Essays in Economic Theory

Essays in Economic Theory

a dissertation presented
by
Jetlir Duraj
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
March 2020

To my father.

Dedikuar babait tim Lefter Duraj, sacrifica e të cilit në përkujdesjen ndaj mamasë time, bëri të mundur të arrija deri këtu.

Dissertation Advisors:
Professor Tomasz Strzalecki
Professor Eric Maskin

Author:
Jetlir Duraj

Essays in Economic Theory

Abstract

This dissertation consists of three independent essays in microeconomic theory, focusing on aspects of learning in varied economic settings, both single and multi-agent.

In chapter 1, I study a discrete-time dynamic bargaining game in which a buyer can choose to learn privately about her value of the good. I assume information generation takes time and is endogenous, and that verifiable disclosure of evidence is possible. These assumptions result in a folk-theorem type of result about the delay. Moreover, near the high-frequency limit, all stationary equilibria feature non-extreme prices and non-extreme payoffs.

In chapter 2 (co-authored with Kevin He), we study how a benevolent sender communicates non-instrumental information over time to a Bayesian receiver who experiences news utility and exhibits diminishing sensitivity. We show that one-shot resolution of uncertainty is strictly suboptimal under many commonly used functional forms. We identify additional conditions that imply the sender optimally releases good news in small pieces but bad news in one clump and show how diminishing sensitivity may lead to commitment problems for the sender.

In chapter 3 (co-authored with Yi-Hsuan Lin), we consider an agent who privately learns information about a payoff-relevant uncertain state of the world

through a sequential experiment. Suppose the analyst observes the joint distribution over chosen action and decision time. We show that such data uniquely identify costs of information in two popular canonical cases. Moreover, we show how an outside observer with access to such data can conduct welfare analysis despite being oblivious of the sequential experiment of the agent.

Contents

Listing of figures

Acknowledgments

I AM GRATEFUL TO Tomasz Strzalecki, Eric Maskin and Jerry Green for their guidance and continuous support, including the support during the academic job market. I am also grateful to Matthew Rabin for his support, guidance and discussions related to the material of the second chapter of the dissertation. Finally, I am also grateful to Drew Fudenberg for valuable discussions and feedback over the years related to most of the projects I worked on while at Harvard. Unfortunately, I am unable to thank him for support during the academic job market.

I also thank my coauthors in Economics Sichao He, Yi-Hsuan Lin for being nice and conscientious co-authors. Finally, I also thank Artur Ananin, Anna Avdulaj, Arjada Bardhi, Sichao He, Jonathan Libgober, Rinald Murataj, Gianluca Rinaldi, Nicola Rosaia and Vitali Wachtel for their friendship during the last year of my PhD studies.

0
Introduction

This dissertation presents three independent essays in microeconomic theory, focusing on aspects and implications of learning, both in strategic as well as in single-agent settings. The three chapters deal with learning in different economic settings, ranging from bargaining to optimal disclosure of non-instrumental information. They also vary in the methodology used, from dynamic game theory to axiomatic decision theory, thus validating the title of the thesis.

In chapter 1, I study a discrete-time dynamic bargaining game in which a buyer can choose to learn privately about her value of the good. Learning is about information generation, rather than related to information processing costs. I assume that information sources take time to arrive and that the amount of their exploitation is endogenous. After learning, the buyer can disclose verifiable evidence of her valuation to the seller. The seller has the bargaining power, i.e. he is the only party making price offers. These features make for a more realistic model of negotiations. Examples include venture capital negotiations or procurement of new technologies, which sometimes feature significant delay due to endogenous costly learning. They also include certain M&A activities, negotiations at the book stage between a buyer and a seller, or negotiations between a real estate developer and a property owner looking to sell. I look for Perfect Bayesian Equilibria (PBEs) in this dynamic game. I study both the case where exploitation of information sources and search for information sources is costless (a benchmark), as well as the case where both of these components of information generation are costly.[1]

The model analysis delivers two sets of results. The first set are general results about PBEs. For any set of fixed parameters of the game, the buyer receives informational rents for any period-length only if learning is costly. In particular, the mere possibility of learning new information does not always ensure informational rents for the buyer, unless information generation is costly. Moreover, the disclosure choice of the buyer is 'uniform' across all model versions: the buyer who receives good news is indifferent between disclosing or not good news; the buyer with good news never discloses, thus pooling with the buyer who has not learned yet. This strategic option value derived from learning is the mechanism with which the buyer ensures positive payoff in equilibria. Another implication of the disclosure choice is that in equilibria in which the buyer with bad news discloses immediately, there is no scope for Coasian dynamics: the positive selection effect (buyers with good news don't disclose) from non-disclosure is strong enough to (at least) neutralize the negative selection effect from the rejection of prices.

The second set of results are derived under stationarity equilibrium refinements and require the qualification that the frequency of interaction is large enough. This is a popular technical approach in dynamic bargaining games.

[1]Learning is still associated with economic costs in the 'costless' model, because information sources take time to arrive and agents are impatient.

Additionally, I focus the analysis on equilibria in which the buyer with bad news discloses immediately. The rationale behind this requirement is that equilibria in which a buyer stays in the game despite a zero continuation payoff, are not robust to the introduction of minimal bargaining frictions like overhead costs or storage costs for the good.

The high-frequency limits of stationary equilibria result in a folk-theorem type of result about the delay until agreement. Maximal delay is achieved in equilibria with mixed pricing. Near the high-frequency limit, all stationary equilibria feature non-extreme prices and non-extreme payoffs. This is in stark contrast with respect to two more 'extreme' situations: the case in which the buyer cannot learn (ultimatum game), and the case in which the buyer does not need to learn about her value of the good (the classical seller-offer game with initial private information of the buyer, see e.g. Fudenberg et al. [1985] and Gul et al. [1986]). The intuition for the result is simple: in my model the buyer can ensure positive payoff because she *can* learn, but learning is *costly and takes time* and this is common knowledge. The latter fact allows the seller to extract a non-trivial amount of rent. The main technical innovation in the analysis is that the model set up allows for closed-form solutions and for comparative statics results in stationary equilibria with respect to the information choice of the buyer. This is despite the overall complexity of the dynamic game.

Chapter 2 is devoted to a model of dynamic information design. In our model, a benevolent sender communicates non-instrumental information over time to a Bayesian receiver who experiences gain-loss utility over changes in beliefs ("news utility"). News utility is characterized by two features: loss aversion and diminishing sensitivity over the magnitude of news. It is well-known in the literature that in the same set up, with loss aversion but without diminishing sensitivity, one-shot resolution of uncertainty is the optimal disclosure mechanism.

We show how to compute the optimal dynamic information structure for arbitrary news-utility functions. With diminishing sensitivity over the magnitude of news, one-shot resolution of uncertainty is strictly suboptimal under commonly used functional forms. Information structures that deliver bad news gradually are strictly dominated by one-shot resolution of uncertainty. We identify additional conditions that imply the sender optimally releases good news in small pieces but bad news in one clump. We give results and numerical examples in the same spirit, both for a mean-based version of news utility with various popular parametric specifications, as well as for the more 'traditional' version of news utility from Kőszegi and Rabin [2009]. In a quadratic specification of news utility and for a horizon of two periods, we can verify analytically the conditions for the optimality of the one-shot-bad-partial-good news information structure. These analytic conditions say that diminishing sensitivity is strong enough when compared to loss aversion near zero.

When the sender lacks commitment power, diminishing sensitivity leads to a credibility problem for good-news messages. Due to diminishing sensitivity there are utility gains from false hope and the costs associated to the revelation of the true state of the world at the end of the horizon are low, whenever loss aversion is not high enough. This intuition delivers the result that without loss aversion, the babbling equilibrium is essentially unique. Higher loss aversion enables equilibria with higher utility, so that more loss-averse receivers may enjoy higher equilibrium news-utility. This is in stark contrast to the case with commitment, and shows explicitly that in the presence of diminishing sensitivity, the ability to commit has value for the sender.

We discuss applications of our results. News utility and diminishing sensitivity lead to endogenous informational preferences. We show how this may explain the co-existence of different editorial policies in the media industry and justify the prevalence of the sudden-death format for game shows.

Chapter 3 (co-authored with Yi-Hsuan Lin) is devoted to the study of random choice induced by learning through a sequential experiment. Motivated by a sizable empirical and experimental literature about choice data and response times, we postulate that an analyst has access to random choice data showing the joint frequencies of choice and also decision time (RCDT). We assume the analyst lacks knowledge of the sequential experiment that the agent is using to learn about a payoff-relevant state of the world. Our aim is to show what can be identified through this type of data without making any parametric assumptions about the form of the sequential experiment.[2] We focus on two distinct cases of costly experimentation. In the first case, the agent discounts future payoffs geometrically (impatience). In the second, they pay a constant flow cost of time (price per unit of time). We show that such random choice data uniquely identify the discount factor in the first case and the cost of time in the second case, besides identifying the agent's prior and taste. The method works for essentially arbitrary sequential experiment.

As a first step towards identification, we show how an outside observer with access to this random choice data can conduct welfare analysis despite being oblivious of the technology of sequential experiments the agents are using. Technically, the results use an envelope theorem argument (Milgrom and Segal [2002]) and adaptations of insights from Lu [2016]. The output are welfare identities which express the welfare of an agent as a function of the observable RCDT. We leverage this approach to prove several typical welfare comparative statics in our setting, e.g. how can RCDT be used to estimate welfare change due to a lump-sum transfer or due to a subsidy towards the duration of the experimentation.

Once the welfare identities are proven, we can use these to identify parameters of the agent. The identification of her taste uses canonical procedures. We take variations of decision problems, the associated welfare identities and combine them to identities which only depend on the RCDT of the agent. From these identities we can solve for the prior of the agent and for the learning costs in both cases, discount factor in the case of impatience and flow costs of time in the remaining case. Our results illustrate the power of decision time data for identification in learning models.

All three chapters revolve around learning, albeit in different micro-theoretic models and setups. The first chapter characterizes optimal learning and its

[2]In particular, we allow for arbitrary correlation between experimental outcomes at different points in time.

consequences in a novel bargaining model with several realistic features, the second shows how to design a learning protocol for non-instrumental information for a behavioral agent who experiences news utility, whereas the third focuses on what can be identified from random choice and decision time data originating from an agent who is learning through an unknown sequential experiment.

1

Bargaining with endogenous learning

This chapter is based on my job market paper with the same title as the chapter. I am indebted to Eric Maskin, Jerry Green, Drew Fudenberg and Tomasz Strzalecki for their continuous support and encouragement. I thank Arjada Bardhi, Daniel Barron, Daniel Clark, Krishna Dasaratha, Oliver Giesecke, Ed Glaeser, Brett Green, Kevin He, Giacomo Lanzani, Robin Lee, Shengwu Li, Jonathan Libgober, Niccolò Lomys, Anh Nguyen, Indira Puri, Matthew Rabin, Sarah Ridout, Marco Schwarz, Kremena Valkanova and Alex Wolitzky for helpful comments in different stages of this project, as well as seminar audiences at Harvard and MIT for their feedback during presentations. I also thank Artur Ananin, Anna Avdulaj, Krishna Dasaratha, Tomasz Grzelak, Minella Kalluci, Rinald Murataj and Pauline Rueckerl for helpful conversations on business situations in which costly endogenous learning leads to delay in agreement. Any errors are mine.

1.1 Introduction

In bargaining situations often one party can learn privately during the negotiations. For instance, before investing in a start-up, an institutional investor typically conducts market surveys or seeks expert advice about proprietary technology. Learning takes time and resources, but creates an informational advantage. The newly-informed party gains strategic option value, because she can influence the negotiations by disclosing the information.

Appraising the market potential of innovative products is in fact of prime importance in negotiations in the venture capital (VC) industry. To quote the legendary Silicon Valley engineer and venture capitalist Eugene Kleiner: *No matter how ground-breaking a new technology, how large a potential market, make certain customers actually want it.*[1] This investigative process is typically arduous and may lead to significant delays in negotiations. In this spirit, the recent survey study Gompers et al. [2019] finds that closing a deal in the VC industry in the U.S. takes on average 83 days.[2] They also document significant variance in delay until agreement depending on industry, firm characteristics and location; e.g. average delay in California is 65 days, for late-stage firms it is 106 days.

Considerations of market appraisal also afflict venture capital negotiations in emerging markets where initial uncertainty may be related to regulatory trends or the evolution of cultural tastes (see e.g. the case study 'Sula Vineyards' in Zeisberger et al. [2017a] about private equity investment in the nascent wine industry in India).

Endogenous learning pervades negotiations beyond those in the VC industry. A publisher usually prospects the market while bargaining with a new author for the copyrights on her book. A real-estate developer hires lawyers to perform title

[1] See Kleiner's laws in `http://entrepreneurhalloffame.com/eugene-kleiner/` or `www.economist.com/obituary/2003/12/04/eugene-kleiner`.

[2] Deal selection and closing can take even longer in private equity situations not involving innovative products. See Zeisberger et al. [2017b] and Gompers et al. [2019] for case studies that exemplify this.

exams and employs market analysts to forecast the value of a property in a year's time. In government procurement situations, private companies often conduct studies on the benefits of a particular federal contract while at the same time negotiating the terms of agreement.

Motivated by the examples above, this paper studies an abstract dynamic bargaining game in which a buyer can endogenously generate information over time about her valuation of a unit good offered for sale by a seller.

The analysis sheds light on the following questions.

i. When may it be without loss to assume exogenous information, rather than endogenous acquisition of information?

ii. What is the buyer's optimal information acquisition strategy? How can one compare information acquisition across different bargaining environments (comparative statics)?

iii. How can one quantify the efficiency loss and delay in agreement when information takes time to arrive, is costly and endogenously acquired?

The model I study is a modified version of the standard one-buyer/one-seller dynamic bargaining game as introduced in Fudenberg et al. [1985] and Gul et al. [1986].[3] Henceforth the buyer is called Buyer (female pronouns) and the seller is called Seller (male pronouns). Seller has zero value for the good he initially owns (a normalization). The good may be of high or low (non-negative) value to Buyer. Initially the two parties share the same prior belief about Buyer's value.

An informational asymmetry arises over time as Buyer learns endogenously and privately about her valuation. Buyer's learning is stochastic and occurs over time: I consider both the case in which she is able to influence the rate of learning and the one in which she is not. Once the opportunity to learn arrives, Buyer chooses the extent to which she wants to exploit this opportunity, i.e. how accurate a signal to acquire. I assume new information is verifiable and hence

[3]In bargaining theory parlance I consider a seller-offer, (weak) gap case game. See Ausubel et al. [2002] for a survey of the classical dynamic bargaining literature with incomplete information.

cannot be misrepresented. Once new information is acquired, Buyer decides whether and when to disclose it to Seller.

This model departs from traditional bargaining models in two crucial ways. First, players are initially symmetrically informed and one party (Buyer) may become more informed over time. Second, disclosure of new information is feasible and any information disclosed is verifiable. Therefore, Seller learns from two different channels about the current valuation of Buyer: from the information disclosure decision, as well as from the rejection of past prices. The latter is the traditional learning channel in the standard dynamic bargaining game, whereas the former is a new channel.

There are two main takeaways from the analysis. First, learning creates option value for Buyer which is strategic in nature. The possibility to learn typically ensures Buyer a surplus. Second, because learning takes time and is costly, Seller also receives some surplus from the negotiations. This contrasts with the outcomes in the standard game of Fudenberg et al. [1985] and Gul et al. [1986] in which prices and Seller equilibrium payoff are generically extreme because Buyer possesses initial private information for free: Seller typically receives minimal rent in the traditional seller-offer game.

Without the possibility to learn, Seller extracts all surplus immediately. With the possibility to learn, Buyer can wait until she learns about her valuation, disclose if her valuation is low and otherwise keep silent to profit from the private information by pretending she has not learned yet. Thus, the fact that learning is possible, but associated with costs, typically implies non-extreme equilibrium payoffs. Non-extreme payoffs imply trade at non-extreme prices.

One can draw an analogy of the *strategic option value* from learning with the actions of an investor in a capital market. Without the possibility to learn the payoff structure of Buyer at the moment of trade is akin to the investor acquiring a stock at the fair price. She receives zero profit from the trade. The possibility to learn is akin to the investor owning a call option on the stock at the beginning of time: if there is good news the option is exercised at a profit at the moment of trade, if there is bad news there is no downside at the moment of trade.[4] Acquiring and exercising the strategic option in my model is inefficient because learning leads to delay, as well as other costs associated to it. Nevertheless, learning happens in all equilibria and typically results in positive Buyer payoff.

I focus first on the case in which the exploitation of an information source is costless, a useful benchmark for comparison to the often more realistic case of costly information.[5] The costless learning benchmark may also have independent merit in some applications, e.g. when information may arrive through a social network, say of friends or colleagues, rather than through employing internal or external resources in an advisory capacity.

Next, I consider extensively the more realistic case of costly learning. Some of the results change when compared to the case of costless learning; for example, Buyer may receive zero surplus in the costless learning case, so costly learning is needed to guarantee non-extreme outcomes. Learning costs can arise in two distinct dimensions in my model: *exploration/search costs* for sources of information and *exploitation costs* of a newly found learning source. For instance, the manager of a pharmaceutical lab may hire additional scientists to search for new test ideas on the effectiveness of a new drug compared to more traditional ones (exploration/search phase), or once a new test idea is available, she may decide on the scale of the study that implements it (exploitation phase). I endogenize both of these aspects of Buyer learning and show that they behave differently from each other.

The next subsection explains the results of the paper in more detail.

[4]The analogy to a perfect capital market is not complete, because Buyer receives a positive payoff overall from exercising the strategic option value. This is because she is the only party who can learn in my model, which corresponds to a capital market imperfection in the investor analogy.

[5]Strictly speaking, whenever Buyer has to wait for information to arrive there are opportunity costs from waiting, which are reflected in the fact that future utility is discounted. With *costless learning* I mean that there are no *direct, physical* costs related to acquiring and exploiting information sources.

1.1.1 Preview of results

I focus on perfect Bayesian equilibria (PBE): both players follow optimal contingent plans given their beliefs about Buyer's value, and their beliefs are derived from equilibrium strategies and Bayes' rule whenever possible.[6] Discrete-time dynamic bargaining games are notoriously difficult to study. Therefore, many (but not all) results hold under the qualification that the period-length is small enough (near the high-frequency limit, henceforth *near the HFL*), and some hold in the limit as period-length shrinks to zero and time becomes continuous (henceforth *in the HFL*).[7] When period-length shrinks I assume the players discount payoffs within that period less and that the probability that Buyer receives an opportunity to learn diminishes. In particular, the probability of learning within a period vanishes as period-length goes to zero.

The efficient outcome in the game is for trade to happen immediately so that learning does not start.[8] This is because learning always leads to delay, apart from any additional costs that may be associated with it. All equilibria with/without costs feature positive delay and thus are inefficient. This is because Buyer accepts the delay in order to capture the strategic option value from learning new information: if she learns bad news about her valuation, she may disclose this to get a lower price, whereas if she learns good news she can keep silent and pretend no new information has arrived to buy the good at a price below its actual value. Nevertheless, *approximate* efficiency is a feasible equilibrium outcome near the HFL, even if learning is costly. A detailed explanation of the results follows.

