been studied in a number of subsequent papers, including Ben-Porath and Kahneman (2003), Fudenberg and Levine (2007b), and Escobar and Toikka (2011). Public communication has also been used as a stepping stone to results for games where communication is not allowed (Hörner and Olszewski (2006), Hörner and Olszewski (2009), and Sugaya (2011)).

⁸When the unlagged signals are imperfect, the signals in our auxiliary games are not almost common knowledge in the sense of Mailath and Morris (2002), so the Hörner and Olszewski (2009) construction does not apply.

⁹In addition, the folk theorem has been extended to recurrent stochastic games with perfectly or imperfectly observed actions (Dutta (1995), Fudenberg and Yamamoto (2011), and Hörner *et al.* (2011)).

Bhaskar and Obara (2011) studied lags that were either deterministic or stochastic with length at most 1. Both papers considered “short lags” and also restricted to the case of a single long-run player facing a sequence of short-run opponents; this paper allows fairly general stochastic lags and considers the case of all long-run players.

Several papers in the stochastic games literature studied deterministic lags of perfect signals (e.g., Lagziel and Lehrer (2012), Levy (2009), and Yao *et al.* (2011)); this sort of lag does not introduce private information and so is quite different from the lags we study. In Abreu *et al.* (1991) consecutive signals are grouped together and delivered at once, so the delay does not introduce private information.

2.2 General Model

This section introduces a general model that encompasses all the settings discussed subsequently. We consider a repeated game with n players $i \in I \equiv \{1, \dots, n\}$, each of whom has a finite action space A_i . In each period $t = 0, 1, 2, \dots$, each player i chooses a possibly mixed action a_i^t ; this generates a sequence of pure action profiles $\{a^t\}_{t=0}^\infty$. Each player i has a finite signal space Ω_i , and there is a private signal structure π over $\Omega \equiv \prod_{i \in I} \Omega_i$; at each time t , a private signal profile is generated by π according to the conditional probability $\pi(\omega_1, \dots, \omega_n \mid a^t)$.

Thus far, the repeated game has the structure of a standard repeated game with private monitoring. We now relax the assumption that players receive signals of period- t play immediately after period t by replacing it with the assumption that the monitoring structure is *private with stochastic lags*. As in the usual model, upon the choice of a period- t action profile a^t , a private signal profile ω^t is generated according to the conditional distribution $\pi(\omega^t \mid a^t)$. However, the players need not immediately observe their components of the signal profile. Instead, we assume that each player i observes his private signal of period- t play, ω_i^t , at a stochastic time $t + L_i^t$, where $\{L^t \equiv (L_1^t, \dots, L_n^t)\}_t$ is a collection of random variables that take values in $(\mathbb{N} \cup \{\infty\})^n$. We assume that the vectors $\{L^t\}$ are distributed identically and independently across t , with probability density function $\lambda :$

$(\mathbb{N} \cup \{\infty\})^n \rightarrow [0, 1]$. We denote by λ_i the density of the marginal distribution of observation lags of player i , L_i . (The case $L_i^t = \infty$ is interpreted as the event in which player i never receives any information about the period- t private signal.) We let Λ_i denote the cumulative marginal distribution function of player i 's observation lags, i.e. $\Lambda_i(\ell) = \sum_{m=0}^{\ell} \lambda_i(m)$, for $\ell \in (\mathbb{N} \cup \{\infty\})$.

Observation of ω_i^t takes place in period $t + L_i^t$ after the choice of that period's actions.¹⁰ When player i observes ω_i^t , he also observes a "timestamp" indicating that ω_i^t is associated with play in period t . That is, for example, when a player observes that player j played "C" in a prisoner's dilemma, she is informed about the period to which the observation applies, rather than just getting a signal that "player j played C sometime in the past."¹¹

As one concrete example, consider a repeated public goods game in which every period two friends must decide whether or not to exert effort to provide benefits (or gifts) for each other. The friends live far apart, so the benefits must be shared via postal mail. This induces a lag in observation of the realized signals of the friend's action. Furthermore the postmark dates serve as natural timestamps.

Alternatively suppose that n coauthors who write numerous papers together and suppose that the quality of the paper is determined by the sum of the authors' efforts. Each period they complete a paper and submit it to a journal. The editor then makes a decision and mails a letter to each of the authors. Here the decision reveals the project's quality and so provides evidence about partners' efforts; in a two-player game if the effort \rightarrow quality \rightarrow editor's letter map is deterministic and monotone, the letter perfectly reveals partners' efforts, but more typically letters have a stochastic component. Here the project itself serves as a natural timestamp.

Players have perfect recall and receive no further information.

¹⁰Thus player i cannot respond to the period- $(t + L_i^t)$ observation information until time $t + L_i^t + 1$.

¹¹The assumption of timestamps renders our model a smaller departure from the usual repeated game monitoring structure than a model in which players observe only an aggregate measure of the frequencies with which opponents took various actions. Note that it is not clear how players would interpret signals received without timestamps when the expected path of play is not constant over time.

In one part of the paper we allow for communication in every period. Thus, we include message spaces M_i in the general model; when we want to rule out communication we set $M_i = \emptyset$ for each i . After the realization of private signal profile ω^t and after the observation of all private information $\omega_i^{t'}$ for which $t' + L_i^{t'} \leq t$, at each time $t = 0, 1, \dots$, each player i reports a message m_i chosen from the message space M_i . After all of these reports are (simultaneously) submitted, all players immediately observe the message profile $m = (m_1, \dots, m_n)$.

We let H^t denote the set of t -period histories. For a given $h^t \in H^t$ and any $t' \leq t$, we denote by $h^{t,t'}$ the profile of information about the t' -period signal that has been observed by each player. If player i has not yet observed the k -th component of his private signal, $\omega_{i,k}$ in time t' then we specify that $h_{i,k}^{t,t'} = \infty$.

Finally we describe the payoff structure. A sequence of action profiles $\{a^t\}$ chosen by the players generates a total payoff

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t g_i(a^t).$$

In Section 2.5, we prove a Nash threat folk theorem rather than a full folk theorem. To facilitate this, we fix a Nash equilibrium α^* of the stage game and normalize payoffs of players so that $g_i(\alpha^*) = 0$ for all i . We let V denote the convex hull of the feasible set of payoffs, and let V_{α^*} be the convex hull of the set consisting of $g(\alpha^*) = 0$ and the payoff vectors Pareto-dominating $g(\alpha^*) = 0$: $V_{\alpha^*} \equiv \{v \in V \mid v \geq 0\}$. We assume that $\text{int}(V_{\alpha^*})$ is non-empty.

In contrast, the theorems of Section 2.3 and 2.4 concern full folk theorems; thus, we define V^* to be the set of individually rational payoffs of V . With this notation, we are ready to discuss our folk theorems.

We let $G(\delta, \pi, \lambda)$ be the repeated game with discount factor δ , lag distribution λ , and monitoring structure π , and let $E(\delta, \pi, \lambda)$ denote the set of sequential equilibrium payoffs of $G(\delta, \pi, \lambda)$. We let $G(\delta, \pi) \equiv G(\delta, \pi, \text{imm})$, where imm is the (degenerate) distribution which puts full weight on immediate observation, and define $E(\delta, \pi) \equiv E(\delta, \pi, \text{imm})$ similarly.

Finally we introduce the concept of delayed-response strategies, which are used throughout the remainder of the paper to prove our folk theorems. We call σ a *delayed-response strategy profile* in the repeated game if there exists some K such that the repeated game can be divided into K “threads”, with the ℓ -th thread consisting of periods $\ell, K + \ell, 2K + \ell, \dots$, so that at any period t players condition their strategies only on messages and signal information generated within the thread containing period t .

2.3 Bounded Lags

We first present a simple analysis of a repeated game with observation lags in which the lag is certain to be no more than some finite bound.

Assumption 2.3.1. There exists some $K < \infty$ such that $\Pr(\max_i L_i \leq K) = 1$.

With this assumption, it is common knowledge that all players will have seen the signal generated in period t by period $t + K$. This restriction allows us to show that every equilibrium payoff attainable for sufficiently large discount factors in the repeated game *without observation lags* with any private monitoring structure π can also be attained in the associated repeated game *with observation lags* for sufficiently patient players. We show this using delayed-response strategies. Note that the following result does not impose any restrictions on π ; we use such conditions for our folk theorems later but they are not needed here.

Theorem 2.3.2. *Suppose Assumption 2.3.1 holds. Furthermore suppose that $v \in E(\delta, \pi)$ for all $\delta \in (\underline{\delta}, 1)$ where $0 < \underline{\delta} < 1$. Then there exists some $\delta^* \in (0, 1)$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta \in (\delta^*, 1)$.*

Proof. We divide the periods of the repeated game into $K + 1$ threads, with the ℓ -th thread consisting of periods $\ell, (K + 1) + \ell, 2(K + 1) + \ell, \dots$. Now, we suppose that $v \in E(\delta, \pi)$ is generated by the strategy profile σ in the game without lags.

As the information lag has an upper bound of K , the signals generated in periods $\ell, (K + 1) + \ell, \dots, (j - 1)(K + 1) + \ell$ are observed by all players by period $j(K + 1) + \ell$.

Thus, we may define a delayed-response strategy profile σ^K by specifying that in period $t = j(K + 1) + \ell$ ($0 \leq \ell \leq K$), players play according to $\sigma(h^{t,\ell}, h^{t,(K+1)+\ell}, \dots, h^{t,(j-1)(K+1)+\ell})$.

It is clear that the delayed-response strategy profile σ^K generates a payoff profile of v . Moreover, it is an equilibrium for discount factor δ^{K+1} . Thus, taking $\delta^* = \delta^{K+1}$ gives the result. \square

The proof of Theorem 2.3.2 relies heavily on Assumption 2.3.1. For example, if the support of λ were concentrated on $(0, \dots, 0), (1, \dots, 1), \dots, (K, \dots, K)$, and (∞, \dots, ∞) , then the proof above would not work, since each of the threads that it constructs would be a repeated game with a private monitoring structure $\tilde{\pi}$ that is different from π . More problematically, if $\lambda_i(k) > 0$ for all i and $k \in \mathbb{N}$, so that all lag lengths have positive probability for all players, then no matter how far apart the threads are spaced, there is always a positive probability that a realized lag will be longer than this chosen spacing, and the threads considered in the proof above cannot be identified with a private monitoring game at all. In the next two sections, we study and demonstrate how these issues can be resolved when additional assumptions are placed on the monitoring structure π . Therefore for the remainder of the paper, we dispense with Assumption 2.3.1 and allow λ to be an arbitrary probability distribution on $(\mathbb{N} \cup \{\infty\})^n$.

2.4 Lagged Almost-Perfect Monitoring with Two Players

In this section, we extend an approach of HO2006 to obtain a folk theorem for two-player games with lagged almost-perfect monitoring. We focus on the two-player case since the techniques of HO2006 extend naturally to this setting.¹²

¹²We do not know whether our folk theorem extends to games with n players; we discuss related issues in Section 2.6.

2.4.1 Model

We restrict the general monitoring structure introduced above. First, we assume that there are only two players. We assume the monitoring structure to be that of *lagged ε -perfect monitoring*: We allow a general lag structure here, but restrict the private signal space of each player i to be $\Omega_i = A_j$ and furthermore assume that π is ε -perfect in the sense of HO2006. We assume that the private signal space of $\Omega_i = A_j$ so that we may extend the techniques of HO2006.¹³

Definition 2.4.1. A private monitoring structure π is ε -perfect if for every action profile $a \in A$, $\pi(a_2, a_1 \mid a_1, a_2) > 1 - \varepsilon$.

The Folk Theorem

We now prove the following folk theorem.

Theorem 2.4.2. Suppose that $v \in \text{int}(V^*)$. Then there exists some $\bar{\varepsilon} \in (0, 1)$ such that for all lag distributions λ for which $\lambda_i(\infty) < \bar{\varepsilon}$ (for $i = 1, 2$), there exists some $\bar{\delta}$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \bar{\delta}$, and all private monitoring structures π that are $\bar{\varepsilon}$ -perfect.¹⁴

To prove Theorem 2.4.2, we first analyze an auxiliary repeated game with rare observation lags, in which the probability of instantaneous observation of the private signal is very close to 1. We show that the HO2006 approach to repeated games with almost-perfect monitoring can be extended to lagged repeated games with almost-perfect monitoring, so long as positive lags are sufficiently rare, and use this to obtain a folk theorem in the auxiliary game. We then convert the associated auxiliary-game strategies to delayed-response strategies by multithreading the game with lags. A positive lag in a particular thread corresponds to a lag that exceeds the number of threads, so by taking the delay long enough we can shrink

¹³Note that the work of HO2006 does contain a section that extends the analysis to more general private signal spaces where $\Omega_i \neq A_j$. However, as Yuichi Yamamoto pointed out to us, that argument contains an error so we cannot use it here.

¹⁴We thank Yuichi Yamamoto for pointing out a problem with our earlier proof of this result and then suggesting the approach we use now.

the probability of a positive lag close to 0. We thus obtain a folk theorem in the game with stochastic lags.

2.4.2 Auxiliary Repeated Game with Rare Observation Lags

This subsection establishes the following folk theorem for the game where $\lambda_i(0)$ is close to 1 for all players i :

Theorem 2.4.3. *Let $v \in \text{int}(V^*)$. Then there exist $\bar{\varepsilon}, \underline{\delta} \in (0, 1)$ such that if $\lambda_i(0) > 1 - \bar{\varepsilon}$ (for $i = 1, 2$), π is $\bar{\varepsilon}$ -perfect, and $\delta > \underline{\delta}$, then $v \in E(\delta, \pi, \lambda)$.*

Our proof of this theorem adapts a technique of HO2006 to the environment with small observation lags. The HO2006 construction for the case of almost-perfect monitoring uses the same strategies as in the perfect-monitoring construction at histories that are on the equilibrium path of that equilibrium (the “regular histories”), and then uses standard full-rank arguments to show there are continuation payoffs (at the end of the review phase) that preserve the belief-freeness property at the “erroneous” histories—those which are off the path of play under perfect monitoring. When the monitoring is close to perfect, the additional variation introduced in these continuation payoffs converges to zero. We use a similar argument, grouping histories together by treating a delayed observation as one that never arrives, and classifying as “erroneous” any history in which *some* player observes the opponent’s action with a strictly positive lag. We then construct continuation payoffs associated to these histories by applying full-rank arguments to the “immediate observation structure” defined below.

Note first that because the information lag is not bounded, it is possible that information about some past event arrives very late in the repeated game. Such possibilities cannot be ignored—even though they happen with very low probability—since they may potentially affect a player’s beliefs about his opponent’s continuation play. Our extension deals with this problem by constructing equilibria that are belief-free every T periods for the repeated

game with the probability of lagged observation sufficiently small.¹⁵ This means that only information about the past T periods is relevant for computing best replies. Thus, we can ensure that effects on beliefs due to observation lags lasting more than T periods are unimportant.

Note next that lags of length less than T do affect players' on-path beliefs, so the HO2006 arguments do not directly apply. We extend them to lags with $\lambda(0)$ close to 1 by adding the histories where observations arrive with a positive lag to the set of "erroneous" histories.

Preliminaries

We let H_i^t be the set of t -period histories in the repeated game *with observation lags*, with elements denoted in the form

$$h_i^t = (a_i^0, a_i^1, \dots, a_i^{t-1}, h_i^{1,o}, h_i^{2,o}, \dots, h_i^{t,o}).$$

Here, $h_i^{t,o}$ denotes all of the *new information* about the past play of player $-i$ that player i receives in period t . Furthermore denote by S_i^T the set of strategies in the T -times repeated game with information lags. Let \tilde{H}_i^t be the set of t -period histories in the repeated game *without observation lags and with perfect monitoring* with a typical element of \tilde{H}_i^t denoted by \tilde{h}_i^t . Also denote the set of strategies in the T -times repeated game with perfect monitoring and no observation lags by \tilde{S}_i^T .

Now we partition the set of private histories in the T -times-repeated stage game into H_i^R and H_i^E , the *regular* and *erroneous* histories. To do this we first define restricted strategy sets \tilde{S}_i and \tilde{S}_i^o for $i = 1, 2$ in the T -times repeated game with *perfect monitoring*. Partition the set A_i into two subsets, denoted G and B. We call an instance of the T -times repeated game with perfect monitoring a *block*, and say that a player i *sends message* $M \in \{G, B\}$ if he picks an action in M in the first period of a block. As in HO2006, we fix a payoff vector v to be achieved in equilibrium and pick four action profiles $a^{X,Y}$ for $(X, Y) \in \{G, B\}^2$ with

¹⁵A strategy is *belief-free at time t* if the continuation strategy at time t , $s_i | h_i^{t-1}$, is a best response against $s_{-i} | h_{-i}^{t-1}$ for all pairs of histories $(h_i^{t-1}, h_{-i}^{t-1})$. (Here, as we define formally below, " $|$ " indicates the restriction of a strategy to a given history set.)

$w_i^{X,Y} = g_i(a^{X,Y})$, $X, Y \in \{G, B\}$, where $w_i^{G,G} > v_i > w_i^{B,B}$, and

$$w_1^{G,B} > v_1 > w_1^{B,G}, \quad w_2^{B,G} > v_2 > w_2^{G,B}.^{16}$$

Choose $\underline{v}_i < \bar{v}_i$ with $v_i^* < \underline{v}_i < v_i < \bar{v}_i$ —where v_i^* is player i 's minmax payoff—such that

$$[\underline{v}_1, \bar{v}_1] \times [\underline{v}_2, \bar{v}_2] \subset \text{int}(\text{co}\{w^{G,G}, w^{B,B}, w^{B,G}, w^{G,B}\}).$$

We let \tilde{S}_i^T be the set of *block strategies* for player i , i.e. the set of strategies for the T -period perfect monitoring repeated game. We let \tilde{S}_i be the set of strategies $\tilde{s}_i \in \tilde{S}_i^T$ such that $\tilde{s}_i[\tilde{h}_i^t] = a_i^{M_2, M_1}$ for all $\tilde{h}_i^t = (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}), \dots, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1})$ with $a \in \{M_i\} \times G$ ($t \geq 1$). We then let

$$\begin{aligned} \tilde{A}_i(\tilde{h}_i^t) &\equiv \{a_i \in A_i : \exists \tilde{s}_i \in \tilde{S}_i \text{ such that } \tilde{s}_i[\tilde{h}_i^t](a_i) > 0\}, \\ \tilde{S}_i^\rho &\equiv \{\tilde{s}_i \in \tilde{S}_i : \tilde{s}_i[\tilde{h}_i^t](a_i) > \rho \text{ for all } \tilde{h}_i^t \text{ and } a_i \in \tilde{A}_i(\tilde{h}_i^t)\}.^{17} \end{aligned}$$

Define $\tilde{H}_i^{R,t}$ to be the set of period- t private histories of player i in the T -times-repeated game with perfect monitoring that are on the equilibrium path for some (and therefore, every) strategy profile in $\tilde{S}_1^\rho \times \tilde{S}_2^\rho$. Then we identify each $\tilde{h}_i^t \in \tilde{H}_i^{R,t}$ with the unique element of $h_i^t \in H_i^t$ such that h_i^t and \tilde{h}_i^t report exactly the same observations about the play of player $-i$ at all times and h_i^t contains no observations with a positive lag (all observations are observed instantaneously). Define $H_i^{R,t}$ as the image of $\tilde{H}_i^{R,t}$ under this identification, and denote this identification by $\tilde{h}_i^t \simeq h_i^t$ for $\tilde{h}_i^t \in \tilde{H}_i^{R,t}$ and $h_i^t \in H_i^t$. Also define the set of erroneous histories to be $H_i^{E,t} = H_i^t \setminus H_i^{R,t}$. This means that $H_i^{E,t}$ includes any private histories in which player i did not immediately observe the period- t' play of player $-i$ for some $t' < t$.

Additionally define the set of strategies $\mathcal{S}_i \subseteq S_i^T$ in the repeated game *with observation*

¹⁶These action profiles can be assumed to be pure, either with the use of a public randomization device or by picking a quadruple of sequences of action profiles such that the average payoff of each of the sequences satisfy the above properties.

¹⁷As in the HO2006 constructions, given any history h_i^t , the set \tilde{S}_i imposes either no restrictions on $s_i[h_i^t]$ or restricts $s_i[h_i^t]$ to a single action. In particular any strategy $\tilde{s}_i \in \tilde{S}_i^\rho$ puts positive weight on all actions after any erroneous history.

lags as the set

$$\mathcal{S}_i \equiv \{s_i \in \mathcal{S}_i^T : \exists \tilde{s}_i \in \tilde{\mathcal{S}}_i \text{ such that } \tilde{s}_i[\tilde{h}_i^t] = s_i[h_i^t] \text{ for all } \tilde{h}_i^t \in \tilde{H}_i^t \text{ where } \tilde{h}_i^t \simeq h_i^t\}.$$

Additionally, define

$$\mathcal{A}_i(h_i^t) \equiv \{a_i \in A_i : \exists s_i \in \mathcal{S}_i \text{ such that } s_i[h_i^t](a_i) > 0\},$$

$$\mathcal{S}_i^\rho \equiv \{s_i \in \mathcal{S}_i : s_i[h_i^t](a_i) > \rho \text{ for all } h_i^t \in H_i^t \text{ and } a_i \in \mathcal{A}_i(h_i^t)\}^{18}$$

Finally we define strategies $s_i^B, s_i^G \in \mathcal{S}_i^\rho$ by mapping the strategies \tilde{s}_i^B and \tilde{s}_i^G defined by HO2006 in a perfect monitoring repeated game to strategies in our environment with private monitoring and observation lags in a natural way. (The details of this definition are included in Appendix B.1.)

Proof of Theorem 2.4.3

The proof of Theorem 2.4.3 follows from three key lemmata; once these lemmata have been established, the remainder of the proof follows exactly as in HO2006. The first lemma adapts Lemma 1 of HO2006 to our setting of repeated games with information lags. Because the proof requires some nontrivial modifications, we include the argument here. As we show in the Appendix, analogous modifications can be made to the proofs of Lemmata 2 and 3 of HO2006; Theorem 2.4.3 then follows.

We write $s_i \mid H_i$ for the restriction of strategy s_i to history set H_i . We let \tilde{U}_i^T be the payoff of player i in the T -times repeated game *with perfect monitoring and no observation lags*. Analogously define U_i^T to be the ex-ante payoff of player i in the T -times repeated game *with private monitoring structure π and observation lags*. We consider a version of the T -times repeated game (with observation lags) which is augmented with a *transfer* $\zeta_{-i} : H_i^T \rightarrow \mathbb{R}$ at the end of the T -th period. In this *auxiliary scenario*, the payoff of i under strategy profile s is

¹⁸Just as in the case of $\tilde{\mathcal{S}}_i^\rho$ in Footnote 17, $s_i^\rho \in \mathcal{S}_i^\rho$ puts positive weight on all actions after any erroneous history.

taken to be

$$U_i^A(s, \xi_i) \equiv U_i^T(s) + (1 - \delta)\delta^T \mathbb{E}(\xi_i | s).$$

The set of best responses of player i in the auxiliary scenario with opponent's strategy s_{-i} and own transfer ξ_i is denoted $B_i(s_{-i}, \xi_i)$.

With these notations, we have the following lemma that defines the transfer ξ_i^B received after "bad" messages.

Lemma 2.4.4. *For every strategy profile $\bar{s} | H^E$, there exists $\bar{\varepsilon} > 0$ such that whenever $\lambda_i(0) > 1 - \bar{\varepsilon}$ for $i = 1, 2$ and π is $\bar{\varepsilon}$ -perfect, then there exists a nonnegative transfer $\xi_i^B : H_{-i}^T \rightarrow \mathbb{R}_+$ such that $S_i^T = B_i(\bar{s}_{-i}^B, \xi_i^B)$ where $\bar{s}_{-i}^B | H_{-i}^R = s_{-i}^B | H_{-i}^R$ and $\bar{s}_{-i}^B | H_{-i}^E = \bar{s}_{-i} | H_{-i}^E$, and for every $s_i \in B_i(\bar{s}_{-i}^B, \xi_i^B)$,*

$$\lim_{\varepsilon \rightarrow 0} U_i^A(s_i, \bar{s}_{-i}^B, \xi_i^B) = \max_{\bar{s}_i \in \bar{S}_i^T} \tilde{U}_i^T(\bar{s}_i, \bar{s}_{-i}^B).$$

This generalizes Lemma 1 of HO2006 to a repeated game in which information does not arrive instantaneously. To do so, we must contend with the fact that H_{-i}^T contains many more histories than in their private monitoring environment because information may arrive with lag, so that it is not immediately clear how to construct the ξ_i^B . We handle this issue by partitioning the set of histories into sets which are *past-observation equivalent*, in the sense that for any two time- t histories h^t and $h^{t'}$ in the same set, the $(t - 1)$ -period truncations of h^t and $h^{t'}$ are equal. We then identify each of the elements of this partition with a particular history in a private monitoring repeated game with the *immediate monitoring structure* μ induced by λ and π defined over the space of signal profiles $(\Omega_1 \cup \{\infty\}) \times (\Omega_2 \cup \{\infty\})$:

$$\mu(\hat{\omega}_1, \hat{\omega}_2 | a_1, a_2) = \begin{cases} \sum_{\tau_1=1}^{\infty} \sum_{\tau_2=1}^{\infty} \lambda(\tau_1, \tau_2) & \hat{\omega}_1, \hat{\omega}_2 = \infty \\ (\sum_{\tau=1}^{\infty} \lambda(\tau, 0)) \left(\sum_{\omega'_1 \in A_2} \pi(\omega'_1, \hat{\omega}_2 | a_1, a_2) \right) & \hat{\omega}_1 = \infty \text{ and } \hat{\omega}_2 \neq \infty \\ (\sum_{\tau=1}^{\infty} \lambda(0, \tau)) \left(\sum_{\omega'_2 \in A_1} \pi(\hat{\omega}_1, \omega'_2 | a_1, a_2) \right) & \hat{\omega}_1 \neq \infty \text{ and } \hat{\omega}_2 = \infty \\ \lambda(0, 0) \pi(\hat{\omega}_1, \hat{\omega}_2 | a_1, a_2) & \hat{\omega}_1, \hat{\omega}_2 \neq \infty. \end{cases}$$

This monitoring structure represents the information about the period- t action of the opponent that is available in period- $t + 1$, treating positive lags as the null signal. With this

identification, we can extend the arguments of HO2006 to arrive at our desired conclusion.

We construct the transfers ζ_i^G received after a “good” message in the repeated game with rare lags in a fashion closely similar to those specified in Lemma 2.4.4.

Lemma 2.4.5. *For every strategy profile $\bar{s} \mid H^E$, there exists $\bar{\varepsilon} > 0$ such that, whenever $\Pr(L > 0) < \bar{\varepsilon}$ and π is $\bar{\varepsilon}$ -perfect, there exists a nonpositive transfer $\zeta_i^G : H_{-i}^T \rightarrow \mathbb{R}_-$ such that*

$$\{s_i \in s_i^T : s_i \mid H_i^R = \hat{s}_i \mid H_i^R \text{ for some } \hat{s}_i \in \mathcal{S}_i \text{ and } s_i \mid H_i^E = \bar{s}_i \mid H_i^E\} \subseteq B_i(\bar{s}_{-i}^G, \zeta_i^G \mid \bar{s}_i)$$

where $\bar{s}_{-i}^G \mid H_{-i}^R = \bar{s}_{-i}^G \mid H_{-i}^R$ and $\bar{s}_{-i}^G \mid H_{-i}^E = \bar{s}_{-i} \mid H_{-i}^E$. Furthermore $\zeta_i^G : H_{-i}^T \rightarrow \mathbb{R}_-$ can be chosen so that, for every $s_i \in B_i(\bar{s}_{-i}^G, \zeta_i^G \mid \bar{s}_i)$, we have

$$\lim_{\varepsilon \rightarrow 0} U_i^A(s_i, \bar{s}_{-i}^G, \zeta_i^G) = \min_{\bar{s}_i \in \bar{\mathcal{S}}_i} \tilde{U}_i^T(\bar{s}_i, \bar{s}_{-i}^G),$$

ζ_i^G depends continuously on \bar{s} , and ζ_i^G is bounded away from $-\infty$.

We relegate the proof to the Appendix. The remainder of the proof of Theorem 2.4.3 follows along the same lines as in HO2006, defining $\bar{s}_{-i}^B \mid H_{-i}^E$ and $\bar{s}_{-i}^G \mid H_{-i}^E$, ζ_i^G and ζ_i^B as the fixed point of the relevant correspondence. The construction works because of Lemma 2.4.5 and the fact that play at periods $T, 2T, \dots$, is belief free (by Lemma 2.4.4). Thus for example if player i receives information about the play of player $-i$ in period $T - m$ at some time $T + l$, this does not have any effect on his best response calculation since player i 's strategy only depends on the history of information about the events occurring after period T .

2.4.3 The Repeated Game with Frequent Observation Lags

In the previous section, we required that the probability of a positive lag be small. In this section, we show that even if the lags are frequent and possibly very long, the folk theorem still obtains when $\lambda_i(\infty)$ is sufficiently small for $i = 1, 2$.

The following lemma employs a technique similar to that used in the proof of Theorem 2.3.2, using delayed-response strategies to relate the equilibrium payoffs in the game with rare observation lags to those with possibly long lags.

Lemma 2.4.6. *Suppose $v \in E(\delta, \pi, \hat{\lambda})$ for all lag distributions $\hat{\lambda}$ such that $\hat{\lambda}(0) > 1 - \bar{\epsilon}$ and all $\delta \in (\underline{\delta}, 1)$. Then for all lag distributions λ such that $\lambda_i(\infty) < \bar{\epsilon}/2$ for $i = 1, 2$, there exists some $\delta^* \in (0, 1)$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \delta^*$.*

Proof. Choose $K \in \mathbb{N}$ such that $(1 - \Lambda_i(K - 1)) < \bar{\epsilon}$ for $i = 1, 2$. and set $\delta^* = \underline{\delta}^{\frac{1}{K}}$. Then there exists a positive integer $K^* \geq K + 1$, such that $\delta^{K^*} \in (\underline{\delta}, 1)$ for every $\delta > \delta^*$.

Now divide the repeated game $G(\delta, \pi, \lambda)$ into K^* distinct repeated game “threads,” the ℓ -th ($1 \leq \ell \leq K^*$) of which is played in periods $\ell, K^* + \ell, 2K^* + \ell, \dots$. Because $K^* \geq K + 1$, each of these separate repeated games is equivalent to $G(\delta^{K^*}, \pi, \hat{\lambda})$ for some $\hat{\lambda}$ such that $\hat{\lambda}(0) > 1 - \bar{\epsilon}$, and each repeated game thread can be treated independently, as players never condition their play in the ℓ -th thread on information received about play in the ℓ' -th repeated games ($\ell' \neq \ell$). Because $v \in E(\delta^{K^*}, \pi, \hat{\lambda})$, it is then clear that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \delta^*$. \square

Theorem 2.4.2 follows directly from Lemma 2.4.6 and Theorem 2.4.3.

Remark. Ellison (1994) used threading primarily to lower the discount factor. By contrast, we use threading to ensure that the probability of lags being longer than the thread length remains low, so that players (with high probability) observe signals of play within a thread before choosing new actions within that thread. Thus the number of threads required in our proof is independent of the discount factor, while Ellison (1994) required the number of threads to become arbitrarily large as the discount factor approaches 1.