Exogenous intensity, costless accuracy. I first consider the case where the opportunity to learn new information arrives privately to Buyer at a positive Poisson rate of λ and Buyer can freely choose how much information to acquire. In this case, it is a weakly dominant strategy for Buyer to learn perfectly whenever

[6] See part IV of Fudenberg and Tirole [1991] for a formal definition.

[7] Period-length is inversely proportional to the frequency of interaction. Thus, *near the HFL* equivalently means *for a high enough frequency of interaction*.

[8] I consider the case of negative good values as an extension (see below). With negative values learning can be welfare-enhancing.

she gets the chance. In fact, acquiring full information whenever possible is strictly optimal in a PBE in which Buyer receives a positive equilibrium payoff. Inuitively, free information never hurts, and so is always consumed by an agent without commitment. Thus, whenever information is costless but arrives at a random date, the assumption of endogenous choice of informativeness is equivalent to the assumption that information is exogenous, one-shot and delivers conclusive evidence to Buyer about her value.

Second, near the HFL there are generically no PBEs in which Buyer with bad news discloses immediately and Seller screens sequentially the Buyer for the other valuations, as long as there is no disclosure. Thus, the usual logic of Seller 'screening down the demand curve' that appears in many standard dynamic bargaining models fails in this setup.[9] The reason is that the disclosure decision leads to an *interim* update of Seller about the private information of Buyer. If Buyer with bad news always discloses, no disclosure is interpreted as indication of higher valuations. This counteracts the belief update after a price rejection, which is interpreted by Seller as an indication of lower values. Overall, the first effect is strong enough to overpower the second.

I also show that for high enough λ, but fixed period-length, there are equilibria in which Buyer gets zero payoff, because she cannot succesfully pretend to be Buyer who has not received any news. Therefore the possibility of private endogenous learning *per se* does not necessarily ensure informational rents for Buyer. In these equilibria, Seller always charges the highest price that may be accepted by some Buyer type he deems feasible, given the information he has about Buyer's valuation *(high-price equilibria)*. High-price equilibria do not survive costly learning. Moreover, there cannot be any equilibria with zero Buyer payoff whenever period-length is small enough. This is because of Seller's lack of commitment: the arrival of information in the first period is very unlikely so that asking for a high price is suboptimal already in the first period. This results in

[9] See e.g. Fudenberg et al. [1985], Gul et al. [1986] for classical and many recent papers, e.g. Fuchs and Skrzypacz [2010] or Hwang and Li [2017] for results featuring equilibrium dynamics in which seller *screens down the demand curve*.

informational rents for Buyer.

I focus then the analysis on stationary equilibria in which Seller mixes between at most two prices after non-disclosure, Buyer who has learned good news trades without delay, whereas Buyer who has learned bad news discloses immediately and accepts the revised price of Seller. Henceforth these are called *strongly stationary equilibria*. These always exist near the HFL and besides strongly stationary equilibria with pure pricing by Seller, there are also strongly stationary equilibria near the HFL that feature mixed pricing upon non-disclosure. In particular, there is equilibrium multiplicity.

In the HFL, any mixing by Seller upon non-disclosure disappears and prices converge to a single number. Thus, Seller asks for a flat price in the HFL, unless he sees evidence of bad news at which point he revises the price down and the game ends. In the HFL of strongly stationary equilibria, both players have positive payoffs. The amount of expected real time delay is indeterminate and varies from zero to $\frac{1}{\lambda}$: for each value strictly between zero and $1/\lambda$, there is a corresponding equilibrium with expected real time delay equal to said value. In particular, near the HFL there are strongly stationary equilibria with pure pricing whose outcome is arbitrarily close to efficiency. In these equilibria, Buyer who has not learned yet is made indifferent on path between buying now and continuing and chooses to buy now with high probability. This leads to vanishing delay near the HFL and to low probability of learning on path. Overall, this results in payoffs close to efficiency.

The indeterminacy in expected delay and payoffs occurs because the price charged in the HFL makes Buyer who has not learned yet indifferent between waiting to capture the strategic option value and stopping immediately. It is this possibility to wait that in turn creates the strategic option value from learning. This is because it allows Buyer who has learned good news to 'pool' with Buyer who has not learned yet and thus possibly get the high-value good at a bargain. On the other hand, because information is not readily available to Buyer at the start of time, Buyer never captures all surplus in any strongly stationary equilibria.

Overall, the combination of endogenous learning and disclosure of bad news ensures the existence of equilibria which feature neither extreme prices, nor extreme payoffs.

Exogenous intensity, costly accuracy. I consider two distinct costs of exploiting an information source, conditional on its arrival. In the first case, costs are deterministic and variable and more accurate information costs more. In the second case, costs are stochastic but lump-sum: whenever the opportunity to learn arrives, a cost is drawn from a distribution. If Buyer pays the cost she can exploit the learning source at no additional marginal cost; otherwise, she may wait for future opportunities to learn and more favorable draws of the lump-sum exploitation cost.

Surprisingly, the same set of results can be proven for both models of accuracy costs. First, *for any period-length*, all PBEs feature a positive payoff for Buyer. This is in contrast to the costless case. Common knowledge of Buyer's costs of information acts as an insurance device for Buyer's payoff. The intuition is the following. Because of the strategic option value of learning, all PBEs feature some non-trivial amount of learning from Buyer. Because learning is private, Seller has to compensate Buyer for the costs of learning *on average*. This results in positive ex-ante surplus for Buyer because of the discretionary nature of her disclosure decision. Recall that in the case of costless accuracy there exist PBEs with zero Buyer payoff.

Second, just as in the case of costless accuracy, there are generically no PBEs near the HFL in which Buyer with bad news discloses and which feature deterministically falling prices on path (Seller does not 'screen down a demand curve'). The intuition is the same as before.

Subsequently, I focus again on strongly stationary equilibria. Near the HFL there exist such equilibria featuring pure pricing on path from Seller, for any parameter values of intensity and accuracy costs. Under a condition postulating that the arrival rate of the opportunity to learn λ is not too high compared to the players' impatience level, there are additional strongly stationary equilibria with mixed Seller pricing near the HFL.

In all strongly stationary equilibria with pure pricing Seller incentivizes the information acquisition of Buyer by offering her the opportunity to buy the good at a low price in case of good news. In contrast to the case of costless learning, the price spread does not disappear in the HFL of mixed pricing equilibria. Near the HFL of equilibria with mixed pricing Seller charges most of the time the reservation price of Buyer with good news as long as there is no disclosure. To incentivize learning in such equilibria Seller occasionally charges a low price so that Buyer has incentives to learn, whenever the opportunity arrives. But since the probability to learn in a single period is very low when period-length is very small, Seller promises the low price within a period with probability declining to zero, as period-length goes to zero. This leads to maximal delay in real time for such equilibria, because Buyer has no choice but to wait for the opportunity to learn to realize the strategic option value.

In the HFL of strongly stationary equilibria with pure pricing expected delay is again indeterminate. Moreover, near the HFL there are strongly stationary equilibria that are almost efficient, and all such equilibria exhibit pure pricing. This is despite costly learning. The reasoning for these results is the same as in the costless case.

In the HFL, the payoffs of Buyer and Seller give insight into the inefficiency sources: the deviation from full efficiency is a weighted sum of the ex-ante surplus and the learning costs incurred on path. The inefficiency becomes larger the more impatient the players are or the lower the arrival rate λ.

Extension: endogenous intensity. I extend the model to allow for endogenous costly choice of the learning intensity λ. In this most general version of the model, both aspects of learning, intensity and accuracy, are endogenous.[10] Costs of intensity are deterministic and variable and they are incurred at the beginning of every period, as long as bargaining continues. All general results from the model with costly accuracy but exogenous intensity carry over in this more general framework. In particular, there are again strongly stationary

[10] Results about costless choice of λ are contained in section A.4 of the appendix.

equilibria near the HFL, which are uniquely parametrized by the average price quoted by Seller upon non-disclosure. Compared to the case of exogenous intensity, there is now an additional potential source of inefficiency coming from the intensity costs. This is quantified in the HFL, just as the previous sources of inefficiency.

Once the rate of opportunities to learn is a choice variable for Buyer, comparative statics for information choice are possible. Across all strongly stationary equilibria and in the HFL, the endogenous and stationary level of intensity increases in the ex-ante level of optimism and decreases in the patience level of the players. In contrast, accuracy choice is broadly speaking 'reverse-U-shaped' in the level of ex-ante optimism and independent of the impatience level.

Extension: pre-learning negotiations. The reasons why learning is inefficient in my main model are two-fold. First, trade is ex-ante efficient and therefore learning does not add to the social welfare. Second, I assume that Buyer starts to learn before Seller can make the first price offer. This is a natural feature of many real-world negotiations: the party who becomes first interested in the trade may naturally start to gather information privately before she actually shows her interest to the other side of the market.

If Buyer could commit to not start learning before she approaches Seller and instead allows him to make a first, pre-learning offer, then the inefficiency would disappear. This is true independently of period-length. I show this by adding an ex-ante stage to the bargaining game. In this stage, Seller can make a first price offer before the learning from Buyer's side can start. In equilibrium, Buyer accepts the price offer immediately and learning does not happen on path. This extension suggests, that there is scope for a more systematic study of the design of bargaining institutions for markets in which parties typically engage in costly private learning before and during negotiations.

Extension: possible negative Buyer value. Suppose that the lowest value of Buyer is negative but that in expectation trade is efficient ex-ante (i.e. positive Buyer value occurs with high enough probability to compensate for the possibility of negative value). Assume in addition that Buyer is free to walk away from the bargaining at any moment. In this situation learning can strictly improve welfare, despite being costly. This is because Buyer can always walk away when she learns bad news that lead to a negative valuation for the good.

In this set up Buyer's information acquisition choice is very similar to the case of non-negative values and in strongly stationary equilibria bargaining ends whenever Buyer receives news about the good. When she receives good news she trades without delay. When she receives bad news, she either discloses immediately whenever it leads to a positive valuation or she walks away immediately whenever it leads to a negative valuation.

If learning were impossible, the efficient outcome under imperfect information about the value of the good would again be to trade immediately. This would require trade in *both* states of the world, i.e. Buyer would incur a loss from trade ex-post, whenever value of good is negative. I illustrate that learning limits this downside by lowering the probability of trade in the case of a negative valuation.

Outline of the rest of the paper. The next subsection discusses related literature. Section 1.2 introduces the basic model, states auxiliary results which are valid for all model versions and studies the case of costless learning. Section 1.3 introduces costs for accuracy and studies their implications. Section 1.4.1 discusses endogenous intensity and contains some comparative statics results. Section 1.4.2 shows how efficiency can be restored through pre-learning negotiations. Section 1.4.3 discusses negative values. Section 1.5 concludes. Formal proofs are contained in the appendix. Results that are not central to the main takeaways of this paper are contained in section A.4 of the appendix.

1.1.2 Related literature

The study of bargaining games with informational asymmetries has a long tradition in economic theory. This literature begins with the seminal papers Fudenberg and Tirole [1983], Sobel and Takahashi [1983], Cramton [1984], Fudenberg et al. [1985], Fudenberg et al. [1985] and Gul et al. [1986]. These papers focus on the case in which one or both bargaining parties have *initial* private information about their valuations and study how bargaining parties learn about the private information of their strategic opponent from price offers and rejections. A focal point of the analysis is the validity of the Coase conjecture. This conjecture prescribes that, as bargaining parties interact more and more frequently, delay until agreement vanishes and the informed party's rent is maximal.

Almost since its beginning, the literature with asymmetric information has focused on understanding the economic forces behind inefficient delay and whether the Coase conjecture survives, in one form or another, more complicated economic environments. To mention a few seminal contributions, Cramton [1984] and Chaterjee and Samuelson [1987] find that two-sided private initial information may lead to costly delay, Rubinstein [1985] that delay is possible whenever a player is uncertain about the time preferences of her bargaining counterpart, whereas Deneckere and Liang [2006] find that the same may happen if the parties have interdependent values.[11]

Starting from Fudenberg et al. [1987], the literature has studied bargaining under the existence of other potential trading partners, or more generally outside options. Board and Pycia [2014] shows that Seller may get significant surplus (and thus the Coase conjecture fails), even though agreement is immediate, whenever Buyer has an outside option at the beginning of the game. Fuchs and Skrzypacz [2010] as well as Hwang and Li [2017], Hwang [2018a] and Lomys [2017] focus on the effects of stochastic arrivals of outside options. Hwang and

[11] This list is by far incomplete. Other models of delay include: Abreu and Gul [2000] due to irrational players and reputation building, Feinberg and Skrzypacz [2005] due to higher-order beliefs and Yildiz [2004] due to ex-ante optimism (non-common priors).

Li [2017] and Hwang [2018a] in particular look at the case where the exogenous arrival of the outside option leads to a FOSD-shift upwards of the valuation of Buyer.[12] They consider private arrival of the outside option and find similar equilibrium dynamics as the strongly stationary equilibria in this paper, in addition to verifying the Coasian conjecture under some parameter restrictions.

While the existence or arrival of an outside option might also be interpreted as additional private information, none of the above mentioned papers models information acquisition explicitly. Learning new information leads to a *mean-preserving-spread (MPS)* in the distribution of Buyer valuations and in this paper, this MPS turns into a FOSD-shift upwards in the equilibrium dynamics of strongly stationary equilibria through the equilibrium disclosure choice of Buyer who receives bad news. Thus, the resulting FOSD-shift in valuations is an equilibrium property, rather than assumed in the model primitives. Additionally, results here show that the existence of sequential screening dynamics near the HFL hinges upon the assumption of private initial information, which is absent in the model of this paper. Finally, in the set up considered in this paper Coase conjecture fails and equilibria typically exhibit non-extreme prices and payoffs for *all* parameter values of the game considered.

This paper also connects to the literature on evolving valuations/roles in dynamic situations. Ortner [2017] and Ortner [2019] consider bargaining situations in which the change is exogenous, whereas Bergemann and Välimaki [2019] offer a review of the recent related literature on dynamic mechanism design, in which the standing assumption is commitment from the part of one of the players. More related to this work, Ravid [2019] studies a seller-offer game, in which Seller has the private initial information about the quality of the good and Buyer is rationally inattentive to past prices and the product's quality. Similarly to this work, Buyer's valuation changes endogenously, there may be delay until agreement and Buyer obtains positive surplus in the equilibria characterized. In contrast to this work, there is no disclosure decision for information, because the costs of information are not due to *information generation/production* but rather to

[12] FOSD stands for *First-order stochastic dominance.*

information-processing. Maximal delay in my model is achieved by equilibria with mixed pricing, whereas Ravid [2019] focuses on equilibria with pure pricing. Moreover, positive costs of information are not necessary to ensure positive Buyer payoff in the model of this paper, whereas they are in Ravid [2019]. Finally, in the setting of this paper sequential screening dynamics do not play a role, whereas they play an important one in Ravid [2019].

Daley and Green [2019] considers a model of bargaining in which Seller has initial private information about the value of the good and public, exogenous news delivers garbled information about the value of the good. In their setting interdependence of values is necessary but not sufficient for delay. Esö and Wallace [2019] considers a game with interdependent values in which both Buyer and Seller can learn their value of the good exogenously at a random date, privately and independently of each other. They assume each player can disclose verifiably their updated valuation to the other party. The interplay between interdependent values and two-sided exogenous learning implies there is no scope for equilibria with mixed prices and that with infinite horizon there is no delay near the high-frequency limit.

In this paper only Buyer learns endogenously and privately and I model explicitly the costly generation process of new information. In terms of results, I characterize equilibria which feature delay in real time, even though values of Buyer and Seller are independent and trade is efficient in every state of the world. Moreover, I show that equilibria with price mixing on path, which are absent from the above mentioned papers, are not knife-edge cases but a robust prediction of endogenous private learning.

Crémer and Khalil [1992] and Crémer et al. [1998] are classical works on information acquisition before the signing of a contract. More recently, Shi [2012] and Li [2019] consider auction settings in which Buyer types can acquire costly signals about their valuation before bidding in the auction.[13] Kirpalani and

[13]Shi [2012], Hwang [2018b] and Esö and Wallace [2019] are to the best of my knowledge the only other papers which consider an incomplete information game between a set of buyers and a seller without the assumption of initial private information.

Madsen [2019] studies how private information acquisition and social learning through public investments affect investment timing in settings in which investments are non-rival.

In contrast, this paper studies a bilateral dynamic bargaining model and focuses on the delay caused by endogenous learning, besides studying the optimal information acquisition of Buyer in both of its dimensions: intensity and accuracy.

Finally, this paper relates to the classic literature on strategic information transmission of verifiable information without commitment, beginning with Grossman [1981] and Milgrom [1981].[14] In particular, the strongly stationary equilibria of this paper feature partial unravelling analogous to Dye [1985]: Buyer with bad news presents evidence if she has it, while Buyer with good news pools with Buyer without evidence. Esö and Wallace [2014] consider a static bargaining model with verifiable disclosure and two-sided private information and focus on the value of verifiability of private information. Most of the work in the literature on verifiable disclosure has been on static disclosure; notable exceptions are Acharya et al. [2011] and Guttman et al. [2014], which focus on dynamic verifiable disclosure to a market rather than a strategic audience. Most recently, DeMarzo et al. [2019] studies a model in which a seller designs tests she can use to certify product quality to a market. Similar to this paper, it considers the case of costly design and shows the existence of equilibria with partial revelation. However, the audience of the disclosure is non-strategic and the model is static; thus any study of potential inefficient delay due to costly test design is moot.

1.2 The model

There are two players: Buyer and Seller. Seller owns an indivisible good whose value for Seller is zero (a normalization). Buyer has potential value $\theta \in \{\bar{v}, \underline{v}\}$ for the good with $\bar{v} > \underline{v} \geq 0$.

[14] See Milgrom [2008] and Dranove and Jin [2010] for surveys.

Time is discrete, denoted by $t = \Delta, 2\Delta, \ldots$ with period-length given by $\Delta > 0$. I use the expression *near the HFL* to mean *for all* $\Delta > 0$ *small enough* and *in the HFL* to mean *in the limit as* $\Delta \to 0$.

At the beginning of the game Buyer does not know her value for the good and both Buyer and Seller share a common prior of high value π_0, which is strictly between zero and one. Denote by $\hat{v} = \pi_0 \bar{v} + (1 - \pi_0)\underline{v}$ the ex-ante common knowledge valuation of Buyer. Players are impatient with common discount factor $\delta = e^{-r\Delta}$, for some $r > 0$.

Learning. Every period starting from $t = 1$ with probability $\mu = 1 - e^{-\lambda\Delta}, \lambda > 0$ opportunities to learn about θ arrives for the *Buyer*. For tractability, I make the major simplifying assumption that Buyer can exploit only one opportunity to learn, i.e. she can learn additional information only once. Its arrival is private information of Buyer. Whenever the opportunity to learn arrives, Buyer can pick an experiment of the form[15]

$$\mathcal{E}_a : \{\underline{v}, \bar{v}\} \to \mathcal{P}(\{H, L\}) \text{ with } \mathcal{E}_a(\bar{v})(H) = \mathcal{E}_a(\underline{v})(L) = a \in \left[\frac{1}{2}, 1\right].$$

The accuracy chosen by Buyer is also private. a is the accuracy and it is a choice variable of Buyer. H stands for a signal which gives (possibly partial) evidence of high value and L for a signal which gives (possibly partial) evidence of low value. The informativeness of the experiment $\mathcal{E}_a$ is denoted by $I(a)$ and is given by $I(a) = \frac{a}{1-a}$.

Results would go through with other parametric forms of experiments, as long as they lead to concave value of information. If one allows for general experiments, this requires introducing two informativeness parameters, since there are two states of the world. It turns out that value of information can be non-concave in this model, when general two-parametric experiments are allowed.[16]

[15]Henceforth $\mathcal{P}(X)$ for a metric space X denotes the set of Borel probability measures over X.

[16]See section A.4 of the appendix for this. The possibility of non-concavity of the value of information in information acquisition models is a well-known phenomenon since the seminal work of

The parameter λ is called *intensity* throughout, whereas a is called *accuracy*. Intensity is exogenously given and time-invariant until section 1.4.1. In analogy to the terminology of the experimentation literature, intensity describes the *exploration/search* rate of the learning process, whereas accuracy choice is akin to the choice of how much to *exploit* an already available learning source.

Observing signal L or H after performing experiment $\mathcal{E}_a$ with accuracy $a \in [\frac{1}{2}, 1]$ leads to an updated valuation for Buyer (expected value given the signal and the accuracy a chosen). The possible valuations in $[\underline{v}, \bar{v}]$ that Buyer may have at each moment in time are called the **Buyer's type**. These consist of $\hat{v}$, the valuation of Buyer before she learns, as well as the updated valuations after learning. For all purposes of the analysis, Seller's belief after a public history about the private information of Buyer may be summarized through a probability distribution over possible valuations in $[\underline{v}, \bar{v}]$.

Disclosure choice. In every period, after opportunity to learn and learning, Buyer can choose to disclose verifiably her updated valuation, or choose to not disclose her updated valuation. Buyer can delay disclosure.

Under this assumption, non-disclosure can be due to two reasons only: Buyer has not learned yet or she has learned and chosen not to disclose until that point in time. The timeline of the game within a period is as in Figure 1.2.1.

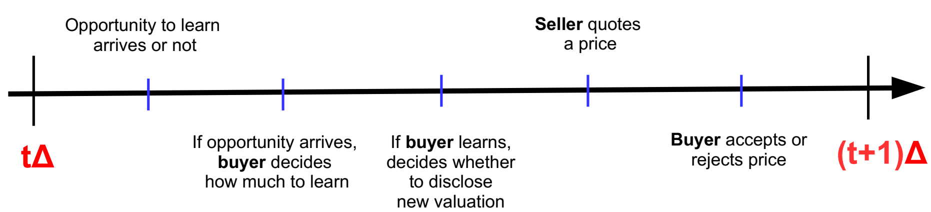

Figure 1.2.1: Timeline within a period.

Radner and Stiglitz [1984]. See Chade and Schlee [2002] for a modern treatment of this issue.

Histories, strategies and equilibrium. This game has two types of histories: *private and public*. A public history consists of a sequence of disclosure or non-disclosure events, as well as of rejected prices. Only Buyer has access to private histories, which include, in addition to publicly available information, both the occurrence of the arrival of the opportunity to learn as well as the learning outcome. A strategy for Buyer after a private history prescribes her choice of a, if that history ends with the arrival of the opportunity to learn, her choice of a probability of disclosure, if that history prescribes probability of disclosure at the end and there has been an opportunity to learn in the past. Finally, it prescribes an acceptance probability for a price quoted by Seller, if the history ends with that price quoted by Seller. A strategy for Seller prescribes a (possibly mixed) price offer after every public history ending with a disclosure/non-disclosure by Buyer.

Throughout, an equilibrium is a *perfect Bayesian equilibrium (PBE)*.[17] Buyer's strategy prescribes an optimal move after every private history, given Seller's strategy and Bayes updating about her value of the good. Seller's strategy prescribes an optimal mixing over prices after every public history in which he is called upon to quote a price, given Buyer strategy and Bayes' updating about the evolution of Buyer's valuation (whenever possible using Buyer's strategy).

Introduction of the assumptions on costly learning is deferred to sections 1.3 and 1.4.1.

Before delving into the analysis of the main results, I state two auxiliary results, which are very helpful in simplifying the analysis of PBEs throughout.[18] The first one mirrors similar results in the classical works Fudenberg et al. [1985] and Gul et al. [1986] on the seller-offer game with initial private information. It is valid for all model versions considered in this work. Before stating it, I define the concept of *reservation prices*.

Definition 1. Fix a PBE, a Buyer type w and a private history h which ends just

[17] See part IV of Fudenberg and Tirole [1991] for a formal definition.

[18] Both Lemma 1 and Lemma 2 here are also valid word-for-word for the models with costs. Thus, I do not restate them in sections 1.3 and 1.4.1.

before Seller has the possibility to quote a price. The reservation price of type w after h is the highest price that type w is willing to pay after h given continuation play in the PBE.

Lemma 1 describes general properties of PBEs of dynamic bargaining games in which Seller is the only player who makes price offers.

Lemma 1. *In any PBE the following hold true.*

1) Fix a public history h after which Seller is asked to quote a price and let $\underline{w}$ be the lowest possible Buyer valuation according to Seller's belief distribution over Buyer-types after h. Seller asks for at least $\underline{w}$ after h.

2) In any PBE, after every public history in which it is Seller's turn to move, Seller asks for prices among all reservation prices (given continuation play) of Buyer types she thinks are feasible right after that history.

3) After every private history, the Buyer type with the highest reservation price that has positive probability after that history, accepts an offer equal to that reservation price with positive probability.

4) After any disclosure event, Seller quotes a price equal to the disclosed valuation with probability one.

The strategies of a player in a PBE may be history-dependent and look back at more than just the preceding periods. Therefore, there is the theoretical chance that Buyer can reward Seller for prices lower than $\underline{v}$ by using history-dependent continuation play. Part 1) shows that this can never happen in a PBE.

The proof of part 1) also establishes the *skimming property*. It says that, after every public history, if Buyer of type w accepts a price p then so does every type with a strictly higher valuation $w' > w$.

Part 2) follows immediately from part 1) and the skimming property: if Seller wouldn't charge reservation prices, he would be leaving surplus on the table at no future benefit. Part 3) holds necessarily in every PBE to ensure that best responses of Seller are well-defined. Finally, part 4) is a direct implication of the assumption that learning is one-shot, disclosure is verifiable and that this is common knowledge: if Buyer discloses, Seller knows her valuation will not change in the future and so asks for the full surplus. Parts 1) and 4) imply that Buyer with bad news has zero surplus in any PBE.

Lemma 1 has several important economic implications whose proofs are contained in the appendix. First, it implies that in every PBE the reservation prices are strictly increasing in the type of Buyer. Second, it implies that there are no quiet periods in any PBE, i.e. after any history in which Seller is called upon to play, the probability of agreement is positive.

The second auxiliary result concerns the disclosure decision.

Lemma 2. *It holds true in all PBEs:*

- Buyer has strict incentives not to disclose good news on path, whenever the PBE features a positive Buyer payoff

- Buyer is indifferent between disclosing or not disclosing bad news on path,

- there are no strict incentives to delay disclosure of bad news.

Intuitively, if Buyer has received good news she cannot have strict incentives to disclose because she hopes to get a price lower than her reservation price. If she receives bad news, she knows her continuation payoff is zero and thus is indifferent between disclosing and not disclosing.

Equilibrium refinements

I close this subsection by introducing several equilibrium refinements which are used in the rest of the paper.

Refinement with respect to the disclosure decision. Equilibria in which Buyer with bad news does not disclose with positive probability are not 'robust' to the introduction of some slightly more realistic features into the model. For instance, Seller may have very small inventory costs for the good. Alternatively, he or Buyer may have small but positive overhead costs for continuing the bargaining, e.g. paperwork costs or costs for intermediaries for the communication or other meeting costs. In all of these cases the game ends once Buyer receives bad news. If Seller has small overhead costs for bargaining or inventory costs for the good, she 'bribes' Buyer with bad news into disclosing immediately by offering her a negligible surplus. If Buyer has small overhead costs for continuing the bargaining, she discloses bad news immediately to avoid future overhead costs.

In the following, I call an equilibrium a *disclosure equilibrium,* if it prescribes that Buyer who receives bad news on path discloses immediately and accepts with probability one the price offered subsequently by Seller.

STATIONARITY. As is usual for many dynamic bargaining games, I often focus the analysis on *stationary equilibria*.

Stationary equilibria satisfy the following properties.

i. Buyer's on-path actions depend only on her current type and Seller's current belief over Buyer types,

ii. Seller's on-path actions depend only on his belief distribution over Buyer types,

iii. If off-path play leads to a Seller-belief that happens with positive probability on path, ensuing play of Seller follows his on-path strategies.

Under a very mild technical requirement, section A.4 of the appendix shows that for every sequence of disclosure equilibria as $\Delta \to 0$ the beliefs at the start of each period converge to the degenerate distribution on $\hat{v}$, as long as bargaining goes on. This motivates requiring this property near the HFL as well.

A stationary equilibrium is called **strongly stationary**, if as long as bargaining goes on, Seller starts each period on path with belief concentrated on Buyer type $\hat{v}$.[19] Strongly stationary equilibria are analytically tractable and have intuitive closed-form solutions.

[19]Proposition 3 from subsection 1.2.1 below shows that there exist stationary equilibria which are not strongly stationary. Theorem 1, Propositions 6 and 7 show that strongly stationary equilibria exist near the HFL for both costless and costly learning.

Note that Buyer with good news in a strongly stationary equilibrium never rejects her reservation price on path. In particular, under strong stationarity the game never continues past the period of the arrival of information. This bounds the delay across all strongly stationary equilibria with disclosure because it takes in expectation $\frac{1}{\lambda}$ in real time for the opportunity to learn new information to arrive.

Finally, in the rest of the paper I call an equilibrium a **strongly stationary equilibrium with mixed pricing** if it satisfies

i. Seller mixes after on-path histories,

ii. it is a disclosure equilibrium,

iii. it is strongly stationary.

I call an equilibrium a **strongly stationary equilibrium with pure pricing** it it satisfies ii. and iii. above and i. is replaced with

i'. Seller does not mix after on-path histories.

The main economic property of the strongly stationary equilibria characterized in the rest of the paper is the strategic option value of Buyer from learning. This results in non-extreme payoffs for both Buyer and Seller and non-extreme prices.

1.2.1 A benchmark: costless learning

I now consider the benchmark case of costless learning. Thus, intensity is fixed $\lambda > 0$ throughout and Buyer can acquire any level of accuracy a for free, whenever the opportunity arises. The first major implication of costless choice of accuracy is that in every PBE it is a (weakly) best response for Buyer to learn conclusively, whenever the opportunity to learn arises. Learning conclusively is necessarily a strict best response in any PBE in which Buyer payoff is positive. The proof of the following Proposition is in section A.4 of the appendix.

Proposition 1. *There does not exist any PBE with a private history h (either on- or off-path), such that all of the following conditions are fulfilled* *i. Buyer is uncertain of the value of good,*

ii. at the end of h Buyer has an opportunity to learn,

iii. with positive probability after h and after opportunity to learn, Buyer picks $a < 1$,

iv. Buyer has positive continuation payoff after picking $a < 1$.

This result gives a micro-foundation for the assumption of

exogenous arrival of a one-shot and conclusive learning opportunity,

thus answering Question 1 in the introduction. If the learning is costless, then this assumption is without loss of generality for any PBE in which Buyer has a positive payoff. Moreover, this result simplifies the analysis of the costless case in that, there are only two possibilities: either Buyer receives zero equilibrium payoff or her information acquisition decision on the equilibrium path is trivial, because she chooses to learn perfectly whenever she can.

The next result formalizes the idea, that the usual logic of sequential screening and the related intuition of the optimality of 'screening down the demand curve' (so-called Coasian dynamics) fail in this model. First, I define the concept of sequential screening of valuations in this set up.

Definition 2. Say that a PBE features sequential screening of valuations if and only if on path

- as long as there is no disclosure, Seller quotes a decreasing sequence of deterministic prices $\{r_l, l \leq K\}$ $(K \leq \infty)$ with r_1 a reservation price of Buyer with good news
- the sequence of Seller-beliefs $\gamma_l \in \mathcal{P}([\underline{v}, \bar{v}])$, $2 \leq l \leq K$ that Seller entertains at the beginning of every period starting from the second, is strictly decreasing over time in the FOSD-sense.[20]

[20] At the start of the game there is no initial private information, i.e. $\gamma_1(\hat{v}) = 1$.

I first show that the logic of Seller pricing is enough to exclude disclosure PBEs when $\underline{v} > 0$ if the period-length is small enough. The case $\underline{v} = 0$ is more difficult to treat and I introduce additional assumptions.[21] It is treated in detail in the appendix.[22]

Proposition 2. *Suppose accuracy is costless and that $\underline{v} > 0$. Then there are no disclosure equilibria near the HFL in which Seller screens the valuations sequentially.*

Recall that the economic rationale for equilibria in which Seller 'screens down the demand curve' is that of screening for the initial private information that Buyer might possess at the start of the game. When there is common knowledge of an initial distribution of valuations and Buyer cannot learn, the fact that bargaining continues can only be interpreted as indication of lower valuations.[23] When the demand curve is endogenous and there is a disclosure decision, under the assumption that Buyer with bad news discloses immediately, Seller updates twice within a period as long as bargaining continues. He updates once from non-disclosure (an indication of higher valuations) and once from the rejection of a price (an indication of lower valuations). The first movement in beliefs is large enough as to neutralize the effect of the second, so that overall the classical sequential screening result fails.

Positive Buyer payoff only near the HFL

This subsection shows that when accuracy is costly, the mere possibility to learn does not ensure informational rents for Buyer, unless the frequency of interaction with Seller is high enough.

I first look for stationary equilibria in which Seller quotes $p_H = \bar{v}$ with probability one every period on path, as long as there is no disclosure. Let this

[21] In particular, I restrict to PBEs which are stationary and satisfy an equilibrium refinement called *'divinity in bargaining'*. See the appendix for more.

[22] In section A.4 of the appendix Proposition 2 is generalized to its 'real-time' counterpart: the sequential screening dynamics are allowed to start at some date $T(\Delta) \geq 1$ and so that $T(\Delta)\Delta \to 0$ as $\Delta \to 0$.

[23] This belief updating logic underlies the traditional 'Coasian dynamics' result. See chapter 10 of Fudenberg and Tirole [1991] for more on Coasian dynamics.

type of equilibrium be called a *stationary high-price equilibrium*. Buyer of type $\bar{v}$ accepts $p_t = \bar{v}$ with some probability $q_1 \in (0, 1)$ at $t = 1$. This results upon rejection of p_H in positive probability of type $\bar{v}$ at the beginning of period $t = 2$. Let this probability be γ and let $q(\gamma)$ be the probability with which the type $\bar{v}$ rejects p_H in $t \geq 2$.

Upon non-disclosure of $\underline{v}$ within a period, Seller updates her belief of Buyer with good news from γ to

$$U(\gamma) = \frac{\gamma + (1 - \gamma)\mu\pi_0}{1 - (1 - \gamma)\mu(1 - \pi_0)}. \tag{1.1}$$

This *interim* update is higher than γ: there is a *positive selection effect*, because no disclosure is stronger indication that Buyer may have learned good news.

If Buyer of type $\theta = \bar{v}$ accepts p_H with probability q, Seller updates the belief of $\bar{v}$ from $U(\gamma)$ to

$$B(U(\gamma), q) = \frac{U(\gamma)(1 - q)}{U(\gamma)(1 - q) + 1 - U(\gamma)}. \tag{1.2}$$

This is also the belief with which Seller starts the new period. It is strictly lower than the interim update $U(\gamma)$, because of a *negative selection effect*: rejection of a price is indication of lower valuations.

The condition for stationary beliefs on path from $t = 2$ on is given by

$$B(U(\gamma), q(\gamma)) = \gamma. \tag{1.3}$$

Thus, the positive and negative selection effects balance out at γ, whenever the type $\bar{v}$ rejects the price $p_H = \bar{v}$ with probability $1 - q(\gamma)$.

In the following let $W(\gamma)$ denote the stationary payoff of Seller from $t = 2$ onwards in the stationary high-price equilibrium. This aggregates over time the profit Seller makes from the arrival of the type $\bar{v}$, conditional on her accepting the price and the profit she makes from the type $\underline{v}$ who discloses immediately.

Given that type $\underline{v}$ discloses, the only possibly viable deviation for Seller is to ask for the reservation price of Buyer who has not learned yet. The following necessary conditions need to be satisfied for the stationary equilibrium.

$$U(0) \cdot q_1 \cdot \bar{v} + (1 - U(0)q_1)\, \delta W(\gamma) \geq \hat{v}, \qquad \text{Seller-optimality at } t = 1, \quad (1.4)$$

and

$$U(\gamma)q(\gamma) \cdot \bar{v} + (1 - U(\gamma)q(\gamma))\delta W(\gamma) \geq \hat{v}, \qquad \text{Seller-optimality at, } t \geq 2 \quad (1.5)$$

with

$$\gamma = B(U(0), q_1).$$

γ here gives the stationary belief at the beginning of periods $t \geq 2$. In the appendix I show the following result.

Proposition 3. *Let* $U(0) = \frac{\mu\pi_0}{1-\mu+\mu\pi_0}$ *with* $\mu = 1 - e^{-\lambda\Delta}$*, be the probability on the type* $\bar{v}$ *in the first period after no disclosure and* $W = \frac{\mu}{1-\delta+\delta\mu}(U(0)\bar{v} + (1 - U(0))\hat{v})$*. Then whenever the parameters satisfy*

$$(C - high) \qquad U(0)\bar{v} + (1 - U(0))\delta W > \hat{v},$$

there exists a stationary high-price equilibrium. In this equilibrium Buyer payoff is zero and Seller asks with probability one for $\bar{v}$ *as long as the bargaining continues and there is no disclosure.*

For fixed other parameters of the game, $(C - high)$ *is always satisfied when* λ *is high enough.*

The intuition for the existence of the stationary high-price equilibria is simple. When λ is large Seller knows that with high probability Buyer will know the value of the good very soon after bargaining starts. Non-disclosure is a strong indicator of good news whenever λ is large. Thus, Buyer with good news cannot successfully pool with Buyer who has not learned yet. But the strategic option value from learning comes precisely from being able to pool with type $\hat{v}$!

Stationary high-price equilibria from Proposition 3 do not survive near the HFL. The intuition is that as $\Delta \to 0$ the probability that Buyer has learned before any fixed period K goes to zero as well, at a speed of Δ. Thus, Seller cannot ask for the highest price already at $t = 1$. The price should be lower than $\bar{v}$ with positive probability in the first period, whenever Δ is small enough. But this implies that Buyer receives a positive information rent with positive probability already in the first period. This intuition does not depend on the assumption of stationarity. Therefore, more generally Buyer always ensures a positive payoff near the HFL.