2.5 Lagged Public Monitoring

2.5.1 Model

In this section, we consider an n -player repeated game in which the monitoring structure of the repeated game is *public with stochastic lags*: There is a set of public signals, denoted Y , and we set $\Omega_i = Y$ for all players $i \in I$. Furthermore we assume that π is supported on the set

$$\{(y_1, \dots, y_n) \in Y^n : y_1 = y_2 = \dots = y_n\}.$$

That is, the monitoring structure of the underlying repeated game without lags is public. With a slight abuse of notation, we then write $\pi(y | a)$ as shorthand for $\pi((y, \dots, y) | a)$.

We place a mild restriction on the support of the monitoring structure π .

Assumption 2.5.1. For every pure action profile $a \in A$, there exist $y, y' \in Y$ with $y \neq y'$ such that $\pi(y | a), \pi(y' | a) > 0$.

Note that the argument used for the case of lagged perfect monitoring does not work here because the analogous auxiliary game does not have almost-perfect monitoring. Moreover, an extension of the Hörner and Olszewski (2009) construction to repeated games with rare observation lags is not possible, because that construction assumes that each player assigns high probability to the event that all players observe the same signal as in the setting of Mailath and Morris (2002); this condition is possibly violated when a player observes the low-probability “null” signal.¹⁹ The possibility of receiving an uninformative signal also prevents the application of the folk theorem of Sugaya (2011), because the necessary full rank condition fails. Thus, instead of invoking or adapting existing results for general private monitoring games, we allow for the possibility of communication that is perfectly and publicly observed at the end of every period, i.e. $M_i \neq \emptyset$. We assume that, unlike signals, messages are observed without delay. In the context of our “joint coauthorship” example of Section 3.2, the authors can quickly reach each other by phone or email after the reports arrive. We show that as long as $|M_i| \geq |Y| + 1$ for all i , a folk theorem can be established.

¹⁹We believe that threading combined with Hörner and Olszewski (2009) yields a folk theorem when lags are sufficiently positively correlated because the auxiliary repeated game corresponding to a thread can be treated as an almost-public monitoring game with the possibility of an uninformative null signal. The techniques we develop in this section are more novel.

2.5.2 Structure of the Observation Lags

In this section, we allow for the possibility that lags may be correlated (across agents). Define:²⁰

$$\Lambda(T) \equiv \Pr\left(\min_i \{L_i\} \leq T\right) \quad \text{and set} \quad \gamma_i \equiv \lim_{T \rightarrow \infty} \left[\frac{1 - \Lambda(T)}{1 - \Lambda_i(T)} \right].$$

The quantity γ_i represents the limiting conditional probability that player i assigns to the event that players $j \neq i$ have not received signals about period- t play within T periods, when he himself has also not received any signal about period- t play within T periods. Note that if $\Pr(\{L : L_i = \infty\}) > 0$, then $\gamma_i = \frac{\lambda((\infty, \dots, \infty))}{\lambda_i(\infty)}$.

For our results in the section, we assume that $\lambda_i(\infty)$ and γ_i are both small for all i . It is easy to see what kind of lags satisfy the first condition. The second condition is a bit more subtle and so we illustrate it through some concrete examples in the Appendix.

2.5.3 The Folk Theorem

We begin our analysis with the simple observation that the repeated play of α^* is an equilibrium of the game with observation lags.²¹ We use this fact along with techniques from Abreu *et al.* (1990) and FLM to construct equilibria that generate any payoff profile $v \in \text{int}(V_{\alpha^*})$.

To use the techniques of FLM, we need to impose some additional assumptions on the public monitoring structure π . Recall the following definition from FLM.

Definition 2.5.2. Let π be a public monitoring structure. Then a mixed action profile α has pairwise full rank for a pair $i, j \in I$ if the $((|A_i| + |A_j|) \times |Y|)$ matrix

$$\begin{pmatrix} (\pi(\cdot \mid a_i, \alpha_{-i}))_{a_i \in A_i} \\ (\pi(\cdot \mid a_j, \alpha_{-j}))_{a_j \in A_j} \end{pmatrix}$$

²⁰Although we assume that $\Lambda_i(T) < 1$ for all $T \in \mathbb{N}$, our results extend to the case in which there exists some player i and some T^* such that $\Lambda_i(T^*) = 1$. In that case, we can simply take the number of threads to be larger than T^* , so that player i 's signal structure in the auxiliary game need not contain ∞ as one of its elements.

²¹Note that for such play, the communication strategies are irrelevant, so we need not specify them.

has rank $|A_i| + |A_j| - 1$.

We will maintain the following restriction on π throughout the rest of this section.

Assumption 2.5.3. For all pairs i, j , there exists a profile α that has pairwise full rank for that pair.

We can now state our folk theorem for repeated games with public monitoring and stochastic lags with communication.

Theorem 2.5.4. *Let $v \in \text{int}(V_{\alpha^*})$ and suppose that π satisfies Assumptions 2.5.1 and 2.5.3. Furthermore suppose that $|M_i| \geq |Y| + 1$ for all i . Then there exist some $\varepsilon^* \in (0, 1)$ such that for every λ such that $\gamma_i < \varepsilon^*$ and $\lambda_i(\infty) < \varepsilon^*$ for all i , there exists $\delta^* \in (0, 1)$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \delta^*$.*

As a preview of our proof, it is important that both γ_i and $\lambda_i(\infty)$ are small for all i . The need of the latter condition should be intuitive. For example, in a two-player game, if it is likely that one player never observes any information, the other player would have an incentive to play myopically. As we will see, the former condition is important for establishing truthful communication of signals. The remainder of the section proves Theorem 2.5.4.

2.5.4 Private Monitoring Game with Communication

Incentives for Truthful Communication

We first consider a private monitoring game with communication (in every period) and *no observation lags* for which each player's message space is $M_i = \tilde{Y} \equiv Y \cup \{\infty\}$. The results of this section are of stand-alone interest: the case where players might sometimes not see the signal seems plausible and it leads to a form of private monitoring that does not appear to be covered by past results. Let us first define some notation. For a vector $\tilde{y} \in \tilde{Y}^n$, define $\mathcal{I}(\tilde{y}) = \{i : \tilde{y}_i \neq \infty\}$ and $|\tilde{y}| \equiv |\mathcal{I}(\tilde{y})|$. Define the following set

$$\mathcal{V} \equiv \{(\tilde{y}_1, \dots, \tilde{y}_n) \in \tilde{Y}^n : |(\tilde{y}_1, \dots, \tilde{y}_n)| > 0 \text{ and } \tilde{y}_j = \tilde{y}_k \forall j, k \text{ such that } \tilde{y}_j, \tilde{y}_k \neq \infty\}.$$

The monitoring structure is then supported on the set $\mathcal{Y} \cup \{(\infty, \dots, \infty)\}$. For any $\tilde{y} \in \mathcal{Y}$, we define $\vec{y} \in Y$ to be the $y \in Y$ such that $\tilde{y}_j = y$ for all j such that $\tilde{y}_j \neq \infty$.

Now consider a private monitoring structure π^{Pr} that is supported on the set $\mathcal{Y} \cup \{(\infty, \dots, \infty)\}$ with the following additional features.

Assumption 2.5.5. $\pi^{\text{Pr}}((\infty, \dots, \infty) \mid a)$ is constant across all $a \in A$.

Assumption 2.5.6. $\sum_{\vec{y}=\vec{y}} \pi^{\text{Pr}}(\tilde{y} \mid a) = (1 - \pi^{\text{Pr}}((\infty, \dots, \infty) \mid a))\pi(y \mid a)$ for all $y \in Y$.

The reasons for these restrictions become clear when we relate this game to the repeated game with observation lags. We say that this monitoring structure is ε -close to π if π^{Pr} is such that

$$\pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid a, \tilde{y}_i) < \varepsilon \quad (2.1)$$

for all $a \in A$ and all $\tilde{y}_i \in A \cup \{\infty\}$. Note that this definition of ε -closeness to a public monitoring structure is quite different from the one used by Hörner and Olszewski (2009). The key difference is in the conditional probability $\pi^{\text{Pr}}((\infty, \dots, \infty) \mid a, \tilde{y}_i = \infty)$. Hörner and Olszewski (2009) assumed this conditional probability to be close to 1. Here, we assume that it is very small.

We denote by $G^{\text{Pr}}(\delta, \pi^{\text{Pr}})$ the private monitoring game with discount factor δ and private monitoring structure π^{Pr} (and communication) and let $E^{\text{Pr}}(\delta, \pi^{\text{Pr}})$ be the set of sequential equilibrium payoffs of $G^{\text{Pr}}(\delta, \pi^{\text{Pr}})$. We now show the following.

Theorem 2.5.7. *Let $v \in \text{int}(V_{a^*})$. Then there exist $\underline{\delta}, \bar{\delta} \in (0, 1)$ with $\underline{\delta} < \bar{\delta}$ and $\bar{\varepsilon} \in (0, 1)$ such that $v \in E^{\text{Pr}}(\delta, \pi^{\text{Pr}})$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all private monitoring structures π^{Pr} that are $\bar{\varepsilon}$ -close to π .*

To prove this theorem we construct strategies that generate a payoff profile of v , and are public perfect in the sense of Kandori and Matsushima (1998): strategies in the non-communication stages of the game depend only on the sequence of message profiles reported in the history. These strategies use a form of grim-trigger reversion to static Nash equilibrium when the messages disagree, in order to provide incentives for truthful reporting. We prove the theorem in two parts. We first prove a lemma demonstrating that

truth-telling is *incentive compatible* (i.e. that each player i should report message $m_i = y$ upon seeing signal $y \in Y$) when ε is sufficiently small given strategies with this grim-trigger property.

Lemma 2.5.8. *Let W be a convex, compact set that is a subset of $\text{int}(V_{a^*})$. Consider a collection of public perfect strategy profiles $\{\sigma^{\delta, \pi^{\text{Pr}}}\}$, indexed by δ and π^{Pr} , for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all private monitoring structures π^{Pr} that are $\bar{\varepsilon}$ -close to π and have the following properties.*

1. *In period t , each player i (truthfully) communicates the signals $\tilde{y}_i^t \in \tilde{Y} = M_i$ he observes in period t .*
2. *If there exists some t such that $m^t \notin \mathcal{Y}$, then all players i play α_i^* .*
3. *Strategies are such that $\sigma^{\delta, \pi^{\text{Pr}}}(m^0, \dots, m^t) = \sigma^{\delta, \pi^{\text{Pr}}}(\bar{m}^0, \dots, \bar{m}^t)$ whenever $\bar{m}^\tau = \vec{m}^\tau$ for all $\tau = 0, \dots, t$.*
4. *Expected continuation values are always contained in W for play of $\sigma^{\delta, \pi^{\text{Pr}}}$ in the game $G^{\text{Pr}}(\delta, \pi^{\text{Pr}})$ whenever the message history contains only elements in the set \mathcal{Y} .*

Then there exists $\varepsilon^ \leq \bar{\varepsilon}$ such that for all private monitoring structures π^{Pr} that are ε^* -public except at infinity and all $\delta \in [\underline{\delta}, \bar{\delta}]$, truthful communication is incentive compatible at any private history in $G^{\text{Pr}}(\delta, \pi^{\text{Pr}})$ given continuation play determined by $\sigma^{\delta, \pi^{\text{Pr}}}$ and truthful communication by all other players.*

Proof. We check that there are no profitable one-stage deviations in which a player misreports once and then follows the continuation strategy prescribed by $\sigma_i^{\delta, \pi^{\text{Pr}}}$. First note that if the player is at a history in which there exists some t at which $m^t \notin \mathcal{Y}$, then all players play α_i^* forever from that point on. Since then continuation play does not depend on the message being sent, all players are indifferent to the message that they send after such a history. Thus it is incentive compatible.

So it remains to analyze incentives for truth-telling after histories in which $m^t \in \mathcal{Y}$ for all t . Suppose first that player i sees the null signal. Then by reporting ∞ , player i obtains

an expected payoff of

$$\sum_{\vec{y} \in \mathcal{Y}} \pi^{\text{Pr}}(\vec{y} \mid \alpha, \tilde{y}_i = \infty) w_i(\vec{y})$$

for some $\alpha \in \prod_{i=1}^n \Delta(A_i)$ and some expected continuation value function $w : Y \rightarrow W$.²²

If instead player i reports $y' \in Y$, he obtains a payoff of

$$\pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha, \tilde{y}_i = \infty) w_i(y') + \sum_{\vec{y} = y'} \pi^{\text{Pr}}(\vec{y} \mid \alpha, \tilde{y}_i = \infty) w_i(y')$$

Thus, to show that truth-telling is incentive compatible after all histories in which a player observes the null signal, it suffices to show that there exists ε^* sufficiently small so that

$$\sum_{y \in Y \setminus \{y'\}} \pi^{\text{Pr}}(\{\vec{y} : \vec{y} = y\} \mid \alpha, \tilde{y}_i = \infty) w_i(y) > \pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha, \tilde{y}_i = \infty) w_i(y') \quad (2.2)$$

for all $y' \in Y$, all $\alpha \in \prod_{i=1}^n \Delta(A_i)$, all $w : Y \rightarrow W$, $i = 1, \dots, n$, and all π^{Pr} ε^* -close to π . Assumptions 2.5.1 and 2.5.6 imply:

$$\begin{aligned} M(\pi^{\text{Pr}}, y', \alpha) &\equiv \sum_{y \in Y \setminus \{y'\}} \pi^{\text{Pr}}(\{\vec{y} : \vec{y} = y\} \mid \alpha, \tilde{y}_i = \infty) \\ &= (1 - \pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha)) \sum_{y' \in Y \setminus \{y'\}} \pi(y \mid \alpha) \\ &= (1 - \pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha))(1 - \pi(y' \mid \alpha)) \\ &> 0. \end{aligned}$$

Note that for a fixed map, $w : Y \rightarrow W$ and $\alpha \in \prod_{i=1}^n \Delta(A_i)$, (2.2) holds for all $i = 1, \dots, n$ and all $y' \in Y$ if and only if

$$\sum_{y \in Y \setminus \{y'\}} \frac{\pi^{\text{Pr}}(\{\vec{y} : \vec{y} = y\} \mid \alpha, \tilde{y}_i = \infty)}{M(\pi^{\text{Pr}}, y', \alpha)} w_i(y) > \frac{\pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha, \tilde{y}_i = \infty)}{M(\pi^{\text{Pr}}, y', \alpha)} w_i(y'). \quad (2.3)$$

Now let π^{Pr} be ε -close to π . As $\varepsilon \rightarrow 0$, $M(\pi^{\text{Pr}}, y') \rightarrow 1 - \pi(y \mid \alpha)$ and thus because π^{Pr} is ε -

²²Note that in any sequential equilibrium if a player observes signal ∞ , he still believes that all other players played according to their prescribed actions, i.e. that there have been no “unexpected” events.

close to π ,

$$\frac{\pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha, \tilde{y}_i = \infty)}{M(\pi^{\text{Pr}}, y', \alpha)} < \frac{\varepsilon}{M(\pi^{\text{Pr}}, y', \alpha)} \rightarrow \frac{0}{1 - \pi(y' \mid \alpha)} = 0.$$

Note that for any value of $\varepsilon > 0$, because W is convex, the left hand side of inequality (2.3) is an element of W .

Therefore because W is compact and contained in the interior of V_{α^*} , there is some ε^* such that inequality (2.3) holds for all π^{Pr} ε^* -close to π . Moreover $\prod_{i=1}^n \Delta(A_i)$ and the set of all maps $w : Y \rightarrow W$ are both compact. Therefore such an ε^* can be taken uniformly across all α and all maps $w : Y \rightarrow W$. This shows that all players will report the null signal truthfully when π^{Pr} is ε^* -close to π .

Now suppose that player i observes $y \in Y$. By reporting truthfully, player i obtains a payoff of $w_i(y)$ for some map $w : Y \rightarrow W$. However by reporting $y' \in Y$ with $y' \neq y$, player i obtains a payoff of

$$\pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid \alpha, \tilde{y}_i = y) w_i(y') \quad (2.4)$$

while reporting ∞ yields a payoff of

$$(1 - \pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid \alpha, \tilde{y}_i = y)) w_i(y). \quad (2.5)$$

Clearly $w_i(y)$ is at least the expression in (2.5) for any π^{Pr} since $w_i(y) \geq 0$. Furthermore we can take ε^* sufficiently small so that $w_i(y) > \max_{y' \neq y} \{\varepsilon w_i(y')\}$ for all $y \in Y$, all maps $w : Y \rightarrow W$, all $i = 1, \dots, n$, and all $\varepsilon < \varepsilon^*$. Then all players have an incentive to report truthfully upon observing an informative signal when π^{Pr} is ε^* -close to π since

$$w_i(y) > \max_{y' \neq y} \{\varepsilon w_i(y')\} \geq \max_{y' \neq y} \{\pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid \alpha, \tilde{y}_i = y) w_i(y')\}$$

and

$$w_i(y) \geq (1 - \pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid \alpha, \tilde{y}_i = y)) w_i(y)$$

trivially. This concludes the proof. \square

Remark. The fact that (2.1) is small for all $\tilde{y}_i \in A_i \cup \{\infty\}$ is crucial. Otherwise, because a

message profile of (∞, \dots, ∞) results in reversion to the static Nash equilibrium, player i upon observation of the null signal may have an incentive to deviate and report some signal $y \in Y$.

Remark. Because the set W in the lemma does not depend on delta, neither does ε^* . This is important for our folk theorem as we must establish a claim about all games with a private monitoring structure that is ε^* -close to π and all discount factors in an interval.

Remark. Players have no incentive to either communicate or respond to a signal that arrives late, since in equilibrium players do not respond to such communication. This is similar to the way in which we treat late signals in Section 2.4, where the belief-free property of the equilibrium construction allows us to show that players do not have an incentive to respond to late signals.

Non-Communication Stages

Lemma 2.5.8 provides a sufficient condition for truth-telling to be incentive compatible. We now show that given truthful communication by all players at all histories, we can construct a collection of strategies $\{\sigma^\delta, \pi^{\text{PF}}\}$ that satisfy the necessary properties of Lemma 2.5.8 for truthful communication and in which all players are also playing best-responses in the non-communication stages of the game.

To construct such strategies $\sigma^\delta, \pi^{\text{PF}}$, we first specify that players play α^* whenever in the history there exists some t such that $m^t \notin \mathcal{Y}$. Then it is trivial that playing α_i^* is a best response at such a history since opponents play α_{-i}^* forever. It remains to specify play after histories in which all messages in the history are elements of \mathcal{Y} . We do this by considering public strategies that only depend on the history of messages.

Given strategies that satisfy conditions 1, 2, and 3 of Lemma 2.5.8 we can simplify the analysis to that of an auxiliary public monitoring game defined in the following discussion. The auxiliary game is one of standard simultaneous moves in which public signals arise according to the conditional probability distribution π every period. We then modify this repeated game so that at the beginning of periods $1, 2, \dots$, the game ends with probability ε

and each player receives flow payoffs of $0 = g_i(\alpha^*)$ thereafter. This corresponds exactly to the event in which all players report the null signal, triggering all players to play according to α^* forever.²³

In the modified game, payoffs are given by

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t (1 - \varepsilon)^t g_i(a^t). \quad (2.6)$$

We denote this public monitoring game by $G^{\text{pu}}(\delta, \varepsilon)$ and let $EP^{\text{u}}(\delta, \varepsilon)$ be the set of sequential equilibrium payoffs of $G^{\text{pu}}(\delta, \varepsilon)$. Note that in this game the feasible payoff set is not constant in δ and ε , and in particular for any fixed $\varepsilon > 0$, as $\delta \rightarrow 1$, the feasible payoff set converges to $\{0\}$, just as the payoffs to grim trigger strategies converge to those of static Nash equilibrium as $\delta \rightarrow 1$ in a repeated game with imperfect public monitoring. However for any fixed δ , as $\varepsilon \rightarrow 0$, the feasible payoff set converges to V , the feasible payoff set of the original public monitoring game. Our analysis takes care in addressing this issue.

In order to extend the arguments of FLM to this modified repeated game, we first renormalize payoffs so that the feasible payoff set is indeed equal to V . We do this by multiplying the payoffs by a factor of $(1 - \delta(1 - \varepsilon))/(1 - \delta)$ to get payoff structure

$$(1 - \delta(1 - \varepsilon)) \sum_{t=0}^{\infty} \delta^t (1 - \varepsilon)^t g_i(a^t). \quad (2.7)$$

Now, our modified game corresponds to a repeated game with discount factor given by $\delta(1 - \varepsilon)$, hence all of the conclusions of FLM can be applied to this game, with the appropriate assumptions on the (original) public monitoring structure.

Before we proceed with the analysis of the game, recall the definition of self-generation (Abreu *et al.* (1990)).

Definition 2.5.9. For $W \subset \mathbb{R}^n$, define the sets $B(W, \delta, \varepsilon)$ and $\hat{B}(W, \delta, \varepsilon)$ as follows. Let $B(W, \delta, \varepsilon)$ be the set of $v \in \mathbb{R}^n$ such that there exists some mixed action profile α and a map

²³Because all players report truthfully at all histories, message profiles $m \in \tilde{Y}^n \setminus \{\mathcal{Y} \cup (\infty, \dots, \infty)\}$ never occur on the equilibrium path. Thus the “grim phase” of playing α^* forever is only triggered in the event of message profile $m = (\infty, \dots, \infty)$; this happens with probability ε .

$w : Y \rightarrow W$ such that

$$v = (1 - \delta)g(\alpha) + \delta(1 - \varepsilon) \sum_{y \in Y} w(y)\pi(y|\alpha), \quad \text{and}$$

$$v_i \geq (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta(1 - \varepsilon) \sum_{y \in Y} w_i(y)\pi(y | a_i, \alpha_{-i}),$$

for all $a_i \in A_i$ and all i . Analogously define $\hat{B}(W, \delta, \varepsilon)$ to be the set of $v \in \mathbb{R}^n$ such that there exists some mixed action profile α and a map $w : Y \rightarrow W$ such that

$$v = (1 - \delta(1 - \varepsilon))g(\alpha) + \delta(1 - \varepsilon) \sum_{y \in Y} w(y)\pi(y|\alpha), \quad \text{and}$$

$$v_i \geq (1 - \delta(1 - \varepsilon))g_i(a_i, \alpha_{-i}) + \delta(1 - \varepsilon) \sum_{y \in Y} w_i(y)\pi(y | a_i, \alpha_{-i})$$

for all i . We say that W is *self-generating in the repeated game with payoff structure (2.6) with discount factor δ and absorption probability ε* if $W \subseteq B(W, \delta, \varepsilon)$. Similarly, W is *self-generating in the repeated game with payoff structure (2.7) with discount factor δ and absorption probability ε* if $W \subseteq \hat{B}(W, \delta, \varepsilon)$.

Because the public monitoring game $G^{\text{Pr}}(\delta, \varepsilon)$ has a slightly different structure from that of a standard public monitoring game, the consequences of self-generation are not immediate from past theorems, but the same ideas apply as shown in the next lemma.

Lemma 2.5.10. *Suppose W is compact and that $W \subseteq B(W, \delta, \varepsilon)$. Then $W \subseteq EP^u(\delta, \varepsilon)$.*

The proof of Lemma 2.5.10 is completely standard, so we omit it. FLM applied to the repeated game with discount factor $\delta(1 - \varepsilon)$ yields the following lemma.

Lemma 2.5.11. *Suppose that Assumption 2.5.3 holds. Let \hat{W} be a smooth, compact, convex set in the $\text{int}(V_{a^*})$. Then there exists $\bar{\delta} \in (0, 1)$ and $\bar{\varepsilon} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ and all $\varepsilon < \bar{\varepsilon}$, $\hat{W} \subseteq \hat{B}(\hat{W}, \delta, \varepsilon)$, that is, \hat{W} is self-generating in the repeated game with payoff structure (2.7) with discount factor δ and absorption probability ε .*

Next, we translate the payoff set used in Lemma 2.5.11 back into payoffs without the renormalization. To do this, we define (for a set \hat{W}) a set W under the payoff normalization

given by (2.6):

$$W = \frac{1 - \delta}{1 - \delta(1 - \varepsilon)} \hat{W}. \quad (2.8)$$

Of course for any fixed ε and a fixed set \hat{W} , as $\delta \rightarrow 1$, W shrinks (setwise) towards the point-set $\{0\}$. Thus for any choice of $v \in \text{int}(V_{\alpha^*})$, v will necessarily lie outside of W for δ close to 1, so it is not immediate from that for any discount factor δ , one can construct a self-generating set containing v according to the B operator rather than the \hat{B} operator. The next lemma shows that this can be done for a non-empty interval of discount factors.

Lemma 2.5.12. *Let $v \in \text{int}(V_{\alpha^*})$ and suppose that Assumption 2.5.3 holds. Consider the repeated game with payoffs given by (2.6). Then there exist $\underline{\delta}, \bar{\delta} \in (0, 1)$ with $\underline{\delta} < \bar{\delta}$ and $\bar{\varepsilon} \in (0, 1)$ such that $v \in E^{\text{pu}}(\delta, \varepsilon)$ for all $\varepsilon < \bar{\varepsilon}$ and all $\delta \in [\underline{\delta}, \bar{\delta}]$. Furthermore there exists some compact set $W \subseteq \text{int}(V_{\alpha^*})$ such that the equilibrium corresponding to payoff v can be taken to have continuation values that always lie in W for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all $\varepsilon < \bar{\varepsilon}$.*

Proof. Fix some $v \in \text{int}(V_{\alpha^*})$. Then choose a compact, smooth, convex set $\hat{W} \subseteq \text{int}(V_{\alpha^*})$ such that $v \in \text{int}(\hat{W})$. Since \hat{W} is bounded away from 0 and contains v , there exists some $\eta < 1$ and compact set W such that $v \in \eta' \hat{W} \subseteq W \subseteq \text{int}(V_{\alpha^*})$ for all $\eta' \in [\eta, 1]$. By Lemma 2.5.11, there exists some $\underline{\delta}$ and ε^* such that $\hat{W} \subseteq \hat{B}(\hat{W}, \delta, \varepsilon)$ for all $\delta \geq \underline{\delta}$ and all $\varepsilon < \bar{\varepsilon}$.

Now choose $\bar{\delta} \in (\underline{\delta}, 1)$ arbitrarily. Then choose

$$\bar{\varepsilon} = \min \left\{ \frac{(1 - \eta)(1 - \bar{\delta})}{\bar{\delta}\eta}, \varepsilon^* \right\}.$$

This then implies that for all $\varepsilon < \bar{\varepsilon}$ and all $\delta \in [\underline{\delta}, \bar{\delta}]$,

$$v \in W_{\delta, \varepsilon} \equiv \frac{1 - \delta}{1 - \delta(1 - \varepsilon)} \hat{W} \subseteq W \subseteq \text{int}(V_{\alpha^*}).$$

Furthermore $\hat{W} \subseteq \hat{B}(\hat{W}, \delta, \varepsilon)$ for all $\varepsilon < \bar{\varepsilon}$ and all $\delta \in [\underline{\delta}, \bar{\delta}]$.

This observation allows us to establish the claims of the lemma. To see this, we note that for every $\delta \in [\underline{\delta}, \bar{\delta}]$ and all $\varepsilon < \bar{\varepsilon}$, every $\check{w} \in \hat{W}$ can be written in the form

$$\check{w}_i = (1 - \delta(1 - \varepsilon))g_i(\alpha) + \delta(1 - \varepsilon) \sum_{y \in Y} \hat{w}_i(y) \pi(y|\alpha)$$

for all i for some α and some $\hat{w} : Y \rightarrow \hat{W}$ so that α_i is a best response given the expected continuation payoff \hat{w}_i and opponents' current mixed action profile α_{-i} . Translating payoffs into the original normalization under (2.6), yields

$$\frac{1 - \delta}{1 - \delta(1 - \varepsilon)} \check{w}_i = (1 - \delta)g_i(\alpha) + \delta(1 - \varepsilon) \sum_{y \in Y} \frac{1 - \delta}{1 - \delta(1 - \varepsilon)} \hat{w}_i(y) \pi(y | \alpha).$$

We then see that

$$\frac{1 - \delta}{1 - \delta(1 - \varepsilon)} \hat{w}_i(y) \in W_{\delta, \varepsilon}$$

for all $y \in Y$ and all i . Thus $v \in W_{\delta, \varepsilon} \subseteq B(W_{\delta, \varepsilon}, \delta, \varepsilon)$ and $W_{\delta, \varepsilon} \subseteq W$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all $\varepsilon < \bar{\varepsilon}$. Then from Lemma 2.5.10, if $v \in W \subseteq B(W_{\delta, \varepsilon}, \delta, \varepsilon)$ then $v \in E^{\text{pu}}(\delta, \varepsilon)$. Therefore $v \in E^{\text{pu}}(\delta, \varepsilon)$ for all $\varepsilon < \bar{\varepsilon}$ and all $\delta \in [\underline{\delta}, \bar{\delta}]$. \square

Then we relate the auxiliary game $G^{\text{pu}}(\delta, \varepsilon)$ back to the original private monitoring game $G^{\text{pr}}(\delta, \pi^{\text{pr}})$ as follows. We let $\varepsilon = \pi^{\text{pr}}((\infty, \dots, \infty) | a)$.²⁴ Furthermore when constructing strategies that satisfy condition 3 of Lemma 2.5.8, players play as if they are observing a public signal structure over $Y \cup \{\infty\}$ with $\pi^{\text{pu}}(\infty | a) = \varepsilon$ and $\pi^{\text{pu}}(y | a) = \sum_{\vec{y}=y} \pi^{\text{pr}}(\vec{y} | a) = (1 - \varepsilon)\pi(y | a)$ by Assumption 2.5.5. With these observations, lemmas 2.5.8 and 2.5.12 together prove Theorem 2.5.7.

2.5.5 The Repeated Game with Observation Lags

We now prove Theorem 2.5.4. To this end, let us first link the private monitoring game with communication, $G^{\text{pr}}(\delta, \pi^{\text{pr}})$, to the original repeated game with public monitoring and observation lags: For a given lag distribution λ and some $T \in \mathbb{N}$, we define the induced private monitoring structure π^{pr} in the following way:

$$\pi^{\text{pr}}(\vec{y} | a) = \begin{cases} \Pr(\{L : L_i \geq T \forall i \in \mathcal{I}(\vec{y}), L_i < T \forall i \notin \mathcal{I}(\vec{y})\}) \pi(y | a) & \text{if } \vec{y} \in \mathcal{Y}, \vec{y} = y \\ \Pr(\{L : L_i \geq T \forall i\}) & \text{if } \vec{y} = (\infty, \dots, \infty) \\ 0 & \text{otherwise.} \end{cases}$$

²⁴Here we use Assumption 2.5.5 so that ε does not depend on $a \in A$.

Note that π^{Pr} satisfies Assumptions 2.5.5 and 2.5.6. Then given this monitoring structure, we define the game $\tilde{G}(\delta, \lambda, T) = G^{\text{Pr}}(\delta, \pi^{\text{Pr}})$, and let $\tilde{E}(\delta, \lambda, T)$ be the set of sequential equilibrium payoffs of $\tilde{G}(\delta, \lambda, T)$ for which equilibrium play depends only on the message histories.

In constructing an equilibrium for the repeated game with observation lags, we suppose that the message spaces in each period are $M_i = \tilde{Y}$. Henceforth $G(\delta, \lambda)$ and $E(\delta, \lambda)$ specifically refer to the repeated game with observation lag distribution λ , discount factor δ , and message spaces $M_i = \tilde{Y}$.