Proposition 4. *Fix all parameters of the game except for the period-length Δ. There are no equilibria with zero Buyer payoff if Δ is small enough.*

If an equilibrium has zero Buyer payoff it is necessarily a high-price equilibrium (albeit maybe not stationary): Seller quotes $\bar{v}$ as long as there is no disclosure and bargaining goes on. Otherwise Buyer waits until the first period that Seller quotes with positive probability a price lower than $\bar{v}$ to realize informational rents with positive probability. The same argument as above for stationary equilibria shows that Seller would do better by offering some positive information rent already in the first period, as period-length shrinks to zero.

Intuitively, Seller would like to commit to quoting prices less often as period-length shrinks, so that he can become relatively certain that Buyer has learned in the meanwhile and her willingness to pay has increased. When there is no commitment across periods Seller quotes a price every period and thus allows Buyer to achieve positive information rent already in the first period with positive probability.[24]

Strongly stationary equilibria.

Next, I focus on *strongly stationary equilibria* in which Seller may or may not mix between prices on path upon non-disclosure. If he mixes on path, then he puts

[24] Incidentally, Proposition 4 also shows that costs of information are not necessary for Buyer to ensure a positive equilibrium payoff, unless there are rational-inattention costs of processing information, as Ravid [2019] shows.

positive probability on two prices $p_H > p_L$. p_H is the reservation price of Buyer who has learned good news, whereas p_L of Buyer who has not learned yet. If Seller does not mix on path after non-disclosure, he necessarily quotes the price p_L. This follows from the requirements of strong stationarity and the refinement with respect to the disclosure decision.

Let $p \in [0, 1)$ be the probability with which p_H is quoted upon non-disclosure. Suppose in the following, that in equilibrium Buyer of type $\hat{v}$, who has the option value, accepts her reservation price p_L with some stationary probability $q \in [0, 1]$.

I look extensively at the HFL of sequences of strongly stationary equilibria.

Definition 3. Say that a sequence of strongly stationary equilibria indexed by period-length $\Delta > 0$ with $\Delta \to 0$ *converges in the HFL*, if as $\Delta \to 0$

i. the average price quoted by Seller upon non-disclosure $\hat{p}(\Delta)$ converges,

ii. the sequence $q(\Delta)$ of acceptance probabilities of Buyer type $\hat{v}$ satisfies

$$\frac{q(\Delta)}{\Delta} \to \varkappa,$$

for some $\varkappa \in [0, \infty]$.

Say that a HFL corresponds to some $\varkappa$ if there exists a sequence of strongly stationary equilibria such that $\frac{q(\Delta)}{\Delta}$ converges to $\varkappa$ as $\Delta \to 0$.

Let $U(0) = \frac{\mu\pi_0}{1-\mu+\mu\pi_0}$. This is the (stationary) on-path probability that Seller puts on the type $\bar{v}$ after non-disclosure. Denote by $V_\Delta(q, p)$ the stationary Seller-payoff in the equilibrium, if the period-length is Δ.

Seller optimality upon non-disclosure in strongly stationary equilibria with pure pricing is ensured if

$$(U(0)+(1-U(0))q)p_L+(1-U(0))(1-q)\delta V_\Delta(q, 0) \geq U(0)p_H+(1-U(0))\delta V_\Delta(q, 0). \tag{1.6}$$

(1.6) ensures that deviating to p_H is not profitable for Seller. Lemma 1 ensures that these are the only 'relevant' deviations for Seller on path.

Seller indifference upon non-disclosure in strongly stationary equilibria with mixed pricing is equivalent to

$$U(0)p_H + (1 - U(0))\delta V_\Delta(q,p) = p_L(U(0) + (1 - U(0))q) + (1 - U(0))(1 - q)\delta V_\Delta(q,p). \tag{1.7}$$

The left-hand side is the payoff from charging p_H upon non-disclosure, whereas the right-hand side is the payoff from charging p_L upon non-disclosure.

Part 2) of Lemma 1 implies the following relations for the reservation pricing of types $\bar{v}, \hat{v}$ in strongly stationary equilibria with mixing probability of Seller given by $p \in [0, 1)$.

$$\bar{v} - p_H = \delta(p(\bar{v} - p_H) + (1 - p)(\bar{v} - p_L)), \tag{1.8}$$

and

$$\hat{v} - p_L = \frac{\mu\pi_0(\bar{v} - p_H) + \delta(1 - \mu)(1 - p)(\hat{v} - p_L)}{1 - \delta p(1 - \mu)}. \tag{1.9}$$

On the left-hand side of (1.8) and (1.9) is the payoff of the respective type if she decides to buy now when facing her reservation price, and on the right-hand side is the payoff if she decides to continue. Denote $\hat{p} = pp_H + (1 - p)p_L$ the average price quoted on path upon non-disclosure. (1.9) can be re-written with use of (1.8) as

$$\hat{v} - p_L = \frac{\delta\mu}{1 - \delta + \delta\mu}\pi_0(\bar{v} - \hat{p}).$$

This depicts the option value of the type $\hat{v}$. Buyer has the option to wait for the opportunity to learn at which case she gets a payoff of $\bar{v} - \hat{p}$, if she learns good news and of zero, if she learns bad news. The payoff from exercising the option value is discounted by $\frac{\delta\mu}{1-\delta+\delta\mu}$. This is the 'effective discount rate' for the option value, because the type $\hat{v}$ sometimes stops at price p_L (payoff has weight $1 - \delta$) and otherwise continues next period in the hopes of getting the chance to exercise the option value (payoff has weight $\delta\mu$).

The next result characterizes strongly stationary equilibria near the HFL and gives a complete characterization of their convergence in HFL.

Theorem 1. *1) [Existence near the HFL] Strongly stationary equilibria with both pure and mixed pricing exist near the HFL.*

2) [Uniqueness near the HFL] Near the HFL, every strongly stationary equilibrium with mixed pricing is unique up to the mixing probability $p \in (0, 1)$ of Seller.

Near the HFL, every strongly stationary equilibrium with pure pricing is unique up to the acceptance probability q of Buyer of type $\hat{v}$.

3) [Delay in the HFL] There exists HFL of strongly stationary equilibria corresponding to any $\varkappa \in [0, \infty]$.

In any HFL of strongly stationary equilibria with mixed pricing $\varkappa$ is 0 and positive expected delay is $\frac{1}{\lambda}$.

In any HFL of strongly stationary equilibria with pure pricing $\varkappa$ is in $(0, \infty]$ and expected delay is given by

$$\begin{cases} \frac{1}{\lambda+\varkappa}, & \text{if } \varkappa \in (0, \infty), \\ 0, & \text{if } \varkappa = \infty. \end{cases}$$

4) [Pricing in the HFL] In any HFL of strongly stationary equilibria prices converge to

$$\psi = \frac{r\hat{v} + \lambda(1 - \pi_0)\underline{v}}{r + \lambda(1 - \pi_0)}.$$

In particular, there is no price spread in the HFL of equilibria with mixed pricing.

5) [Payoff and efficiency properties in the HFL] Buyer and Seller payoffs in any converging sequence of strongly stationary equilibria with $\frac{q(\Delta)}{\Delta} \to \varkappa$ are unique. Buyer and Seller payoffs lie in $(\underline{v}, \hat{v})$ for all $\varkappa \in [0, \infty]$.

The efficiency loss in the HFL of equilibrium sequences with $\varkappa \in [0, \infty)$ is given by

$$\frac{r}{r + \lambda + \varkappa}\hat{v}.$$

There is no efficiency loss in the HFL of sequences with $\varkappa = \infty$.

Near the HFL, the mixing probability p of Seller upon non-disclosure is a sufficient statistic for the construction of the equilibria with mixed pricing: whenever two strongly stationary equilibria with mixed pricing share the same mixing probability of Seller, they prescribe identical play on path.[25] The Seller indifference condition pins down a unique q. Strongly stationary equilibria with pure pricing allow for a unique Seller mixing probability $p = 0$, but the acceptance probability for the reservation price of type $\hat{v}$ is determined only up to a lower bound.

The pricing in the HFL has a simple structure: Seller asks for a flat price ψ, unless he sees evidence that $\theta = \underline{v}$ and subsequently revises price down to $\underline{v}$. Buyer waits with some probability until she gets the information to end the game with either of the two prices. This is optimal due to two reasons. First, because ψ is lower than $\hat{v}$, but not too low as to compensate for the option value from learning. Second, as $\Delta \rightarrow 0$ the loss due to impatience from waiting an additional period is small, whereas the option value from learning remains strictly positive as $\Delta \rightarrow 0$.

The price spread $p_H - p_L$ in equilibria with mixed pricing, which is due to Seller attempting to screen the types $\{\hat{v}, \bar{v}\}$, disappears as Δ becomes smaller and smaller. This shows that the inefficient delay near the HFL of such equilibria originates mostly from Buyer of type $\hat{v}$ waiting to realize her option value from learning, rather than Seller trying to screen the types $\{\hat{v}, \bar{v}\}$ upon non-disclosure.

ψ corresponds to the HFL of the reservation pricing of type $\hat{v}$. This leads to the indeterminacy in expected delay in the HFL and implies that equilibrium multiplicity survives HFL. Depending on the acceptance probability, expected delay in the HFL can be any number in $[0, \frac{1}{\lambda}]$. The degree of inefficiency in the HFL of a strongly stationary equilibrium is characterized by the difference between the ex-ante surplus $\hat{v}$ and the sum of Buyer and Seller payoff in the HFL. The last part of Theorem 1 shows that near the HFL there are strongly stationary

[25]Whenever discussing strongly stationary equilibria, I use the words *a collection of variables are a sufficient statistic for the equilibrium* in the sense they are used here, i.e. the equilibrium is unique within the class of strongly stationary equilibria, once the value of the variables in consideration is fixed.

equilibria with (necessarily) pure pricing that are arbitrarily close to efficiency. These correspond to cases in which $\varkappa$ is 'very large', i.e. the sequence of acceptance probabilities $q(\Delta)$ falls relatively slowly on the scale of Δ.

In the HFL Seller pricing, and Buyer and Seller payoffs are non-extreme. In particular, even after accounting for losses from delay due to learning, Seller's payoff is not minimized as in the classical seller-offer game of Fudenberg et al. [1985] and Gul et al. [1986].

Finally, the potential inefficiency from learning is ceteris paribus decreasing in the patience level of the players. This is because delay hurts more, the more impatient players are.

1.2.2 Discussion of alternative assumptions

To get a better sense of the economic factors driving the results in the costless case and beyond, it is instructive to consider variations in the model assumptions with the classical game from Fudenberg et al. [1985] and Gul et al. [1986] in mind.

Suppose for a moment that Buyer knows her true valuation before she approaches Seller and that this would be common knowledge. Thus, Buyer has *initial* private information. Suppose that Buyer can disclose her valuation verifiably. Similar to the results in the model of this paper, type $\underline{v}$ is indifferent in her disclosure decision. In equilibria in which she discloses immediately, the average prices will be approximately $\bar{v}$ near the HFL unless there is disclosure. In equilibria in which she never discloses, the average prices are near $\underline{v}$ near the HFL. In either case, one bargaining party receives all the surplus and prices are extreme near the HFL, just as in the traditional game. Therefore, in my model it is the assumption of private arrival of information that ensures non-extreme equilibrium payoffs *and* prices, whenever the assumption of verifiable disclosure is maintained.

Suppose alternatively, that types are endogenous as in the model of this paper, but that communication between Buyer and Seller is impossible. I conjecture that in this case the equilibria are again extreme near the HFL, in that they either have extreme payoffs or equilibrium play ultimately exhibits extreme prices. In contrast, the combination of verifiable disclosure and stochastic evolution of types in my model enables equilibria near the HFL which exhibit neither extremeness of payoffs nor extremeness of prices.

1.3 Costly learning

This section introduces costs for accuracy. I consider two distinct models of accuracy costs and always assume that parametric forms on costs are common knowledge. In the first case, costs are deterministic and marginal costs of picking a higher accuracy are positive. I assume that acquiring an experiment of very low accuracy costs very little. In the second case, the costs are stochastic and independent of accuracy. Thus, the second model is one of fixed costs of learning. I assume in this second case, that arbitrarily low costs have positive probability. Formally, the assumptions are as follows.

Deterministic variable costs of accuracy. The experiment $\mathcal{E}_a$ with informativeness $I(a) = \frac{a}{1-a}$ costs $c(I(a))$ with $c : [1, \infty) \to \mathbb{R}_+$ satisfying

- $c(1) = 0$ and $c'(1) = 0$
- c is strictly convex and increasing
- $\lim_{I \to \infty} c'(I) = +\infty$.

It is easy to see that in the case of deterministic variable costs, in any PBE, Buyer always learns whenever she gets the chance, be it even by a bit, provided the option value from learning is strictly positive. This follows from the assumption that $c'(1) = 0$, i.e. an experiment close to uninformative costs almost nothing. It follows that the *learning rate* and the intensity λ are the same for deterministic variable costs of accuracy, in any equilibrium in which Buyer has positive option value from learning.

Stochastic fixed costs of accuracy. Conditional on an opportunity to learn having arrived and independent of everything else, a fixed cost $c \in (0, \infty)$ is drawn, which is distributed according to a distribution F. If Buyer pays c, she can verify the state (equivalently: can pick $a = 1$) at no additional cost. Otherwise she can wait for future draws. F satisfies the following requirements.

- F is continuous and has finite first moment[26]
- (Possibility of arbitrarily low costs) F puts positive probability to a neighborhood of zero.

Because of the lump-sum nature of the stochastic costs, and the fact that Buyer can always wait for lower cost draws, whenever the current cost draw is too high, the *rate at which the agent learns* becomes endogenous and distinct from the rate of arrival of opportunities to learn, given by the intensity parameter λ.[27] To avoid confusion in notation, I use the definition $\mu_0 = 1 - e^{-\lambda\Delta}$ for the case of stochastic fixed costs only, for the probability of the arrival of the opportunity to learn within a period (intensity) and keep the notation μ for the probability with which Buyer actually learns within a period in a stationary equilibrium.[28] The latter is now endogenous.

Despite their significant differences, broadly speaking the same set of results turns out to be true for both of the models of costly learning. The first result establishes that costs ensure a positive Buyer payoff in every PBE. This is true for any Δ, in contrast to the case of costless learning, in which, for every $\Delta > 0$ one can construct PBEs with zero Buyer payoff.

Theorem 2. *Pick any $\Delta > 0$. If learning is costly, every PBE has a positive Buyer payoff.*

Proof-sketch. Fix a $\Delta > 0$.

[26] An alternative and for-all-purposes-equivalent assumption is that F is continuous, has bounded support contained in $[0, +\infty)$.

[27] Recall that in the case of deterministic variable costs these are the same.

[28] It is easily established that this object is time-stationary in a stationary equilibrium.

A PBE with zero Buyer payoff can only happen if on path Seller erases the option value from learning. He can only do this by quoting, after every public history when it is his turn to move, a price equal to the reservation price of the highest Buyer type she deems feasible at that moment in the game. But if there is no option value from learning Buyer strictly prefers not learn on path, because learning is costly. If Buyer does not to learn on path, then the best response of Seller is to ask for $\hat{v}$, after every public history, as long as there is no agreement.

Suppose this is the case and consider first the model with deterministic variable costs. After every private history which ends with the arrival of an opportunity to learn, Buyer does want to learn. This is because learning very little costs very little (recall $c'(1) = 0$), whereas the benefit from learning is an order of magnitude larger than the increase in marginal costs. Since it can happen with positive probability that Buyer receives the opportunity to learn in every period that the game goes on and she has not learned before (in particular, also in the first period), such a strategy would give Buyer positive payoff with positive probability. This is a contradiction.

Consider next the model with stochastic fixed costs. In this case, the option value from learning is strictly positive (at least as large as $\pi_0(\bar{v} - \hat{v})$). Again, since the opportunity to learn arrives with positive probability every period that the bargaining goes on, and the probability that the cost draw is below $\pi_0(\bar{v} - \hat{v})$ is strictly positive (due to the assumption of the possibility of arbitrarily low costs), the same argument as in the case of deterministic variable costs leads to a contradiction. □

The intuition for this result is surprisingly simple. The surplus may change only through Buyer-learning. Because the learning is private information, Seller can only gives incentives to Buyer to learn *on average*, and not conditional on every realized learning outcome. Because of the discretionary nature of the information disclosure decision, this creates informational rents for Buyer. That the proof works *for any* $\Delta > 0$, depends crucially on the assumption of zero marginal costs for experiments close to uninformative for the case of deterministic variable costs and the assumption of the possibility of arbitrarily low costs in the case of stochastic fixed costs.

Theorem 2 complements Proposition 4 in Ravid [2019], because it exhibits another situation in which costs of information (in this case of production, rather than information processing costs) ensure a positive payoff across PBEs. In both models learning creates surplus because it creates private information for Buyer. Positive Buyer payoff does not come from initial private information as in the classical setting of Fudenberg et al. [1985] and Gul et al. [1986], because Coasian forces are absent. The positive buyer payoff comes instead from the fact that learning is private and costly.

Finally, another implication of the proof of Theorem 2 is that there are no PBEs with costly learning in which Buyer chooses not to learn with probability one, whenever the opportunity to learn comes.[29] This implies that the no-sequential-screening-of-valuations result from the costless case extends to the case of deterministic variable costs. This is true without additional assumptions even when $\underline{v} = 0$, because any learning event in the case of deterministic variable costs leads to an updated valuation $\underline{w} > 0$ (bad news is never conclusive). Moreover, by adapting the proof of Proposition 2 the same result can be shown to hold for the case of stochastic fixed costs with $\underline{v} > 0$. Summarizing, one has the following extension of Proposition 2.

Proposition 5. *There are no disclosure equilibria near the HFL in which Seller screens the valuations sequentially, in the case of*

- deterministic variable costs on accuracy

- stochastic fixed costs of accuracy with $\underline{v} > 0$.

This extension is not surprising, because the proof of Proposition 2 only relies on the logic of Seller-pricing: as the period-length shrinks, the probability that Buyer has learned before a fixed finite date vanishes. Thus, as the length of the period shrinks, there is no common knowledge of a date in which the private information of Buyer is present and Seller can start the sequential screening.

[29]See Corollary 9 in the appendix.

1.3.1 Strongly stationary equilibria with accuracy costs

I construct the same type of strongly stationary equilibria as in the costless case. For strongly stationary equilibria with mixed pricing and costly learning, the sufficient statistic for the construction of the equilibria is the average price upon non-disclosure $\hat{p} = pp_H + (1-p)p_L$.

The case of deterministic variable costs of accuracy. Suppose the stationary valuation of Buyer with good news is given by $\bar{w}$, whereas the stationary valuation of Buyer with bad news is given by $\underline{w}$. It holds

$$\underline{v} < \underline{w} < \hat{v} < \bar{w} < \bar{v}.$$

The *option value* from information acquisition is a function of the pair $(a, \hat{p})$:

$$\begin{aligned} V_A(a, \hat{p}) &= \pi_0 a\bar{v} + (1-\pi_0)(1-a)\underline{v} - (\pi_0 a + (1-\pi_0)(1-a))\hat{p} \\ &= (\pi_0 a + (1-\pi_0)(1-a))\left(\bar{w} - \hat{p}\right). \end{aligned}$$

When the opportunity to learn arrives, Buyer learns and ends the bargaining in the same period. She discloses bad news to get the lower price $\underline{w}$, and does not disclose good news in which case she pays in expectation $\hat{p}$ to Seller.

Let the accuracy chosen on path be a. The valuations $\bar{w}, \underline{w}$ are given by

$$\bar{w}(a) = \frac{a\pi_0\bar{v} + (1-a)(1-\pi_0)\underline{v}}{a\pi_0 + (1-a)(1-\pi_0)}, \quad \underline{w}(a) = \frac{(1-a)\pi_0\bar{v} + a(1-\pi_0)\underline{v}}{(1-a)\pi_0 + a(1-\pi_0)}.$$

Whenever the opportunity to learn arrives, optimal learning results in the following two incentive constraints.

$$(OL - intensive) \quad a \in \arg\max_{\bar{a}}\{V_A(\bar{a}, \hat{p}) - c(I(\bar{a}))\}, \tag{1.10}$$

and

$$(OL - extensive) \quad V_A(a, \hat{p}) - c(I(a)) \geq \hat{v} - p_L. \tag{1.11}$$

OL stands for 'optimal learning'. Incentive constraint (1.10) refers to the *intensive margin* of the learning decision (i.e. how accurate a signal to acquire), whereas (1.11) to the *extensive margin* of the learning decision (i.e. whether to acquire a costly signal).

The reservation prices of Buyer with good news $\bar{w}$ and of Buyer who has not learned yet are given by

$$\bar{w}(a) - p_H = \delta(\bar{w}(a) - \hat{p}),$$

and

$$\hat{v} - p_L = \frac{\delta\mu}{1 - \delta + \delta\mu}\left(V_A(a, \hat{p}) - c(I(a))\right).$$[30]

I denote by

$$BL(\hat{p}) = V_A(a(\hat{p}), \hat{p}) - c(I(a(\hat{p}))), \tag{1.13}$$

the endogenous *benefit from learning* in the stationary equilibrium with sufficient statistic $\hat{p}$.

Seller indifference condition remains the same as in (1.7). Note that because $\hat{v} \geq p_L$, *OL-intensive* and the reservation pricing for Buyer of type $\hat{v}$ imply immediately that *OL-extensive* is satisfied. Therefore, this constraint can be dropped in the following w.l.o.g.[31] *OL-intensive* leads to the first-order condition

$$\pi_0(\bar{v} - \hat{p}) + (1 - \pi_0)(\hat{p} - \underline{v}) = c'\left(\frac{a}{1-a}\right)\frac{1}{(1-a)^2}.$$

This determines uniquely the optimal accuracy $a(\hat{p})$ and with it, also the rest of the variables of the equilibrium, except for the acceptance probability q of the type $\hat{v}$ for the price p_L. q is determined uniquely by Seller's indifference condition

[30] This follows from some algebra, starting with the reservation price relation for type $\hat{v}$:

$$\hat{v} - p_L = \delta\left(\mu(\pi_0 a + (1-\pi_0)(1-a))(\bar{w} - pp_H - (1-p)p_L) - \mu c(I(a)) + (1-\mu)p\delta V_L + (1-\mu)(1-p)(\hat{v} - p_L)\right),$$

where V_L is the continuation payoff of the type $\hat{v}$ when she starts a new period in the stationary equilibrium. Due to reservation pricing it holds $\delta V_L = \hat{v} - p_L$.

[31] Another way to see that *OL-extensive is redundant* is to combine Corollary 9 from the appendix and use the assumed stationarity of the PBE.

in equilibria with mixed pricing and is determined only up to a lower bound in equilibria with pure pricing.

The case of stochastic fixed costs of accuracy. When costs of accuracy are lump-sum but stochastic, there is no intensive margin for the learning decision: Buyer learns perfectly whenever she pays the costs. The extensive margin decision is explained as follows.

Let V_N be the continuation utility for Buyer of type $\hat{v}$. Upon non-disclosure, with probability p she faces a price of p_H which she rejects with probability one and gets the continuation payoff $\delta \hat{V}$, where $\hat{V}$ is the continuation payoff of starting a period in the stationary equilibrium as type $\hat{v}$. Due to reservation pricing and Seller's belief dynamics on path, it holds

$$\delta \hat{V} = \hat{v} - p_L.$$

On the other hand, with probability $1 - p$ type $\hat{v}$ faces price p_L and has continuation payoff $\hat{v} - p_L$. Overall, it follows $V_N = \hat{v} - p_L$. Let V_A be the continuation utility if Buyer learns, with the learning costs not yet subtracted. With probability π_0 Buyer becomes the high type $\bar{v}$ and so receives continuation utility $p(\bar{v} - p_H) + (1 - p)(\bar{v} - p_L) = \bar{v} - \hat{p}$. With probability $1 - \pi_0$ Buyer becomes type $\underline{v}$, discloses immediately, receives a payoff of zero and the game ends. It follows $V_A = \pi_0(\bar{v} - \hat{p})$.

The costs c are worth paying if and only if

$$c \leq V_A - V_N,$$

that is, if and only if they are low enough. In particular, the stationary probability μ that Buyer of type $\hat{v}$ learns within a period is given by

$$\mu = \mu_0 F\left(\pi_0(\bar{v} - \hat{p}) - \hat{v} + p_L\right). \tag{1.14}$$

In the following let $\bar{\mu}(\hat{p}, \Delta) = F(\pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L(\hat{p}, \Delta)))$ be the probability of incurring the costs, conditional on the opportunity to learn having arrived. Denote also $\bar{\mu}(\hat{p}) = F(\pi_0(\bar{v} - \hat{p}) - (\hat{v} - \bar{p}_L(\hat{p})))$ for any HFL of $\bar{\mu}(\hat{p}, \Delta)$ as $\Delta \to 0$. In difference to the case of deterministic variable costs, $\bar{\mu}$ is equilibrium-dependent and different from the probability of learning μ.

Seller indifference condition and the reservation price relation for type $\bar{v}$ are the same as in the case of costless learning (namely, formally the same as in (1.7) and (1.8)), with the major difference that now μ is endogenously determined in equilibrium. Reservation pricing for the type $\hat{v}$ leads to

$$\hat{v} - p_L = \frac{\delta\mu}{1 - \delta + \delta\mu} BL(\hat{p}),$$

with the endogenous *benefit of learning* $BL(\hat{p})$ given by

$$BL(\hat{p}) = \pi_0(\bar{v} - \hat{p}) - \mathbb{E}[c | c \leq V_A - V_N]. \tag{1.15}$$

The option value from the costless case, given by $\pi_0(\bar{v} - \hat{p})$ is reduced in (1.15) by the expected costs of learning, conditional on the event that learning occurs.

The following Proposition establishes existence of strongly stationary equilibria near the HFL.

Proposition 6. *Pick any $\pi_0, \underline{v}, \bar{v}$ and λ, r. In both cases of accuracy costs the following holds.*

1) [Existence near the HFL] Strongly stationary equilibria with pure pricing always exist near the HFL.

There exists an open neighborhood $\mathcal{N}$ of $\hat{v}$ such that strongly stationary equilibria with mixed pricing and average price upon non-disclosure $\hat{p} \in \mathcal{N}$ exist near the HFL, whenever the following condition is satisfied.

$$(P) \quad r > \lambda \text{ if } \pi_0 \leq \frac{1}{2} \text{ or } r > \sqrt{2}\lambda \text{ if } \pi_0 > \frac{1}{2}.$$

2) [Uniqueness near the HFL] For any fixed average price $\hat{p} \in \mathcal{N}$ the quantities $p_L(\hat{p}, \Delta), p_H(\hat{p}, \Delta), p(\hat{p}, \Delta), q(\Delta, \hat{p})$ are uniquely determined in every strongly stationary equilibrium with mixed pricing.

Pure pricing equilibria are unique up to the acceptance probability q of Buyer of type $\hat{v}$.

In the case of deterministic variable costs, $a(\hat{p})$ is unique and decreasing in $\hat{p}$ if $\pi_0 > \frac{1}{2}$, increasing in $\hat{p}$ if $\pi_0 < \frac{1}{2}$ and independent of $\hat{p}$ if $\pi_0 = \frac{1}{2}$.

In the case of mixed pricing the acceptance probability $q(\Delta, \hat{p})$ is unique.

The condition (P) $r > \lambda$ if $\pi_0 \leq \frac{1}{2}$ or $\sqrt{2} < \frac{r}{\lambda}$ if $\pi_0 > \frac{1}{2}$ for the existence of mixed pricing equilibria are used in the proof to show existence of the mixing probability $q(\hat{p}, \Delta)$ of the type $\hat{v}$, whenever Δ is small enough. (P) is not minimal (see the proof of Proposition 6 in the appendix for more on this), but it does not depend on the precise parametric specification of costs. Namely, it ensures existence of mixed pricing equilibria for any cost function c in the deterministic variable case and for any F in the stochastic fixed case, as long as these satisfy the original assumptions at the beginning of this section. (P) requires that the intensity is not too high compared to the impatience level of the players.[32]

The sufficient statistic for the construction of strongly stationary equilibria with mixed pricing is the average price $\hat{p}$. In the case of pure pricing there is one additional degree of freedom: the acceptance probability q for type $\hat{v}$.

The next result gives the HFL characterization of the strongly stationary equilibria with accuracy costs.

Theorem 3. *Pick any $\pi_0, \underline{v}, \bar{v}$ and r, λ.*

1) [Existence in the HFL] There exists HFL of strongly stationary equilibria with pure pricing corresponding to any $\varkappa \in [0, \infty]$.

Let condition (P) from Proposition 6 be satisfied and $\mathcal{N}$ as in Proposition 6. For every $\hat{p} \in \mathcal{N}$ with $\mathcal{N}$ as in Proposition 6 there exists a sequence of strongly stationary equilibria with mixed pricing such that the sequence of average prices $\hat{p}(\Delta)$ along the sequence converges to $\hat{p}$.

2) [Delay in the HFL] In any HFL of strongly stationary equilibria expected delay

[32] In section 1.4.1 the choice of λ is endogenous and costly, so that these parametric assumptions can be transferred to corresponding assumptions on the cost of choosing the intensity λ.

is given by

$$\begin{cases} \frac{1}{\lambda+\varkappa}, \ \textit{if } \varkappa \in [0,\infty) \textit{ and accuracy costs are deterministic and variable}, \\ \frac{1}{\lambda\bar{\mu}(\hat{p})+\varkappa}, \ \textit{if } \varkappa \in [0,\infty) \textit{ and accuracy costs are stochastic and fixed}, \\ 0, \ \textit{if } \varkappa = \infty. \end{cases}$$

3) [Pricing in the HFL] In both cases of accuracy costs the price spread in the HFL of a sequence of strongly stationary equilibria with mixed pricing is bounded away from zero, and the low price is charged with vanishingly small probability.

4) [Payoff and efficiency properties in the HFL] Buyer and Seller payoffs in any converging sequence of strongly stationary equilibria with $\frac{q(\Delta)}{\Delta} \to \varkappa$ *are unique. Buyer and Seller payoffs lie in* $(\underline{v}, \hat{v})$ *for all* $\varkappa \in [0,\infty]$.

The efficiency loss in the HFL of equilibrium sequences with $\varkappa \in [0,\infty)$ *is given by*

$$\frac{r}{r+\lambda+\varkappa}\hat{v} + \frac{\lambda}{r+\lambda+\varkappa}c(I(a(\hat{p}))), \tag{1.16}$$

in the case of deterministic variable costs of accuracy. It is given by

$$\frac{r}{r+\lambda\bar{\mu}(\hat{p})+\varkappa}\hat{v} + \frac{\lambda\bar{\mu}(\hat{p})}{r+\lambda\bar{\mu}(\hat{p})+\varkappa}\mathbb{E}_F[c|c \le \pi_0(\bar{v}-\hat{p}) - (\hat{v} - \bar{p}_L(\hat{p}))], \tag{1.17}$$

in the case of stochastic fixed costs of accuracy.

There is no efficiency loss in the HFL of sequences with $\varkappa = \infty$.

Theorem 3 showcases the differences as well as the commonalities between the cases of costless and costly learning.

First, the expected delay in the HFL of all $\varkappa$ but $\varkappa = \infty$ is positive in both cases. It equals that of the case of costless information in the case of deterministic variable accuracy costs. It becomes equilibrium-dependent in the case of stochastic fixed accuracy costs, because the rate of learning diverges from the rate of arrival of opportunities to learn.

Second, in the case of costly learning the price spread does not disappear for equilibria with mixed pricing. The reason for this is that it is necessary even in the limit to subsidize the information costs incurred with positive probability for any $\Delta > 0$. For Δ positive but small, Seller promises to occasionally charge a low price so that Buyer has incentives to learn whenever the opportunity arrives (gets the high value good at a bargain). Since the probability to learn in a single period is very low when Δ is very small, Seller promises the low price within a period less and less often as Δ vanishes. Overall, in the HFL of mixed pricing equilibria the reason for the delay is the same as in the costless case: Buyer waits to realize the option value associated with learning new information.[33]

Third, the price distribution of Seller in mixed pricing equilibria converges in the HFL to the reservation price of Buyer with good news, so that in the HFL of mixed pricing equilibria Buyer who has not learned yet waits until she can acquire new information. This leads to maximal expected delay in the HFL. In the HFL of pure pricing equilibria Seller incentivizes learning by offering the reservation price of Buyer who has not learned yet. Just as in the case of costless learning this this allows for multiplicity in expected delay in the HFL. In particular, near the HFL there are strongly stationary equilibria with (necessarily) pure pricing that are arbitrarily close to efficiency. This holds true despite the fact that in every PBE with costly learning Buyer learns with positive probability on path! The reason is that along a sequence of strongly stationary equilibria which converge to efficiency as Δ shrinks to zero, learning happens less and less often and Buyer of type $\hat{v}$ accepts her reservation price with higher and higher probability.

Fourth, the first part of both (1.16) and (1.17) quantifies how much of the ex-ante surplus is wasted whenever there is delay in the HFL. The second term of the sum in (1.16) and (1.17) quantifies the loss due to costly learning. The share of the 'pie' lost due to delay does not depend on Seller pricing in the model with deterministic variable costs, but it does so in the model with stochastic fixed costs. In the costless case, less patient players waste more of the overall gains from

[33]The proof shows that in general, the HFL of mixed pricing equilibria cannot correspond to HFL of pure pricing equilibria. Therefore, the two cases must be treated separately also in the HFL.

trade because waiting hurts more. If learning is costly, there are two effects whenever players become more impatient and equilibrium features inefficient delay. First, delay until agreement hurts more and so leads to higher inefficiency. This is the same effect as in the costless case. Second, a more impatient Buyer discounts learning costs more, which *ceteris paribus* should lower inefficiency. Since the costs of learning incurred in any equilibrium are less than the ex-ante surplus $\hat{v}$, the first effect dominates, so that inefficiency falls with the patience level even when learning is costly.[34]

Fifth, the inefficiency in the HFL is always decreasing in $\varkappa > 0$. This is intuitive: a larger and positive $\varkappa$ means that Buyer of type $\hat{v}$ waits longer in real time to realize her strategic option, even though she may not be worse off by accepting the current price offer of Seller. In the case of mixed pricing it holds $\varkappa = 0$, because Seller quotes the reservation price of Buyer with good news as long as there is no disclosure.

Overall, just as in the costless case in the HFL Seller pricing, and Buyer and Seller payoffs are non-extreme.

1.4 Extensions

1.4.1 Costly choice of intensity

I now introduce costly choice of intensity and thus endogenize the intensity choice of Buyer. Exploration/search and exploitation are very different aspects of endogenous learning. Typically, an agent decides first how actively she is going to search for new information sources and only after, how extensively she is going to exploit the ones she has found. This distinction bears out in the results of this section because these two aspects of learning behave differently in the HFL of strongly stationary equilibria. Here I focus on the case of costly choice of

[34]To see that costs of learning on path are necessarily lower than $\hat{v}$, note that they are lower than the option value from learning, which in turn is always less than the ex-ante surplus $\hat{v}$. This is straightforward to see in (1.15) for the case of stochastic fixed costs and follows from proof arguments of Proposition 6 in the case of deterministic variable costs.

intensity. In section A.4 of the appendix I also consider the case of costless choice of intensity.[35] Formally, costs on intensity are modeled as follows.

Costs on intensity. At the beginning of every period $t \geq 1$ Buyer picks the probability μ that a learning opportunity arrives within that period at the cost $\mathcal{C}(\Delta, \mu)$. The following is satisfied for $\mathcal{C}$:

- the function $\mathcal{C} : (0, \infty) \times [0, 1) \to \mathbb{R}_+$ is differentiable
- it satisfies

$$\lim_{\Delta \to 0} \frac{\mathcal{C}(\Delta, \lambda\Delta)}{\Delta} = f(\lambda), \qquad \lim_{\Delta \to 0} \frac{\partial}{\partial \mu} \mathcal{C}(\Delta, \Delta\lambda) = f'(\lambda) \tag{1.18}$$

  uniformly on $\lambda > 0$, with $f : [0, \infty) \to \mathbb{R}_+$ differentiable, strictly increasing and convex with

  – $f(0) = f'(0) = 0$ and
  – $\lim_{\lambda \to \infty} f'(\lambda) = +\infty$.

The second requirement states that $\mathcal{C}$ scales appropriately with time: as period-length goes to zero, both absolute and marginal costs of picking intensity go to zero on the same scale as the period-length. f is therefore the cost of intensity in real time. An example that satisfies the conditions is $\mathcal{C}(\Delta, \mu) = \Delta \cdot f(\mu)$ for f satisfying the properties above.

All general insights from section 1.3 apply to the set up with costs on intensity. In particular, all equilibria feature positive Buyer payoff for every $\Delta > 0$ and there exist strongly stationary equilibria near the HFL. In the HFL of strongly stationary equilibria with mixed pricing the price spread does not disappear, but the low price quoted upon non-disclosure is quoted only with vanishing

[35] The results on costless choice of intensity have some implications for the classical result of generic uniqueness of sequential equilibria in the traditional seller-offer, weak gap game from Fudenberg et al. [1985] and Gul et al. [1986]. Namely, the generic uniqueness result for Coasian dynamics established in Fudenberg et al. [1985] and Gul et al. [1986] depends crucially on the assumption that Buyer *can commit* to having private initial information, before she approaches Seller.

probability. Moreover, despite costly learning, there are near the HFL strongly stationary equilibria (necessarily with pure pricing) that are arbitrarily close to efficiency. The intuitions are the same as in the case of exogenous intensity.