Lemma 2.5.13. *Suppose that $v \in \tilde{E}(\delta, \lambda, T)$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ for some fixed λ and all $T \geq T^*$, where $0 < \underline{\delta} < \bar{\delta} < 1$. Then there exists some $\delta^* \in (0, 1)$ such that $v \in E(\delta, \lambda)$ for all $\delta > \delta^*$.*

As in the proof of Lemma 2.4.6, the proof here also divides the repeated game into threads, mapping each thread to an auxiliary game of the form described in the preceding sections. However because the lemma here additionally allows for communication, care in defining the communication strategies is necessary in order to appropriately construct the map from threads to auxiliary games.

Proof. We set $\delta^* = (\underline{\delta}/\bar{\delta})^{\frac{1}{T^*+1}}$, so that for every $\delta > \delta^*$, there exists a positive integer multiple of $T^* + 1$, $N(\delta)$, such that $\delta^{N(\delta)} \in [\underline{\delta}, \bar{\delta}]$.

Now we divide the repeated game $G(\delta, \lambda)$ into $N(\delta)$ distinct repeated game threads, the ℓ -th ($1 \leq \ell \leq N(\delta)$) of which is played in periods $\ell, N(\delta) + \ell, 2N(\delta) + \ell, \dots$. In our construction, players communicate the public signal generated at the end of period $(k - 1)N(\delta) + m$ at the end of period $kN(\delta) + (m - 1)$. If they have not yet seen the signal of that period's play they report the null signal. Then each repeated game thread is equivalent to a private monitoring game of the form described in the previous section.

As in the proof of Lemma 2.4.6, each repeated game can be treated independently, as players never condition their play in the ℓ -th repeated game on information received about play in the ℓ' -th repeated games ($\ell' \neq \ell$). Moreover, any equilibrium of $\tilde{G}(\delta^{N(\delta)}, \lambda, N(\delta))$ where play depends only on the message history can be embedded into an equilibrium of

one of the repeated game threads. But since $N(\delta) > T^* + 1$, we have $v \in \tilde{E}(\delta^{N(\delta)}, \lambda, N(\delta))$, so it is then clear that $v \in E(\delta, \lambda)$ for all $\delta > \delta^*$. \square

We can now finish the proof of Theorem 2.5.4.

Proof of Theorem 2.5.4. By Theorem 2.5.7, there exist $\underline{\delta}, \bar{\delta} \in (0, 1)$ with $\underline{\delta} < \bar{\delta}$ and $\varepsilon^* \in (0, 1)$ such that $v \in E^{\text{Pr}}(\delta, \pi^{\text{Pr}})$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all π^{Pr} that is ε^* -close to π .

Then choose $\bar{\varepsilon} > 0$ such that $\frac{\varepsilon}{1-\varepsilon} < \varepsilon^*$ for all $\varepsilon \leq \bar{\varepsilon}$. Now suppose that $\lambda_i(\infty) < \bar{\varepsilon}$ and $\gamma_i < \bar{\varepsilon}$ for all i . Then there exists a (finite) K^* such that

$$\frac{\Pr(L_i \leq K, L_j > K \forall j \neq i)}{\Lambda_i(K)} \leq \frac{\Pr(L_j > K \forall j \neq i)}{\Lambda_i(K)} < \varepsilon^* \quad \text{and} \quad \frac{1 - \Lambda(K)}{1 - \Lambda_i(K)} < \varepsilon^*$$

for all i and $K \geq K^*$. Thus $v \in \tilde{E}(\delta, \lambda, T)$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all $T \geq K^*$ since it is easy to show that the π^{Pr} induced by λ and T is in fact ε^* -close to π for all $T \geq K^*$. This however means—by Lemma 2.5.13—that there exists some $\delta^* \in (0, 1)$ such that $v \in E(\delta, \lambda)$ for all $\delta > \delta^*$; this concludes the proof. \square

Remark. Note that the proof of this theorem uses delayed-response strategies in three ways: to ensure that in each thread there is very low probability of all players' lags being longer than the thread length; so that even after not observing any signal from the previous period in a thread, players believe with high probability that others have observed an informative signal; *and* to map discount factors near 1 in the game $G(\delta, \lambda)$ to intermediate discount factors in the auxiliary games. The first feature is also present in the proof of Lemma 2.4.6. The second and third features are specific to the proof here. The second, ensured by the assumption that γ_i is small, is key to establishing incentives for truthful communication, so that π^{Pr} in the auxiliary game can be shown to satisfy condition (2.1). The third feature is closely analogous to the use of threads in the work of Ellison (1994).

Remark. Note that an important part of the proof of Theorem 2.5.4 is that messages are instantaneously observed. As in the literature on private monitoring games with communication, this is important since the messages serve to make private information public.

However, it is straightforward to extend our argument to settings in which messages are observed with a bounded lag. To see this, suppose that lags arrive within \check{K} periods with probability 1. We separate the game into $\hat{K} \equiv \max\{K, \check{K}\} + 1$ threads, where K is as in the proof of Theorem 2.5.4 (page 95). Each thread is further subdivided into a pair of subthreads, respectively played in “even” and “odd” thread periods; players communicate information observed in the even (resp. odd) subthread in periods of the odd (resp. even) subthread. Since the gap between thread periods is at least \check{K} , all messages sent in the even (resp. odd) subthread arrive with probability 1 before the next period of the odd (resp. even) subthread. Thus messages about play in the odd (resp. even) subthread arrive in time for the next round of play in that subthread. More formally, the ℓ -th thread is separated into two subthreads so that:

1. In periods $(2k)\hat{K} + \ell$, the players send messages about the signals generated in period $(2k - 1)\hat{K} + \ell$, and in periods $(2k + 1)\hat{K} + \ell$, the players play the appropriate responses to the messages sent in periods $2\hat{K} + \ell, 4\hat{K} + \ell, \dots, (2k)\hat{K} + \ell$.
2. In periods $(2k + 1)\hat{K} + \ell$, the players send messages about the signals generated in period $(2k)\hat{K} + \ell$, and in periods $(2k + 2)\hat{K} + \ell$, the players play the appropriate responses to the messages sent in periods $\hat{K} + \ell, 3\hat{K} + \ell, \dots, (2k + 1)\hat{K} + \ell$.

Under this construction, with the number of threads larger than $\hat{K} = \max\{K, \check{K}\} + 1$, messages sent in period $(2k)\hat{K} + \ell$ (resp. period $(2k + 1)\hat{K} + \ell$) are observed with probability 1 by the time at which players must act on them—period $(2k + 1)\hat{K} + \ell$ (resp. period $(2k + 2)\hat{K} + \ell$). However, we do not know whether a folk theorem would obtain if lags of message transmission are possibly unbounded.²⁵

²⁵In any event, as motivated in footnote 6, it seems reasonable to assume that message delays are much shorter than signal lags.

2.6 Discussion and Conclusion

As we argued in the introduction, the key role of the repeated games model makes it important to understand which of its many simplifications are essential for the folk theorem. We have extended this result to two settings in which players' information about others' play arrives with stochastic lags. In both of the settings we consider, there is a special but natural form of private information, as players do not know whether and when their opponents observe signals.

Our proof in the case of almost-perfect monitoring (and no communication) depends on the methods of HO2006. Unfortunately, our proof technique does not extend to repeated games with n players. We could attempt to classify any history containing the null signal as an erroneous history and follow the approach of HO2006 for n -player games, but this approach is invalid because of the HO2006 n -player proof's requirement of *communication phases*. For repeated games with observation lags having finite support (possibly including ∞), it may seem that the discussion in Remark 4 of HO2006 regarding almost-perfect monitoring private monitoring games with general signal spaces could be useful. This is due to the fact that as long as the lag distribution has finite support, we can take the K chosen in Lemma 2.4.6 to be sufficiently large so that each thread corresponds to a private monitoring game.²⁶ However the conjecture in Remark 4 of HO2006 regarding the partition of signals contains an error and thus cannot be applied.²⁷ Instead, we conjecture that the set of all belief-free equilibrium payoffs in n -player games without communication can be attained in the analogous games with lags. Using results from Yamamoto (2009), one could then obtain a lower bound on the limit equilibrium payoff sets for n -player repeated games with almost-perfect monitoring structures and observation lags.

²⁶Note that this is not the case if the lag distribution's support is not finite.

²⁷Specifically, Remark 4 suggests that one can find a partition of the private signals to restore the invertibility of the appropriate information matrix so that their results go through, with the elements of the partition treated as the set of private signals. However inference about others' private histories is different across different signals within the same element of the partition so that it is not clear whether the appropriate incentive compatibility conditions would hold.

A more substantial extension of our results would be to the case in which the lag distribution varies with the discount factor. It seems likely that our results would extend to settings in which longer lags become somewhat more likely as players become more patient, but we do not know how rapid an increase can be accommodated.

Chapter 3

Contagious Commitment via Unknown Patience

3.1 Introduction

The literature on reputation in repeated interactions demonstrates how the introduction of a small amount of uncertainty regarding the “rationality” of the long run player can generate large benefits.¹ Specifically whenever there is some small probability that the long run player is some behavioral Stackelberg type who always plays the Stackelberg action, then a sufficiently patient player can guarantee payoffs close to the Stackelberg payoff in all equilibria of the game.² This literature typically fixes the type space of player 1 and shows that a sufficiently patient long run rational player can guarantee payoffs arbitrarily close to the Stackelberg payoff in all equilibria. However when we instead fix the payoff function of the long run player (in particular fixing the patience of the long run player), if monitoring of the long run player’s action is noisy, then whenever the type space assigns sufficiently small probability to the behavioral types, there typically exist equilibria in which the long

¹See for example Kreps *et al.* (1982), Fudenberg and Levine (1989), and Mailath and Samuelson (2006).

²Recall that the Stackelberg action is the action that he would prefer to choose in the stage game if player 2 could observe player 1’s choice of action before choosing his action. The payoff in the corresponding strategy profile is then called the Stackelberg payoff.

run rational player obtains payoffs much lower than the Stackelberg payoff.³ In other words, reputation effects vanish in these models when the probability of the behavioral types is sufficiently small for a long run player with a fixed level of patience.

This paper demonstrates that if uncertainty over patience of the long run player is introduced along with the typical uncertainty regarding the rationality of the long run player strong reputation effects emerge even when the probability assigned to the behavioral types is arbitrarily small. We show that in a game where the long run player must choose an action that he must commit to for the entirety of the game, the introduction of uncertainty regarding patience together with uncertainty about the rationality of the long run player generates very strong reputation effects that do not depend on the probability of the behavioral types. We show more strongly that the required amount of uncertainty regarding the patience of the rational long run player necessary to generate these strong reputation effects is quite small.⁴

With uncertainty over patience, these strong reputation effects are the result of a contagion effect initiated by very patient types. Whenever there is some positive probability that the long run player is a behavioral Stackelberg type, regardless of the size of this probability, some positive mass of the most patient types have a strict incentive to play the Stackelberg action in all equilibria. This then reinforces the reputation effect by effectively increasing the probability of types that play the Stackelberg action in all equilibria. This increase in turn incentivizes those less patient types to also play the Stackelberg action, leading to a contagion of types who play the Stackelberg action in all equilibria. Our main result shows that even when very patient types are extremely small in probability, this contagion effect is very strong so that types with intermediate levels of patience (bounded away from no discounting) obtain high payoffs in all equilibria even when the probability of the behavioral types becomes arbitrarily small.

To prove our main theorem, we show that there must exist some $x^* > 0$ such that

³See Section 3.5 for a precise statement.

⁴See Assumption 3.4.1 for a precise statement of the requirements imposed on the distribution over patience.

whenever the type space assigns positive probability to a Stackelberg commitment type, the probability with which the long run player plays the Stackelberg action must be at least x^* in all equilibria.⁵ Consider the hypothetical scenario in which the probability that the long run player plays the Stackelberg action is given by $x > 0$. Now given such an x , using the arguments from Fudenberg and Levine (1989), we can show that the most patient types will strictly prefer to play the Stackelberg action. Thus there exists some lower bound X such that at least a mass $X(x) > 0$ must strictly prefer to play the Stackelberg action given x . Thus an equilibrium must be such that $x \geq X(x)$. If X is such that there exists some x^* so that for all $x \in (0, x^*)$, $x < X(x)$ then we have shown that any equilibrium with $x > 0$ must also satisfy $x > x^*$, which is the desired conclusion. Our proof technique will establish that X must have this property under suitable conditions on the type space.

The reason why a very weak condition such as that imposed in Assumption 3.4.1 suffices for the reputation theorem is due to the fact that the benefits of reputation erode very slowly as x decreases. More precisely, the lower bound (derived by Gossner (2011)) to playing the Stackelberg action for a rational type with discount factor δ takes the form $w(-(1-\delta)\log(x))$ for some continuous function w .⁶ Importantly $-\log(x)$ increases to infinity as $x \rightarrow 0$ at a very slow rate and so for a rational type with discount factor δ , the value to playing the Stackelberg action does not fall very rapidly as x approaches 0.⁷ This means that even for small values of x , a large mass of rational types are still willing to play the Stackelberg action in all equilibria, providing a strong kick-start to the contagion effect highlighted above. In contrast, if the lower bound on the payoff to playing the Stackelberg action took the form $w(-(1-\delta)x^{-1})$, then Assumption 3.4.1 would no longer suffice for a result in the spirit of our main theorem. Thus the fact that the lower bound takes the particular form derived in Gossner (2011) is crucial for the main theorem.

In Section 3.6, we explicitly calculate the exact size of the contagion effect and compute

⁵Note importantly that x^* does not depend on the probability the type space assigns to the Stackelberg type.

⁶The fact that the lower bound can be expressed as a function of $-(1-\delta)\log(x)$ can also be derived from Fudenberg and Levine (1992).

⁷In fact for any $\rho > 0$, there exists some x^* such that whenever $x < x^*$, $x^{-\rho} > -\log(x)$.

the payoffs that each type gets in equilibria as the size of the Stackelberg type becomes very small. We perform these calculations in a continuous time model in which there is one behavioral type who always plays the Stackelberg action and a uniform distribution of patience levels of the rational types. Consistent with the predictions of Section 3.4, we show that as the probability of the behavioral type converges to zero, the mass of types who play the Stackelberg action converges to large masses. We show that this limit is unique and that even moderately patient types obtain very high payoffs even in the limit.

3.1.1 Literature Review

As discussed in the introduction, this paper contributes to the literature that studies reputation effects in repeated games, which began with Kreps *et al.* (1982) and Fudenberg and Levine (1989). Fudenberg and Levine (1992) extended their results to settings in which the long run player's actions are observed imperfectly. These papers show that when some uncertainty about the rationality of the long run player is introduced, a sufficiently patient player can guarantee payoffs close to the Stackelberg payoff in all equilibria of the game.

In contrast, Faingold and Sannikov (2011) study reputation effects in a continuous time model when the discount factor of the long run player is not necessarily close to one. As a result, the players do not necessarily obtain payoffs close to the Stackelberg payoff. Their continuous time approach allows them to analyze reputation effects of impatient long run players through the study of ordinary differential equations. As a result their model can study the evolution of beliefs over time in equilibrium that the standard discrete time model cannot.

Methodologically, the techniques in this paper rely on bounds to reputation that are developed in Gossner (2011), which improve on the payoff bounds obtained in Fudenberg and Levine (1992). These bounds prove useful for our purposes as it allows us to obtain sufficient estimates on the mass of types who have strict incentives to play certain actions. Similarly Faingold (2013) illustrates the usefulness of Gossner's bounds in studying reputation effects when the interactions between the long run and short run player become very

frequent.

There is also a string of recent papers that model reputation effects without behavioral commitment types. For example, Board and Meyer-ter Vehn (2013) study an alternative reputation model where the actions of the long run player have long-lasting but transitory effect on “quality”. As a result, the long run player has an incentive to provide effort to improve the quality of his product, yielding reputation-like effects. Relatedly Dilmé (2012) study a model that also generates rich reputation dynamics in a model where the long run player faces costs to switching their actions. Additionally Bohren (2011) studies more general stochastic games where actions have a persistent effect on an evolving state variable. This influence of the long run player’s action on the state variable then gives rise to effects that resemble the traditional reputation effects. Finally Weinstein and Yildiz (2012) shows that in finitely repeated games, any arbitrary commitment type that is programmed to play a certain strategy in the can be constructed in a standard finitely repeated game using only incomplete information about the stage game payoffs. By introducing higher order uncertainty about the stage game payoffs of the long run player, Weinstein and Yildiz (2012) construct types whose unique rationalizable action in the incomplete information game is to play the strategy that the commitment type is programmed to play.

Finally some recent papers study reputation effects in models that relax some of the restrictive assumptions that the standard reputation models imposes. Liu and Skrzypacz (2014) study a model in which the short run players can only observed a limit number of observations regarding the reputation builder’s action. Liu (2011) studies a similar reputation model where the short run player must pay a cost to acquire information about the long run player’s past chosen actions. As a result, both of these papers exhibit reputation dynamics where play switches between phases of reputation building and reputation exploitation and spending. Finally Jehiel and Samuelson (2012) study reputation building in a model with short run players who form beliefs about the long run player’s intended course of play according to a simpler rule than that required in sequential equilibrium. As a consequence, the long run player can guarantee payoffs that are strictly higher than the

Stackelberg payoff.

The remainder of the paper is organized as follows. Section 3.2 describes the model. Section 3.3 proves the existence of a Nash equilibrium of the described game so that our results are not vacuous. Section 3.5 illustrates the necessity of both behavioral commitment types and arbitrarily patient rational types for reputation results. Section 3.4 presents our main reputation theorem and its proof illustrating how uncertainty about discount factors help generate strong reputation effects even for small probabilities of the commitment types. Section 3.6 performs some numerical computations in a continuous time modification of the main game to illustrate the exact size of the contagion effect. Finally Section 3.7 concludes.

3.2 Model

There are two players $i = 1, 2$, each with a finite action space A_i . Player 1 moves only once at time 0 and picks an action $a_1 \in A_1$.⁸ There is an infinite sequence of short run player 2's who each picks an action $a_2 \in A_2$ at times $t = 0, 1, 2, \dots$. At the end of each period, the players observe a stochastic outcome $y \in Y$ which is drawn independently and identically from a finite set Y according to the probability density function $\pi(\cdot | a_1) \in \Delta(Y)$ where a_1 is the action chosen by player 1 in period 0.

Note importantly that the public signal distribution $\pi(\cdot | a_1)$ potentially depends on the action a_1 . This dependence is used by player 2 to make inferences about the actually chosen action a_1 . We impose the following standard assumption on the public signal structure.

Assumption 3.2.1. For all $y \in Y$ and all $a_1 \in A_1$, $\pi(y | a_1) > 0$. Furthermore $(\pi(\cdot | a_1))_{a_1 \in A_1}$ forms a matrix that has full row rank.

The first part of the assumption states that all signals are possible regardless of the action a_1 . The second part assumes that the long run player's action is statistically identifiable so that eventually if enough observations of Y are observed, player 2 will learn the true

⁸This is a significant departure from the classical reputation literature where both players choose actions at each time $t = 0, 1, 2, \dots$. As discussed in the introduction, we want to study how uncertainty regarding patience can alleviate the cost of influencing the short run players' beliefs about his own action.

action. Note that Assumption 3.2.1 holds generically if we regard the public signal structure $(\pi(\cdot | a_1))_{a_1 \in A_1}$ as a vector in $\mathbb{R}^{|A_1| \cdot |Y|}$.

We now introduce two sources of private information regarding player 1: uncertainty over *rationality* and *patience*. The first is standard in the literature on reputation building but the second is new. To model these two sources of uncertainty, we construct a type space (Ω, μ) where each $\omega \in \Omega$ represents a type of player 1. In our model we assume that there is no uncertainty regarding the type of player 2 and thus there are no types for player 2.

We define a type space to be a pair (Ω, μ) of a measurable space Ω together with a probability measure μ on Ω . In this paper we consider a specific form for Ω : we partition Ω into two sets Ω_r and Ω_b (so that $\Omega = \Omega_r \cup \Omega_b$) where Ω_r denotes the set of all rational types and Ω_b denotes the set of all behavioral types. Each $\omega \in \Omega_b$ is associated with a pure strategy $a_1 \in A_1$ and is programmed to always play his associated strategy at time 0. Let us denote the type $\omega \in \Omega_b$ that corresponds to the action a_1 as ω_{a_1} . The set of rational types Ω_r is defined on the interval $[0, 1)$. We interpret a type in Ω_r as $\delta \in [0, 1)$ where δ represents the discount factor of the type. We let μ_r denote the conditional distribution over Ω_r and assume that this probability measure admits a density f in $L^1([0, 1])$ with cumulative density function F .

As a concrete example of an economic setting that we have in mind, consider a firm servicing a group of customers that needs to decide on the hire of one of two candidates to run its operations. The firm is able to perfectly verify the abilities of the two candidates however this information is not available to its consumers. The better candidate is able to produce better products attracting more demand from the consumers. However the better candidate comes at a cost as the firm must pay him higher wages. Because information about the ability level of the hired candidate arrives to the customers imperfectly, the firm may not be able to convince the market immediately that he has hired the better candidate.

Having described the type space, let us define the strategies of the players. A pure strategy for player 1 is a measurable map s_1 from $\Omega_r \rightarrow A_1$. Note that such a map can be identified with a function \hat{s}_1 in $L^\infty([0, 1]^{m-1})$ where each coordinate takes a value in the

binary set $\{0,1\}$ and $\hat{s}_1^\ell(\delta) = 1$ if and only if $\delta \in \Omega_r$ and $s_1(\delta) = a_1^\ell$. Let S_1 be the set of pure strategies in this space. In the Appendix we show that S_1 can be endowed with the appropriate topology to make it a nonempty, compact, metric space. We then define the set of player 1 mixed strategies Σ_1 as the set of Borel probability measures over S_1 . With a slight abuse of notation, given any $\sigma_1 \in \Sigma_1$, $a_1 \in A_1$, and a type space (Ω, μ) , we define $\sigma_1[a_1]$ as the total probability that player 1 plays a_1 :

$$\sigma_1 [a_1^\ell] = \mu [\omega_{a_1^\ell}] + \mu[\Omega_r] \int_{s_1 \in S_1} \int_{\delta \in \Omega_r} s_1^\ell(\delta) dF(\delta) d\sigma_1[s_1].$$

Similarly define $\sigma_1^r [a_1^\ell] \equiv \sigma_1 [a_1^\ell] - \mu [\omega_{a_1^\ell}]$ to be the total probability that a rational type plays a_1^ℓ .

To define player 2's strategy, we first need to define histories. Let H^t be the set of t -period public histories and define $H = \bigcup_{t=0}^{\infty} H^t$. Then we define player 2's strategy to be a map $\sigma_2 : H \rightarrow \Delta(A_2)$. Note that we restrict to strategies of player 2 that are public and do not depend on player 2's private history.⁹

Let us now specify the payoff functions of the players. Player 1 derives ex-ante utility of

$$V(\sigma_1, \sigma_2) = \int_{S_1} \int_0^1 \mathbb{E}_{s_1(\delta)} [U_1(s_1(\delta), \sigma_2, \delta)] dF(\delta) d\sigma_1[s_1]$$

from playing a mixed strategy $\sigma_1 \in \Sigma_1$ against $\sigma_2 \in \Sigma_2$ where

$$\mathbb{E}_{a_1} U_1(a_1, \sigma_2, \delta) = \mathbb{E}_{a_1} (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_1(a_1, \sigma_2^t(h^t)).$$

Player 1's objective is to maximize $V(\sigma_1, \sigma_2)$ against σ_2 . This means that $V(s_1, \sigma_2) \geq V(s_1', \sigma_2)$ for all $s_1' \in S_1$ and all $s_1 \in S_1$ σ_1 -almost everywhere. Note that this in turn means that for σ_1 -almost every $s_1 \in S_1$,

$$\mathbb{E}_{s_1(\delta)} U_1(s_1(\delta), \sigma_2, \delta) = \max_{a_1' \in A_1} \mathbb{E}_{a_1'} U_1(a_1', \sigma_2, \delta)$$

at F -almost every $\delta \in [0, 1]$.

⁹This is without loss of generality for equilibrium since player 1's action is only chosen in period 0 and thus does not change based on the history.

Player 2 simply maximizes expected payoffs at every history $h \in H$ conditional on available information:

$$\sigma_2(h) \in \arg \max_{\sigma'_2} \mathbb{E}_{\sigma_1} [u_2(a_1, \sigma'_2) \mid h].$$

In other words player 2 is completely myopic and only cares about how $\sigma_2(h)$ performs against his beliefs about player 1's played action in the current period.¹⁰ The solution concept we use is Nash equilibrium (σ_1, σ_2) where σ_1 is an ex-ante best response against σ_2 according to the payoff function V and σ_2 is a best response at every history $h \in H$ according to beliefs consistent with Bayesian updating.¹¹

Finally we impose the following mild assumptions on the stage game utility functions u_1 and u_2 :

Assumption 3.2.2. Given any pure strategy $a_1 \in A_1$ of player 1, player 2 has a unique strict best response which we denote $B(a_1)$.

Assumption 3.2.3. Given any two distinct actions $a_1, a'_1 \in A_1$, $u_1(a_1, a_2(a_1)) \neq u_1(a'_1, a_2(a'_1))$.

Note that these assumptions hold for generic payoff functions. Let us also extend B to the space of mixed actions $\alpha_1 \in \Delta(A_1)$ by defining $B(\alpha_1)$ as the set of best responses in the stage game for player 2 against α_1 . Let $m = |A_1|$ and without loss of generality, order the actions of player 1 so that

$$u_1(a_1^1, B(a_1^1)) > u_1(a_1^2, B(a_1^2)) > \dots > u_1(a_1^m, B(a_1^m)).$$

We shorten notation further by defining $a_2^k = B(a_1^k)$ and $u_1(a_1^k, a_2^k) = v_1^k$. Recall in the literature that a_1^1 is commonly called the Stackelberg action of the stage game. As a final piece of notational simplification, define $\omega_S = \omega_{a_1^1}$.

¹⁰This is without loss of generality since even if player 1 faced a single non-myopic player 2, player 2 would have no incentive to play a non-myopic best response at any history because his actions do not influence the signal process.

¹¹Note that because the public signal structure has full support, Bayes' rule pins down unique beliefs at all public histories.

3.3 Existence of Equilibrium

Due to the fact that the type space is a continuum, there is no pre-existing general existence theorem that we can immediately apply here and we must explicitly prove the existence of an equilibrium. We first find that when a_1^1 is a Nash equilibrium of the stage game, an equilibrium trivially exists if the behavioral type probability is sufficiently small. This implies the following:

Theorem 3.3.1. *Suppose that a_1^1 is a Nash equilibrium of the stage game. Then there exists some $v^* > 0$ such that if $\mu[\Omega_b] < v^*$, then there exists a Nash equilibrium in which all rational types play a_1^1 .*

Obviously there are two limitations to the statement above. First a_1^1 may not be a Nash equilibrium of the stage game and secondly we may be interested in analyzing games in which the probability of a behavioral type is not necessarily small. Demonstrating existence in such games requires a different proof technique as the construction of trivial equilibria is no longer possible. For example consider the strategy profile in which all players play a_1^1 with probability one. This cannot be an equilibrium of the repeated game if $\mu[\Omega_b]$ is very small and a_1^1 is not a Nash equilibrium. The reason is that because almost all types are playing a_1^1 (with the exception of the behavioral types), player 2 places large probability on player 1 playing a_1^1 for a long time regardless of the history of realized public signals. However the less patient types of player 1 would then have an incentive to cheat because player 2's posterior does not decrease until very late in the game even if he were to play a non-Stackelberg action. Therefore in any equilibrium, at least a positive fraction of the types must necessarily play an action besides a_1^S with positive probability when $\mu[\Omega_b]$ is sufficiently small. Nevertheless, we can still prove an existence theorem using non-constructive methods.

Theorem 3.3.2. *Suppose μ_r is absolutely continuous with respect to the Lebesgue measure. Then a Nash equilibrium exists.*

Proof. See Appendix. □

3.4 Reputation and Contagion

We now show that when even a moderate amount of uncertainty regarding patience (in a sense to be made precise) exists, a non-vanishing fraction of rational types must play a_1^1 with probability 1 even when $\mu[\omega_S]$ vanishes to zero. Thus uncertainty regarding patience interacts with uncertainty about the rationality of the long run player to substantially lower the cost of convincing the short run player of his intended action, which in turn guarantees payoffs close to what one would obtain under perfect monitoring for even a moderately patient long run player. We will show in Section 3.5 Theorem 3.5.2 that some restrictions on μ_r are necessary in order to prove such a theorem. In particular, Theorem 3.4.2 requires at the very least that types of arbitrarily high levels of patience exist with positive probability. The condition that we impose however is arguably slightly stronger and is summarized in the following assumption.

Assumption 3.4.1. μ_r is an absolutely continuous measure with respect to the Lebesgue measure on $[0, 1)$ and let F be its cumulative distribution function (cdf). There exists some $\underline{\delta} < 1$ and some $k \geq 1$ such that F is continuously k -times differentiable at all points $\delta \geq \underline{\delta}$ and

$$\lim_{\delta \rightarrow 1} D^k F(\delta) \neq 0.$$

In words, the assumption states that the density does not vanish too rapidly near $\delta = 1$. This precludes for example the distributions studied in Section 3.5 where $F(\delta^*) = 1$ for some $\delta^* < 1$. However note that the assumption does not require the density of arbitrarily patient types to be non-vanishing. A trivial example of a probability measure that satisfies this assumption is the uniform distribution over $[0, 1)$. Similarly any distribution whose density converges to some strictly positive number as $\delta \rightarrow 1$ also satisfies the assumption. Moreover any cdf that has *any* derivative that converges to a number other than zero as $\delta \rightarrow 1$ also satisfies the assumption. Thus the assumption is quite weak.¹²

¹²In fact if F can be extended to a function \tilde{F} that is analytic on an open set $H \supseteq (0, 1]$, then if F violates Assumption 3.4.1, F must be identically zero on H which is a contradiction of the fact that F is a cdf. Therefore if F has an analytic extension to an open set $H \supseteq (0, 1]$, then F must satisfy Assumption 3.4.1.

With this let us now state the main theorem. Due to our negative results from the previous section, we must impose further restrictions on the type space: $\mu[\omega^S] > 0$ and $\mu[\Omega_r]$ is bounded away from zero. For a given $\rho > 0$, let us define the set \mathcal{F}_ρ in the following manner:

$$\mathcal{F}_\rho \equiv \{(\Omega, \mu) : \mu[\omega^S] > 0, \mu[\Omega_r] > \rho\}.$$

Note importantly that \mathcal{F} places no restrictions on the size of $\mu[\omega_S]$ other than it being positive. Finally given a cumulative distribution function F over $\Omega_r = [0, 1)$, let $\mathcal{F}_\rho(F)$ be the set of elements of \mathcal{F}_ρ such that the conditional distribution over rational types coincides precisely with F :

$$\mathcal{F}_\rho(F) \equiv \{(\Omega, \mu) \in \mathcal{F}_\rho : \mu_r = F\}.$$

Theorem 3.4.2. *Suppose that F satisfies Assumption 3.4.1 and let $\rho > 0$. Then under Assumptions 3.2.1, 3.2.2, and 3.2.3, for every $\varepsilon > 0$, there exists some δ^* such that all types with $\delta \geq \delta^*$ obtain a payoff of at least $v_1^1 - \varepsilon$ in every Nash equilibrium of $G(\Omega, \mu)$ for all type spaces $(\Omega, \mu) \in \mathcal{F}_\rho(F)$.*

Importantly note the order of quantifiers. The classical reputation results allow the threshold δ^* to vary across type spaces (Ω, μ) . What distinguishes this theorem is that δ^* holds uniformly across all type spaces in $\mathcal{F}_\rho(F)$. Thus δ^* depends only on F and ρ and not on the specific manner in which μ weights the elements of Ω_b versus Ω_r . To prove the theorem, we show that regardless of type spaces in $\mathcal{F}_\rho(F)$, there exists some v^* such that $\sigma_1[a_1^1] > v^*$ in all Nash equilibria. Having established this argument, then the theorem above is an immediate application of FL modified to the game we analyze.