The condition (P) for the existence of strongly stationary equilibria with mixed pricing in the statement of Proposition 6 can now be transferred to assumptions that involve the real-time cost of intensity f. The new existence conditions look as follows.

$$(P') \quad \text{if } \pi_0 \leq \frac{1}{2}, \text{ then } f'(r) > \frac{1}{2}\pi_0 \bar{v}, \quad \text{if } \pi_0 > \frac{1}{2}, \text{ then } f'\left(\frac{1}{\sqrt{2}}r\right) > \frac{\sqrt{2}}{\sqrt{2}+1}\pi_0 \bar{v}. \tag{1.19}$$

(P') assumes that acquiring intensity is not too cheap.[36] Just as in the case of exogenous intensity, the average price upon non-disclosure $\hat{p}$ is a sufficient statistic for the construction of the equilibria. Analogously, in the case of pure pricing, equilibria are are unique up to the price quoted by seller upon non-disclosure and the acceptance probability of Buyer of type $\hat{v}$.

At the beginning of every period in a strongly stationary equilibrium, Buyer picks the intensity λ. She trades off the physical costs with the benefit from learning. Because the benefit of learning is stationary and not influenced by the acceptance probability q of type $\hat{v}$, so is the intensity choice in any strongly stationary equilibrium. This implies that the average price upon non-disclosure is a sufficient statistic for the information acquisition choice of Buyer in strongly stationary equilibria.

The first-order condition for the choice of intensity in the HFL of deterministic variable costs is given by

$$f'(\lambda)\frac{r+\lambda}{r} = BL(\hat{p}), \tag{1.20}$$

where $BL(\hat{p})$ is the stationary benefit from learning defined in (1.13). For the

[36]The proof of existence is contained in section (A.3) of the appendix. Here I focus only on the analysis in the HFL.

case of stochastic fixed costs the respective intensity choice is characterized by the first-order condition

$$f'(\lambda)\frac{r+\lambda\bar{\mu}(\hat{p})}{r\bar{\mu}(\hat{p})} = BL(\hat{p}), \tag{1.21}$$

where $BL(\hat{p})$ is the stationary benefit from learning defined in (1.15) and $\bar{\mu}(\hat{p}) = F(\pi_0(\bar{v}-\hat{p}) - (\hat{v} - \bar{p}_L(\hat{p})))$ is the HFL of the stationary probability that the state is verified by Buyer, conditional on the opportunity to learn having arrived. Here, $\bar{p}_L(\hat{p}) = \lim_{\Delta\to 0} p_L(\hat{p},\Delta)$ is the HFL of the reservation price of the type $\hat{v}$.

In both cases of costs on accuracy the endogenous intensity $\lambda(\hat{p})$ is strictly decreasing in $\hat{p}$. In the HFL of the case of deterministic variable costs, the marginal costs of picking the intensity $f'(\lambda)\frac{r+\lambda}{\lambda}$ do not depend on the pricing of Seller, whereas in the model with stochastic fixed costs the marginal costs of picking the intensity $f'(\lambda)\frac{r+\lambda\bar{\mu}(\hat{p})}{r\bar{\mu}(\hat{p})}$ depend on the average price $\hat{p}$. This is because the rate of arrival of opportunities to learn and the rate of learning coincide in the model with deterministic variable costs, but they diverge in the model with stochastic fixed costs. Recall that in the latter case, the probability of learning conditional on having the opportunity to learn is an endogenous object.

The following Proposition summarizes some of the interesting results from the analysis in the HFL. I focus only on delay and payoff properties in the HFL for brevity's sake.

Proposition 7. *Let $\lambda(\hat{p})$ satisfy (1.20) in the case of deterministic variable costs and (1.21) in the case of stochastic fixed costs.*

1) [Delay in the HFL] In any HFL of strongly stationary equilibria expected delay is given by

$$\begin{cases} \frac{1}{\lambda(\hat{p})+\varkappa}, \text{ if } \varkappa \in [0,\infty), \hat{p}\in\mathcal{N} \text{ and accuracy costs are deterministic and variable}, \\ \frac{1}{\lambda(\hat{p})\bar{\mu}(\hat{p})+\varkappa}, \text{ if } \varkappa \in [0,\infty), \hat{p}\in\mathcal{N} \text{ and accuracy costs are stochastic and fixed}, \\ 0, \text{ if } \varkappa = \infty. \end{cases}$$

2) [Payoff and efficiency properties in the HFL] Buyer and Seller payoffs lie in $(\underline{v}, \hat{v})$ *for all* $\varkappa \in [0, \infty]$.

The efficiency loss in the HFL of equilibrium sequences with $\varkappa \in [0, \infty)$ *is given by*

$$\frac{r}{r+\lambda(\hat{p})+\varkappa}\hat{v} + \frac{\lambda(\hat{p})}{r+\lambda(\hat{p})+\varkappa}c(I(a(\hat{p}))) + \frac{f(\lambda(\hat{p}))}{r+\lambda(\hat{p})+\varkappa}, \tag{1.22}$$

in the case of deterministic variable costs of accuracy. It is given by

$$\frac{r}{r+\lambda(\hat{p})\bar{\mu}(\hat{p})+\varkappa}\hat{v} + \frac{\lambda(\hat{p})\bar{\mu}(\hat{p})}{r+\lambda(\hat{p})\bar{\mu}(\hat{p})+\varkappa}\mathbb{E}_F[c|c \leq \pi_0(\bar{v}-\hat{p})-(\hat{v}-\bar{p}_L(\hat{p}))] + \frac{f(\lambda(\hat{p}))}{r+\lambda(\hat{p})+\varkappa}, \tag{1.23}$$

in the case of stochastic fixed costs of accuracy.

There is no efficiency loss in the HFL of sequences with $\varkappa = \infty$.

There are two main differences between (1.22) and (1.16), as well as (1.23) and (1.17). First, for fixed $\varkappa < \infty$ the share of the ex-ante surplus that is lost due to the delay until agreement is endogenous in the case of deterministic variable accuracy costs. It also depends on Seller-pricing. This is because the rate of learning is now a strategic variable for Buyer in the case of deterministic variable costs. Second, there is a third term in the sum of inefficiencies that quantifies the additional welfare loss due to the fact that intensity is costly for Buyer. In the HFL, the decision on accuracy a in the case of deterministic variable costs and the decision whether to verify the state in the case of stochastic fixed costs do not depend directly on the arrival rate λ. Therefore, for fixed deterministic variable cost c or distribution of stochastic lump-sum cost F, as costs of intensity f converge to zero uniformly the inefficiency due to delay and due to costly choice of intensity disappears, because $\lambda(\hat{p})$ becomes arbitrarily large. But the inefficiency due to costly accuracy persists.

Comparative statics for information acquisition

In a learning model in which intensity and accuracy are endogenous it is natural to ask the question of how these two distinct aspects of learning behave across ex-ante different environments. The Proposition in this subsection delivers an answer to this question for the quantities a, λ in the HFL. Despite the equilibrium multiplicity the comparative static comparisons are clean, if one compares stationary equilibria with the same average price $\hat{p}$ upon non-disclosure. Comparing equilibria with the same average price upon non-disclosure is natural. $\hat{p}$ is potentially empirically observable and at the same time it is a sufficient statistic for the information acquisition choice of Buyer in strongly stationary equilibria.[37]

The result of this subsection is as follows.[38]

Proposition 8. *[Comparative statics in the HFL.]*

1) Suppose there are two strongly stationary equilibria in the HFL with the same average price $\hat{p}$ and all parameters the same, except for the discount rates $r_1 > r_2$. Then the equilibrium intensity is higher for r_1 than r_2.

2) Suppose there are two strongly stationary equilibria with the same average price $\hat{p}$ and all parameters the same, except for prior of high value $\pi_0^1 > \pi_0^2$. Then the equilibrium intensity is higher for π_0^1.

3) (deterministic variable accuracy costs). *Equilibrium accuracy is independent of the discount rate. Suppose there are two strongly stationary equilibria with the same average price $\hat{p}$ and all parameters the same except for $\pi_0^1 > \pi_0^2$. Equilibrium accuracy is higher for π_0^1 if $\frac{\bar{v}+\underline{v}}{2} > \hat{p}$, and it is higher for π_0^2 if $\frac{\bar{v}+\underline{v}}{2} < \hat{p}$.*

4) (stochastic fixed accuracy costs). *Suppose in the case of stochastic fixed costs there are two strongly stationary equilibria with the same average price $\hat{p}$ and all*

[37] For small changes of parameters like ex-ante prior of high value π_0 and discount rate r the existence neighborhoods $\mathcal{N}$ in Proposition 7 for mixed pricing equilibria overlap and so comparative statics for stationary equilibria with the same $\hat{p}$ are possible even if they are of the mixed pricing sort, at least for marginal changes in π_0, r or of the cost distribution F.

[38] More complete results can be found in subsection A.3.2 of the appendix. I focus here on a, λ for brevity's sake.

parameters the same, except for distribution of lump-sum costs $F_1 >_{FOSD} F_2$.[39] *Then λ_1 is lower than λ_2.*

The comparative statics of λ are all intuitive: the more impatient Buyer and/or the more optimistic at the outset of the bargaining she is, the higher the incentives to explore for information sources. The higher the accuracy costs, the lower the incentives to explore in the first place.

Equilibrium accuracy in the case of deterministic variable costs is independent of the impatience level. Whenever Buyer has the chance to exploit an opportunity to learn, the option value from learning is independent of the discount rate and so are the incurred accuracy costs. This is because, as long as there is no agreement, Seller uses the same stationary pricing strategy upon non-disclosure. To get an intuition for the rest of the result about accuracy, suppose $\frac{\bar{v}+\underline{v}}{2} > \hat{p}$ and $\hat{p}$ is close to $\hat{v}$. Then Buyer is relatively pessimistic ex-ante about the value of the good. As long as she becomes more optimistic $(\frac{1}{2} > \pi_0^1 > \pi_0^2)$, she also becomes more uncertain about the value of the good and as a consequence has stronger incentives to exploit the learning opportunity. If instead $\frac{\bar{v}+\underline{v}}{2} < \hat{p}$ and $\hat{p}$ is close to $\hat{v}$, then Buyer is relatively optimistic ex-ante about the value of the good. The marginal benefit of additional learning when she becomes even more optimistic $(\pi_0^1 > \pi_0^2 > \frac{1}{2})$ is low in this case, so that the equilibrium accuracy $a(\hat{p})$ is lower for the higher prior π_0^1.

Other comparative results can be proven, and their proofs are skipped here for the sake of length. E.g. for the case of deterministic variable costs, one can easily show that, whenever accuracy costs fall pointwise $(c_1(a) \leq c_2(a)$ for all $a \in [\frac{1}{2}, 1))$ equilibrium intensity $\lambda(\hat{p})$ weakly increases. A slightly more convoluted argument is needed to show that the probability of learning conditional on an opportunity to learn $\bar{\mu}(\hat{p}) = F(\pi_0(\bar{v} - \hat{p}) - (\hat{v} - \bar{p}_L(\hat{p})))$ falls, if F increases in the FOSD-sense.[40]

[39] This means that F_1 dominates F_2 in the sense of first-order stochastic dominance. Formally, $F_1(x) \leq F_2(x)$ for all $x > 0$ with the inequality strict for some $x > 0$.

[40] There is a direct effect, as can be seen from the definition of $\bar{\mu}(\hat{p})$, but also an indirect one counteracting it, because the HFL of $p_L(\hat{p})$ given by $\bar{p}_L(\hat{p})$, increases when F increases in the FOSD-sense. The direct effect is stronger.

1.4.2 Pre-learning negotiations

In the game introduced in section 1.2 learning is a socially wasteful activity for every $\Delta > 0$, because it does not raise the ex-ante surplus, nor does it eliminate any pre-game informational asymmetry. Therefore, a third party who has social welfare in mind would prohibit learning altogether, if that were feasible. Another way to avoid the inefficiency from learning is to give Buyer the opportunity to commit to pre-learning negotiations. In fact, such commitment would lead to efficient outcomes.

To see this formally, consider an extensive-form expansion of the basic game from section 1.2 in which there is a first stage at a period $t = 0$ in which Seller can make an offer to Buyer, before learning can start. If the offer at $t = 0$ is accepted the game ends with agreement, whereas if the offer is rejected by Buyer the bargaining game from section 1.2 is played.

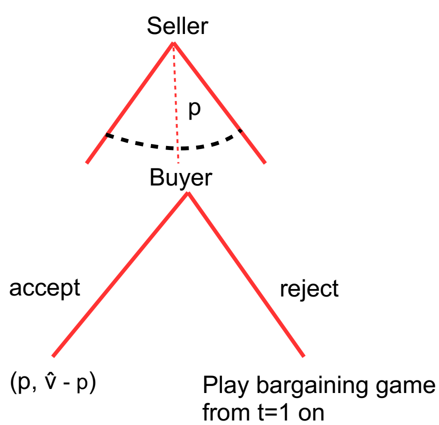

Figure 1.4.1: Game with pre-learning negotiations.

The following Proposition is straightforward.

Proposition 9. *Pick any $\Delta > 0$. All perfect Bayesian equilibria of the game with pre-learning negotiations are efficient. In particular, all perfect Bayesian equilibria feature agreement at time $t = 0$.*

Note that Buyer's payoff in the game with pre-learning negotiations quantifies the value of learning from an ex-ante perspective. Suppose that the PBE has a price p accepted with probability one in $t = 0$. Then it holds $\hat{v} - p = \delta V_B(p)$ with $V_B(p)$ the payoff of Buyer in the continuation bargaining game, if she rejects p. Then the opportunity to learn is worth precisely $\hat{v} - p$. Proposition 3 implies that the value of learning can be zero whenever learning is costless but not immediate, so that Buyer may not get an informational rent from the opportunity to learn. In contrast, Theorem 2 shows that Buyer always receives an informational rent whenever learning is costly.

Proposition 9 gives a rationale for the intervention of a third party with social welfare in mind or for the introduction of commitment devices that allow Buyer to commit to pre-learning negotiations. This intervention is the more valuable, the larger the period-length $\Delta > 0$ is, or equivalently the lower the frequency of interaction between the bargaining parties.

The crucial assumption that allows efficiency to be restored for any $\Delta > 0$ is that trade is ex-ante efficient. In other situations, e.g. Buyer may also experience disutility from acquiring the good, learning is not necessarily wasteful and delay might be efficiency-enhancing. This is illustrated in the next subsection.

1.4.3 Negative lowest Buyer value

Consider a variation of the game introduced in section 1.2 with two differences. First, assume that $\underline{v} < 0 < \hat{v} = \pi_0 \bar{v} + (1 - \pi_0)\underline{v}$. Thus, trade is not efficient in every state of the world, even though it is efficient to trade immediately, if learning is impossible. Second, assume that Buyer can walk away from negotiations whenever she chooses. Many negotiation settings share these features. E.g. a significant share of VC, or more generally, private equity negotiations break down after the investors learn new information and conclude that the transaction is not worthwhile.

The analysis of behavior for this variation of the model follows closely the results of the main model.[41] This holds throughout except for payoff and efficiency considerations.

For brevity's sake, here I illustrate only that learning can be welfare-enhancing, even in strongly stationary equilibria with $\varkappa = 0$ (recall that these are the strongly stationary equilibria with the largest expected delay in the baseline model). I look at the case of costly learning to give inefficiency its 'best shot'. For simplicity of exposition I assume intensity $\lambda > 0$ is exogenously given and focus on the mixed pricing equilibria.

Proposition 10. *Suppose that $\underline{v} < 0 < \hat{v}$.*

1) If learning is impossible the efficient outcome with respect to the available information is to trade with probability one.

2) [Deterministic variable costs of accuracy] Consider the HFL of strongly stationary equilibria with mixed pricing and $\hat{p} = \hat{v}$ and the following parameter restrictions:

$\pi_0 > \frac{1}{2}$, $\sqrt{2}\lambda + 2\varepsilon > r > \sqrt{2}\lambda + \varepsilon$ *with some* $\varepsilon > 0$ *and*

$$\frac{\sqrt{2}}{\sqrt{2}+1}\frac{\pi_0}{1-\pi_0} < -\frac{\underline{v}}{\hat{v}} < \frac{\pi_0}{1-\pi_0}.$$

Then learning is welfare enhancing whenever the accuracy costs c are near enough to zero with respect to the topology of uniform convergence on compact sets and ε is small enough.

3) [Stochastic variable costs of accuracy] Consider the HFL of strongly stationary equilibria with mixed pricing and $\hat{p} = \hat{v}$ and the following parameter restrictions:

$$\frac{r}{r+\lambda}\frac{\pi_0}{1-\pi_0} < -\frac{\underline{v}}{\hat{v}} < \frac{\pi_0}{1-\pi_0}.$$

Then learning is welfare enhancing whenever the distribution of fixed accuracy costs F is near enough to zero with respect to the topology of weak convergence of probability distributions.

[41] The changes are minor. Details are available upon request.

Learning is not always welfare enhancing when $\underline{v} < 0$. Intuitively, this holds true in case the costs F or c are 'high enough' with respect to the other parameters of the game.

1.5 Concluding Discussion

In most real-life bargaining situations a bargaining party has the chance to learn privately about the terms of trade during the negotiation process. She will typically take up this opportunity, even if it is costly and leads to delay, because additional information may give her a strategic advantage in negotiations. This paper considers such a situation explicitly and shows that this may lead to delay, even in situations in which ex-ante trade is efficient.

Moreover, this paper models the learning process explicitly in a bargaining environment. Endogenous information acquisition and disclosure of new information with the aim of influencing bargaining positions is a realistic feature of many real-world market interactions, from merger and acquisition negotiations to government leases of natural resources. This paper is a first attempt to introduce this realistic feature into the dynamic bargaining theory literature.

Many extensions and variations to this work are natural topics for future research. First, many real-world bargaining situations feature initial information asymmetries *in addition to* the possibility of endogenous learning as negotiations progress. Examples include mergers and acquisitions involving industry leaders or management buyouts. It is natural to expect that sequential screening dynamics reappear in situations in which, besides the possibility of endogenous learning and selective disclosure of information, initial private information is a prominent feature.

Second, in my model learning is wasteful because, besides being costly, it creates information asymmetry between the two bargaining parties. In situations in which there is initial private information and the values of Buyer and Seller are interdependent, learning and the possibility to communicate learning outcomes may lead to improvements in efficiency.[42] A careful study of the effects of endogenous learning in settings with interdependent values is left for future research.

Third, a serious study of the interaction of competitive pressure and endogenous learning is missing from this paper, even though competitive pressure is a significant factor for negotiations in many business situations, e.g. in private equity or procurement of investment goods.[43] A better understanding of the interaction between initial information asymmetries, competitive pressure and endogenous information acquisition in dynamic bargaining environments is an exciting topic left for future research.

Finally, my model implies that one-sided endogenous private learning in bargaining situations is compatible with a wide variety of welfare outcomes, from approximate efficiency to significant inefficiency. This suggests studying the role of limited commitment and other institutional forms with the aim of *designing* bargaining outcomes. One hopes that deepening our understanding of endogenous costly learning in bargaining situations will lead to valuable insights on how to better devise institutions that facilitate negotiations in the real world.

[42]These are typically situations in which the lemons problem is prevalent (Akerlof [1970]).