3.4.1 Proof Sketch

Let us first highlight the essence of the arguments used in proving Theorem 3.4.2 in the context of a game in which there are two *rational* types: one type ω_1 who does not discount payoffs and is a time-average payoff maximizer and another type ω_δ with discount factor $\delta \in (0, 1)$. Then it is easy to show that the time-average payoff maximizer must always choose the Stackelberg action in any Nash equilibrium as long as $\mu[\omega_S] > 0$ since playing

a_1^1 results in exactly the Stackelberg payoff for this type whereas any other action yields a strictly lower payoff.

Then the fact that ω_1 plays a_1^1 in all Nash equilibria reinforces reputation effects for type ω_δ of lower patience even when $\mu[\omega_S]$ is small since there is now a measure $\mu[\omega_S] + \mu[\omega_1]$ of types who play a_1^1 in all equilibria. Regardless of the size of $\mu[\omega_S]$, there must always be a measure $\mu[\omega_1]$ of types who play a_1^1 in all equilibria as long as $\mu[\omega_S]$ is positive. As a consequence if $\mu[\omega_S] > 0$, $\varepsilon > 0$, and $\eta > 0$, there exists some δ^* such that if $\delta > \delta^*$, then type ω_δ is able to guarantee himself a payoff of at least $v_1^1 - \varepsilon$ in all equilibria and all type spaces with $\mu[\omega_1] > \eta$.

The argument above relies on a very restrictive assumption: the type space places strictly positive probability on a type (the ω_1 type) whose uniquely rationalizable action is the Stackelberg action in any type space in which $\mu[\omega_S] > 0$. The remainder of this section will illustrate how this simple argument can indeed be extended to quite general type spaces in which such a type ω_1 may not exist.

We illustrate heuristically how the proof can indeed be extended. The intuition can be seen most clearly in the two action case and so suppose that A_1 consists of the Stackelberg action a_1^1 and an action $a_1^2 \neq a_1^1$. Consider a type space (Ω, μ) with $\mu[\omega_S] > 0$ and $\Omega = \Omega_r \cup \{\omega_S\}$. Let σ be a Nash equilibrium. Note that if $\sigma_1[a_1^1] > 1/2$, then reputation effects are already strong and so let us assume that $\sigma_1[a_1^1] \leq 1/2$. This then means that $\sigma_1[a_1^2] \geq 1/2$.

Now choose $\varepsilon > 0$ such that $v_1^1 - \varepsilon > v_1^2 + \varepsilon$. Then with the use of upper bounds to payoffs obtained in Fudenberg and Levine (1992) or Gossner (2011), we can find $\delta^* < 1$ such that the most that any type $\delta > \delta^*$ can obtain from playing a_1^2 is $v_1^2 + \varepsilon$. Importantly note that this δ^* holds uniformly across all σ_1 with the property that $\sigma_1[a_1^1] \leq 1/2$ and is in particular independent of the size of $\mu[\omega^S]$.

At the same time because $\sigma_1[a_1^1] \geq \mu[\omega^S] > 0$, we can also find $\delta_0 \geq \delta^*$ such that all types $\delta > \delta_0$ obtains at least $v^1 - \varepsilon = u_1(a_1^1, a_1^2) - \varepsilon$ from playing a_1^1 .¹³ This then implies that

¹³Unlike δ^* , note that δ_0 cannot be taken to be independent of neither $\mu[\omega^S]$ nor σ_1 .

for any type $\delta > \delta_0$, $\sigma_1(\delta)$ must assign probability one to a_1^1 and therefore the probability which σ_1 assigns to a_1^1 must at the very least be:

$$\sigma_1[a_1^1] \geq X_0 \equiv \mu_r[(\delta_0, 1)]\mu[\Omega_r] + \mu[\omega_S].$$

Now note that because δ_0 depends on $\mu[\omega_S]$ (and may a priori converge to one as $\mu[\omega^S] \rightarrow 0$) the argument above is insufficient for establishing a uniform lower bound on $\sigma_1[a_1^1]$ across all equilibria and all type spaces with $\mu[\omega_S] > 0$.

However we can iterate this argument using the new lower bound X_0 on the probability with which a_1^1 must be played. More generally, given a probability $\sigma_1[a_1^1] > 0$, using the same arguments above, we can find a decreasing function $\hat{\delta}(\sigma_1[a_1^1])$ (as a function of $\sigma_1[a_1^1]$) such that all types $\delta > \hat{\delta}(\sigma_1[a_1^1])$ must play a_1^1 with probability one. This then implies that for all σ_1 such that $\sigma_1[a_1^1] \leq 1/2$,

$$\sigma_1[a_1^1] \geq \mu_r[(\hat{\delta}(\sigma_1[a_1^1]), 1)]\mu[\Omega_r] + \mu[\omega_S] \geq X(\sigma_1[a_1^1]) \equiv \mu_r[(\hat{\delta}(\sigma_1[a_1^1]), 1)]\mu[\Omega_r]. \quad (3.1)$$

This function X , under fairly mild conditions on the measure μ_r , turns out to possess the nice property that the derivative near zero is strictly bigger than 1.

As evidenced by Figure 3.1, this property implies that X must lie strictly above the 45-degree line in a neighborhood around zero. Furthermore because $\mu[\omega_S] > 0$ and consequently $\sigma_1[a_1^1] > 0$, inequality (3.1) illustrates that any equilibrium must have the property that $\sigma_1[a_1^1]$ lies strictly to the right of the point at which F crosses the 45-degree line. This therefore shows that no matter how small $\mu[\omega^S]$ is, as long as it is positive, $\sigma_1[a_1^1]$ is at least x^* . At this point, one may wonder how we can conclude the particular shape of the function X that lies at the heart of the conclusion just established. Broadly speaking, this observation is a consequence of the fact that the value of a reputation effect emerging from the existence of ω_S erodes very slowly as $\mu[\omega_S]$.¹⁴ We now shed light on this fact in the next subsection.

¹⁴In particular see Lemmata 3.4.3 and 3.4.14 as well as Theorem 3.4.7 for the details.




Figure 3.1: Lower Bound on $\sigma_1[a_1^1]$

3.4.2 Details of the Proof

Mathematical Preliminaries

Having illustrated the main ideas of the proof technique, we proceed to the proof in more detail. Given any $\varepsilon > 0$, define the function $G_\nu : (0, 1] \rightarrow \mathbb{R}$:

$$G_\nu(x) = \left(1 - F \left(1 + \frac{\nu}{\log x} \right) \right).$$

Before relating our discussion to the derivation of our key lemma, we first note a purely mathematical lemma that will prove useful for the estimates we wish to obtain.

Lemma 3.4.3. *The following two statements hold.*

1. Let $p : \mathbb{R} \rightarrow \mathbb{R}$ be a polynomial over the reals. Then $xp(\log x) \rightarrow 0$ as $x \rightarrow 0$.
2. Suppose h is some function $h : (-\infty, 0] \rightarrow \mathbb{R}$ such that there exists $k \geq 1$ such that

$\lim_{x \rightarrow 0} D^k h(x) \neq 0$. Let $q(x) = h((\log x)^{-1})$. Then $\lim_{\mu \rightarrow 0} Dq(\mu) = +\infty$.

Proof. We now prove the first statement. To prove the statement, it is sufficient to show that $\lim_{x \rightarrow 0} x(\log x)^k = 0$ for all $k \geq 0$. This is trivial for $k = 0$. Suppose that the statement holds for $k - 1$. Then using L'Hospital's rule,

$$\lim_{x \rightarrow 0} x(\log x)^k = \lim_{x \rightarrow 0} \frac{(\log x)^k}{1/x} = \lim_{x \rightarrow 0} -\frac{k(\log x)^{k-1} x^{-1}}{x^{-2}} = \lim_{x \rightarrow 0} -k(\log x)^{k-1} x = 0.$$

This proves the first claim.

Then the second claim is a consequence of the first.

$$Dq(x) = -\frac{Dh((\log x)^{-1})}{x(\log x)^2}$$

Let k^* be the minimum positive integer $k > 0$ for which $\lim_{x \rightarrow 0} D^k h(x) \neq 0$. If $k^* = 1$, then clearly $\lim_{x \rightarrow 0} Dq(x) = +\infty$ since $\lim_{x \rightarrow 0} x(\log x)^2 = 0$.

To prove the result for $k^* > 1$, we show that if $k^* \geq k \geq 1$, then

$$\lim_{x \rightarrow 0} Dq(x) = \lim_{x \rightarrow 0} (-1)^k \frac{D^k h((\log x)^{-1})}{xp(\log x)}$$

for some $p : \mathbb{R} \rightarrow \mathbb{R}$ a polynomial. Clearly this holds for $k = 1$. Now suppose it holds for $k < k^*$. Then by L'Hospital's rule, we have:

$$\begin{aligned} \lim_{x \rightarrow 0} Dq(x) &= \lim_{x \rightarrow 0} (-1)^k \frac{D^k h((\log x)^{-1})}{xp(\log x)} \\ &= \lim_{x \rightarrow 0} (-1)^{k+1} \frac{D^{k+1} h((\log x)^{-1})}{p(\log x) + Dp(\log x)} \frac{1}{x(\log x)^2} \\ &= \lim_{x \rightarrow 0} (-1)^{k+1} \frac{D^{k+1} h((\log x)^{-1})}{x(\log x)^2 (p(\log x) + Dp(\log x))} \end{aligned}$$

Then note that $(\log x)^2 (p(\log x) + Dp(\log x))$ is a polynomial in $\log x$. Thus this completes the induction.

As a consequence, we have:

$$\lim_{x \rightarrow 0} Dq(x) = \lim_{x \rightarrow 0} (-1)^{k^*} \frac{D^{k^*} h((\log x)^{-1})}{xp(\log x)} = +\infty,$$

where the last equality uses statement 1 of the lemma. This concludes the proof. \square

Remark. The role that $\log x$ plays will become apparent once we begin to study how beliefs are updated by the short run player. However the above lemma illustrates a remarkable property of the function $(\log x)^{-1}$ near $x = 0$, which is essential in guaranteeing that the contagion effect is strong enough to ensure that reputation effects persist even when the probability of behavioral types becomes arbitrarily small.

Now with the help of Lemma 3.4.3, Assumption 3.4.1 directly translates to a useful mathematical property regarding the function G_ν , which we summarize in the following corollary.

Corollary 3.4.4. *Suppose that Assumption 3.4.1 holds and let $\nu > 0$. Then*

$$\lim_{x \rightarrow 0} DG_\nu(x) = +\infty.$$

Relating this property back to the arguments in Section 3.4.1 outlining the proof of Theorem 3.4.2, Corollary 3.4.4 will be used to show that the lower bound function F takes the shape illustrated in Figure 3.1. The fact that the curve F in Figure 3.1 is very steep near 0 in the figure is not a mere coincidence and rather corresponds exactly to the property demonstrated above that the function $G_{\bar{c}, \varepsilon}$ has infinite derivative at exactly zero.¹⁵

Proof of Corollary 3.4.4: The proof is a direct consequence of Lemma 3.4.3. Define the function: $\beta(x) = 1 + \nu x$. Then we can rewrite G_ν :

$$G_\nu(x) = 1 - F(\beta((\log x)^{-1})).$$

Now define $h(x) = 1 - F(\beta(x))$. By Lemma 3.4.3 it is sufficient to show that there exists some $k \geq 1$ such that $\lim_{x \rightarrow 0} D^k h(x) \neq 0$.

Then note that

$$D^k h(x) = -\nu^k D^k F(\beta(x)).$$

¹⁵Note that an infinite slope near 0 is not necessary for the arguments. Rather it is sufficient that the slope near 0 is strictly greater than 1. Nevertheless the strong conclusion of Corollary 3.4.4 is the result of a very mild assumption.

Thus

$$\lim_{x \rightarrow 0} D^k h(x) \neq 0 \Leftrightarrow \lim_{x \rightarrow 0} D^k F(\beta(x)) \neq 0.$$

But Assumption 3.4.1 implies that there exists $k \geq 1$ such that $\lim_{x \rightarrow 0} D^k F(\beta(x)) \neq 0$. \square

A Review of Gossner (2011)

To establish that the function G defined above is indeed a lower bound on $\sigma_1[a^S]$ we first review some tools developed in the classical reputation literature on upper and lower bounds on equilibrium payoffs. To derive these upper and lower bounds we specifically use the tools developed and studied by Gossner (2011). Recall the following definitions. Given any two measures $p, q \in \Delta(Y)$, the relative entropy is defined as

$$d(p \parallel q) = \sum_{y \in Y} p(y) \log \left(\frac{p(y)}{q(y)} \right).$$

Note that in this paper due to Assumption 3.2.1, the measures of interest always have full support on Y and so $d(p \parallel q)$ is always well-defined.

Definition 3.4.5. $\alpha_2 \in S_2$ is an ε -entropy-confirming best response to $\alpha_1 \in S_1$ if there exists α'_1 such that the following conditions hold:

- α_2 is a best response to α'_1 .
- $d(\pi(\cdot \mid \alpha_1, \alpha_2) \parallel \pi(\cdot \mid \alpha'_1, \alpha_2)) \leq \varepsilon$.

$B_{\alpha_1}^d(\varepsilon)$ denotes the set of ε -entropy-confirming best responses to α_1 .

We now define the following functions.

$$\underline{v}_{\alpha_1}(\varepsilon) = \min_{\alpha_2 \in B_{\alpha_1}^d(\varepsilon)} u_1(\alpha_1, \alpha_2).$$

Similarly define

$$\bar{v}_{\alpha_1}(\varepsilon) = \max_{\alpha_2 \in B_{\alpha_1}^d(\varepsilon)} u_1(\alpha_1, \alpha_2).$$

We denote the supremum of all convex functions that lie below \underline{v}_{α_1} as \underline{w}_{α_1} and the infimum of all concave functions above \bar{v}_{α_1} as \bar{w}_{α_1} .¹⁶

Finally recall the following facts from Gossner (2011).

- \underline{v}_{α_1} and \underline{w}_{α_1} are nonincreasing and both continuous at 0.
- \underline{w}_{α_1} and \bar{w}_{α_1} are nondecreasing and both continuous at 0.
- $\underline{v}_{\alpha_1}(0) = \underline{w}_{\alpha_1}(0)$ and $\underline{w}_{\alpha_1}(0) = \bar{w}_{\alpha_1}(0)$.

With these facts and statistical identifiability of player 1's action due to Assumption 3.2.1, we immediately obtain the following lemma.

Lemma 3.4.6. *Let $\eta > 0$. Then there exists some $v^* > 0$ such that*

$$\begin{aligned} |\underline{w}_{a_1}(v) - u_1(a_1, a_2(a_1))| &< \eta, \\ |\bar{w}_{a_1}(v) - u_1(a_1, a_2(a_1))| &< \eta \end{aligned}$$

for all $v \leq v^*$ and all $a_1 \in A_1$.

With all of these definitions, we have the following payoff bound due to Gossner (2011) adapted to the specific game we analyze.

Theorem 3.4.7. *Let σ be any Nash equilibrium of the Bayesian game. Then given any $a_1 \in A_1$ such that $\sigma_1[a_1] > 0$,*

$$\underline{w}_{a_1}(- (1 - \delta) \log[\sigma_1[a_1]]) \leq U_1(a_1, \sigma_2) \leq \bar{w}_{a_1}(- (1 - \delta) \log[\sigma_1[a_1]]).$$

The following section will now use these payoff bounds and the property of the function G_v illustrated in Lemma 3.4.4 to prove Theorem 3.4.2.

3.4.3 Concluding the Proof of Theorem 3.4.2

Because of Gossner (2011) and Theorem 3.4.7, the following lemma immediately implies the main theorem and the remainder of the section will be devoted to proving it.

¹⁶Technically, Gossner (2011) did not define the function \bar{v}_{α_1} nor \bar{w}_{α_1} but the properties of these functions are essentially the same as the function \bar{w} defined there.

Lemma 3.4.8. *Let $\rho > 0$ and suppose F satisfies Assumption 3.4.1. Then there exists some $v^* > 0$ such that $\sigma_1[a_1^1] > v^*$ for all $\sigma \in E(\Omega, \mu)$ and all $(\Omega, \mu) \in \mathcal{F}_\rho(F)$.*

Sequence of Games

Consider a sequence of games with type spaces given by (Ω, μ^n) where we fix the conditional probability measure on Ω_r with corresponding cumulative distribution function F . We assume that $\mu^n[\omega^S] \rightarrow 0$ but that $\mu^n[\omega^S] > 0$ for all n . We now study the properties of the limit of equilibria of this sequence of games. More formally, denote by $E(\Omega, \mu_n)$ to be the set of equilibrium strategy profiles of the game $G(\Omega, \mu_n)$. To further simplify the exposition, we introduce the following definition.

Definition 3.4.9. Given $\lambda \in \Delta(A_1)$, we say that a sequence $\{(\sigma^n, \mu^n)\}_n$ approximates λ if

1. $(\Omega, \mu^n) \in \mathcal{F}_\rho(F)$ for all n ,
2. $\sigma^n \in E(\Omega, \mu^n)$ for all n ,
3. $\mu^n[\omega^S] \rightarrow 0$,
4. and $\sigma_1^n[a] \rightarrow \lambda[a]$ for all $a \in A_1$.

Then let us define the following set of distributions over A_1 :

$$\Pi_\rho(F) \equiv \{\lambda \in \Delta(A_1) : \exists \text{ a sequence } \{(\sigma^n, \mu^n)\}_n \text{ that approximates } \lambda\}.$$

In words, the set $\Pi_\rho(F)$ denotes the set of distributions over actions that can be approximated by equilibrium distributions over player 1 actions of games that place small probability on the Stackelberg commitment type ω^S . Recall that when F is a point mass distribution at some $\delta^* < 1$, then λ that places probability one on a player 1 strict Nash action is indeed approximated by some sequence (σ^n, μ^n) .

To arrive at our main theorem, we study properties of the set $\Pi_\rho(F)$. Toward the end of proving Lemma 3.4.8, we show first that it is sufficient to demonstrate the existence of some $v > 0$ such that $\lambda[a^S] > v$ for all $\lambda \in \Pi_\rho(F)$. This greatly simplifies the analysis by allowing

us to study the limit of equilibria as the probability of the behavioral types vanish rather than studying the sets of equilibria of all type spaces in $\mathcal{F}_\rho(F)$, which is undoubtedly a very large space.

Lemma 3.4.10. *Suppose that there exists $\nu > 0$ such that $\lambda[a^S] > \nu$ for all $\lambda \in \Pi_\rho(F)$. Then there exists $\nu' > 0$ such that $\sigma_1[a^S] > \nu'$ for all $\sigma \in E(\Omega, \mu)$ for all $(\Omega, \mu) \in \mathcal{F}_\rho(F)$.*

Proof. Suppose such a $\nu' > 0$ did not exist. Then for every n , we can find some $(\Omega, \mu^n) \in \mathcal{F}_\rho(F)$ with $\sigma^n \in E(\Omega, \mu^n)$ with $\sigma_1^n[a^S] < \frac{1}{n}$. But then this implies that $\sigma_1^n[a^S] \rightarrow 0$ as $n \rightarrow \infty$ and in particular $\mu^n(\omega^S) \rightarrow 0$.

Because $[0, 1]^{m-1}$ is a compact subset of \mathbb{R}^m we can find a convergent subsequence $\sigma_1^{n_\ell}$. But then note that we have found a sequence such that $\sigma^{n_\ell} \in E(\Omega, \mu^{n_\ell})$ for all ℓ and $(\Omega, \mu^{n_\ell}) \in \mathcal{F}_\rho(F)$ such that

$$\sigma_1^{n_\ell}[a] \rightarrow \lambda[a] \quad \forall a \in A$$

for some $\lambda \in \Delta(A_1)$ and $\mu^{n_\ell}(\omega^S) \rightarrow 0$. This then implies that $\lambda \in \Pi_\rho(F)$. But by construction we also had $\lambda[a^S] = 0$. This then of course contradicts the main assumption of the lemma. \square

Properties of $\Pi_\rho(F)$

We now turn our attention to the study of the set $\Pi_\rho(F)$. With the help of Lemma 3.4.10, it is sufficient to prove that there exists some $\nu > 0$ such that $\lambda[a^S] > \nu$ for all $\lambda \in \Pi_\rho$.

Lemma 3.4.11. *$\Pi_\rho(F)$ is a closed set.*

Proof. The proof follows from a standard diagonalization argument. \square

An immediate corollary is the following observation which we will use to prove Theorem 3.4.2 by way of contradiction.

Corollary 3.4.12. *Suppose there exists no $\nu > 0$ such that $\lambda[a^S] > \nu$ for all $\lambda \in \Pi_\rho(F)$. Then there exists $\lambda' \in \Pi_\rho(F)$ such that $\lambda'[a^S] = 0$.*

Now define $\mathcal{L}(\lambda)$ as the support of actions given a distribution $\lambda \in \Pi_\rho(F)$:

$$\mathcal{L}(\lambda) = \{\ell = 1, \dots, m : \lambda[a_1^\ell] > 0\}.$$

Using this language due to the corollary just stated, it remains to show that $1 \in \mathcal{L}(\lambda)$ for all $\lambda \in \Pi_\rho(F)$. We first

Lemma 3.4.13. *Suppose that $\lambda \in \Pi_\rho(F)$ and let $\ell^* = \min \mathcal{L}(\lambda)$. Suppose further that (σ^n, μ^n) approximates λ . Then there exist n^* and δ^* such that $\sigma_1^n(\delta)[a_1^\ell] = 0$ for all $n > n^*$, $\delta > \delta^*$ and all $\ell \in \mathcal{L}(\lambda)$ such that $\ell > \ell^*$.*

Proof. We use Theorem 3.4.7. Fix an $\varepsilon > 0$ sufficiently small so that

$$v_1^{\ell^*+1} < v_1^{\ell^*} - \varepsilon < v_1^{\ell^*} + \varepsilon < v_1^{\ell^*-1}.$$

By Lemma 3.4.6 we can choose $\nu^* > 0$ such that

$$\begin{aligned} v_1^{\ell^*} + \varepsilon &< \underline{w}_{a_1^\ell}(\nu) \text{ for all } \ell < \ell^* \\ v_1^{\ell^*} - \varepsilon &< \underline{w}_{a_1^{\ell^*}}(\nu) < \overline{w}_{a_1^{\ell^*}}(\nu) < v_1^{\ell^*} + \varepsilon, \text{ and} \\ \overline{w}_{a_1^\ell}(\nu) &< v_1^{\ell^*} - \varepsilon \text{ for all } \ell > \ell^* \end{aligned}$$

for all $\nu \leq \nu^*$. Then because $\zeta_n^\ell \rightarrow \zeta^\ell > 0$ for all $\ell \in \mathcal{L}$, there exists some n^* and δ^* such that $-(1 - \delta) \log(\zeta_n^\ell) \leq \nu^*$ for all $\ell \in \mathcal{L}$, all $n \geq n^*$, and all $\delta \geq \delta^*$.

Then by Theorem 3.4.7, we obtain the following inequalities regarding the payoff to playing a_1^ℓ for $\ell \in \mathcal{L}$ in the equilibrium σ^n for $n \geq n^*$ and for any type $\delta \geq \delta^*$:

$$\begin{aligned} v_1^{\ell^*} - \varepsilon &< U_1(a_1^{\ell^*}, \sigma_2^n, \delta) < v_1^{\ell^*} + \varepsilon, \\ U_1(a_1^\ell, \sigma_2^n, \delta) &< v_1^{\ell^*} - \varepsilon. \end{aligned}$$

This immediately implies that all types $\delta \geq \delta^*$ must play a mixed strategy with support contained in $\mathcal{L}^c \cup \{\ell^*\}$ in the equilibrium σ^n for all $n \geq n^*$. \square

The lemma above is very intuitive. Given a sequence (σ^n, μ^n) that approximates λ , because all actions $\ell \in \mathcal{L}$ are played with positive probability in the limit, tight payoff

bounds for each action $\ell \in \mathcal{L}$ hold *uniformly* for all n sufficiently large and all sufficiently patient types.¹⁷ Then from these payoff bounds, sufficiently patient players must obtain payoffs close to v_1^ℓ for playing an action $\ell \in \mathcal{L}$ at all n sufficiently large. Clearly for these sufficiently patient players and these values of n , this means that the payoff from playing ℓ^* must dominate the payoff for playing any other action $\ell \in \mathcal{L}$ for these sufficiently patient players, which yields our desired conclusion.

Now we can conclude the proof. We have yet to invoke Corollary 3.4.4 and we will do so in this final step of the proof. In this final step, we study the upper tail of the distribution of types in Ω_r and the actions that these types can play in equilibria that approximate a distribution $\lambda \in \Pi_\rho(F)$. Invoking Lemma 3.4.13, we know that the upper tail must concentrate their play on the set $\{\ell^*\} \cup \mathcal{L}^c$ where $\mathcal{L}^c \equiv \{1, \dots, m\} \setminus \mathcal{L}$. Unlike in Lemma 3.4.13 however, payoff bounds to playing an action $\ell \in \mathcal{L}^c$ no longer hold *uniformly* for all n sufficiently large and all sufficiently patient types. So ruling out the play of certain actions in $\{\ell^*\} \cup \mathcal{L}^c$ by these types is far more difficult. We use Corollary 3.4.4 precisely to deal with this challenge. As discussed before, Corollary 3.4.4 will establish the existence of a strong contagion effect that rules out the possibility that $\sigma_1^n[a_1^1]$ vanishes to zero.

Lemma 3.4.14. *Suppose F satisfies Assumption 3.4.1 and let $\rho > 0$. Then there exists $\nu > 0$ such that $\lambda[a^S] > \nu$ for all $\lambda \in \Pi_\rho(F)$.*

Proof. Suppose that the lemma is false so that by Corollary 3.4.12, we can find $\lambda \in \Pi_\rho(F)$ such that $\lambda[a^S] = 0$. With a slight abuse of notation, let $\mathcal{L} = \mathcal{L}(\lambda)$. Note by definition that $1 \notin \mathcal{L}$ since $\lambda[a_1^1] = \lambda[a^S] = 0$.

Choose some sequence (σ^n, μ^n) that approximates λ . Using Lemma 3.4.13, there exists some $\delta^* < 1$ and some n^* such that $\sigma_1^n(\delta)[a_1^\ell] = 0$ for all $\delta > \delta^*$, $n > n^*$, and $\ell \in \mathcal{L} \setminus \{\ell^*\}$. So without loss of generality, by replacing the original sequence with the subsequence starting at $n^* + 1$, we can assume (σ^n, μ^n) to be a sequence that approximates λ and at the same time $\sigma_1^n(\delta)[a_1^\ell] = 0$ for all $\ell \in \mathcal{L} \setminus \{\ell^*\}$ and all $\delta > \delta^*$. Furthermore because

¹⁷This would not be true if instead $\ell \notin \mathcal{L}$, because $\sigma_1^n[a_1^\ell] \rightarrow 0$ for $\ell \notin \mathcal{L}$.

$\sigma_1^n[a_1^{\ell^*}] \rightarrow \lambda[a_1^{\ell^*}] > 0$, we can replace the sequence (σ^n, μ^n) with a further subsequence such that $\sigma_1^n[a_1^{\ell^*}] > \lambda[a_1^{\ell^*}]/2 > 0$ for all n .

Fix an $\varepsilon > 0$ sufficiently small so that

$$v_1^{\ell^*+1} < v_1^{\ell^*} - \varepsilon < v_1^{\ell^*} + \varepsilon < v_1^{\ell^*-1}.$$

Again as in the proof of Lemma 3.4.13, due to Lemma 3.4.6, we can choose $v^* > 0$ such that for all $v < v^*$,

$$\begin{aligned} v_1^{\ell^*} + \varepsilon &< \underline{w}_{a_1^\ell}(v) \text{ for all } \ell < \ell^* \\ v_1^{\ell^*} - \varepsilon &< \underline{w}_{a_1^{\ell^*}}(v) < \overline{w}_{a_1^{\ell^*}}(v) < v_1^{\ell^*} + \varepsilon, \text{ and} \\ \overline{w}_{a_1^\ell}(v) &< v_1^{\ell^*} - \varepsilon \text{ for all } \ell > \ell^*. \end{aligned}$$

Now because $\sigma_1^n[a_1^{\ell^*}] > \lambda[a_1^{\ell^*}]/2 > 0$ for all n , there exists some $\bar{\delta} \geq \delta^*$ such that

$$-(1 - \delta) \log \left(\sigma_1^n[a_1^{\ell^*}] \right) < v^*$$

for all $\delta > \bar{\delta}$ and all n . Of course this implies that for all $\delta > \bar{\delta}$ and all n ,

$$v_1^{\ell^*} + \varepsilon > \overline{w}_{a_1^{\ell^*}} \left(-(1 - \delta) \log \left(\sigma_1^n[a_1^{\ell^*}] \right) \right). \quad (3.2)$$

Now by Corollary 3.4.4, we can choose $\bar{\varepsilon} \in (0, \varepsilon)$ such that

$$\lim_{\bar{\zeta} \rightarrow 0} \frac{\rho}{\bar{\zeta} + (m-1)\bar{\varepsilon}} \min \left\{ 1 - F \left(1 + \frac{v^*}{\log \bar{\zeta}} \right), 1 - F(\bar{\delta}) \right\} > 1.$$

Then we can choose $\varepsilon^* \in (0, \bar{\varepsilon})$ such that

$$\frac{\rho}{\bar{\zeta} + (m-1)\varepsilon^*} \min \left\{ 1 - F \left(1 + \frac{v^*}{\log \bar{\zeta}} \right), 1 - F(\bar{\delta}) \right\} > 1 \quad (3.3)$$

for all $\bar{\zeta} < \varepsilon^*$.

We now establish an upper bound on the discount factor of a rational type who plays an action in \mathcal{L} with positive probability in a Nash equilibrium. Choose any $\ell' < \ell^*$. Due to

inequality (3.2), a rational type with discount factor $\delta > \bar{\delta}$ such that

$$\underline{w}_{a_1^{\ell'}} \left(-(1-\delta) \log \left(\sigma_1^n[a_1^{\ell'}] \right) \right) > v_1^{\ell^*} + \varepsilon$$

cannot play $a_1^{\ell^*}$ with positive probability. Thus any type who plays $a_1^{\ell^*}$ with positive probability must have a discount factor δ such that either $\delta \leq \bar{\delta}$ or

$$-(1-\delta) \log \left(\sigma_1^n[a_1^{\ell'}] \right) > v^* \Leftrightarrow \delta < 1 + \frac{v^*}{\log \left(\sigma_1^n[a_1^{\ell'}] \right)}.$$

This means that for any n , all types with discount factor $\delta > \max \left\{ 1 + v^* \left(\log \left(\sigma_1^n[a_1^{\ell'}] \right) \right)^{-1}, \bar{\delta} \right\}$ must play an action supported in $\mathcal{L}^c \equiv \{1, \dots, m\} \setminus \mathcal{L}$, which implies that for all n ,

$$\begin{aligned} \sum_{\ell \in \mathcal{L}^c} \sigma_1^n[a_1^\ell] &\geq \mu^n(\Omega_r) \min \left\{ 1 - F \left(1 + \frac{v^*}{\log \left(\sigma_1^n[a_1^{\ell'}] \right)} \right), 1 - F(\bar{\delta}) \right\} \\ &> \rho \min \left\{ 1 - F \left(1 + \frac{v^*}{\log \left(\sigma_1^n[a_1^{\ell'}] \right)} \right), 1 - F(\bar{\delta}) \right\}. \end{aligned}$$

Recall by the definition of \mathcal{L} that for all $\ell \in \mathcal{L}^c$, $\sigma_1^n[a_1^\ell] \rightarrow 0$ and therefore we can choose n^* such that $\sigma_1^n[a_1^\ell] < \varepsilon^*$ for all $n > n^*$ and all $\ell \in \mathcal{L}^c$. Fix some $n > n^*$. We now divide the inequality above on both sides by $\sum_{\ell \in \mathcal{L}^c} \sigma_1^n[a_1^\ell]$ yielding:¹⁸

$$\begin{aligned} 1 &\geq \left(\sum_{\ell \in \mathcal{L}^c} \sigma_1^n[a_1^\ell] \right)^{-1} \rho \min \left\{ 1 - F \left(1 + \frac{v^*}{\log \left(\sigma_1^n[a_1^{\ell'}] \right)} \right), 1 - F(\bar{\delta}) \right\} \\ &\geq \frac{\rho}{\sigma_1^n[a_1^{\ell'}] + (m-1)\varepsilon^*} \min \left\{ 1 - F \left(1 + \frac{v^*}{\log \left(\sigma_1^n[a_1^{\ell'}] \right)} \right), 1 - F(\bar{\delta}) \right\}. \end{aligned}$$

But this directly contradicts inequality (3.3) and concludes the proof. \square

Then Lemma 3.4.14 together with Lemma 3.4.10 immediately imply Lemma 3.4.8, proving Theorem 3.4.2.

¹⁸Recall that $\sum_{\ell \in \mathcal{L}^c} \sigma_1^n[a_1^\ell] > 0$ because $\sigma_1^n[a_1^1] > 0$ for all n and $1 \notin \mathcal{L}$.

3.5 Negative Results

In this section, we present two types of negative results that arise as the consequence of particular specifications of the type space. Unlike the theorems analyzed in the classical reputation literature, we fix the payoff functions of all players and vary the type space of the long run player in order to study its' effects on the set of equilibria.

3.5.1 Necessity of Behavioral Commitment Types

First we illustrate why behavioral commitment types are generally necessary for reputation building.

Theorem 3.5.1. *Suppose that the stage game has a pure Nash equilibrium a^* and suppose that $\mu[\Omega_b] = 0$. Then there exists a Nash equilibrium in which all rational types play a_1^* with probability one.*

Proof. The proof is simple. Consider the strategy profile σ in which all types in Ω_r play a_1^* . Then a best response for player 2 is to play a_2^* at all histories since player 2 assigns probability one to player 1 playing a_1^* at every history regardless of the signals observed. But then it is clear that a_1^* is a best response for all player 1 types. \square

This shows that uncertainty regarding discount factors is not sufficient by itself to generate reputation effects. The reason is due to the fact that in the standard complete information model when the long run player is known to be rational and his discount factor is common knowledge, the repetition of a pure static Nash equilibrium is always a Nash equilibrium of the repeated game regardless of the discount factor.

Note however the above theorem simply notes that $\mu[\Omega_b]$ must be strictly positive and does not necessarily require $\mu[\Omega_b]$ to be larger than some pre-specified value $\xi > 0$. We now study the structure of equilibria when $\mu[\Omega_b]$ is strictly positive but small. We show that even when this is the case, the set of equilibrium payoffs of the long run player generally includes payoffs far below the long run player's Stackelberg payoff.

3.5.2 Necessity of Arbitrarily Patient Rational Types

Suppose that the discount factor of the rational type is known to be at most $\delta^* < 1$ i.e. $F(\delta^*) = 1$. Then it is easy to show the following:

Theorem 3.5.2. *Suppose $F(\delta^*) = 1$ for some $\delta^* < 1$ and the stage game has a strict Nash equilibrium a^* . Then there exists some ζ^* such that if $\mu[\Omega_b] < \zeta^*$, there exists a Nash equilibrium in which all rational types play a_1^* with probability one.*

First note that this includes the case of the standard model in which the discount factor of player 1 is perfectly known. The argument is intuitive. When all rational types play a_1^* and the probability of non-rationality is sufficiently small, by deviating to a_1^1 , player 1 cannot convince player 2 that he indeed played a_1^1 until very late in the game. Because of discounting, these benefits arrive too late in the game for such a deviation to be desirable.

Proof. Consider the strategy in which all rational types play a_1^* . Then given any type space (Ω, μ) , because of Assumption 3.2.1, the beliefs of player 2 about player 1's action at any public history h^t , which we denote $\mu(h^t)$, is uniquely defined due to Bayes' rule. Given any type space (Ω, μ) , define player 2's strategy to be such that $\sigma_2(h^t) \in B(\mu(h^t))$. This now defines a strategy profile σ . We will now show that given any (Ω, μ) , σ constructed in the above manner is indeed a Nash equilibrium whenever $\mu[\Omega_b]$ is sufficiently small.

Let $\bar{u} = \max_{a \in A} u_1(a)$. Because a^* is a strict Nash equilibrium, we can find t^* such that

$$\max_{a_1 \neq a_1^*} \left(1 - (\delta^*)^{t^*}\right) u_1(a_1, a_2^*) + (\delta^*)^{t^*} \bar{u} < u_1(a^*).$$

Then let $\bar{\pi}$ be the maximum possible likelihood ratio among all actions and all public signals when comparing against a_1^* :

$$\bar{\pi} \equiv \max_{y, a} \frac{\pi(y | a)}{\pi(y | a_1^*)}.$$

Note then that given any t period history of signals $(y^0, y^1, \dots, y^{t-1})$, the probability that

player 2 assigns to player 1 playing a_1^* is at least:

$$\begin{aligned} & \frac{(1 - \mu[\Omega_b])}{\mu[\Omega_b] \frac{\max_{a_1 \in A_1} \pi(y^0|a_1)}{\pi(y^0|a_1^*)} \dots \frac{\max_{a_1 \in A_1} \pi(y^{t-1}|a_1)}{\pi(y^{t-1}|a_1^*)} + (1 - \mu(\Omega_b))} \\ \geq & \frac{1 - \mu[\Omega_b]}{\mu[\Omega_b] \bar{\pi}^{t-1} + (1 - \mu[\Omega_b])}. \end{aligned}$$

Now let $\nu \in (0, 1)$ be such that whenever player 2 assigns probability at least ν to player 1 playing a_1^* , player 2's unique best response is to play a_2^* .

Then choose there choose κ^* such that

$$\frac{1 - \kappa^*}{\kappa^* \bar{\pi}^{t^*-1} + (1 - \kappa^*)} > \nu.$$

Then it is easy to see that as long as $\mu[\Omega_b] \leq \kappa^*$,

$$\frac{1 - \mu[\Omega_b]}{\mu[\Omega_b] \bar{\pi}^{t-1} + (1 - \mu[\Omega_b])} > \nu$$

for all $t \leq t^*$. This then means that at any history h^t with $t \leq t^*$, player 2 must play a_2^* with certainty. With this an upper bound on player 1's payoff to deviating to another action a_1 is

$$\max_{a_1 \neq a_1^*} \left(1 - (\delta^*)^{t^*}\right) u_1(a_1, a_2^*) + (\delta^*)^{t^*} \bar{u}.$$

By definition of t^* , the above is strictly less than $u_1(a_1^*)$ and so we have indeed checked that player 1's actions are incentive compatible. This concludes the proof. \square

Bagwell (1995) proves essentially the same theorem in the context of a two-stage game. Additionally the point here corresponds to the theorems found in Cripps *et al.* (2004) who illustrate the same point in a different setting where the player can change his action every period. Note that the above result depends importantly on the assumption that the action of player 1 is observed imperfectly. For example, if the action was perfectly observed, then even without any behavioral types, a sufficiently patient player 1 would always choose a_1^1 in all Nash equilibria. In contrast, the imperfection in observation of the long run player's action (and more precisely the full support assumption of the public monitoring structure) eliminates the possibility of large jumps in player 2's posterior regarding the played action

of player 1 especially when $\sigma_1[a_1]$ is close to one. This then implies that when $(a_1, a_2(a_1))$ is a strict Nash equilibrium of the stage game, a deviation to an action $a'_1 \neq a_1$ is very costly since player 2's belief about the played action being a_1 remains close to one for a long time. Unlike in Section 3.4, we can make $\sigma_1[a_1]$ very close to one because the highest patient types are no longer available to rectify and reinforce reputation effects.

3.6 Numerical Simulations

3.6.1 Strictly Dominated Action: Product Choice Game

Let the stage game be given by the following two player game.

Table 3.1: *Stage Game*

	B	D
H	α, x	$-\beta, 0$
L	$\chi, -y$	$0, 0$

Player 1 chooses a technology of either H or L in period 0. Playing H incurs a cost to player 1 ($\alpha < \xi$, $\beta > 0$). However if player 1 played H and player 2 was convinced that he did so, player 2 would prefer to play B whereas if player 2 were to become convinced that player 1 played L , then he would play D . Let ω_* be the cutoff likelihood ratio above which the short run player plays B . To simplify notation, let $\kappa = (\lambda_H - \lambda_L)^2$.

We simplify the model such that the type space of discount factors (the amount of discounting over a unit interval of time) is given by the uniform distribution over $(0, 1)$. Furthermore assume that only rational types exist in the model. Because the model is in continuous time, it is convenient to transform the discount factor into a discount rate: $r \equiv -\log \delta$. Furthermore conditional on the firm choosing technology $\theta \in \{H, L\}$, the stochastic process of public signals observed by the consumers is given by:

$$dX_t = \lambda_\theta dt + dZ_t$$

where $\lambda_H > \lambda_L$. With this simple setup, we can now perform a rich analysis of the product

choice game described above.

Bayesian Inference

We first solve the filtering problem. Conditional on the long run player playing $\theta \in \{H, L\}$, we can calculate the process dp_t^θ :

$$\begin{aligned} dp_{t,H} &= \kappa p_{t,H}(1 - p_{t,H})^2 dt + \sqrt{\kappa} p_{t,H}(1 - p_{t,H}) dZ_t, \\ dp_{t,L} &= -\kappa p_{t,L}^2(1 - p_{t,L}) dt + \sqrt{\kappa} p_{t,L}(1 - p_{t,L}) dZ_t. \end{aligned}$$

It is computationally much more convenient to express all terms in terms of the likelihood ratio:

$$\omega_t^H = \frac{p_t^H}{1 - p_t^H}, \omega_t^L = \frac{p_t^L}{1 - p_t^L}.$$

We obtain the following SDE:

$$\begin{aligned} d\omega_t^H &= \kappa \omega_t^H dt + \sqrt{\kappa} \omega_t^H dZ_t \\ d\omega_t^L &= \sqrt{\kappa} \omega_t^L dZ_t. \end{aligned}$$

Value Functions

With the above, we can now solve for the value function of a type r that plays H whenever $\omega \leq \omega_*$:

$$V_H(r, \omega_t^H) = -r\beta dt + (1 - r)dt \left(V_H(r, \omega_t^H) + \mathbb{E}[dV_H(r, \omega_t^H)] \right).$$

Rearranging we obtain:

$$r \left(V_H(r, \omega_t^H) + \beta \right) dt = \mathbb{E}[dV_H(r, \omega_t^H)];$$

We can derive a similar equation at any $\omega > \omega_*$ which implies the following ODE for V_H :

$$rV_H(r, \omega) = \begin{cases} -r\beta + (\lambda_H - \lambda_L)^2 \omega \left(\frac{\partial V_H}{\partial \omega} + \frac{\omega}{2} \frac{\partial^2 V_H}{\partial \omega^2} \right) & \text{if } \omega < \omega_*, \\ r\alpha + (\lambda_H - \lambda_L)^2 \omega \left(\frac{\partial V_H}{\partial \omega} + \frac{\omega}{2} \frac{\partial^2 V_H}{\partial \omega^2} \right) & \text{if } \omega > \omega_*. \end{cases}$$

Solving the above ODE piecewise, we obtain the following:

$$V_H(r, \omega) = \begin{cases} -\beta + C_1 \omega^{-\frac{1}{2}} \left(1 + \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) + C_2 \omega^{\frac{1}{2}} \left(-1 + \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) & \text{if } \omega < \omega_* \\ \alpha + C_3 \omega^{-\frac{1}{2}} \left(1 + \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) + C_4 \omega^{\frac{1}{2}} \left(-1 + \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) & \text{if } \omega > \omega_*. \end{cases}$$

Because of the boundary conditions, $V^H(r, 0) = -\beta$ and $\lim_{\omega \rightarrow +\infty} V^H(r, \omega) = \alpha$, we have $C_1 = 0$ and $C_4 = 0$. Thus the above simplifies to:

$$V_H(r, \omega) = \begin{cases} -\beta + C_2 \omega^{-\frac{1}{2}} \left(1 - \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) & \text{if } \omega < \omega_* \\ \alpha + C_3 \omega^{-\frac{1}{2}} \left(1 + \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) & \text{if } \omega > \omega_*. \end{cases}$$

We then solve for C_3 and C_4 explicitly by employing the smooth pasting and value matching conditions:

$$\begin{aligned} \lim_{\omega \uparrow \omega_*} V_H(\omega) &= \lim_{\omega \downarrow \omega_*} V_H(\omega) \\ \lim_{\omega \uparrow \omega_*} V'_H(\omega) &= \lim_{\omega \downarrow \omega_*} V'_H(\omega). \end{aligned}$$

After evaluation and algebraic manipulation, we get the following closed-form solution to the differential equation:

$$V_H(r, \omega) = \begin{cases} -\beta + \frac{\alpha + \beta}{2} \left(\frac{\sqrt{\kappa}}{\sqrt{8r+\kappa}} + 1 \right) (\omega/\omega_*)^{-\frac{1}{2}} \left(1 - \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) & \text{if } \omega < \omega_* \\ \alpha + \frac{\alpha + \beta}{2} \left(\frac{\sqrt{\kappa}}{\sqrt{8r+\kappa}} - 1 \right) (\omega/\omega_*)^{-\frac{1}{2}} \left(1 + \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) & \text{if } \omega > \omega_*. \end{cases}$$

Similarly we calculate the value functions of a type r that plays L :

$$rV_L(r, \omega) = \begin{cases} r\chi + \kappa \frac{\omega^2}{2} \frac{\partial^2 V_L}{\partial \omega^2} & \text{if } \omega > \omega_*, \\ \kappa \frac{\omega^2}{2} \frac{d^2 V_L}{d\omega^2} & \text{if } \omega < \omega_*. \end{cases}$$

This then yields the following solution to the differential equation:

$$V_L(r, \omega) = \begin{cases} \frac{\chi}{2} \left(1 - \frac{\sqrt{\kappa}}{\sqrt{8r+\kappa}}\right) (\omega/\omega_*)^{\frac{1}{2}} \left(1 + \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) & \text{if } \omega < \omega_*, \\ \chi - \frac{\chi}{2} \left(1 + \frac{\sqrt{\kappa}}{\sqrt{8r+\kappa}}\right) (\omega/\omega_*)^{\frac{1}{2}} \left(1 - \frac{\sqrt{8r+\kappa}}{\sqrt{\kappa}}\right) & \text{if } \omega > \omega_*. \end{cases}$$

Equilibrium

First note that for every ω , $V_H(r, \omega) - V_L(r, \omega)$ is strictly decreasing in r whenever $\omega < \omega_*$. Thus there exists a function $r^* : [0, \infty) \rightarrow [0, \infty)$ such that for all ω , $V_H(r, \omega) > V_L(r, \omega)$ for all $r < r^*(\omega)$ and $V_H(r, \omega) < V_L(r, \omega)$ for all $r > r^*(\omega)$.

This then implies that every equilibrium is uniquely determined by a cutoff discount rate r^* such that the equilibrium likelihood ratio is

$$\bar{\omega} = \frac{1 - e^{-r^*(\bar{\omega})}}{e^{-r^*(\bar{\omega})}} = e^{r^*(\bar{\omega})} - 1$$

and $V_L(r^*, \bar{\omega}) = V_H(r^*, \bar{\omega})$. All of this yields the equation:

$$V_L(r^*, e^{r^*} - 1) = V_H(r^*, e^{r^*} - 1). \quad (3.4)$$

Thus in any equilibrium, r^* is a fixed point of Equation (3.4).

Note that thus far we have not introduced any behavioral commitment types into the game and thus $r^* = 0$ is a trivial equilibrium. However we now introduce the idea of *robustness to the introduction of the Stackelberg type* ω_S .

Definition 3.6.1. Let σ be an equilibrium of the continuous time game with only rational types. Then σ is robust to the introduction of ω_S if for any sequence of games $(\Omega_r \cup \{\omega_S\}, \mu^n)$ with $\mu^n[\omega_S] \rightarrow 0$, there exist $\sigma^n \in E(\Omega_r \cup \{\omega_S\}, \mu^n)$ such that $\sigma^n[a_1^1] \rightarrow \sigma[a_1^1]$.

It is an easy modification of the arguments in Section 3.4 to prove that the only equilibrium that is robust to the introduction of ω_S (the Stackelberg commitment type) is when $r^* > 0$. Furthermore r^* is the unique point not equal to zero at which Equation (3.4) holds.

Theorem 3.6.2. *There exists a unique equilibrium that is robust to the introduction of ω_S . In this equilibrium, the cutoff discount rate r^* satisfies Equation (3.4) and is strictly positive.*

Size of the Contagion Effect

Using this framework we can now explicitly compute the size of the contagion effect. Consider the following specification of payoffs: $\alpha = 1, \beta = 1, \xi = 2, x = 1, y = 1$. In this

scenario $\omega^* = 1$ so that player 2 plays B at time t if and only if $p_t \geq \frac{1}{2}$. Furthermore let $\kappa = (\lambda_H - \lambda_L)^2 = 1.5$.

Table 3.2: *Stage Game*

	B	D
H	1, 1	-1, 0
L	2, -1	0, 0

Given these parameters we can compute the mass of types playing the In the equilibrium robust to the introduction of the Stackelberg type, $\delta^* \equiv e^{-r^*} \approx 0.617$. Thus $\sigma_1[H] \approx 0.383$ and so approximately 38 percent of all rational types play H in equilibrium. Given this, in Figure 3.2, we plot the payoffs of all rational types in the equilibrium robust to the introduction of the Stackelberg type. As a consequence this means that even when the probability of the Stackelberg type is extremely small, there will be approximately 38 percent of rational types who play the Stackelberg action in any equilibrium. This leads to high payoffs even for types with moderate levels of patience. For example, a type with a discount factor of 0.95 obtains a payoff of approximately 0.81 in the robust equilibrium.

If instead $\kappa = 2$ so that the signal structure is more informative, then $\delta^* \approx 0.476$ and thus over half of the population plays the Stackelberg action in the unique robust equilibrium. As a result the payoffs of all types are at least 0.57 as illustrated in Figure 3.3. Furthermore the same type $\delta = 0.95$ from the previous example now obtains a payoff of approximately 0.92 in the robust equilibrium while a much less patient type $\delta = 0.8$ also obtains a relatively high payoff of 0.76.

3.7 Conclusion and Discussion

To conclude, the paper leaves open natural questions for further research. It is unclear whether private information regarding discount factors can generate reputation effects with small probability of behavioral commitment types in the standard reputation model where the long run player can continuously change his action. This problem is substantially




Figure 3.2: *The payoff of type δ in the robust equilibrium.*




Figure 3.3: *The payoff of type δ in the robust equilibrium.*

complicated by two factors. First the lack of commitment to an action means that it may be a best response for a very patient player to pick a non-Stackelberg action at some point in the future. Perfect commitment to an action avoids this complication and thus allows for strong contagion effects starting with these very patient types. Secondly, the long run player's decision at every point in time will depend on the short run player's posterior about not just whether he is a rational or behavioral type but also on the short run player's beliefs about the long run player's discount factor. Because this is changing over time, it is important to understand how the short run player's perceived distribution over the long run player's patience evolves in order to study whether reputation can be maintained.

Another potentially interesting open question is whether one can obtain similar results if incomplete information concerns the monitoring structure of player 2 rather than the patience level of player 1. Consider player 1 who has private information about the technology $\pi_\theta(\cdot | a_1)$ governing the mapping from intended action to public consequences Y . Intuitively those types that know that $\pi(\cdot | a_1)$ is essentially one of perfect monitoring will have a strict incentive to play the Stackelberg action. This may lead to a similar contagion effect for those who have less perfect technologies. However to study this question, new techniques must be developed to obtain bounds on how the payoff estimates depend on the monitoring structure, which is not directly addressed in Gossner (2011).

References

- ABREU, D., MILGROM, P. and PEARCE, D. (1991). Information and timing in repeated partnerships. *Econometrica*, **59** (6), 1713–1733.
- , PEARCE, D. and STACCHETTI, E. (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, **58**, 1041–1064.
- AHN, T. K., OSTROM, E., SCHMIDT, D. and WALKER, J. (2003). Trust in two-person games: Game structures and linkages. In E. Ostrom and J. Walker (eds.), *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, Russell Sage, New York, pp. 323–351.
- AUMANN, R. J. and SHAPLEY, L. S. (1976). Long-term competition—A game theoretic analysis. In N. Megiddo (ed.), *Essays in Game Theory in Honor of Michael Maschler*, New York: Springer-Verlag, pp. 1–15.
- BAGWELL, K. (1995). Commitment and observability in games. *Games and Economic Behavior*, **8** (2), 271–280.
- BANDIERA, O. and RASUL, I. (2006). Social networks and technology adoption in northern mozambique. *Economic Journal*, **116** (514), 869–902.
- BAPTISTA, R. (1999). The diffusion of process innovations: a selective review. *International Journal of the Economics of Business*, **6** (1), 107–129.
- BEN-PORATH, E. and KAHNEMAN, M. (2003). Communication in repeated games with costly monitoring. *Games and Economic Behavior*, **44** (2), 227–250.
- BERGEMANN, D. and VÄLIMÄKI, J. (1997). Market diffusion with two-sided learning. *RAND Journal of Economics*, pp. 773–795.
- and VÄLIMÄKI, J. (2000). Experimentation in markets. *Review of Economic Studies*, **67** (2), 213–234.
- BERGIN, J. and MACLEOD, W. B. (1993). Continuous time repeated games. *International Economic Review*, **34**, 21–37.
- BHASKAR, V. and OBARA, I. (2011). Reputation and information lag, mimeo.
- BOARD, S. and MEYER-TER-VEHN, M. (2013). Reputation for quality. *Econometrica*, forthcoming.
- and MEYER-TER-VEHN, M. (2013). Reputation for quality. *Econometrica*, **81** (6), 2381–2462.

- BOHREN, J. A. (2011). Stochastic games in continuous time: Persistent actions in long-run relationships, mimeo.
- BOLTON, P. and HARRIS, C. (1999). Strategic experimentation. *Econometrica*, **67** (2), 349–374.
- CHAMLEY, C. and GALE, D. (1994). Information revelation and strategic delay in a model of investment. *Econometrica*, pp. 1065–1085.
- CHE, Y.-K. and HÖRNER, J. (2013). Optimal design for social learning, mimeo.
- COMPTE, O. (1998). Communication in repeated games with imperfect private monitoring. *Econometrica*, **66** (3), 597–626.
- CONLEY, T. G. and UDRY, C. R. (2010). Learning about a new technology: Pineapple in Ghana. *American Economic Review*, **100** (1), 35–69.
- CRIPPS, M. W., MAILATH, G. J. and SAMUELSON, L. (2004). Imperfect monitoring and impermanent reputations. *Econometrica*, **72** (2), 407–432.
- DAVIES, S. (1979). *The Diffusion of Process Innovations*. Cambridge: CUP Archive.
- DILMÉ, F. (2012). Building and milking trust: Reputation as a moral hazard phenomenon, mimeo.
- DUTTA, P. (1995). A folk theorem for stochastic games. *Journal of Economic Theory*, **66** (1), 1–32.
- ELLISON, G. (1994). Cooperation in the prisoner’s dilemma with anonymous random matching. *Review of Economic Studies*, **61**, 567–588.
- and FUDENBERG, D. (1993). Rules of thumb for social learning. *Journal of Political Economy*, **101** (4), 612–643.
- ESCOBAR, J. and TOIKKA, J. (2011). A folk theorem with Markovian private information, mimeo.
- FAINGOLD, E. (2013). Reputation and the flow of information in repeated games, mimeo.
- and SANNIKOV, Y. (2011). Reputation in continuous-time games. *Econometrica*, **79** (3), 773–876.
- FOSTER, A. D. and ROSENZWEIG, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, pp. 1176–1209.
- FRIEDMAN, J. W. (1971). A noncooperative equilibrium for supergames. *Review of Economic Studies*, **38** (1), 1–12.
- FUDENBERG, D. and LEVINE, D. (2007a). Continuous time limits of repeated games with imperfect public monitoring. *Review of Economic Dynamics*, **10**, 173–192.

- and — (2007b). The Nash-threats folk theorem with communication and approximate common knowledge in two player games. *Journal of Economic Theory*, **132** (1), 461–473.
- and — (2009). Repeated games with frequent signals. *Quarterly Journal of Economics*, **124** (1), 233–265.
- and LEVINE, D. K. (1989). Reputation and equilibrium selection in games with a patient player. *Econometrica*, **57** (4), 759–778.
- and — (1992). Maintaining a reputation when strategies are imperfectly observed. *Review of Economic Studies*, **59** (3), 561–579.
- , — and MASKIN, E. (1994). The folk theorem with imperfect public information. *Econometrica*, **62** (5), 997–1031.
- and MASKIN, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, **54** (3), 533–554.
- and OLSZEWSKI, W. (2011). Repeated games with asynchronous monitoring of an imperfect signal. *Games and Economic Behavior*, **72**, 86–99.
- and YAMAMOTO, Y. (2011). The folk theorem for irreducible stochastic games with imperfect public monitoring. *Journal of Economic Theory*, **146**, 1664–1683.
- GACHTER, S., HERRMANN, B. and THONI, C. (2004). Trust, voluntary cooperation, and socio-economic background: Survey and experimental evidence. *Journal of Economic Behavior & Organization*, **55**, 505–531.
- GEROSKI, P. A. (2000). Models of technology diffusion. *Research policy*, **29** (4), 603–625.
- GORT, M. and KLEPPER, S. (1982). Time paths in the diffusion of product innovations. *Economic Journal*, **92** (367), 630–653.
- GOSSNER, O. (2011). Simple bounds on the value of a reputation. *Econometrica*, **79** (5), 1627–1641.
- GRILICHES, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, **25** (4), 501–522.
- HALAC, M., LIU, Q. and KARTIK, N. (2013). Optimal contracts for experimentation, mimeo.
- HÖRNER, J. and OLSZEWSKI, W. (2006). The folk theorem for games with private almost-perfect monitoring. *Econometrica*, **74** (6), 1499–1544.
- HÖRNER, J. and OLSZEWSKI, W. (2009). How robust is the folk theorem? *Quarterly Journal of Economics*, **124** (4), 1773–1814.
- HÖRNER, J., SUGAYA, T., TAKAHASHI, S. and VIEILLE, N. (2011). Recursive methods in discounted stochastic games: An algorithm for $\delta < 1$ and a folk theorem. *Econometrica*, **79**, 1277–1318.

- JEHIEL, P. and SAMUELSON, L. (2012). Reputation with analogical reasoning. *The Quarterly Journal of Economics*, **127** (4), 1927–1969.
- JENSEN, R. (1982). Adoption and diffusion of an innovation of uncertain profitability. *Journal of Economic Theory*, **27** (1), 182–193.
- KANDORI, M. and MATSUSHIMA, H. (1998). Private observation, communication and collusion. *Econometrica*, **66** (3), 627–652.
- KAPUR, S. (1995). Technological diffusion with social learning. *Journal of Industrial Economics*, **43** (2), 173–195.
- KEILLOR, B. D. (2007). *Marketing in the 21st Century*, vol. 4. Greenwood Publishing Group.
- KELLER, G. and RADY, S. (2010). Strategic experimentation with poisson bandits. *Theoretical Economics*, **5** (2), 275–311.
- and — (2013). Breakdowns, mimeo.
- , — and CRIPPS, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, **73** (1), 39–68.
- KREPS, D. M., MILGROM, P. R., ROBERTS, J. and WILSON, R. J. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, **27** (2), 245–252.
- LAGZIEL, D. and LEHRER, E. (2012). Regret-free strategy with delayed information, mimeo.
- LEVY, Y. (2009). Stochastic games with information lag, mimeo.
- LIU, Q. (2011). Information acquisition and reputation dynamics. *Review of Economic Studies*, **78**, 1400 – 1425.
- and SKRZYPACZ, A. (2014). Limited records and reputation bubbles. *Journal of Economic Theory*, **151**, 2 – 29.
- MACLEOD, W. B. (2007). Reputations, relationships, and contract enforcement. *Journal of Economic Literature*, pp. 595–628.
- MAILATH, G. J. and MORRIS, S. (2002). Repeated games with almost-public monitoring. *Journal of Economic theory*, **102** (1), 189–228.
- and SAMUELSON, L. (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford: Oxford University Press.
- MANSFIELD, E. (1961). Technical change and the rate of imitation. *Econometrica*, **29** (4), 741–766.
- (1968). *Industrial Research and Technological Innovation: An Econometric Analysis*. New York: Norton.
- MCLAREN, C. H. and DEPAOLO, C. A. (2009). Movie data. *Journal of Statistics Education*, **17** (1).

- MUNSHI, K. (2004). Social learning in a heterogeneous population: Technology diffusion in the indian green revolution. *Journal of Development Economics*, **73** (1), 185–213.
- MURTO, P. and VÄLIMÄKI, J. (2011). Learning and information aggregation in an exit game. *Review of Economic Studies*, **78** (4), 1426–1461.
- NARAYANAN, V. K. and O’CONNOR, G. C. (2010). *Encyclopedia of Technology and Innovation Management*. Wiley-Blackwell.
- OBARA, I. (2009). Folk theorem with communication. *Journal of Economic Theory*, **144**, 120 – 134.
- PERSONS, J. C. and WARTHER, V. A. (1997). Boom and bust patterns in the adoption of financial innovations. *Review of Financial Studies*, **10** (4), 939–967.
- RENY, P. J. (2008). Non-cooperative games (equilibrium existence). In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*, Basingstoke: Palgrave Macmillan.
- ROSENBERG, D., SOLAN, E. and VIEILLE, N. (2007). Social learning in one-arm bandit problems. *Econometrica*, **75** (6), 1591–1611.
- RUBINSTEIN, A. (1994). Equilibrium in supergames. In N. Megiddo (ed.), *Essays in Game Theory in Honor of Michael Maschler*, New York: Springer-Verlag, pp. 17–28, master’s thesis, Hebrew Univeristy of Jerusalem, 1977.
- SANNIKOV, Y. (2007). Games with imperfectly observed actions in continuous time. *Econometrica*, **75**, 1285–1329.
- and SKRZYPACZ, A. (2007). Impossibility of collusion under imperfect monitoring with flexible production. *American Economic Review*, **97**, 1794–1823.
- and — (2010). The role of information in repeated games with frequent actions. *Econometrica*, **78**, 847–882.
- SEKIGUCHI, T. (1997). Efficiency in repeated prisoner’s dilemma with private monitoring. *Journal of Economic Theory*, **76** (2), 345–361.
- STOKEY, N. L. (1979). Intertemporal price discrimination. *The Quarterly Journal of Economics*, **93** (3), 355–371.
- SUGAYA, T. (2011). Folk theorem in repeated games with private monitoring, mimeo.
- WEINSTEIN, J. and YILDIZ, M. (2012). Reputation without commitment, mimeo.
- YAMAMOTO, Y. (2009). A limit characterization of belief-free equilibrium payoffs in repeated games. *Journal of Economic Theory*, **144** (2), 802–824.
- YAO, T., XU, S. H. and JIANG, B. (2011). Trigger strategy equilibriums in stochastic differential games with information time lags: an analysis of cooperative production strategies. *SIAM Journal of Control and Optimization*, **49** (3), 1197–1220.
- YOUNG, H. P. (2009). Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American Economic Review*, **99** (5), 1899–1924.