[43]One *ad hoc* way to incorporate ideas of competition into the model of this paper is to assume that the discount factor of the players reflects, besides impatience, the possibility that negotiations break down due to exogenous reasons, one of them being the arrival of a superior bargaining partner who 'steals away' the chance for a deal between the original players. A more realistic model could have discount factors of the players depend on their respective estimate of the current Buyer valuation.

2

Dynamic Information Design with Diminishing Sensitivity Over News

This chapter is coauthored with Kevin He. We thank Drew Fudenberg, Jerry Green, Jonathan Libgober, Erik Madsen, Pietro Ortoleva, Matthew Rabin, Collin Raymond, the MIT information design reading group, and our seminar participants for insightful comments. We also benefited from conversations with Krishna Dasaratha, Ben Enke, Simone Galperti, David Hagmann, Marina Halac, Johannes Hörner, David Laibson, Shengwu Li, Elliot Lipnowski, Gautam Rao, and Tomasz Strzalecki at an early stage of the project. Any errors are ours.

2.1 Introduction

When people give others news, they are often mindful of the information's psychological impact. For example, this consideration affects the way CEOs announce earnings forecasts to shareholders and organization leaders update their teams about recent developments. While the instrumental value of information also plays a significant role, we analyze the under-studied problem of how the audience's psychological reaction to good and bad news shapes the dynamic communication of information. This problem is even more relevant in situations like designing game shows and other entertainment content, where the audience experiences positive and negative reactions over time to news and developments that have no bearing on their personal decision-making.

We consider an informed, benevolent *sender* communicating non-instrumental information to a *receiver* who experiences gain-loss utility over changes in beliefs ("news utility"). The state of the world, privately known to the sender, determines the receiver's consumption at some future date. The sender communicates this state over multiple periods as to maximize the receiver's expected welfare, knowing that the receiver derives utility based on the nature and the magnitude of news each period — good news elates and bad news disappoints. The receiver will exogenously learn the true state just before future consumption.

We focus on how the receiver's *diminishing sensitivity* over news affects the optimal design of information structures. Kahneman and Tversky [1979]'s original formulation of prospect theory envisioned a gain-loss utility component based on deviations from a reference point, where larger deviations carry smaller marginal effects. This idea of diminishing sensitivity is referenced in virtually all subsequent work on reference-dependent preferences, including Kőszegi and Rabin [2009], who first introduced a model of news utility. In almost all cases, however, researchers then specialize for simplicity to a two-part linear gain-loss utility function that allows for loss aversion but not diminishing sensitivity. Four decades since Kahneman and Tversky [1979]'s publication, O'Donoghue and Sprenger [2018]'s review of the ensuing literature summarizes the situation as follows:

> "Most applications of reference-dependent preferences focus entirely on loss aversion, and ignore the possibility of diminishing sensitivity [...] The literature still needs to develop a better sense of when diminishing sensitivity is important."

We argue that diminishing sensitivity over the magnitude of news generates novel predictions for information design. As Kőszegi and Rabin [2009] point out, the two-part linear news-utility model makes the stark prediction that people prefer resolving all uncertainty in one period ("one-shot resolution") over any other dynamic information structure. We show that diminishing sensitivity over news complicates the sender's problem and leads to a more nuanced optimal information structure. In particular, one-shot resolution is strictly suboptimal for a class of news-utility functions exhibiting diminishing sensitivity. This class includes the commonly used power-function specification. It also includes a tractable quadratic specification, whenever diminishing sensitivity is sufficiently strong relative to the degree of loss aversion. We further identify conditions that imply the optimal information structure treats good news and bad news asymmetrically, disclosing good news gradually but bad news all at once. The direction of this optimal skewness is a central implication of diminishing sensitivity: the "opposite" kind of information structure that divulges all good news at once but doles out bad news in small portions is *never* optimal. In fact, this kind of information structure is even worse than one-shot resolution.

In our model, the receiver knows the sender's strategy and formulates Bayesian beliefs. This framework leads to *cross-state* constraints on the sender's problem. In view of diminishing sensitivity, one might conjecture that the sender should concentrate all bad news in period 1 if the state is bad, and deliver equally-sized pieces of good news in periods 1, 2, 3, ... if the state is good. But these belief paths are infeasible, since a Bayesian audience who knows this strategy and does not receive bad news in period 1 will conclusively infer that the state is good. The receiver should not judge subsequent communication from the sender as further good news or derive positive news utility from them.[1] We show that the sender can nevertheless implement a "gradual good news, one-shot bad news" information structure for a Bayesian receiver by sending a conclusive bad-news signal in a *random* period when the state is bad. In the optimal information structure, conditional on the good state, the receiver may get different amounts of good news in different periods, even though his news-utility function is time-invariant and the sender knows the state from the start.

Another implication of diminishing sensitivity is that people with opposite consumption rankings over states may exhibit opposite informational preferences. In a world with two possible states, A and B, suppose state A realizes if and only if a series of intermediate events all occur successfully. We show that agents who prefer the consumption they get in state A will choose to observe the intermediate events resolve in real-time (gradual information), while agents who prefer the consumption they get in state B will choose to only learn the final state (one-shot information). This prediction distinguishes the news-utility model with diminishing sensitivity from other models of non-instrumental information preference. The result also rationalizes a "sudden death" format often found in game shows, where the contestant must overcome every challenge in a sequence to win the grand prize (as opposed to the grand prize being contingent on beating at least one of several challenges.)

When the sender lacks commitment power, information structures featuring gradual good news encounter a credibility problem. In the bad state, the sender may strictly prefer to lie and convey a positive message intended for the good state. This temptation exists despite the fact that the sender is far-sighted and maximizes the receiver's total news utility over time. The intuition is that the receiver will inevitably feel disappointed upon learning the truth in the future, so his marginal utility of (unwarranted) good news today is larger than his marginal

[1] In the language of information design, these conjectured belief paths violate *Bayesian plausibility*, as they cannot arise from the Bayesian updating of a given prior.

disutility of *heightened* disappointment in the future, thanks to diminishing sensitivity. This perverse incentive to provide false hope in the bad state may preclude all meaningful communication in all states. We show that if the receiver has diminishing sensitivity but not loss aversion (or has low loss aversion), then every equilibrium is payoff-equivalent to the babbling equilibrium. High enough loss aversion, however, can restore the equilibrium credibility of good-news messages by increasing the future disappointment costs associated with inducing false hope today. As a consequence, receivers with higher loss aversion may enjoy higher equilibrium payoffs, which would never happen if the sender had commitment power.

Finally, we characterize the entire family of equilibria featuring gradual good news and study how quickly the receiver learns the state. For a class of news-utility functions that include the square-root and quadratic specifications mentioned before, the sender conveys progressively larger pieces of good news over time, so the receiver's equilibrium belief grows at an increasing rate in the good state. This puts a uniform bound on the number of periods of informative communication across all time horizons and all equilibria.

The rest of the paper is organized as follows. The remainder of Section 2.1 reviews related literature. Section 2.2 defines the sender's problem under the commitment assumption and introduces our model of news utility. Section 2.3 studies the optimal information structure and the relationship between consumption preferences and informational preferences. Section 2.4 focuses on the cheap-talk model when the sender lacks commitment power. Section 2.5 looks at a variant of the model without a deterministic horizon. Section 2.6 discusses other models of preference over non-instrumental information. Section 2.7 concludes.

2.1.1 Related Literature

Since Kőszegi and Rabin [2009], several other authors have analyzed the implications of news utility in such varied settings as asset pricing [Pagel, 2016], life-cycle consumption [Pagel, 2017], portfolio choice [Pagel, 2018], and mechanism design [Duraj, 2018b]. These papers focus on Bayesian agents with two-part linear gain-loss utilities and do not study the role of diminishing sensitivity to news.

Interpreting monetary gains and losses as news about future consumption, experiments that show risk-seeking behavior when choosing between loss lotteries and risk-averse behavior when choosing between gain lotteries provide evidence for diminishing sensitivity over consumption news (see, for example, Rabin and Weizsäcker [2009]). In the same vein, papers in the finance literature that use diminishing sensitivity over monetary gains and losses to explain the disposition effect [Barberis and Xiong, 2012, Henderson, 2012, Kyle, Ou-Yang, and Xiong, 2006, Shefrin and Statman, 1985] also provide indirect evidence for diminishing sensitivity over consumption news.

We are not aware of other work that focuses on how diminishing sensitivity matters for information design with news utility. In fact, except for the work on disposition effect in finance, very few papers deal with diminishing sensitivity in *any* kind of reference-dependent preference. One exception is Bowman, Minehart, and Rabin [1999], who study a consumption-based reference-dependent model with diminishing sensitivity. A critical difference is that their reference points are based on past habits, not rational expectations. In their environment, a consumer who knows their future income optimally concentrates all consumption losses in the first period if income will be low, but spreads out consumption gains across multiple periods if income will be high. As discussed before, the analog of this strategy cannot be implemented in our setting since the receiver derives news utility from changes in rational Bayesian beliefs.

Our model of diminishing sensitivity over the magnitude of news shares the same psychological motivation as Kahneman and Tversky [1979], who base their theory of human responses to monetary gains and losses on human responses to changes in physical attributes like temperature or brightness:

> "Many sensory and perceptual dimensions share the property that the psychological response is a concave function of the

> magnitude of physical change. For example, it is easier to discriminate between a change of 3° and a change of 6° in room temperature, than it is to discriminate between a change of 13° and a change of 16°."

We are not aware of any empirical work designed to measure diminishing sensitivity over news, but will highlight some testable predictions of the model later on.

While some of our results apply to Kőszegi and Rabin [2009]'s model of news utility or to a more general class of such models (e.g., Proposition 11, Proposition 12, Proposition 13, Corollary 11), we mostly focus on the simplest model of news utility where the agent derives gain-loss utility from changes in *expected* future consumption utility. This mean-based model lets us concentrate on the implications of diminishing sensitivity, but differs from Kőszegi and Rabin [2009]'s model where agents make a *percentile-by-percentile* comparison between old and new beliefs. Fully characterizing the optimal information structure using this percentile-based model is out of reach for us, but our numerical simulations in Appendix B.2.2 suggest the answers would be very similar.

Parallel to the recent literature on the applications of news utility discussed above, Dillenberger and Raymond [2018] axiomatize a general class of additive belief-based preferences in the domain of two-stage lotteries by suitably weakening the independence axiom of expected utility. In the case of $T = 2$, our news-utility model belongs to the class they characterize. Under this specialization, our work may be thought of as studying the information design problem, with and without commitment, using some of Dillenberger and Raymond [2018]'s additive belief-based preferences. Dillenberger and Raymond [2018] also provide high-level conditions for additive belief-based preferences to exhibit preference for one-shot resolution. We are able to find more interpretable and easy-to-verify conditions for the sub-optimality of one-shot resolution, working with a specific sub-class of their preferences.

In general, papers on belief-based utility have highlighted two sources of felicity: *levels* of belief about future consumption utility ("anticipatory utility," e.g., Eliaz and Spiegler [2006], Kőszegi [2006], Schweizer and Szech [2018]) and *changes* in belief about future consumption utility ("news utility"). News utility is a function of both the prior belief and the posterior belief, while a given posterior belief brings the same anticipatory utility for all priors [Eliaz and Spiegler, 2006]. As we discuss in Section 2.6, the rich dynamics of the optimal information structure are a unique feature of the news-utility model (with diminishing sensitivity).

Brunnermeier and Parker [2005] and Macera [2014] study the optimal design of beliefs for agents with belief-based utilities that differ from the news-utility setup we consider. Another important distinction is that we focus on the design of *information*: changes in the receiver's belief derive from Bayesian updating an exogenous prior, using the information conveyed by the sender. Macera [2014] considers a non-Bayesian agent who freely chooses a path of beliefs, while knowing the actual state of the world. Brunnermeier and Parker [2005] study the "opposite" problem to ours, where the agent freely chooses a prior belief (over the sequence of state realizations) at the start of the game, then updates belief about future states through an exogenously given information structure.

Our emphasis on information is shared by Ely, Frankel, and Kamenica [2015], who study dynamic information design with a Bayesian receiver who derives utility from suspense or surprise. In contrast to these authors who propose and study an original utility function over belief paths where larger belief movements always bring greater felicity, we consider a gain-loss utility function over changes in beliefs. Because our states are associated with different consumption consequences, changes in beliefs may increase or decrease the receiver's utility depending on whether the news is good or bad. While one-shot resolution is suboptimal in both Ely, Frankel, and Kamenica [2015]'s problem and our problem (under some conditions), the optimal information structure differs. The optimal information structure in our problem is asymmetric, a key implication of diminishing sensitivity. Another difference is that information structures featuring gradual bad news, one-shot good news are worse than one-shot resolution in our problem, while one-shot resolution is the worst possible information structure in Ely, Frankel, and Kamenica [2015]'s problem.

Also within the dynamic information design literature but without behavioral preferences, Li and Norman [2018] and Wu [2018] consider a group of senders moving sequentially to persuade a single receiver. The receiver takes an action after observing all signals. This action, together with the true state of the world, determines the payoffs of every player. While these authors study a dynamic environment, only the final belief of the receiver at the end of the last period matters for the players' payoffs. Indeed, every equilibrium in their setting can be converted into a payoff-equivalent "one-step" equilibrium where the first sender sends the joint signal implied by the old equilibrium, while all subsequent senders babble uninformatively. In our setting, the distribution of the receiver's final belief at the end of the last period is already pinned down by the prior belief at the start of the first period. Yet, different sequences of interim beliefs cause the receiver to experience different amounts of total news utility. The stochastic process of these interim beliefs constitutes the object of design. We provide a general procedure for computing the optimal dynamic information structure in this new setting.

Lipnowski and Mathevet [2018] study a static model of information design with a psychological receiver whose welfare depends directly on posterior belief. They discuss an application to a mean-based news-utility model without diminishing sensitivity in their Appendix A, finding that either one-shot resolution or no information is optimal. We focus on the implications of diminishing sensitivity and derive specific characterizations of the optimal information structure. Our work also differs in that we study a dynamic problem, examine equilibria without commitment, and discuss how the rate of releasing good news changes over time.

2.2 Model

2.2.1 Timing of Events

We consider a discrete-time model with periods $0, 1, 2, \ldots, T$, where $T \geq 2$. There are two players, the sender ("she") and the receiver ("he"). There is a finite state space Θ with $|\Theta| = K \geq 2$. In state θ, the receiver will consume c_θ in period T, deriving from it consumption utility $v(c_\theta)$ where v is strictly increasing. Assume that $c_{\theta'} \neq c_{\theta''}$ when $\theta' \neq \theta''$. We may normalize without loss $\min_{\theta \in \Theta}[v(c_\theta)] = 0$, $\max_{\theta \in \Theta}[v(c_\theta)] = 1$. There is no consumption in other periods and neither player can affect period T's consumption.

The players share a common prior belief $\pi_0 \in \Delta(\Theta)$ about the state, where $\pi_0(\theta) > 0$ for all $\theta \in \Theta$. In period 0, the sender commits to a finite message space M and a strategy $\sigma = (\sigma_t)_{t=1}^{T-1}$, where $\sigma_t(\cdot \mid h^{t-1}, \theta) \in \Delta(M)$ is a distribution over messages in period t that depends on the public history $h^{t-1} \in H^{t-1} := (M)^{t-1}$ of messages sent so far, as well as the true state θ. The sender can commit to any *information structure* (M, σ), which becomes common knowledge between the players. At the start of period 1, the sender privately observes the state's realization, then sends a message in each of the periods $1, 2, \ldots, T-1$ according to the strategy σ. (Section 2.4 studies a cheap talk model where the sender lacks commitment power.) Information about θ is *non-instrumental* in that it does not help the receiver make better decisions, but it can change his belief about future welfare.

At the end of period t for $1 \leq t \leq T-1$, the receiver forms the Bayesian posterior belief π_t about the state after the on-path history $h^t \in H^t$ of t messages. This belief is rational and calculated with the knowledge of the information structure (M, σ). In period T, the receiver exogenously and perfectly learns the true state θ, consumes c_θ, and the game ends. (Section 2.5 considers a random-horizon model where the termination date is random and unknown to both parties.)

Since the receiver is Bayesian, the sender faces cross-state constraints in choosing paths of beliefs. For example, if the sender wishes to use some message $m \in M$ to convey positive but inconclusive news in the first period when the state is good, then the same message must also be sent with positive probability when the state is bad – otherwise, receiving this information in the first period would amount to conclusive evidence of the good state. As we later show, these cross-state constraints imply distortions from perfect "consumption smoothing" of good news.

When $K = 2$, we label two states as **G**ood and **B**ad, $\Theta = \{G, B\}$, so that $v(c_G) = 1$, $v(c_B) = 0$. We also abuse the notation π_t to mean $\pi_t(G)$ in the case of binary states.

In this model, the sender has perfect information about the receiver's future consumption level once she observes the state. Appendix B.2 discusses an extension where the sender's information is imperfect, so that there is residual uncertainty about the receiver's consumption even conditional on the state (i.e., conditional on the sender's private information).

2.2.2 News Utility

The receiver derives utility based on changes in his belief about the final period's consumption. Specifically, he has a continuous *news-utility function* $N : \Delta(\Theta) \times \Delta(\Theta) \to \mathbb{R}$, mapping his pair of new and old beliefs about the state into a real-valued felicity.[2] He receives utility $N(\pi_t \mid \pi_{t-1})$ at the end of period $1 \leq t \leq T$. Utility flow is undiscounted and the receiver has the same N in all periods. The sender maximizes the total expected welfare of the receiver, which is the sum of the news utilities in different periods and the final consumption utility, $\sum_{t=1}^{T} N(\pi_t \mid \pi_{t-1}) + v(c)$. We assume for every $\pi \in \Delta(\Theta)$, both $N(\cdot \mid \pi)$ and $N(\pi \mid \cdot)$ are continuously differentiable except possibly at π.

For many of our results, we study a *mean-based* news-utility model. Kőszegi and Rabin [2009] mention this model, but mostly consider a decision-maker who makes a percentile-by-percentile comparison between his old and new

[2]Since different states lead to different levels of consumption, beliefs over states induce beliefs over consumption.

beliefs. We use the mean-based model to focus on the implications of diminishing sensitivity in the simplest setup. The agent applies a gain-loss utility function, $\mu : [-1, 1] \to \mathbb{R}$, to changes in expected consumption utility for period T. That is, $N(\pi_t \mid \pi_{t-1}) = \mu\left(\sum_{\theta \in \Theta}(\pi_t(\theta) - \pi_{t-1}(\theta)) \cdot v(c_\theta)\right)$. Throughout we assume μ is continuous, strictly increasing, twice differentiable except possibly at 0, and $\mu(0) = 0$. We impose further assumptions on μ to reflect diminishing sensitivity and loss aversion.

Definition 4. Say μ satisfies *diminishing sensitivity* if $\mu^{''}(x) < 0$ and $\mu^{''}(-x) > 0$ for all $x > 0$. Say μ satisfies *(weak) loss aversion* if $-\mu(-x) \geq \mu(x)$ for all $x > 0$. There is *strict loss aversion* if $-\mu(-x) > \mu(x)$ for all $x > 0$.