Appendix A

Appendix to Chapter 1

A.1 Preliminary Mathematical Tools for Equilibrium Analysis

A.1.1 General Properties of the Value to Waiting

Throughout this section N denotes an equilibrium adoption flow, with associated value to waiting W_t^N and no-news posterior p^N . We establish some basic mathematical properties of the value to waiting W^N corresponding to any equilibrium adoption flow N .

Lemma A.1.1. *Let N be an equilibrium flow of adopters. Then W_t^N is continuous in t .*

Proof. This is immediate from Lemma 1.2.4. Note that W_t^N can be written as:

$$W_t^N = \int_t^{\infty} h(\tau) d\tau$$

for some $h \in L^1[0, \infty) \cap L^\infty[0, \infty)$. Then it is immediate that W_t^N is continuous in t . \square

Lemma A.1.2. *Suppose that N is an equilibrium and that $W_t^N < 2p_t^N - 1$ for some $t > 0$. Then there exists some $\nu > 0$ such that W_t^N is continuously differentiable in t on the interval $(t - \nu, t + \nu)$ and*

$$\begin{aligned} \dot{W}_t^N = & (r + \rho + (\varepsilon_G + \lambda_G N_t) p_t^N + (\varepsilon_B + \lambda_B N_t)(1 - p_t^N)) W_t^N \\ & - \rho(2p_t^N - 1) - p_t^N (\varepsilon_G + \lambda_G N_t) \frac{\rho}{\rho + r}. \end{aligned}$$

Proof. By Lemma A.1.1, W_t^N must be continuous in t . Because $2p_t^N - 1$ is continuous in t , there exists some $\nu > 0$ such that $W_\tau^N < 2p_\tau^N - 1$ for all $\tau \in (t - \nu, t + \nu)$. This means that $N_\tau = \rho \bar{N}_\tau$ for all $\tau \in (t - \nu, t + \nu)$ and so N_τ must be continuous at all $\tau \in (t - \nu, t + \nu)$. From Lemma 1.2.4, W_τ^N can be rewritten for all $\tau \in (t - \nu, t + \nu)$ as

$$\begin{aligned} W_\tau^N &= \int_\tau^{t+\nu} \rho e^{-(\rho+r)(s-\tau)} \left(p_\tau^N e^{-\int_\tau^s (\varepsilon_G + \lambda_G N_x) dx} - (1 - p_\tau^N) e^{-\int_\tau^s (\varepsilon_B + \lambda_B N_x) dx} \right) ds \\ &\quad + e^{-(r+\rho)(t+\nu-\tau)} \left(p_\tau^N e^{-\int_\tau^{t+\nu} (\varepsilon_G + \lambda_G N_x) dx} + (1 - p_\tau^N) e^{-\int_\tau^{t+\nu} (\varepsilon_B + \lambda_B N_x) dx} \right) W_{t+\nu}^N \\ &\quad + \int_\tau^{t+\nu} \rho e^{-(\rho+r)(s-\tau)} p_\tau^N \left(1 - e^{-\int_\tau^s (\varepsilon_G + \lambda_G N_x) dx} \right) ds \\ &\quad + e^{-(r+\rho)(t+\nu-\tau)} p_\tau^N \left(1 - e^{-\int_\tau^{t+\nu} (\varepsilon_G + \lambda_G N_x) dx} \right) \frac{\rho}{\rho + r}. \end{aligned}$$

From this it is easy to see that W_τ^N is continuously differentiable with respect to τ for all $\tau \in (t - \nu, t + \nu)$.

The derivative can be computed using Ito's Lemma for processes with jumps. Given the perfect Poisson learning structure, the derivation is simple and we provide it here for completeness.

As above, for any $\Delta < t + \nu - \tau$ we can rewrite W_τ^N as

$$\begin{aligned} W_\tau^N &= \int_\tau^{\tau+\Delta} \rho e^{-(\rho+r)(s-\tau)} \left(p_\tau^N e^{-\int_\tau^s (\varepsilon_G + \lambda_G N_x) dx} - (1 - p_\tau^N) e^{-\int_\tau^s (\varepsilon_B + \lambda_B N_x) dx} \right) ds \\ &\quad + e^{-(r+\rho)\Delta} \left(p_\tau^N e^{-\int_\tau^{\tau+\Delta} (\varepsilon_G + \lambda_G N_x) dx} + (1 - p_\tau^N) e^{-\int_\tau^{\tau+\Delta} (\varepsilon_B + \lambda_B N_x) dx} \right) W_{\tau+\Delta}^N \\ &\quad + \int_\tau^{\tau+\Delta} \rho e^{-(r+\rho)(s-\tau)} p_\tau^N \left(1 - e^{-\int_\tau^s (\varepsilon_G + \lambda_G N_x) dx} \right) ds \\ &\quad + e^{-(r+\rho)\Delta} p_\tau^N \left(1 - e^{-\int_\tau^{\tau+\Delta} (\varepsilon_G + \lambda_G N_x) dx} \right) \frac{\rho}{\rho + r}. \end{aligned}$$

Since this is true for all $\Delta \in (0, t + \nu - \tau)$, the right hand side of this identity, which we denote R_Δ , is continuously differentiable with respect to Δ and satisfies $\frac{d}{d\Delta} R_\Delta \equiv 0$. Taking the limit as $\Delta \rightarrow 0$ and since $\dot{W}_\tau^N = \lim_{\Delta \rightarrow 0} \frac{d}{d\Delta} W_{\tau+\Delta}^N$ by continuous differentiability, we then

obtain that

$$\dot{W}_\tau^N = (r + \rho + (\varepsilon_G + \lambda_G N_\tau)p_\tau + (\varepsilon_B + \lambda_B N_\tau)(1 - p_\tau))W_\tau^N - \rho(2p_\tau - 1) - p_\tau(\varepsilon_G + \lambda_G N_\tau)\frac{\rho}{\rho + r},$$

as claimed. \square

We can prove a similar lemma for the case in which the equilibrium value to waiting is strictly above the payoff to adopting today.

Lemma A.1.3. *Suppose that N is an equilibrium and that $W_t^N > 2p_t^N - 1$ for some $t > 0$. Then there exists some $\nu > 0$ such that W_t^N is continuously differentiable in t on the interval $(t - \nu, t + \nu)$ and*

$$\dot{W}_t^N = (r + p_t^N \varepsilon_G + (1 - p_t^N) \varepsilon_B) W_t^N - p_t^N \varepsilon_G \frac{\rho}{\rho + r}.$$

Proof. The proof of continuous differentiability of W_t^N follows along the same lines as in the proof of Lemma A.1.2. Lemma A.1.1 again implies that if $W_t^N > 2p_t^N - 1$, then there exists $\nu > 0$ such that $W_\tau^N > 2p_\tau^N - 1$ for all $\tau \in (t - \nu, t + \nu)$. By the definition of equilibrium, $N_\tau = 0$ for all $\tau \in (t - \nu, t + \nu)$.

Hence, W_τ^N satisfies

$$\begin{aligned} W_\tau^N &= e^{-r(t+\nu-\tau)} \left(p_\tau^N e^{-\varepsilon_G(t+\nu-\tau)} + (1 - p_\tau^N) e^{-\varepsilon_B(t+\nu-\tau)} \right) W_{t+\nu}^N \\ &\quad + p_\tau^N \int_\tau^{t+\nu} \varepsilon_G e^{-(\varepsilon_G+r)s} \frac{\rho}{\rho + r} ds. \end{aligned}$$

From this it is again immediate that W_τ^N is continuously differentiable in τ .

To compute the derivative, we can proceed as above, rewriting W_τ^N as

$$W_\tau^N = e^{-r\Delta} \left(p_\tau^N e^{-\varepsilon_G\Delta} + (1 - p_\tau^N) e^{-\varepsilon_B\Delta} \right) W_{t+\Delta}^N + p_\tau \int_\tau^{\tau+\Delta} \varepsilon_G e^{-(\varepsilon_G+r)s} \frac{\rho}{\rho + r} ds$$

for any $\Delta < t + \nu - \tau$.

Differentiating both sides of the above equality with respect to Δ and taking the limit as $\Delta \rightarrow 0$, we obtain:

$$\dot{W}_\tau^N = (r + p_\tau^N \varepsilon_G + (1 - p_\tau^N) \varepsilon_B) W_\tau^N - p_\tau^N \varepsilon_G \frac{\rho}{\rho + r},$$

as claimed. □

A.1.2 Special Properties of the Value to Waiting under PBN

Here we focus on learning via perfect bad news. By Equation 1.1, an upper bound on the no-news posterior is given by:

$$\mu(\varepsilon, \Lambda_0, p_0) := \begin{cases} 1 & \text{if } \varepsilon > 0, \\ \frac{p_0}{p_0 + (1-p_0)e^{-\Lambda_0}} & \text{if } \varepsilon = 0. \end{cases}$$

We now show that absent breakdowns, this posterior is attained in the limit.

Lemma A.1.4. *Let N be an equilibrium under PBN. Suppose that $\varepsilon > 0$ or $p_0 > 1/2$. Then $p_t^N \rightarrow \mu(\varepsilon, \Lambda_0, p_0)$ and $W_t^N \rightarrow \frac{\rho}{\rho+r}(2\mu(\varepsilon, \Lambda_0, p_0) - 1)$ as $t \rightarrow \infty$.*

Proof. Consider first the case in which $\varepsilon > 0$. Then trivially $p_t^N \rightarrow 1$ as $t \rightarrow \infty$. So for any $\nu > 0$, there exists some t^* such that whenever $t > t^*$, then $1 - p_t^N < \nu$.

Then we can produce upper and lower bounds on W_t^N :

$$\frac{\rho}{\rho+r}(1-\nu) - \frac{\rho}{\rho+r}\nu < \frac{\rho}{\rho+r}(2p_t^N - 1) \leq W_t^N \leq \frac{\rho}{\rho+r}.$$

Since this is true for any $\nu > 0$, it follows that $\lim_{t \rightarrow \infty} W_t^N = \frac{\rho}{\rho+r}$ as claimed.

Now suppose that $\varepsilon = 0$ and $p_0 > 1/2$. Then note that $W_t^N \leq 2p_t^N - 1$ for all t : Indeed, suppose that $W_t^N > 2p_t^N - 1$ for some t . We can't have that $W_s^N > 2p_s^N - 1$ for all $s \geq t$, since otherwise $W_t^N = 0$, contradicting $W_t^N > 2p_t^N - 1 > 0$. But then we can find $s > t$ such that $W_s^N = 2p_s^N - 1$ and $W_{s'}^N > 2p_{s'}^N - 1$ for all $s' \in (t, s)$. This implies $N_{s'}^N = 0$ for all s' , and hence $W_t^N = e^{-r(s-t)}W_s^N = e^{-r(s-t)}(2p_s^N - 1) = e^{-r(s-t)}(2p_t^N - 1)$, again contradicting $W_t^N > 2p_t^N - 1 > 0$.

Let $N^* := \lim_{t \rightarrow \infty} \int_0^t N_s ds = \sup_t \int_0^t N_s ds \leq \bar{N}_0$. Let $p^* := \lim_{t \rightarrow \infty} p_t^N = \sup_t p_t^N$. For any $\nu > 0$ we can find t^* such that whenever $t > t^*$, then $e^{-\lambda \int_{t^*}^t N_s ds} > 1 - \nu$. Because

$2p_t^N - 1 \geq W_t^N$ for all t , we can then rewrite the value to waiting at time t as:

$$\begin{aligned} W_t^N &= \int_t^\infty \rho e^{-(r+\rho)\tau} \left(p_t^N - (1 - p_t^N) e^{-\lambda \int_t^\tau N_s ds} \right) d\tau \\ &\leq \frac{\rho}{r + \rho} \left(p_t^N - (1 - p_t^N)(1 - \nu) \right) \end{aligned}$$

for all $t > t^*$. Moreover, by optimality $W_t^N \geq \frac{\rho}{\rho+r}(2p_t^N - 1)$ for all t , so combining we have

$$\frac{\rho}{\rho+r}(2p^* - 1) \leq \liminf_{t \rightarrow \infty} W_t^N \leq \limsup_{t \rightarrow \infty} W_t^N \leq \frac{\rho}{r + \rho} (p^* - (1 - p^*)(1 - \nu)).$$

Since this is true for all $\nu > 0$, it follows that

$$\lim_{t \rightarrow \infty} W_t^N = \frac{\rho}{r + \rho} (2p^* - 1).$$

But the above is strictly less than $2p^* - 1$, so for all t sufficiently large we must have $2p_t^N - 1 > W_t^N$. Then for all t sufficiently large, we have $N_t = \rho \bar{N}_t$. Thus, $N^* = \bar{N}_0$ and therefore $p^* = \mu(\varepsilon, \Lambda_0, p_0)$. \square

A.2 Quasi-Single Crossing Property for Equilibrium Incentives

A.2.1 Proof of Theorem 1.4.1 under Perfect Good News

From now on we drop the superscript N from W and p .

Proof. The proof consists of two steps. In the first step, we show that whenever $W_t = 2p_t - 1$, then $W_\tau \geq 2p_\tau - 1$ for all $\tau \geq t$. In the second step, we show that whenever $W_t > 2p_t - 1$, then $W_\tau > 2p_\tau - 1$ for all $\tau > t$.

Step 1: Suppose $W_t = 2p_t - 1$ at some time t and suppose for a contradiction that at some time $s' > t$, we have $W_{s'} < 2p_{s'} - 1$. Let

$$s^* = \sup\{s < s' : W_s = 2p_s - 1\}.$$

By continuity, $s^* < s'$, $W_{s^*} = 2p_{s^*} - 1$, and $W_s < 2p_s - 1$ for all $s \in (s^*, s')$. Then by

Lemma A.1.2, the right hand derivative of $W_s - (2p_s - 1)$ at s^* exists and satisfies:

$$\lim_{s \downarrow s^*} \dot{W}_s - 2\dot{p}_s = r(2p_{s^*} - 1) + p_{s^*} (\varepsilon_G + \lambda_G N_{s^*}) \frac{r}{\rho + r} > 0.$$

This implies that for some $s \in (s^*, s')$ sufficiently close to s^* we have $W_s > 2p_s - 1$, which is a contradiction.

Step 2: Assume $W_t > 2p_t - 1$ at some t and suppose for a contradiction that there exists $s' > t$ such that $W_{s'} = 2p_{s'} - 1$. Let

$$s^* = \inf\{s > t : W_s = 2p_s - 1\}.$$

By continuity, $s^* > t$, $W_{s^*} = 2p_{s^*} - 1$, and $W_s > 2p_s - 1$ for all $s \in (t, s^*)$. Then by Lemma A.1.3 the left-hand derivative of $W_s - (2p_s - 1)$ at s^* exists and is given by:

$$\lim_{s \uparrow s^*} \dot{W}_s - 2\dot{p}_s = r(2p_{s^*} - 1) + p_{s^*} \frac{r}{\rho + r} \varepsilon_G > 0.$$

This implies that for some $s \in (t, s^*)$ sufficiently close to s^* , we must have $W_s < 2p_s - 1$, which is a contradiction. \square

A.2.2 Proof of Theorem 1.4.1 under Perfect Bad News

Proof. The proof consists of two steps. In the first step, we show that whenever $W_t = 2p_t - 1$, then $W_\tau \leq 2p_\tau - 1$ for all $\tau \geq t$. In the second step, we show that whenever $W_t < 2p_t - 1$, then $W_\tau < 2p_\tau - 1$ for all $\tau > t$.

Step 1: Suppose $W_t = 2p_t - 1$ at some time t and suppose for a contradiction that at some time $s' > t$ we have $W_{s'} > 2p_{s'} - 1$. Then because $W_t \rightarrow \frac{\rho}{\rho+r} (2\mu(\varepsilon, \Lambda, p_0) - 1) < 2\mu(\varepsilon, \Lambda, p_0) - 1$ by Lemma A.1.4, there exists $\underline{s} < \bar{s}$ such that $W_{\underline{s}} = 2p_{\underline{s}} - 1$, $W_{\bar{s}} = 2p_{\bar{s}} - 1$, and $W_s > 2p_s - 1$ for all $s \in (\underline{s}, \bar{s})$. By Lemma A.1.3, we have the following two limits:

$$\lim_{s \downarrow \underline{s}} \dot{W}_s = (r + (1 - p_{\underline{s}})\varepsilon)(2p_{\underline{s}} - 1).$$

$$\lim_{s \uparrow \bar{s}} \dot{W}_s = (r + (1 - p_{\bar{s}})\varepsilon)(2p_{\bar{s}} - 1).$$

Also, as usual

$$\lim_{s \downarrow \underline{s}} \frac{d}{ds} (2p_s - 1) = 2p_{\underline{s}}(1 - p_{\underline{s}})\varepsilon$$

$$\lim_{s \uparrow \bar{s}} \frac{d}{ds} (2p_s - 1) = 2p_{\bar{s}}(1 - p_{\bar{s}})\varepsilon.$$

In order that $W_s > 2p_s - 1$ for all $s \in (\underline{s}, \bar{s})$, we need:

$$(r + (1 - p_{\underline{s}})\varepsilon)(2p_{\underline{s}} - 1) \geq 2p_{\underline{s}}(1 - p_{\underline{s}})\varepsilon$$

$$(r + (1 - p_{\bar{s}})\varepsilon)(2p_{\bar{s}} - 1) \leq 2p_{\bar{s}}(1 - p_{\bar{s}})\varepsilon.$$

Rearranging we get:

$$r(2p_{\underline{s}} - 1) \geq (1 - p_{\underline{s}})\varepsilon$$

$$r(2p_{\bar{s}} - 1) \leq (1 - p_{\bar{s}})\varepsilon.$$

But this is impossible given that $p_{\bar{s}} > p_{\underline{s}}$. This completes the proof of Step 1.

Step 2: Suppose that $W_t < 2p_t - 1$ and suppose for a contradiction that there exists some $s' > t$ such that $W_{s'} \geq 2p_{s'} - 1$. Define

$$\underline{s} = \inf\{s' > t : W_{s'} \geq 2p_{s'} - 1\}.$$

By continuity, $W_\tau < 2p_\tau - 1$ for all $\tau \in [t, \underline{s})$ and $W_{\underline{s}} = 2p_{\underline{s}} - 1$.

Furthermore, by Lemma A.1.4, there exists some $\bar{s} \geq \underline{s}$ such that $2p_{\bar{s}} - 1 = W_{\bar{s}}$ and $2p_s - 1 > W_s$ for all $s > \bar{s}$. By Lemma A.1.2, we have the following two limits:

$$H_{\underline{s}} \equiv \lim_{s \uparrow \underline{s}} \left(\dot{W}_s - \frac{d}{ds} (2p_s - 1) \right) = r(2p_{\underline{s}} - 1) - (\varepsilon + \lambda\rho\bar{N}_{\underline{s}})(1 - p_{\underline{s}})$$

$$H_{\bar{s}} \equiv \lim_{s \downarrow \bar{s}} \left(\dot{W}_s - \frac{d}{ds} (2p_s - 1) \right) = r(2p_{\bar{s}} - 1) - (\varepsilon + \lambda\rho\bar{N}_{\bar{s}})(1 - p_{\bar{s}}).$$

As usual, because $W_s < 2p_s - 1$ for all $s \in (t, \underline{s})$ and for all $s > \bar{s}$, we must have $H_{\underline{s}} \geq 0$ and $H_{\bar{s}} \leq 0$. But since $p_{\bar{s}} \geq p_{\underline{s}}$, this is only possible if $\underline{s} = \bar{s} =: s^*$ and $H_{s^*} = H_{\underline{s}} = H_{\bar{s}} = 0$.

Thus,

$$r(2p_{s^*} - 1) = (\varepsilon + \lambda\rho\bar{N}_{s^*})(1 - p_{s^*}).$$

Now consider any $s \in [t, s^*]$. Because $p_s \leq p_{s^*}$ we must have

$$r(2p_s - 1) \leq (\varepsilon + \lambda\rho\bar{N}_s)(1 - p_s).$$

Combining this with the fact that $W_s < 2p_s - 1$ and $N_s = \rho\bar{N}_s$ yields

$$rW_s < (\varepsilon + \lambda\rho\bar{N}_s)(1 - p_s) < (2p_s - W_s)(\varepsilon + \lambda\rho\bar{N}_s)(1 - p_s) + \rho(2p_s - 1 - W_s).$$

Rearranging we obtain:

$$0 < -rW_s + \rho(2p_s - 1 - W_s) + (2p_s - W_s)(\varepsilon + \lambda\rho\bar{N}_s)(1 - p_s).$$

But by Lemma A.1.2, the right-hand side is precisely the derivative $\frac{d}{ds}(2p_s - 1) - \dot{W}_s$. This implies that for all $s \in [t, s^*]$, $2p_s - 1 - W_s$ is strictly increasing, contradicting continuity and the fact that $2p_{s^*} - 1 = W_{s^*}$. This concludes the proof of Step 2. \square

A.3 Equilibrium Uniqueness and Characterization

A.3.1 Equilibrium Characterization under Perfect Bad News

In this section, we do not impose Conditions 1.5.2 or 1.5.3. Recall that $p^* := \min\{\bar{p}, p^\#\}$, where

$$\begin{aligned}\bar{p} &:= \frac{\varepsilon + r}{\varepsilon + 2r}, \\ p^\# &:= \frac{\rho + r}{\rho + 2r}.\end{aligned}$$

Recall the definition of $G : [0, 1] \times \mathbb{R}_+ \rightarrow \mathbb{R}$:

$$G(p, \Lambda) := \int_0^\infty \rho e^{-(r+\rho)\tau} \left(p - (1-p)e^{-(\varepsilon\tau + \Lambda(1-e^{-\rho\tau}))} \right) d\tau.$$

We extend the function to the domain $[0, 1] \times (\mathbb{R}_+ \cup \{+\infty\})$ by defining:

$$G(p, +\infty) := \frac{\rho}{\rho + r} p.$$

Finally, recall the definition of $\Lambda^* : (0, 1) \rightarrow \mathbb{R}_+ \cup \{+\infty\}$:

$$\begin{cases} \Lambda^*(p) = 0 & \text{if } p \leq \underline{p}, \\ 2p - 1 = G(p, \Lambda^*(p)) & \text{if } p \in (\underline{p}, p^\sharp) \\ \Lambda^*(p) = +\infty & p \geq p^\sharp. \end{cases}$$

The proof of Theorem 1.5.1 proceeds in three steps. Assuming that N is an equilibrium, we show in Lemma A.3.1 that if $t_1^* < t_2^*$, then the evolution of adoption behavior on (t_1^*, t_2^*) is uniquely pinned down by an ODE. We next prove Lemma 1.5.4, which provides a characterization of t_1^* and t_2^* in terms of (p_t, Λ_t) . Given these two steps uniqueness is clear. Finally, we check feasibility in Lemma A.3.4, proving equilibrium existence.

Characterization of Adoption between t_1^* and t_2^*

Lemma A.3.1. *Suppose N is an equilibrium with no-news posterior p_t . Suppose that $t_1^* < t_2^*$. Then at (almost) all times $t \in (t_1^*, t_2^*)$,*

$$N_t = \frac{r(2p_t - 1)}{\lambda(1 - p_t)} - \frac{\varepsilon}{\lambda}.$$

Proof. Note that because $2p_t - 1 = W_t^N$ at all $t \in (t_1^*, t_2^*)$ and p_t is weakly increasing, W_t^N and p_t are differentiable almost everywhere (with respect to Lebesgue measure).

Using again the fact that $2p_t - 1 = W_t^N$ at all $t \in (t_1^*, t_2^*)$ we obtain for all $t \in (t_1^*, t_2^*)$:

$$\begin{aligned} W_t^N &= e^{-r(t_2^* - t)} \left(p_t + (1 - p_t) e^{-\int_t^{t_2^*} (\varepsilon + \lambda N_s) ds} \right) (2p_{t_2^*} - 1) \\ &= e^{-r(t_2^* - t)} \left(p_t - (1 - p_t) e^{-\int_t^{t_2^*} (\varepsilon + \lambda N_s) ds} \right). \end{aligned}$$

Then for all t at which W_t^N and p_t are differentiable, we obtain:

$$\begin{aligned} \dot{W}_t^N &= (r + (\varepsilon + \lambda N_t)(1 - p_t)) W_t^N \\ 2\dot{p}_t &= 2p_t(1 - p_t)(\varepsilon + \lambda N_t). \end{aligned}$$

Furthermore, because $W_t^N = 2p_t - 1$ for all $t \in (t_1^*, t_2^*)$, we must have for almost all

$t \in (t_1^*, t_2^*)$:

$$\dot{W}_t^N = 2\dot{p}_t.$$

This means that for almost all $t \in (t_1^*, t_2^*)$:

$$N_t = \frac{r(2p_t - 1)}{\lambda(1 - p_t)} - \frac{\varepsilon}{\lambda}.$$

□

A direct corollary of the above lemma is the following:

Corollary A.3.2. *The posterior at all $t \in (t_1^*, t_2^*)$ evolves according to the following ordinary differential equation:*

$$\dot{p}_t = rp_t(2p_t - 1).$$

Given some initial condition $p = p_{t_1^*}$, this ordinary differential equation admits a unique solution, given by:

$$p_t = \frac{p_{t_1^*}}{2p_{t_1^*} - e^{r(t-t_1^*)}(2p_{t_1^*} - 1)}.$$

Proof of Lemma 1.5.4

We now prove a more general version of Lemma 1.5.4 in which we replace \bar{p} in Lemma 1.5.4 with p^* .

Lemma A.3.3. *Let N be an equilibrium with corresponding no-news posterior $p_{t \geq 0}$ and threshold times t_1^* and t_2^* , and let $\Lambda_{t \geq 0} := \lambda \bar{N}_{t \geq 0}$ describe the evolution of the economy's potential for social learning. Then*

1. $t_2^* = \inf\{t : \Lambda_t < \Lambda^*(p_t)\}$; and
2. $t_1^* = \min\{t_2^*, \sup\{t : p_t < p^*\}\}$.¹

Proof. We first prove both bullet points under the assumption that either $\varepsilon > 0$ or $p_0 > \frac{1}{2}$. Note that in this case Lemma A.1.4 implies that $t_2^* < +\infty$ and we must also have that p_t is strictly increasing for all $t > 0$.

¹We impose the convention that if $\{t \geq 0 : p_t < p^* = \frac{1}{2}\} = \emptyset$, then $\sup\{t \geq 0 : p_t < p^* = \frac{1}{2}\} := 0$.

For the first bullet point, note that by definition of t_2^* and by Theorem 1.4.1, we have $2p_t - 1 > W_t = G(p_t, \Lambda_t)$ for all $t > t_2^*$. This implies that $\Lambda_t < \Lambda^*(p_t)$ for all $t > t_2^*$. Moreover, if $0 < t_2^*$, then by continuity we must have $2p_{t_2^*} - 1 = W_{t_2^*} = G(p_{t_2^*}, \Lambda_{t_2^*})$ and so $\Lambda_{t_2^*} = \Lambda^*(p_{t_2^*})$. In this case, because Λ_s is decreasing in s and p_s is strictly increasing in s and $\Lambda^*(p)$ is increasing in p , we must have $\Lambda_s \geq \Lambda^*(p_s)$ for all $s < t_2^*$. This establishes the first bullet point.

For the second bullet point, it suffices to prove the following three claims:

- (a) If $t_2^* > 0$, then $p_{t_2^*} < p^\sharp$.
- (b) If $t_1^* > 0$, then $p_{t_1^*} \leq \bar{p}$.
- (c) If $t_1^* < t_2^*$, then $p_{t_1^*} \geq \bar{p}$.

Indeed, given (a) and (b), we have that if $0 < t_1^* = t_2^*$, then $p_{t_1^*} \leq p^*$. Given (a)-(c), we have that if $0 < t_1^* < t_2^*$, then $p_{t_1^*} = \bar{p} = p^*$. If $0 = t_1^* < t_2^*$, then (c) implies that $p_0 \geq \bar{p} = p^*$. In all three cases (ii) readily follows. Finally, if $0 = t_1^* = t_2^*$, then there is nothing to prove.

For claim (a), recall from the above that if $t_2^* > 0$, then $\Lambda_{t_2^*} = \Lambda^*(p_{t_2^*})$, whence $p_{t_2^*} < p^\sharp$ because $\Lambda^*(p^\sharp) = +\infty$.

For claim (b), note that if $t_1^* > 0$, then for all $t < t_1^*$, we have $W_t > 2p_t - 1$. Then by Lemma A.1.3 and because $W_{t_1^*} = 2p_{t_1^*} - 1$, we must have

$$\begin{aligned} 0 &\geq \lim_{\tau \uparrow t_1^*} \dot{W}_\tau - 2\dot{p}_\tau = (r + (1 - p_{t_1^*})\varepsilon)(2p_{t_1^*} - 1) - 2p_{t_1^*}(1 - p_{t_1^*})\varepsilon \\ &= r(2p_{t_1^*} - 1) - \varepsilon(1 - p_{t_1^*}), \end{aligned}$$

which implies that

$$p_{t_1^*} \leq \frac{\varepsilon + r}{\varepsilon + 2r} =: \bar{p}.$$

Finally, for claim (c), note that if $t_1^* < t_2^*$, then Lemma A.3.1 implies that for all $\tau \in (t_1^*, t_2^*)$,

$$0 \leq N_\tau = \frac{r(2p_\tau - 1)}{\lambda(1 - p_\tau)} - \frac{\varepsilon}{\lambda}.$$

This implies that for all $\tau \in (t_1^*, t_2^*)$,

$$p_\tau \geq \frac{\varepsilon + r}{\varepsilon + 2r} =: \bar{p},$$

and hence by continuity $p_{t_1^*} \geq \bar{p}$ as claimed. This proves the lemma when either $\varepsilon > 0$ or $p_0 > \frac{1}{2}$. Finally, if $\varepsilon = 0$ and $p_0 \leq \frac{1}{2}$, then it is easy to see that $p_t \equiv p_0$ for all t . Thus, $t_2^* = +\infty = \inf\{t : \Lambda_t < \Lambda^*(p_0) = 0\}$. Also, if $p_0 < \frac{1}{2}$, then $t_1^* = +\infty = \sup\{t : p_t < p^* = \frac{1}{2}\}$; and if $p_0 = \frac{1}{2}$, then $t_1^* = 0 =: \sup\{t \geq 0 : p_t < p^* = \frac{1}{2}\}$. \square

With these lemmas, it is immediate that if an equilibrium exists, then it must take the form of the adoption flow given by Equation 1.6 in Theorem 1.5.1. Moreover, it is easy to see that given initial parameters, Equation 1.6 uniquely pins down the times t_1^* and t_2^* as well as the joint evolution of p_t and N_t at all times (we elaborated on this in the main text), and that whenever $t_1^* < t_2^* < +\infty$, then $2p_t - 1 = W_t$ for all $t \in [t_1^*, t_2^*]$. Provided feasibility is satisfied, it is then easy to check that this adoption flow constitutes an equilibrium.

Feasibility

It remains to check feasibility of the adoption flow implied by Equation 1.6 in Theorem 1.5.1. Note that feasibility is non-trivial only at times $t \in (t_1^*, t_2^*)$.

Lemma A.3.4. *Suppose $N_{t \geq 0}$ is an adoption flow satisfying Equation 1.6 in Theorem 1.5.1 such that $t_1^* < t_2^*$. Then for all $t \in (t_1^*, t_2^*)$,*

$$N_t \leq \rho \Lambda_t.$$

Proof. It suffices to show that

$$\lim_{t \uparrow t_2^*} N_t \leq \rho \bar{N}_{t_2^*}.$$

The lemma then follows immediately since $\rho \bar{N}_t - N_t$ is strictly decreasing in t at all times in (t_1^*, t_2^*) .²

²This is only true if either $\varepsilon > 0$ or $p_0 > \frac{1}{2}$. If $\varepsilon = 0$ and $p_0 = \frac{1}{2}$, then $N_t = 0$ for all t and $t_1^* = 0 < t_2^* = +\infty$. But in this case feasibility is immediate.

To see this, suppose by way of contradiction that $\rho \bar{N}_{t_2^*} < \lim_{t \uparrow t_2^*} N_t$. By continuity this means that there exists some $\nu > 0$ such that $\rho \bar{N}_t < N_t$ for all $t \in (t_2^* - \nu, t_2^*)$. Then note that from the indifference condition at t_2^* , we have that $2p_{t_2^*} - 1 = G(p_{t_2^*}, \lambda \bar{N}_{t_2^*})$. Furthermore because $\Lambda^*(p_t)$ is increasing in t , $2p_t - 1 < G(p_t, \lambda \bar{N}_t)$ for all $t < t_2^*$.

Since at all times $t \in (t_2^* - \nu, t_2^*)$ we have $N_t > \rho \bar{N}_t$, this implies that

$$W_t^N > G(p_t, \lambda \bar{N}_t) > 2p_t - 1.$$

But this is a contradiction since we already checked that the described adoption flow satisfies the condition that $W_t^N = 2p_t - 1$ for all $t \in (t_1^*, t_2^*)$. \square

A.3.2 Equilibrium Characterization under Perfect Good News

Theorem 1.6.1 follows readily from Lemma 1.6.2 and Lemma 1.6.3. Lemma 1.6.3 was proved in the text. It remains to prove Lemma 1.6.2.

Proof of Lemma 1.6.2: Suppose for a contradiction that $t_1^* < t_2^*$. From the definition of these cutoffs, we must have $2p_t - 1 = W_t$ for all $t \in [t_1^*, t_2^*]$. Then for all $t \in (t_1^*, t_2^*)$ and $\Delta \in (0, t_2^* - t)$ we have:

$$\begin{aligned} W_t = & p_t \int_t^{t+\Delta} (\varepsilon + \lambda N_\tau) e^{-\int_t^\tau (\varepsilon + \lambda N_s) ds} e^{-r(\tau-t)} \frac{\rho}{\rho + r} d\tau + \\ & \left((1 - p_t) + p_t e^{-\int_t^{t+\Delta} (\varepsilon + \lambda N_s) ds} \right) e^{-r\Delta} (2p_{t+\Delta} - 1), \end{aligned}$$

where the first term represents a breakthrough arriving at some $\tau \in (t, t + \Delta)$ in which case consumers adopt from then on, yielding a payoff of $e^{-r(\tau-t)} \frac{\rho}{\rho+r}$; and the second term represents no breakthrough arriving prior to $t + \Delta$ in which case, due to indifference, consumers' payoff can be written as $e^{-r\Delta} (2p_{t+\Delta} - 1)$.

Note that we must have $p_t \geq \frac{1}{2}$ on (t_1^*, t_2^*) , since W_t is bounded below by 0. Given that we assume that either $\varepsilon = 0$ or $p_0 \neq \frac{1}{2}$, this means that either $\varepsilon > 0$ or $p_t > \frac{1}{2}$ for t sufficiently

close to t_1^* . Then it follows that for sufficiently small Δ

$$\begin{aligned} W_t &< p_t \left(1 - e^{-\int_t^{t+\Delta} (\varepsilon + \lambda N_s) ds}\right) \frac{\rho}{\rho + r} + \left((1 - p_t) + p_t e^{-\int_t^{t+\Delta} (\varepsilon + \lambda N_s) ds}\right) (2p_{t+\Delta} - 1) \\ &\leq p_t \left(1 - e^{-\int_t^{t+\Delta} (\varepsilon + \lambda N_s) ds}\right) \cdot 1 + \left((1 - p_t) + p_t e^{-\int_t^{t+\Delta} (\varepsilon + \lambda N_s) ds}\right) (2p_{t+\Delta} - 1) \\ &= 2p_t - 1, \end{aligned}$$

where the final equality comes from Bayesian updating of beliefs. This contradicts $W_t = 2p_t - 1$. Thus, $t_1^* = t_2^*$. \square

A.4 Violation of Condition 1.5.2 under Perfect Bad News




Figure A.1: Partition of (p_t, Λ_t) when $\varepsilon \geq \rho$

In this section, we discuss the case in which $\rho \geq \varepsilon$. We saw in Theorem 1.5.1 that the characterization theorem holds even when Condition 1.5.2 is violated.

In this case because $\Lambda^*(p) = +\infty$ for all $p > p^*$, we have:

$$N_t = \begin{cases} 0 & \text{if } \Lambda_t > \Lambda^*(p_t), \\ \rho \bar{N}_t & \text{if } \Lambda_t \leq \Lambda^*(p_t). \end{cases}$$

Note that now partial adoption never occurs and the unique equilibrium reduces to all-or-nothing adoption.

As a result the saturation effect discussed in Section 1.5 is no longer present and welfare always strictly increases in response to an increase in opportunities for social learning:

Proposition A.4.1. *Fix $r > 0$ and $p_0 \in (0, 1)$ and suppose that $\varepsilon \geq \rho > 0$. Then W_0 is strictly increasing in Λ_0 .*

A.5 Inefficiency of Equilibria

A.5.1 Inefficiency under PBN

Proof of Proposition 1.5.8: From Proposition 1.3.3, recall that

$$p^s = \frac{K(\Lambda_0)}{K(\Lambda_0) + \frac{\rho}{r+\rho} \frac{r}{r+\varepsilon}},$$

where

$$K(\Lambda_0) = \int_0^{\infty} \rho e^{-(r+\rho)\tau} e^{-\varepsilon\tau - \Lambda_0(1-e^{-\rho\tau})} d\tau.$$

Note also that

$$K(\Lambda_0) < \frac{\rho}{r + \varepsilon + \rho}$$

which then implies that

$$p^s < \frac{(r + \rho)(r + \varepsilon)}{2(\varepsilon + r)(r + \rho) - \varepsilon\rho} = \underline{p}.$$

Finally observe from the proof of Lemma 1.5.4 that $p_{t_1^*} \geq \underline{p}$.

If $\Lambda_0 > \Lambda^*(\bar{p})$, either $t_1^* > 0$ or $t_2^* > 0$. In the first case, adoption begins too late since $p_{t_1^*} \geq \underline{p} > p^s$ and therefore equilibrium is inefficient. If on the other hand, $t_1^* = 0 < t_2^*$, then again because $p^s < p_{t_1^*}$, adoption is too slow initially since consumers only partially adopt

between t_1^* and t_2^* . Thus again equilibrium is inefficient.

On the other hand, if $\Lambda_0 \leq \Lambda^*(p_0)$, then equilibrium is efficient since both the cooperative benchmark and equilibrium prescribe that absent breakdowns all consumers adopt whenever given an opportunity. \square

A.6 Comparative Statics under PBN

A.6.1 Saturation Effect: Proof of Proposition 1.5.9

Throughout Section A.6 we impose Conditions 1.5.2 (so that $p^* = \bar{p}$) and 1.5.3 as in the text. We first prove Lemma 1.5.6.

Proof of Lemma 1.5.6: Let $\bar{\Lambda}_0 := \max\{\Lambda^*(p_0), \Lambda^*(\bar{p})\}$. We show that $t_1^*(\Lambda_0) < t_2^*(\Lambda_0)$ if and only if $\Lambda_0 > \bar{\Lambda}_0$.

Suppose first that $\Lambda_0 > \bar{\Lambda}_0$. Then by the proof of the first part of Lemma 1.5.4, we must have $t_2^* > 0$ and $\Lambda_{t_2^*} = \Lambda^*(p_{t_2^*})$. If $t_1^* = t_2^* =: t^*$, then by claims (a) and (b) in the proof of Lemma 1.5.4, we must have $p_{t^*} \leq \bar{p}$. But combining these statements, we get

$$\Lambda_{t^*} = \Lambda_0 > \Lambda^*(\bar{p}) \geq \Lambda^*(p_{t^*}) = \Lambda_{t^*},$$

which is a contradiction.

Suppose conversely that $t_1^* < t_2^*$. Then by the proof of Lemma 1.5.4, we have that $\Lambda^*(p_{t_1^*}) < \Lambda_{t_1^*} = \Lambda_0$. That proof also implies that if $0 < t_1^* < t_2^*$, then $p_{t_1^*} = \bar{p} \geq p_0$; and if $0 = t_1^* < t_2^*$, then $p_{t_1^*} = p_0 \geq \bar{p}$. Thus, either way $\Lambda_0 > \bar{\Lambda}_0$, as claimed. \square

Proof of Proposition 1.5.9: For the first bullet point, consider any $\Lambda_0^1 < \Lambda_0^2 \leq \bar{\Lambda}_0$ with corresponding threshold times t_1^i and t_2^i , value to waiting W_t^i , and no-news posteriors p_t^i for $i = 1, 2$. By Lemma 1.5.6, we must have $t_1^i = t_2^i =: t^i$. Let $\hat{t} := \min\{t^1, t^2\}$. Then note that for all $t \leq \hat{t}$, $p_t^1 = p_t^2$ and $\Lambda_t^i = \Lambda_0^i$. By Lemma 1.5.4 this implies that either $0 = t^1 = t^2$ or $t^1 < t^2$. If $0 = t^1 = t^2$, then for all $t > 0$, we have $2p_t^i - 1 > W_t^i$ and

$$p_t^i = \frac{p_0}{p_0 + (1 - p_0)e^{-(\epsilon t + (1 - e^{-\rho t})\Lambda_0^i)}}.$$

Thus, $p_t^1 < p_t^2$ for all $t > 0$. Then by Lemma 1.2.4, $W_0^1 < W_0^2$.

If $t^1 < t^2$, then by definition of the cutoff times

$$W_{t^1}^2 > 2p_{t^1}^2 - 1 = 2p_{t^1}^1 - 1 \geq W_{t^1}^1.$$

Since there is no adoption until t^1 , we have

$$W_0^i = e^{-rt^1} \frac{p_{t^1}^i}{p_0} W_{t^1}^i,$$

which again implies that $W_0^1 < W_0^2$. This proves the first bullet point.

To prove the second bullet point, suppose that $\Lambda_0^2 > \Lambda_0^1 > \bar{\Lambda}_0$. By Lemma 1.5.6, we then have $t_1^i < t_2^i$ for $i = 1, 2$. Moreover, by the proof of Lemma 1.5.4, we have $\max\{p_0, \bar{p}\} = p_{t_1^1}^1 = p_{t_1^2}^2$. Because $N_t^i = N_t^i = 0$ for all $t < t_1^i$ for both $i = 1, 2$, this implies that $t_1^1 = t_1^2 = t_1$. Then

$$W_{t_1}^2 = 2p_{t_1}^2 - 1 = 2p_{t_1}^1 - 1 = W_{t_1}^1.$$

But once again,

$$W_0^i = e^{-rt^1} \frac{p_{t_1}^i}{p_0} W_{t_1}^i,$$

for $i = 1, 2$, whence $W_0^1 = W_0^2$. □

A.6.2 Learning Dynamics

Proof of Lemma 1.5.10

Proof of Lemma 1.5.10. Suppose that $\hat{\Lambda}_0 > \Lambda_0 > \Lambda^*(p_0)$. Recall that we are assuming $p^\# > p_0 > \bar{p}$ so that $t_1^*(\Lambda_0) = t_1^*(\hat{\Lambda}) = 0$ and $\bar{\Lambda}_0 = \Lambda^*(p_0)$. Then by Lemma 1.5.6, we have $t_2^*(\hat{\Lambda}_0), t_2^*(\Lambda_0) > 0$. Let $t_2^* = \min\{t_2^*(\hat{\Lambda}_0), t_2^*(\Lambda_0)\}$. Then because $p_0 = p_0^{\hat{\Lambda}_0} = p_0^{\Lambda_0}$, the ODE in Corollary A.3.2 implies that at all times $t < t_2^*$, we have $p_t^{\hat{\Lambda}_0} = p_t^{\Lambda_0} = p_t$. By Lemma A.3.1, this implies that for all $t < t_2^*$,

$$\lambda N_t = \hat{\lambda} \hat{N}_t. \tag{A.1}$$

To prove the first bullet point, note that Equation A.1 implies that

$$\Lambda_{t_2^*} = \Lambda_0 - \int_0^{t_2^*} \lambda N_t dt < \hat{\Lambda}_0 - \int_0^{t_2^*} \hat{\lambda} \hat{N}_t dt = \hat{\Lambda}_{t_2^*}.$$

By Lemma 1.5.4 and because $p_{t_2^*}^{\Lambda_0} = p_{t_2^*}^{\hat{\Lambda}_0}$, this implies that $t_2^* = t_2^*(\Lambda_0) < t_2^*(\hat{\Lambda}_0)$.

From this and Equation A.1, it is then immediate that $\lambda N_t = \hat{\lambda} \hat{N}_t$ for all $t < t_2^* = t_2^*(\Lambda_0)$, proving the second bullet point. \square

Proof of Proposition 1.5.11

Proof. Clearly $p_t^{\Lambda_0}$ is strictly increasing for all $\Lambda_0 \in (0, \Lambda^*(p_0))$ since in this case $t_2^*(\Lambda_0) = 0$ so that

$$p_t^{\Lambda_0} = \frac{p_0}{p_0 + (1 - p_0)e^{-(\epsilon t + (1 - e^{-\rho t})\Lambda_0)}}.$$

Suppose next that $\hat{\Lambda}_0 > \Lambda_0 \geq \Lambda^*(p_0)$. By Lemma 1.5.10, $t^* := t_2^*(\Lambda_0) < t_2^*(\hat{\Lambda}_0)$, $\lambda N_t = \hat{\lambda} \hat{N}_t$, and $p_t^{\Lambda_0} = p_t^{\hat{\Lambda}_0}$ for all $t \leq t^*$, which proves the first bullet point.

To prove the second bullet point, we claim that there exists some $\nu > 0$ such that at all times $t \in (t^*, t^* + \nu)$, we have $p_t^{\Lambda_0} > p_t^{\hat{\Lambda}_0}$. To see this, we prove the following inequality for the equilibrium corresponding to Λ_0 :

$$\lim_{t \uparrow t^*} \lambda N_t < \lim_{t \downarrow t^*} \lambda N_t. \quad (\text{A.2})$$

In other words, there is necessarily a discontinuity in the equilibrium flow of adoption at exactly t^* . Indeed, because $N_t = \rho \bar{N}_t$ for all $t \geq t^*$ and by continuity of \bar{N}_t , feasibility implies that $\lim_{t \uparrow t^*} \lambda N_t \leq \lim_{t \downarrow t^*} \lambda N_t$. Suppose for a contradiction that $\lim_{t \uparrow t^*} \lambda N_t = \lim_{t \downarrow t^*} \lambda N_t := \lambda N_{t^*}$. Then $\lambda N_{t^*} = \hat{\lambda} \hat{N}_{t^*}$. Moreover, for all $t > t^*$, we have $\lambda N_t = \rho \Lambda_{t^*} e^{-\rho(t - t^*)}$, which is strictly decreasing in t . On the other hand, $\hat{\lambda} \hat{N}_t$ satisfies

$$\hat{\lambda} \hat{N}_t = \begin{cases} \frac{r(2p_t - 1)}{(1 - p_t)} - \epsilon & \text{if } t < t_2^*(\hat{\Lambda}_0) \\ \rho \Lambda_{t_2^*(\hat{\Lambda}_0)} e^{-\rho(t - t_2^*(\hat{\Lambda}_0))} & \text{if } t \geq t_2^*(\hat{\Lambda}_0). \end{cases}$$

Thus, for $t \in [t^*, t_2^*(\hat{\Lambda}_0))$, $\hat{\lambda} \hat{N}_t$ is strictly *increasing* in t . This implies that $\hat{\lambda} \hat{N}_t > \lambda N_t$ for all

$t \in [t^*, t_2^*(\hat{\Lambda}_0))$. But then by Equation 1.1,

$$p_{t_2^*(\hat{\Lambda}_0)}^{\hat{\Lambda}_0} > p_{t_2^*(\hat{\Lambda}_0)}^{\Lambda_0},$$

which by Lemma 1.5.4 implies

$$\hat{\Lambda}_{t_2^*(\hat{\Lambda}_0)} = \Lambda^*(p_{t_2^*(\hat{\Lambda}_0)}^{\hat{\Lambda}_0}) > \Lambda^*(p_{t_2^*(\hat{\Lambda}_0)}^{\Lambda_0}) > \Lambda_{t_2^*(\hat{\Lambda}_0)}.$$

This yields that for all $t \geq t_2^*(\hat{\Lambda}_0)$

$$\hat{\lambda}\hat{N}_t = \rho e^{-\rho(t-t_2^*(\hat{\Lambda}_0))} \hat{\Lambda}_{t_2^*(\hat{\Lambda}_0)} > \rho e^{-\rho(t-t_2^*(\hat{\Lambda}_0))} \Lambda_{t_2^*(\hat{\Lambda}_0)} = \lambda N_t.$$

Thus, $\hat{\lambda}\hat{N}_t > \lambda N_t$ for all $t > t^*$ and hence $p_t^{\hat{\Lambda}_0} > p_t^{\Lambda_0}$ for all $t > t^*$. By Lemma 1.2.4, this implies

$$W_{t^*}^{\hat{\Lambda}_0} > W_{t^*}^{\Lambda_0}.$$

But this is a contradiction, because we have that

$$W_{t^*}^{\hat{\Lambda}_0} = 2p_{t^*}^{\hat{\Lambda}_0} - 1 = 2p_{t^*}^{\Lambda_0} - 1 = W_{t^*}^{\Lambda_0}.$$

This proves that $\lim_{t \uparrow t^*} \lambda N_t < \lim_{t \downarrow t^*} \lambda N_t$. But then,

$$\lim_{t \downarrow t^*} \hat{\lambda}\hat{N}_t = \lim_{t \uparrow t^*} \hat{\lambda}\hat{N}_t = \lim_{t \uparrow t^*} \lambda N_t < \lim_{t \downarrow t^*} \lambda N_t.$$

Therefore there must exist some $\nu > 0$ such that $\hat{\lambda}\hat{N}_t < \lambda N_t$ for all $t \in [t^*, t^* + \nu)$. Together with the fact that $p_t^{\Lambda_0} = p_t^{\hat{\Lambda}_0}$, this implies that $p_t^{\Lambda_0} > p_t^{\hat{\Lambda}_0}$ for all $t \in (t^*, t^* + \nu)$, proving the second bullet point of the proposition.

Finally, for the third bullet point, observe first that there must exist some $t > t^*$ such that $p_t^{\Lambda_0} = p_t^{\hat{\Lambda}_0}$. If not, then by continuity of beliefs $p_t^{\Lambda_0} > p_t^{\hat{\Lambda}_0}$ for all $t > t^*$, and we once again get that $W_{t^*}^{\hat{\Lambda}_0} > W_{t^*}^{\Lambda_0}$, which is false. Then $\bar{t} := \sup\{s \in (t^*, t) : p_s^{\Lambda_0} > p_s^{\hat{\Lambda}_0}\}$ exists, with $\bar{t} > t^*$ by the second bullet point. Further, by continuity, $p_{\bar{t}}^{\Lambda_0} = p_{\bar{t}}^{\hat{\Lambda}_0}$, which implies $\int_0^{\bar{t}} \lambda N_s ds = \int_0^{\bar{t}} \hat{\lambda}\hat{N}_s ds$. This yields $\Lambda_{\bar{t}} < \hat{\Lambda}_{\bar{t}}$. But note that this implies that $\hat{\lambda}\hat{N}_t > \lambda N_t$ for all $t > \bar{t}$: Indeed, if $\bar{t} \geq t_2^*(\hat{\Lambda}_0)$, this is obvious. On the other hand, if $\bar{t} \in (t^*, t_2^*(\hat{\Lambda}_0))$, then we must have $\lambda N_s < \hat{\lambda}\hat{N}_s$ for some $s < \bar{t}$, which implies that $\lambda N_{s'} < \hat{\lambda}\hat{N}_{s'}$ for all $s' \in (s, t_2^*(\hat{\Lambda}_0))$,

because N is strictly decreasing and \hat{N} is strictly increasing on this domain. This implies that

$$p_{t_2^*(\hat{\Lambda}_0)}^{\hat{\Lambda}_0} > p_{t_2^*(\hat{\Lambda}_0)}^{\Lambda_0},$$

which as above implies that

$$\hat{\Lambda}_{t_2^*(\hat{\Lambda}_0)} = \Lambda^*(p_{t_2^*(\hat{\Lambda}_0)}^{\hat{\Lambda}_0}) > \Lambda^*(p_{t_2^*(\hat{\Lambda}_0)}^{\Lambda_0}) > \Lambda_{t_2^*(\hat{\Lambda}_0)}.$$

Then it is again obvious that $\hat{\lambda}\hat{N}_t > \lambda N_t$ for all $t > \bar{t}$. Thus, in either case we get that $p_t^{\hat{\Lambda}_0} > p_t^{\Lambda_0}$ for all $t > \bar{t}$, as claimed by the third bullet point. \square

A.6.3 Adoption Behavior

Proof of Proposition 1.5.12: First note that because $p_0 \geq \bar{p}$, $t_1^*(\Lambda_0) = t_1^*(\hat{\Lambda}_0) = 0$.

Then at all $\Lambda_0 < \Lambda^*(p_0)$, the adoption flow absent breakdowns satisfies $N_t = \rho\bar{N}_t$ for all t . Thus, conditional on a good product we get $A_t(\Lambda_0, G) = A_t(\hat{\Lambda}_0, G) = 1 - e^{-\rho t}$ for all t and all pairs $\Lambda_0, \hat{\Lambda}_0 \leq \Lambda^*(p_0)$.

Now suppose that $\hat{\Lambda}_0 > \Lambda_0 > \Lambda^*(p_0)$. Note that $N_t, \hat{N}_t > 0$ for all $t > 0$ (recall Condition 1.5.3). Let $t^* = t_2^*(\Lambda_0)$. By Lemma 1.5.10, $\lambda N_t = \hat{\lambda}\hat{N}_t$ for all $t < t^*$. Then for all $t < t^*$

$$\frac{N_t}{\bar{N}_0} = \frac{\lambda N_t}{\Lambda_0} = \frac{\hat{\lambda}\hat{N}_t}{\Lambda_0} > \frac{\hat{\lambda}\hat{N}_t}{\hat{\Lambda}_0} = \frac{\hat{N}_t}{\hat{\bar{N}}_0}$$

. Therefore for all $t < t^*$, we have $A_t(\Lambda_0, G) > A_t(\hat{\Lambda}_0, G)$.

Finally note that for all $t \geq t^*$, $N_t = \rho\bar{N}_t$ and so:

$$\begin{aligned} A_t(\Lambda_0, G) &= A_{t^*}(\Lambda_0, G) + \left(1 - e^{-\rho(t-t^*)}\right) (1 - A_{t^*}(\Lambda_0, G)) \\ A_t(\hat{\Lambda}_0, G) &\leq A_{t^*}(\hat{\Lambda}_0, G) + \left(1 - e^{-\rho(t-t^*)}\right) (1 - A_{t^*}(\hat{\Lambda}_0, G)) \end{aligned}$$

where the second inequality follows from feasibility. But because $A_{t^*}(\Lambda_0, G) > A_{t^*}(\hat{\Lambda}_0, G)$, $A_t(\Lambda_0, G) > A_t(\hat{\Lambda}_0, G)$ for all $t > 0$. \square

Proof of Proposition 1.5.13. We first prove the proposition when we increase the information structure from λ to $\hat{\lambda} > \lambda$ holding fixed \bar{N}_0 . Given this, proving the theorem for arbitrary

changes from Λ_0 to $\hat{\Lambda}_0$ is straightforward.

Let N and \hat{N} be the equilibrium under λ and $\hat{\lambda}$, respectively. Note that when $\bar{p} \leq p_0$, $N_t > 0$ for all $t > 0$. Given an arbitrary strictly positive adoption flow M and $t > 0$, consider the following map:

$$\lambda \mapsto \int_0^t M_\tau e^{-\int_0^\tau (\varepsilon + \lambda M_s) ds} d\tau.$$

Note that for any $t > 0$, the above is strictly decreasing in λ . This implies that for all $t > 0$,

$$\int_0^t N_\tau e^{-\int_0^\tau (\varepsilon + \lambda N_s) ds} d\tau > \int_0^t N_\tau e^{-\int_0^\tau (\varepsilon + \hat{\lambda} N_s) ds} d\tau. \quad (\text{A.3})$$

We now show that

$$\int_0^t N_\tau e^{-\int_0^\tau (\varepsilon + \hat{\lambda} N_s) ds} d\tau \geq \int_0^t \hat{N}_\tau e^{-\int_0^\tau (\varepsilon + \hat{\lambda} \hat{N}_s) ds} d\tau$$

which together with Inequality (A.3) implies the desired conclusion that $A_t(\hat{\lambda}, \bar{N}_0, B) < A_t(\lambda, \bar{N}_0, B)$ for all $t > 0$.

To prove this, suppose that there exists some $t > 0$ such that

$$\int_0^t N_\tau e^{-\int_0^\tau (\varepsilon + \hat{\lambda} N_s) ds} d\tau < \int_0^t \hat{N}_\tau e^{-\int_0^\tau (\varepsilon + \hat{\lambda} \hat{N}_s) ds} d\tau. \quad (\text{A.4})$$

Note that by Proposition 1.5.12, $\int_0^\tau N_s ds \geq \int_0^\tau \hat{N}_s ds$ for all $\tau \geq 0$ and so

$$\int_0^t \varepsilon e^{-\int_0^\tau (\varepsilon + \hat{\lambda} N_s) ds} d\tau \leq \int_0^t \varepsilon e^{-\int_0^\tau (\varepsilon + \hat{\lambda} \hat{N}_s) ds} d\tau \quad (\text{A.5})$$

for all $t \geq 0$. Inequalities (A.4) and (A.5) together imply:

$$\int_0^t (\varepsilon + \hat{\lambda} N_\tau) e^{-\int_0^\tau (\varepsilon + \hat{\lambda} N_s) ds} d\tau < \int_0^t (\varepsilon + \hat{\lambda} \hat{N}_\tau) e^{-\int_0^\tau (\varepsilon + \hat{\lambda} \hat{N}_s) ds} d\tau.$$

But this is equivalent to

$$\left(1 - e^{-\int_0^t (\varepsilon + \hat{\lambda} N_s) ds}\right) < \left(1 - e^{-\int_0^t (\varepsilon + \hat{\lambda} \hat{N}_s) ds}\right).$$

This contradicts $\int_0^t N_s ds \geq \int_0^t \hat{N}_s ds$ as found in Proposition 1.5.12.

Having shown the above, consider any change from $\Lambda_0 = \lambda \bar{N}_0$ to $\hat{\Lambda}_0 = \hat{\lambda} \hat{N}_0 > \Lambda_0$. Then there exists $\lambda^* > \lambda$ such that $\hat{\Lambda} = \lambda^* \bar{N}_0$. Let N^* be the equilibrium associated with the pair (λ^*, \bar{N}_0) . By Lemma 1.5.5, unique equilibrium for the pair $(\hat{\lambda}, \hat{N}_0)$ satisfies $\hat{N}_t = (\lambda^* / \hat{\lambda}) N_t^*$. But then the above argument implies:

$$\begin{aligned} A_t(\Lambda, B) &= \mathbb{E} \left[\int_0^t \frac{N_s}{\bar{N}_0} ds \right] > \mathbb{E} \left[\int_0^t \frac{N_s^*}{\bar{N}_0} ds \right] = \mathbb{E} \left[\int_0^t \frac{\lambda^*}{\hat{\lambda} \hat{N}_0} \frac{\hat{\lambda} \hat{N}_0}{\lambda^* \bar{N}_0} N_s^* ds \right] \\ &= \mathbb{E} \left[\int_0^t \frac{\hat{N}_s}{\hat{N}_0} ds \right] = A_t(\hat{\Lambda}, B). \end{aligned}$$

□

A.7 Comparative Statics under PGN

A.7.1 Adoption Behavior

The only statement that was not proved in the text is: $A_t(\hat{\Lambda}_0, G) < A_t(\Lambda_0, G)$ for all $t > t^*(\hat{\Lambda}_0)$, as claimed in the first bullet of Proposition 1.6.7.

Proof. When $\varepsilon = 0$, the statement is trivial, so assume that $\varepsilon > 0$. The claim is also obvious for all $t \leq t^*(\Lambda_0)$ since adoption occurs at the maximal rate under Λ_0 whereas under $\hat{\Lambda}_0$, adoption ceases at times $t \in (t^*(\hat{\Lambda}_0), t^*(\Lambda_0))$ absent breakthroughs.

So assume that $t > t^*(\Lambda_0)$. Recall that the cutoff posterior p^* at which adoption ceases is unchanged upon a change from Λ_0 to $\hat{\Lambda}_0$. Then expected adoption up to time t for any $\Gamma \in [\Lambda_0, \hat{\Lambda}_0]$ can be expressed in the following manner:

$$A_t(\Gamma, G) = \pi^* (1 - e^{-\rho t}) + (1 - \pi^*) \left((1 - e^{-\rho t^*(\Gamma)}) + e^{-\rho t^*(\Gamma)} \int_{t^*(\Gamma)}^t \varepsilon e^{-\varepsilon(\tau - t^*(\Gamma))} (1 - e^{-\rho(t - \tau)}) d\tau \right)$$

where

$$(1 - \pi^*) = \frac{1 - p_0}{1 - p^*}.$$

Now for a fixed t , consider the function:

$$t^* \mapsto \pi^* (1 - e^{-\rho t}) + (1 - \pi^*) \left((1 - e^{-\rho t^*}) + e^{-\rho t^*} \int_{t^*}^t \varepsilon e^{-\varepsilon(\tau - t^*)} (1 - e^{-\rho(t - \tau)}) d\tau \right).$$

Then a straightforward computation yields that the derivative of the above map with respect to any $t^* < t$ is $\rho e^{-(\varepsilon - \rho)t^*} e^{-\varepsilon t} > 0$. Thus, the map is strictly increasing in t^* for all $t^* < t$. Because $t^*(\Gamma)$ is strictly decreasing in Γ , it follows that for all $t > t^*(\Gamma)$, $A_t(\Gamma, G)$ is strictly decreasing for all $\Gamma \in [\Lambda_0, \hat{\Lambda}_0]$. This proves the claim. \square

A.8 Heterogeneous Discount Rate Example

First we show the following basic mathematical fact.

Lemma A.8.1. *Let $\bar{t} > t^*$ and suppose that f and g are real-valued functions such that $f(\tau) = g(\tau)$ for all $\tau \leq t^*$, $f(\tau) < g(\tau)$ for $\tau \in (t^*, \bar{t})$, and $f(\tau) > g(\tau)$ for all $\tau > \bar{t}$. Suppose that*

$$\int_0^\infty e^{-r\tau} f(\tau) d\tau = \int_0^\infty e^{-r\tau} g(\tau) d\tau.$$

Then for all $\hat{r} > r$,

$$\int_0^\infty e^{-\hat{r}\tau} f(\tau) d\tau < \int_0^\infty e^{-\hat{r}\tau} g(\tau) d\tau.$$

Proof. We have

$$\begin{aligned} 0 &= \int_0^\infty e^{-r\tau} (g(\tau) - f(\tau)) d\tau = \int_0^{\bar{t}} e^{-\hat{r}\tau} e^{(\hat{r}-r)\tau} (g(\tau) - f(\tau)) d\tau + \int_{\bar{t}}^\infty e^{-\hat{r}\tau} e^{(\hat{r}-r)\tau} (g(\tau) - f(\tau)) d\tau \\ &< e^{(\hat{r}-r)\bar{t}} \left(\int_0^{\bar{t}} e^{-\hat{r}\tau} (g(\tau) - f(\tau)) d\tau + \int_{\bar{t}}^\infty e^{-\hat{r}\tau} (g(\tau) - f(\tau)) d\tau \right) \\ &< e^{(\hat{r}-r)\bar{t}} \int_0^\infty e^{-\hat{r}\tau} (g(\tau) - f(\tau)) d\tau. \end{aligned}$$

This implies that $\int_0^\infty e^{-\hat{r}\tau} f(\tau) d\tau < \int_0^\infty e^{-\hat{r}\tau} g(\tau) d\tau$, as claimed. \square

As in the main text, assume that $\hat{\lambda} M_0^p > \lambda M_0^p > \Lambda_p^*(p_0)$ and that $p_0 > 1/2$ and $\varepsilon = 0$.

Then modifying the arguments from the proof of Theorem 1.5.1, it is easy to show that when M_0^i is sufficiently small, the unique equilibrium under both information processes λ , $\hat{\lambda}$ will be such that the impatient type adopts immediately upon opportunity at all times absent breakdowns and the patient type only partially adopts until some time $t^* > 0$ after which he switches to immediate adoption:³

$$\begin{aligned} \gamma N_t^i &= \rho M_t^i, \\ \gamma N_t^p &= \begin{cases} \frac{r_p(2p_t-1)}{1-p_t} - \gamma \rho M_t^i & \text{if } t < t^*(\gamma) \\ \gamma \rho M_t^p & \text{if } t \geq t^*(\gamma) \end{cases} \end{aligned}$$

for $\gamma \in \{\lambda, \hat{\lambda}\}$.

Then using arguments analogous to those in Lemma 1.5.10, we can show that $t^*(\lambda) < t^*(\hat{\lambda})$. Furthermore an analogue of Proposition 1.5.11 shows that there must exist some $\bar{t} > t^*(\lambda)$ such that

$$p_t^\lambda \begin{cases} = p_t^{\hat{\lambda}} & \text{if } t \leq t^*(\lambda) \\ > p_t^{\hat{\lambda}} & \text{if } t \in (t^*(\lambda), \bar{t}) \\ < p_t^{\hat{\lambda}} & \text{if } t > \bar{t}. \end{cases}$$

Then using Lemma A.8.1, the proof follows along the lines illustrated in the main text. This proves Theorem 1.7.1.

³The full proof of the modification is available upon request. Here we use a standard argument that shows that whenever the impatient type weakly prefers to wait, then the patient type must strictly prefer to wait. Similarly, if the patient type weakly prefers to adopt then the impatient type must strictly prefer to adopt.

A.9 Cooperative Benchmark

A.9.1 Perfect Bad News

To prove the all-or-nothing property of the optimal policy, we write the Hamilton-Jacobi-Bellman (HJB) equation. Note that there are two state variables, p and \bar{N} .

$$rV(p, \bar{N}) = \max_{0 \leq N \leq \rho \bar{N}} (2p - 1)N + D_p V(p, \bar{N})p(1 - p)(\varepsilon + \lambda N) - D_{\bar{N}} V(p, \bar{N})N - (1 - p)(\varepsilon + \lambda N)V(p, \bar{N}).$$

Since the right hand side is linear in N , it is optimal to always choose either $N = 0$ or $N = \rho \bar{N}$.

To see that the optimal policy must be a cutoff strategy, define

$$\Pi(p, \bar{N}) := (2p - 1) + D_p V(p, \bar{N})p(1 - p)\lambda - D_{\bar{N}} V(p, \bar{N}) - \lambda(1 - p)V(p, \bar{N})$$

and note that whenever $\Pi(p, \bar{N}) < 0$, then

$$rV(p, \bar{N}) = D_p V(p, \bar{N})\varepsilon p(1 - p) - (1 - p)\varepsilon V(p, \bar{N}) \quad (\text{A.6})$$

so that this corresponds to the case where setting $N = 0$ is optimal. It then suffices to prove that

$$\Pi(p, \bar{N}) < 0 \Rightarrow \Pi(p', \bar{N}) < 0 \quad \forall p' < p.$$

To prove this, note first that for every p such that $\Pi(p, \bar{N}) < 0$, there must exist some $p' > p$ such that $\Pi(p', \bar{N}) = 0$. (Otherwise $V(p', \bar{N}) = 0$ for all $p' > p$, which is clearly false.) So it suffices to prove that there cannot exist $\underline{p} < \bar{p}$ such that $\Pi(\underline{p}, \bar{N}) = \Pi(\bar{p}, \bar{N}) = 0$ and $\Pi(p, \bar{N}) < 0$ for all $p \in (\underline{p}, \bar{p})$. Suppose for a contradiction that such an interval (\underline{p}, \bar{p}) exists. Then ordinary differential equation (A.6) implies:

$$V(p, \bar{N}) = \left(\frac{p}{\bar{p}}\right)^{\frac{r+\varepsilon}{\varepsilon}} \left(\frac{1-p}{1-\bar{p}}\right)^{-\frac{r}{\varepsilon}} V(\bar{p}, \bar{N})$$

for all $p \in (\underline{p}, \bar{p})$. Then we can rewrite the expression for $\Pi(p, \bar{N})$ for $p \in (\underline{p}, \bar{p})$:

$$\begin{aligned}\Pi(p, \bar{N}) &= (2p - 1) + \frac{r\lambda}{\varepsilon} V(p, \bar{N}) - D_{\bar{N}} V(p, \bar{N}) \\ &= (2p - 1) + \frac{r\lambda}{\varepsilon} V(p, \bar{N}) - \left(\frac{p}{\bar{p}}\right)^{\frac{r+\varepsilon}{\varepsilon}} \left(\frac{1-p}{1-\bar{p}}\right)^{-\frac{r}{\varepsilon}} D_{\bar{N}} V(\bar{p}, \bar{N}) \\ &= (2p - 1) + \left(\frac{p}{\bar{p}}\right)^{\frac{r+\varepsilon}{\varepsilon}} \left(\frac{1-p}{1-\bar{p}}\right)^{-\frac{r}{\varepsilon}} \left(\frac{\lambda r}{\varepsilon} V(\bar{p}, \bar{N}) - D_{\bar{N}} V(\bar{p}, \bar{N})\right),\end{aligned}$$

Note that if $\frac{\lambda r}{\varepsilon} V(\bar{p}, \bar{N}) - D_{\bar{N}} V(\bar{p}, \bar{N}) \geq 0$, the last expression is increasing in p and so $\Pi(p, \bar{N}) < 0$ which is a contradiction.

If instead $\frac{\lambda r}{\varepsilon} V(\bar{p}, \bar{N}) - D_{\bar{N}} V(\bar{p}, \bar{N}) < 0$, then the second term in the last expression is concave. Furthermore, the derivative of the last expression with respect to p at \bar{p} must be weakly positive: If it were strictly negative, then because $\Pi(\bar{p}, \bar{N}) = 0$, there would exist some $p \in (\underline{p}, \bar{p})$ close to \bar{p} such that $\Pi(p, \bar{N}) > 0$. But if the derivative of $\Pi(p, \bar{N})$ is weakly positive at \bar{p} , then by concavity it must be positive throughout (\underline{p}, \bar{p}) . But this again yields the contradiction that $\Pi(p, \bar{N}) < 0$. This completes the proof.

A.9.2 Perfect Good News

As in the perfect bad news case, we again write the Hamilton-Jacobi-Bellman equation:

$$\begin{aligned}rV(p, \bar{N}) &= \max_{0 \leq N \leq \rho \bar{N}} (2p - 1)N + p(\varepsilon + \lambda N) \left(\frac{\rho}{r + \rho} \bar{N} - V(p, \bar{N})\right) \\ &\quad - D_p V(p, \bar{N}) p(1 - p) (\varepsilon + \lambda N) - D_{\bar{N}} V(p, \bar{N}) N.\end{aligned}$$

Again the right hand side is linear in N and thus the optimal policy always chooses either $N = 0$ or $N = \rho \bar{N}$.

The easiest way to check that an optimal policy exists in cutoff strategies is to simply guess and check that the HJB equation is satisfied by such a strategy. This is straightforward from the social planner policy constructed in Section 1.3.1.

Appendix B

Appendix to Chapter 2

B.1 Definitions of s_i^G and s_i^B

We recall the following definitions of $\tilde{s}_i^G, \tilde{s}_i^B \in \tilde{\mathcal{S}}_i^\rho$ of HO2006.

First, define \tilde{s}_i^g as some strategy such that $\tilde{s}_i^g[\emptyset] \in \Delta G$ and

$$\begin{aligned} &\text{for all } \tilde{h}_i^t = \left(a, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}), \dots, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}) \right), a \in M_i \times M_{-i}, t \geq 1 : \\ &\tilde{s}_i^g[\tilde{h}_i^t] = a_i^{M_2, M_1}; \end{aligned}$$

and define \tilde{s}_i^b such that $\tilde{s}_i^b[\emptyset] \in \Delta B$ and

$$\begin{aligned} &\text{for all } \tilde{h}_i^t = \left(a, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}), \dots, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}) \right), a \in M_i \times M_{-i}, t \geq 1 : \\ &\tilde{s}_i^b[\tilde{h}_i^t] = a_i^{M_2, M_1}. \end{aligned}$$

Moreover define $\tilde{s}_i^m[h_i^t] = \alpha_i^m$ for every history \tilde{h}_i^t that is a continuation of a history

$$\begin{aligned} \tilde{h}_i^r &= \left(a, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}), \dots, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}), (a_i^{M_2, M_1}, a'_{-i}) \right), \\ &a \in B \times M_{-i}, a'_{-i} \neq a_{-i}^{M_2, M_1}, t \geq r \geq 1, \end{aligned}$$

where α_i^m is a possibly mixed minmax action against player $-i$. Then \tilde{s}_i^G and \tilde{s}_i^B are defined as small perturbations of \tilde{s}_i^g and \tilde{s}_i^b to obtain a pair of strategies \tilde{s}_i^G and \tilde{s}_i^B in $\tilde{\mathcal{S}}_i^\rho$. HO2006 showed that by choosing ρ sufficiently small, we may ensure that there exists some T such

that

$$\min_{\tilde{S}_i} \tilde{U}_i^T(\tilde{s}_i, \tilde{s}_{-i}^G) > \bar{v}_i > v_i > \underline{v}_i > \max_{\tilde{S}_i^T} \tilde{U}_i^T(\tilde{s}_i, \tilde{s}_{-i}^B)$$

for sufficiently patient players.

Given the above definitions, we can obtain strategies in repeated games *with observation lags and private monitoring structure* π in a natural way by identifying with each $\tilde{h}_i^t \in \tilde{H}_i^t$ the unique element of $h_i^t \in H_i^t$ such that h_i^t and \tilde{h}_i^t report exactly the same observations about the play of player $-i$ at all times and h_i^t contains no observations with a positive lag. We denote this identification by $\tilde{h}_i^t \simeq h_i^t$.

Similarly we identify $s_i \in S_i^T$ to a strategy $\tilde{s}_i \in \tilde{S}_i^T$ in a natural way. We say that $\tilde{s}_i \simeq s_i$ if $s_i[h_i^t] = \tilde{s}_i[\tilde{h}_i^t]$ for all $\tilde{h}_i^t \in \tilde{H}_i^t$ and all $h_i^t \in H_i^t$ such that $\tilde{h}_i^t \simeq h_i^t$. Then we simply define s_i^G and s_i^B to be strategies such that $\tilde{s}_i^G \simeq s_i^G$ and $\tilde{s}_i^B \simeq s_i^B$. It is easy to see that we can appropriately define s_i^G and s_i^B at all histories $h_i^t \in H_i^t$ that are not identified with some $\tilde{h}_i^t \in \tilde{H}_i^t$ to obtain a pair of strategies $s_i^G, s_i^B \in S_i^O$. Moreover when the probability of observation lag is small and π is very close to perfect monitoring, it is clear that $U_i^T(s_i, s_{-i})$ is close to $\tilde{U}_i^T(\tilde{s}_i, \tilde{s}_{-i})$ where $\tilde{s}_i \simeq s_i$ and $\tilde{s}_{-i} \simeq s_{-i}$. Thus we have for sufficiently patient players and ρ sufficiently small,

$$\min_{S_i} U_i^T(s_i, s_{-i}^G) > \bar{v}_i > v_i > \underline{v}_i > \max_{S_i^T} U_i^T(s_i, s_{-i}^B).$$

B.2 Details of the Proof of Theorem 2.4.3

B.2.1 Proof of Lemma 2.4.4

Proof of Lemma 2.4.4: We wish to specify transfers $\zeta_i^B : H_{-i}^T \rightarrow \mathbb{R}_-$ in such a way that players are indifferent between all possible strategies in the T -period repeated game given auxiliary transfers ζ_i^B . To do this, we define equivalence classes over T -period histories in the following way:

$$(h_{-i}^{T-1}, a_i^{t_1}, a_i^{t_2}, \dots, a_i^{t_m}, a_i^T, a_{-i}^T) \sim (\hat{h}_{-i}^{T-1}, \hat{a}_i^{t_1}, \hat{a}_i^{t_2}, \dots, \hat{a}_i^{t_m}, \hat{a}_i^T, \hat{a}_{-i}^T)$$

if and only if $h_{-i}^{T-1} = \hat{h}_{-i}^{T-1}$ and $a_i^T = \hat{a}_i^T$. Here, if player $-i$ does not obtain information about the play of player i in time T , then a_i^T is taken to be ∞ (representing a null signal). Also

notationally, $a_i^{t_1}, \dots, a_i^{t_m}$ are the elements of $h_i^{T,o}$ that are not equal to a_i^T . We may represent this equivalence class of T period histories in the form (h_{-i}^{T-1}, a_i^T) ; note that this indicates that neither

1. the action played by player $-i$ in period T , nor
2. new information gained about past actions

matter for the determination of the equivalence class.

We define equivalence classes over t -period histories similarly, and represent such an equivalence class by (h_{-i}^{t-1}, a_i^t) . We now define a transfer function ζ_i^B as in HO2006 for some functions θ_t defined over equivalence classes of t -period histories:

$$\zeta_i^B(h_{-i}^T) = \frac{1}{\delta^T} \sum_{t=1}^T \delta^{t-1} \theta_t(h_{-i}^{t-1}, a_i^t).$$

Here, h_{-i}^{t-1} is the t -period truncation of h_{-i}^T and a_i^t is the signal that player $-i$ observed of player i 's period- t action in period t . That is, $a_i^t = \infty$ if player $-i$ does not observe i 's play immediately and is otherwise equal to the actual period- t action of player i .¹

Given any h_{-i}^{T-1} , consider the matrix

$$\left(\mu_{-i}(\cdot \mid a_i, \bar{s}_{-i}^B(h_{-i}^{T-1})) \right)_{a_i \in A_i}.$$

Note that the matrix above has full row rank when $\lambda_{-i}(0)$ is sufficiently close to 1 and π is sufficiently close to perfect monitoring. Therefore the sub-matrix obtained by deleting the column corresponding to the " ∞ " signal is invertible.² We then set $\theta_t(h_{-i}^{t-1}, \infty) = 0$ and solve the system of equations defined by

$$\mu_{-i}(\cdot \mid a_i, \bar{s}_{-i}^B(h_{-i}^{T-1})) \cdot \theta_T(h_{-i}^{T-1}, \cdot) = g_i(a_i^*, \bar{s}_{-i}^B(h_{-i}^{T-1})) - g_i(a_i, \bar{s}_{-i}^B(h_{-i}^{T-1})), \quad (\text{B.1})$$

¹According to this definition, if for example the play of player 1's period-1 action is not observed immediately (i.e. in period 1) by player 2, then the observation of player 1's period-1 action in a later period only has an effect on ζ_i^B through its effect on player $-i$'s play.

²In fact, this sub-matrix approaches the identity matrix as $\lambda(0) \rightarrow 1$ and π approaches perfect monitoring.

where a_i^* is the stage game best response to $\bar{s}_{-i}^B(h_{-i}^{T-1})$. Our preceding observations show that system (B.1) has a unique solution when $\lambda_{-i}(0)$ is sufficiently large and π is sufficiently close to perfect monitoring.

Then in period $T - 1$, player i is indifferent between all of his actions given that player $-i$ plays according to the strategy prescribed by \bar{s}_{-i}^B at history h_{-i}^{T-1} and transfers given by $\theta_T(h_{-i}^{T-1}, \cdot)$, as playing any action a_i generates a payoff of

$$\begin{aligned} & (1 - \delta)\delta^{T-1}g_i(a_i, \bar{s}_{-i}^B(h_{-i}^{T-1})) + (1 - \delta)\delta^{T-1} \sum_{\omega_i \in A_i \cup \{\infty\}} \mu_{-i}(\omega_i | a_i, \bar{s}_{-i}^B) \theta_T(h_{-i}^{T-1}, \omega_i) \\ & = (1 - \delta)\delta^{T-1}g_i(a_i^*, \bar{s}_{-i}^B(h_{-i}^{T-1})). \end{aligned}$$

Suppose that all transfers θ_τ for $\tau \geq t$ have been defined so that player i is indifferent across all of his strategies from period $t + 1$ on. Then define $U_{t+1}(h_{-i}^{t-1}, a_i)$ to be the expected continuation payoff given the transfers at period $t + 1$, given that player $-i$'s history in period $t - 1$ is h_{-i}^{t-1} and player i played a_i in period t .

We now define θ_{t-1} in a similar manner. Again we consider any $h_{-i}^T \in H_{-i}^T$ and consider the following expression:

$$\frac{1}{\delta^T} \sum_{s=t}^T \delta^{s-1} \theta_s(h_{-i}^{s-1}, a_i^s).$$

Again define $\theta_{t-1}(h_{-i}^{t-1}, \infty) = 0$ and consider the matrix

$$\left(\mu_{-i}(\cdot | a_i, \bar{s}_{-i}^B(h_{-i}^{t-1})) \right)_{a_i \in A_i}.$$

Let us denote the sub-matrix obtained by deleting the column corresponding to the null signal " ∞ " by $D(h_{-i}^{t-1})$. This is again invertible when $\lambda_{-i}(0)$ is sufficiently close to 1 and π is sufficiently close to perfect monitoring. Now consider the system of equations

$$\begin{aligned} & (1 - \delta)\delta^{t-1} \left(\mu_{-i}(\cdot | a_i, \bar{s}_{-i}^B(h_{-i}^{t-1})) \cdot \theta_t(h_{-i}^{t-1}, \cdot) + g_i(a_i, \bar{s}_{-i}^B(h_{-i}^{t-1})) \right) + (1 - \delta)U_{t+1}(h_{-i}^{t-1}, a_i) \\ & = (1 - \delta)U_{t+1}(h_{-i}^{t-1}, a_i^*(h_{-i}^{t-1})) + (1 - \delta)\delta^{t-1}g_i(a_i^*(h_{-i}^{t-1}), \bar{s}_{-i}^B(h_{-i}^{t-1})) \end{aligned} \quad (\text{B.2})$$

where $a_i^*(h_{-i}^{t-1})$ is the term that maximizes the expression on the right hand side of the equation above.

Because the matrix $D(h_{-i}^{t-1})$ is invertible, the system (B.2) has a unique solution when we set $\theta_t(h_{-i}^{t-1}, \infty) = 0$. Iterating in this manner allows us to obtain the first part of the lemma.

To achieve non-negativity of transfers, we observe that as the square matrices $D(h_{-i}^{t-1})$ converge to the identity matrix, the solutions $\theta_t(h_{-i}^{t-1}, a_i)$ must be non-negative in the limit. Thus we can make all transfers $\theta_t(h_{-i}^{t-1}, a_i)$ non-negative by adding to all of them a positive constant that converges to zero as $\lambda_{-i}(0)$ and π jointly converge to 1 and perfect monitoring respectively.

Finally we define a strategy $r_i^B \in S_i^T$ in the following way. Let $r_i^B(h_{-i}^{t-1})$ be the action $a_i^*(h_{-i}^{t-1})$ as defined above for all histories h_{-i}^{t-1} that do not contain any null signals, where h_{-i}^{t-1} is the history that corresponds to h_{-i}^{t-1} . Define $r_i^B(h_{-i}^{t-1})$ arbitrarily for all other histories. Then note that as monitoring becomes perfect, the expected value of ζ_i^B goes to zero if players play according to r_i^B and \bar{s}_{-i}^B . By the definition of r_i^B , the payoff in the T -times-repeated game without any transfers then approaches $\max_{s_i \in S_i^T} U_i^T(s_i, \bar{s}_{-i}^B)$; this implies that

$$\lim_{\varepsilon \rightarrow 0} U_i^A(s_i, \bar{s}_{-i}^B, \zeta_i^B) = \max_{\tilde{s}_i} U_i^T(\tilde{s}_i, \bar{s}_{-i}^B)$$

for all $s_i \in S_i^T$. □

B.2.2 Proof of Lemma 2.4.5

Proof. Let $\varepsilon > 0$ be such that π is ε -perfect and $\Pr(L_{-i} > 0) < \varepsilon$. For every $\nu > 0$, observe that there exists ε/ρ small enough such that, for any history $h_{-i}^{t-1} \in H_{-i}^{R,t-1}$ and conditional on observing h_{-i}^{t-1} , player i assigns probability at least $1 - \nu$ to the event that player $-i$ observed the corresponding history h_{-i}^{t-1} . Consider for some $h_{-i}^{t-1} \in H_{-i}^{R,t-1}$ and any action $a_i \in A_i$, the row vector consisting of the probabilities assigned by player i , conditional on history h_{-i}^{t-1} and on action a_i taken by player i in period t , to the different equivalence classes of histories (h_{-i}^{t-1}, a_i) observed by player $-i$ in period t . As in HO2006, we construct a matrix D^{t-1} by stacking the row vectors for all regular histories $h_{-i}^{t-1} \in H_{-i}^{R,t-1}$ and actions $a_i \in A_i$. Note that for small enough ε/ρ , the matrix D^{t-1} has full row rank for every t .

With this we can define $\theta(\cdot, \cdot)$ by setting $\theta(h_{-i}^{t-1}, \infty) = 0$ for any $h_{-i}^{t-1} \in H_{-i}^{t-1}$. This is

possible since the number of rows is exactly the same as in HO2006 and the number of columns corresponding to (h_{-i}^{t-1}, a_i) for some $a_i \neq \infty$ is also the same as in HO2006. This proves the lemma. \square

B.3 Conditions Guaranteeing Small γ_i

When the measures L_1, \dots, L_n are independent and identically distributed, we have

$$\frac{\lambda(\infty, \dots, \infty)}{\lambda_i(\infty)} = \frac{\lambda_1(\infty) \cdots \lambda_n(\infty)}{\lambda_i(\infty)}. \quad (\text{B.3})$$

Clearly when $\lambda_1(\infty) = \cdots = \lambda_n(\infty)$ and $\lambda_1(\infty)$ small, (B.3) is close to 0. Even if L_1, \dots, L_n are not identically distributed but are independent, we again have (B.3) small if $\frac{(\max_i \lambda_i(\infty))^n}{\min_i \lambda_i(\infty)}$ is sufficiently small, i.e. when no player i 's probability of never observing a signal is much smaller than some other player's probability of never observing a signal.

Consider lags for which L_i is split into two components, X and A_i : $L_i = X + A_i$. Note that X is common across all players. We assume that A_i is independent and identically distributed across players, and that A_i is independent of X . Let ζ be the density of X and χ denote the density of A_i .

Then γ_i is small if

$$\frac{(1 - \zeta(\infty))(\chi(\infty))^n + \zeta(\infty)}{\zeta(\infty) + (1 - \zeta(\infty))(\chi(\infty))} \quad (\text{B.4})$$

is small—which is true if $\zeta(\infty)$ is much smaller than $\chi(\infty)$. For example, suppose that $\zeta(\infty) = 0$ but $\chi(\infty) > 0$. Then if $\chi(\infty)$ is small, (B.4) is small.³

³By contrast, consider the case in which $\chi(\infty) = \zeta(\infty) = \varepsilon > 0$. Then (B.4) equals

$$\frac{(1 - \varepsilon)\varepsilon^n + \varepsilon}{\varepsilon + (1 - \varepsilon)\varepsilon},$$

which converges to 1/2 as $\varepsilon \rightarrow 0$ —so γ_i cannot be taken to be small even when ε is small.

Appendix C

Appendix to Chapter 3

C.1 Proof of Existence

Proof. First note that due to absolute continuity of μ_r , there exists $f \in L^1([0,1])$ that represents the density of μ_r . This proof follows via Kakutani's fixed point generalized to locally convex topological vector spaces. The trick is to construct the appropriate topologies on the spaces S_1 and Σ_2 to make compactness and continuity as easy to prove as possible.

This is done in the following manner. Note that we define Σ_2^t as the set of t -period maps $\sigma_2^t : H^t \rightarrow \Delta(A_1)$. Σ_2^t is compact with respect to the following norm:

$$\|\sigma_2^t\| = \sup_{h^t \in H^t} |\sigma_2^t(h^t)|.$$

By Tychonoff's theorem, $\Sigma_2 = \prod_{t \geq 0} \Sigma_2^t$ is compact in the product topology denoted \mathcal{T} . Moreover it is easy to see that (Σ_2, \mathcal{T}) is indeed metrizable with the following metric:¹

$$d(\sigma_2, \alpha_2) = \sum_{t=0}^{\infty} 2^{-t} \frac{\|\sigma_2^t - \alpha_2^t\|}{1 + \|\sigma_2^t - \alpha_2^t\|}.$$

In summary we have shown that (Σ_2, \mathcal{T}) is a compact, convex, metric space.

As we stated in Section 3.2, S_1 can be identified with a subset of all measurable maps

¹For a proof, see any functional analysis textbook containing a section on locally convex topological vector spaces.

from $[0, 1] \rightarrow \{0, 1\}^{m-1}$. Then define the following subset of $L^\infty([0, 1])$:

$$\Gamma = \{f \in L^\infty([0, 1]) : |f(x)| \leq 1 \forall x \in [0, 1]\}.$$

Note that Γ is indeed compact in the weak-star topology endowed on $L^\infty([0, 1])$ by the Banach-Alaoglu theorem. Furthermore Γ is metrizable in the weak-star topology. Then $\Gamma^{m-1} \subseteq L^\infty([0, 1])^{m-1}$ is compact and metrizable in the product topology (generated by the weak-star topologies on each coordinate). Call this topology \mathcal{S} . Notice that S_1 is a closed subset of Γ^{m-1} and thus (S_1, \mathcal{S}) is a compact metric space.

With the above observations, we need to check the continuity of the appropriate utility functions with respect to the appropriate topologies. We consider the following two utility functions:

$$\begin{aligned} V_1(s_1, \sigma_2) &= \int_0^1 \mathbb{E}_{s_1(\delta)} u_1(s_1(\delta), \sigma_2) f(\delta) d\delta \\ &= \sum_{a_1 \in A_1} \mathbb{E}_{a_1} u_1(a_1, \sigma_2) \int_0^1 s_1(\delta) [a_1] f(\delta) d\delta. \\ V_2(s_1, \sigma_2) &= \sum_{\tau=0}^{\infty} \lambda^\tau \sum_{h^\tau \in H^\tau} \frac{\mathbb{E}_{s_1} [u_2(a_1, \sigma_2(h^\tau)) \mid h^\tau]}{|H^\tau|} \end{aligned}$$

where $\lambda \in (0, 1)$. Note that for any two $s_1, s'_1 \in S_1$, if

$$\int_0^1 s_1(\delta) [a_1] f(\delta) d\delta = \int_0^1 s'_1(\delta) [a_1] f(\delta) d\delta$$

for all $a_1 \in A_1$, then $V_2(s_1, \sigma_2) = V_2(s'_1, \sigma_2)$ for all $\sigma_2 \in \Sigma_2$. Thus the distribution over A_1 induced by σ_1 is a sufficient statistic for the computation of $V_2(\cdot, \sigma_2)$. Therefore we write $V_2(\mu, \sigma_2)$ for $V_2(s_1, \sigma_2)$ where μ is the distribution over A_1 induced by s_1 .

It is easy to check the continuity of V_1 in (s_1, σ_2) . Because $S_1 \times \Sigma_2$ is a metric space, it is sufficient to consider sequential continuity. Thus suppose that $(s_1^n, \sigma_2^n) \rightarrow (s_1, \sigma_2)$. By the definition of the definition of the weak-star topology on $L^\infty([0, 1])$, since $f \in L^1([0, 1])$,

$$\int_0^1 s_1^n(\delta) [a_1] f(\delta) d\delta \rightarrow \int_0^1 s_1(\delta) [a_1] f(\delta) d\delta$$

for all $a_1 \in A_1$. Clearly $\mathbb{E}_{a_1} u_1(a_1, \sigma_2^n) \rightarrow \mathbb{E}_{a_1} u_1(a_1, \sigma_2)$ for all $a_1 \in A_1$. Thus V_1 is indeed continuous.

Now let us check that V_2 is indeed continuous. Suppose again that $(s_1^n, \sigma_2^n) \rightarrow (s_1, \sigma_2)$. Let μ^n to be the distribution over A_1 induced by s_1^n and similarly let μ to be the distribution over A_1 induced by s_1 . By the definition of the weak-star topology, $\mu^n \rightarrow \mu$. With this it is easy to see that

$$V_2(\mu^n, \sigma_2^n) \rightarrow V_2(\mu, \sigma_2).$$

Then we can use Glicksberg's existence theorem for continuous games to show that the game with strategy spaces S_1, Σ_2 and payoff functions V_1 and V_2 has a Nash equilibrium in mixed strategies.² In other words, there exists $(\sigma_1, \sigma_2) \in \Sigma_1 \times \Delta(\Sigma_2)$ that is a Nash equilibrium of the above game. But because Σ_2 is already convex and V_2 is already linear in σ_2 , every element of $\Delta(\Sigma_2)$ can be uniquely identified with an element of Σ_2 . Thus there exists $(\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2$ that is a Nash equilibrium.

Upon careful inspection of V_2 , this then implies that σ_2 maximizes expected payoff at all histories (under the stage game payoffs). Therefore we have indeed shown that (σ_1, σ_2) is a Nash equilibrium of the original game. □

²See for example Reny (2008).