We now discuss two important functional forms of μ. In Appendix B.2.2, we compare the optimal information structures for this model and for Kőszegi and Rabin [2009]'s percentile-based model, a class of news-utility functions that do not admit a mean-based representation.

Quadratic News Utility

The quadratic news-utility function $\mu : [-1, 1] \to \mathbb{R}$ is given by

$$\mu(x) = \begin{cases} \alpha_p x - \beta_p x^2 & x \geq 0 \\ \alpha_n x + \beta_n x^2 & x < 0 \end{cases}$$

with $\alpha_p, \beta_p, \alpha_n, \beta_n > 0$. So we have

$$\mu^{'}(x) = \begin{cases} \alpha_p - 2\beta_p x & x \geq 0 \\ \alpha_n + 2\beta_n x & x < 0 \end{cases}, \qquad \mu^{''}(x) = \begin{cases} -2\beta_p & x \geq 0 \\ 2\beta_n & x < 0 \end{cases}.$$

The parameters α_p, α_n control the extent of loss aversion near 0, while β_p, β_n determine the amount of curvature — i.e., the second derivative of μ. We only consider quadratic news-utility functions that satisfy the following parametric restrictions.

A. *Monotonicity*: $\alpha_p \geq 2\beta_p$ and $\alpha_n \geq 2\beta_n$. Monotonicity condition holds if and only if $\mu'(x) \geq 0$ for all $x \in [-1, 1]$.

B. *Loss aversion*: $\alpha_n - \alpha_p \geq (\beta_n - \beta_p)z$ for all $z \in [0, 1]$. This condition is equivalent to loss aversion from Definition 4 for this class of news-utility functions.

A family of quadratic news-utility functions that satisfy these two restrictions can be constructed by choosing any $\alpha \geq 2\beta > 0$ and $\lambda \geq 1$, then set $\alpha_p = \alpha$, $\alpha_n = \lambda\alpha$, $\beta_p = \beta$, $\beta_n = \lambda\beta$. Figure 2.2.1 plots some of these news-utility functions for different values of α, β, and λ.

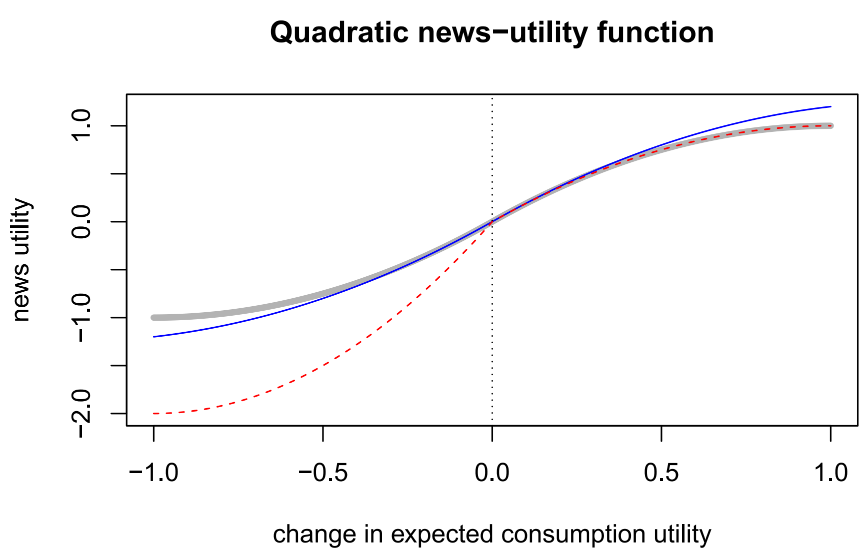

Figure 2.2.1: Examples of quadratic news-utility.
Functions in the family $\alpha_p = \alpha$, $\alpha_n = \lambda\alpha$, $\beta_p = \beta$, $\beta_n = \lambda\beta$. Grey curve: $\alpha = 2$, $\beta = 1$, $\lambda = 1$. Red curve: $\alpha = 2$, $\beta = 1$, $\lambda = 2$. Blue curve: $\alpha = 2$, $\beta = 0.8$, $\lambda = 1$.

Power-Function News Utility

The power-function news-utility $\mu : [-1, 1] \to \mathbb{R}$ is given by

$$\mu(x) = \begin{cases} x^{\alpha} & x \geq 0 \\ -\lambda|x|^{\beta} & x < 0 \end{cases}$$

with $0 < \alpha, \beta < 1$ and $\lambda \geq 1$. Parameters α, β determine the degree of diminishing sensitivity to good news and bad news, while λ controls the extent of loss aversion. This class of functions nests the square-root case when $\alpha = \beta = 0.5$ and is the only class of gain-loss functions to appear in Tversky and Kahneman [1992].

2.3 Optimal Information Structure

In this section, we characterize the optimal information structure that solves the sender's problem. We provide a general inductive procedure to maximize total expected news utility and find an information structure with K messages that achieves this maximum. We show that information structures featuring gradual bad news, one-shot good news are strictly worse than one-shot resolution, then identify sufficient conditions that imply the optimal information structure features gradual good news, one-shot bad news. We illustrate these conditions with the quadratic news-utility specification, finding that the conditions hold whenever diminishing sensitivity is sufficiently strong relative to loss aversion.

We conclude this section by highlighting that agents with opposite consumption preferences over two states of the world can exhibit opposite informational preferences when choosing between one-shot resolution and gradual resolution of uncertainty. This endogenous diversity of information preferences distinguishes news utility with diminishing sensitivity from other models of preference over non-instrumental information in the literature.

2.3.1 A General Backwards-Induction Procedure

For $f : \Delta(\Theta) \to \mathbb{R}$, let cav$f$ be the concavification of f — that is, the smallest concave function that dominates f pointwise. Concavification plays a key role in solving this information design problem, just as in Kamenica and Gentzkow [2011] and Aumann and Maschler [1995].

For $\pi_{T-2}, \pi_{T-1} \in \Delta(\Theta)$ two beliefs about the state, let $U_{T-1}(\pi_{T-1} \mid \pi_{T-2})$ be the sum of the receiver's expected news utilities in periods $T-1$ and T, if he enters period $T-1$ with belief π_{T-2} and updates it to π_{T-1}. More precisely,

$$U_{T-1}(\pi_{T-1} \mid \pi_{T-2}) := N(\pi_{T-1} \mid \pi_{T-2}) + \sum_{\theta \in \Theta} \pi_{T-1}(\theta) \cdot N(1_\theta \mid \pi_{T-1}),$$

where 1_θ is the degenerate belief putting probability 1 on the state θ. Note that by the martingale property of beliefs, if the receiver holds belief π_{T-1} at the end of period $T-1$, then state θ must then realize in period T with probability $\pi_{T-1}(\theta)$.

Let $U^*_{T-1}(\pi_{T-2}) := (\text{cav} U_{T-1}(\cdot \mid \pi_{T-2}))(\pi_{T-2})$. As we will show in the proof of Proposition 11, $U^*_{T-1}(\pi_{T-2})$ is the value function of the sender when the receiver enters period $T-1$ with belief π_{T-2}. It is calculated by evaluating the concavified version of $x \mapsto U_{T-1}(x \mid \pi_{T-2})$ at the point $x = \pi_{T-2}$. By Carathéodory's theorem, there exist weights $w^1, ..., w^K \geq 0$, beliefs $q^1, ..., q^K \in \Delta(\Theta)$, with $\sum_{k=1}^K w^k = 1$, $\sum_{k=1}^K w^k q^k = \pi_{T-2}$, such that $U^*_{T-1}(x) = \sum_{k=1}^K w^k U_{T-1}(q^k \mid x)$. When the receiver enters period $T-1$ with belief π_{T-2}, the sender maximizes his expected payoff using a signaling strategy σ_{T-1} that generates a distribution of posteriors supported on $(q^1, ..., q^K)$ with probabilities $(w^1, ..., w^K)$.

Continuing inductively, using the value function $U^*_{t+1}(x)$ for $t \geq 1$, we may define:

$$U_t(\pi_t \mid \pi_{t-1}) := N(\pi_t \mid \pi_{t-1}) + U^*_{t+1}(\pi_t),$$

which leads to the period t value function $U^*_t(x) := (\text{cav} U_t(\cdot \mid x))(x)$. The maximum expected news utility across all information structures is $U^*_1(\pi_0)$.

Proposition 11 formalizes this discussion. It shows there exists an information structure with K messages that achieves optimality, and the said information structure can be constructed using the sequence of concavifications.

Proposition 11. *The maximum expected news utility across all information* *structures is* $U_1^*(\pi_0)$. *There is an information structure* (M, σ) *with* $|M| = K$ *attaining this maximum, with the property that after each on-path public history* h^{t-1} *associated with belief* π_{t-1}*, the sender's strategy* $\sigma_t(\cdot \mid h^{t-1}, \theta)$ *induces posterior* q^k *at the end of period t with probability* w^k*, for some* $q^1, ..., q^K \in \Delta(\Theta)$, $w^1, ..., w^K \geq 0$, *satisfying* $\sum_{k=1}^K w^k = 1$, $\sum_{k=1}^K w^k q^k = \pi_{t-1}$, *and* $U_t^*(\pi_{t-1}) = \sum_{k=1}^K w^k U_t(q^k \mid \pi_{t-1})$.

A perhaps surprising implication is that the receiver only needs a binary message space if there are two states of the world, regardless of the shape or curvature of the news-utility function N. Figure 2.3.1 illustrates the concavification procedure in an environment with two equally likely states, $T = 5$, and the mean-based news-utility function $\mu(x) = \sqrt{x}$ for $x \geq 0$, $\mu(x) = -1.5\sqrt{-x}$ for $x < 0$. The sender optimally discloses a conclusive bad-news signal in a random period when $\theta = B$, so each period of silence amounts to a small piece of good news. (In Appendix B.2.2, we consider Kőszegi and Rabin [2009]'s percentile-based news utility model in a similar environment with Gaussian distributions of residual consumption uncertainty in the two states. We find a very similar optimal information structure under the same square-root gain-loss function.)

The information-design problem imposes additional constraints relative to a habit-formation model. To see this, consider a "relaxed" version of the sender's problem in the binary-states case where she simply chooses some $x_t \in [0, 1]$ each period for $1 \leq t \leq T - 1$, depending on the realization of θ. The receiver gets $\mu_t(x_t - x_{t-1})$ in period $1 \leq t \leq T$, with the initial condition $x_0 = \pi_0$ and the terminal condition $x_T = 1$ if $\theta = G$, $x_T = 0$ if $\theta = B$. One interpretation of the relaxed problem is that the sender chooses the receiver's sequence of beliefs only subject to the constraint that the initial belief in period 0 is π_0 and the final belief in period T puts probability 1 on the true state. The belief paths do not have to be Bayesian. Another interpretation is that x_t is not a belief, but a consumption level for period t. The receiver's welfare in period t only depends on a gain-loss utility based on how current period's consumption differs from that of period $t - 1$.

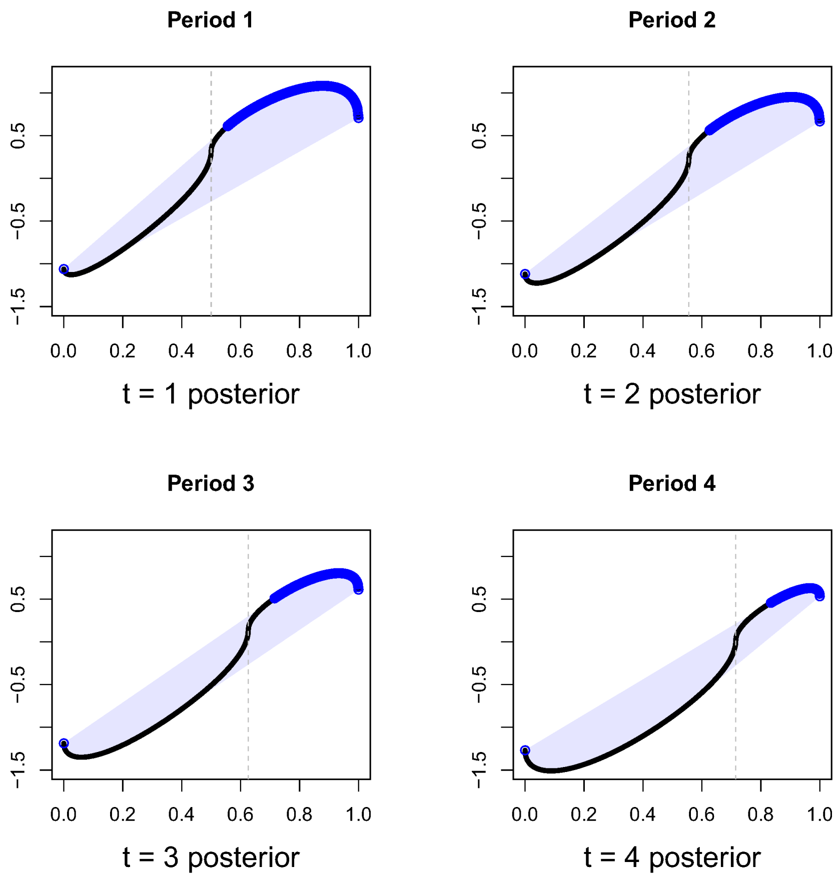

Figure 2.3.1: The concavifications giving the optimal information structure.

Horizon $T = 5$, mean-based news-utility function $\mu(x) = \begin{cases} \sqrt{x} & \text{for } x \geq 0 \\ -1.5\sqrt{-x} & \text{for } x < 0 \end{cases}$, prior $\pi_0 = 0.5$. The dashed vertical line in the t-th graph marks the receiver's belief in $\theta = G$ conditional on not having heard any bad news by the start of period t. The y-axis shows the sum of news utility this period and the value function of entering next period with a certain belief. In the good state of the world, the receiver's belief in $\theta = G$ grows at increasing rates across the periods, $0.5 \to 0.556 \to 0.626 \to 0.715 \to 0.834 \to 1$. In the bad state of the world, the receiver's belief follows the same path as in the good state up until the random period when conclusive bad news arrives.

Provided μ has diminishing sensitivity and exhibits enough loss aversion, the sender maximizes the receiver's utility by choosing $x_t = \pi_0 + \frac{t}{T}(1 - \pi_0)$ in period t when $\theta = G$, and by choosing $x_t = 0$ in every period $t \geq 1$ when $\theta = B$. The belief paths in Figure 2.3.1 differ from these "relaxed" solutions in two ways. First, the receiver gets different amounts of good news (in terms of $\pi_t - \pi_{t-1}$) in different periods when $\theta = G$. Second, the sender sometimes provides false hope in the bad state. These differences come from the Bayesian constraints on beliefs.

2.3.2 Sub-Optimality of One-Shot Resolution

We begin with a sufficient condition on the news-utility function for one-shot resolution to be strictly suboptimal for any T and Θ. Let $\theta_H, \theta_L \in \Theta$ be the states with the highest and lowest consumption utilities. Let $1_H, 1_L \in \Delta(\Theta)$ represent degenerate beliefs in states θ_H and θ_L and let $v_0 := \mathbb{E}_{\theta \sim \pi_0}(v(c_\theta))$ be the ex-ante expected future consumption utility. The symbol $\oplus$ denotes the mixture between two beliefs in $\Delta(\Theta)$.

Proposition 12. *For any T and Θ, one-shot resolution is strictly suboptimal if*

$$\lim_{\varepsilon \to 0^+} \frac{N(1_H \mid (1-\varepsilon)1_H \oplus \varepsilon 1_L)}{\varepsilon} + N(1_H \mid \pi_0) - N(1_L \mid \pi_0)$$
$$> \lim_{\varepsilon \to 0^+} \frac{N(1_H \mid \pi_0) - N((1-\varepsilon)1_H \oplus \varepsilon 1_L \mid \pi_0)}{\varepsilon} - N(1_L \mid 1_H).$$

For the mean-based news-utility model, this condition is equivalent to

$$\mu'(0^+) + \mu(1 - v_0) - \mu(-v_0) > \mu'(1 - v_0) - \mu(-1).$$

In fact, the proof of Proposition 12 shows that whenever its condition is satisfied, some information structure featuring gradual good news and one-shot bad news (to be defined precisely in the next subsection) is strictly better than one-shot resolution.

We can interpret Proposition 12's sufficient condition as "strong enough diminishing sensitivity relative to loss aversion." Evidently, $\mu(1 - v_0) - \mu(-v_0) > 0$, so the condition is satisfied whenever $\mu'(0^+) - \mu'(1 - v_0) \geq -\mu(-1)$. The LHS increases when μ becomes more concave in the positive region, and the RHS decreases when μ is more convex in the negative region. On the other hand, holding fixed the curvature $\mu^{''}(x)$ for $x \neq 0$ and $\mu^{'}(0^{+})$, increasing the amount of loss aversion near 0 (i.e., $\mu^{'}(0^{-}) - \mu^{'}(0^{+})$) increases RHS.

The quadratic news utility provides a clear illustration of this interpretation, as the condition of Proposition 12 holds whenever there is enough curvature relative to the extent of loss aversion.

Corollary 1. *If the receiver has quadratic news utility with $\alpha_n - \alpha_p \leq \beta_n + \beta_p$, then one-shot resolution is strictly suboptimal for any T.*

The difference $\alpha_n - \alpha_p \geq 0$ is the size of the "kink" at 0, that is $\mu^{'}(0^{-}) - \mu^{'}(0^{+})$. On the other side, β_p and β_n control the amounts of curvature in the positive and negative regions, respectively.

The sufficient condition in Proposition 12 is also satisfied by the most commonly used model of diminishing sensitivity, the power function (see, for example, Tversky and Kahneman [1992]). One could think of the power function specification as having "infinite" diminishing sensitivity near 0, as $\mu^{''}(0^{+}) = -\infty$ and $\mu^{''}(0^{-}) = \infty$.

Corollary 2. *Suppose* $\mu(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0 \\ -\lambda \cdot |x|^{\beta} & \text{if } x < 0 \end{cases}$ *for some* $0 < \alpha, \beta < 1$ *and* $\lambda \geq 1$. *Then one-shot resolution is strictly suboptimal for any T.*

While Proposition 12 holds generally, we can find sharper results on the sub-optimality of one-shot resolution for specific news-utility models and environments. Kőszegi and Rabin [2009]'s percentile-based news-utility model stipulates

$$N(\pi_t \mid \pi_{t-1}) = \int_0^1 \mu\left(v(F_{\pi_t}(p)) - v(F_{\pi_{t-1}}(p))\right) dp,$$

where $F_{\pi_t}(p)$ and $F_{\pi_{t-1}}(p)$ are the p-th percentile consumption levels according to beliefs π_t and π_{t-1}, respectively. Whenever μ exhibits diminishing sensitivity to gains and there are at least 3 states, one-shot resolution is suboptimal. This result does not require any assumption about loss aversion or diminishing sensitivity in losses.

Proposition 13. *In Kőszegi and Rabin [2009]'s percentile-based news-utility model, provided the gain-loss utility function satisfies $\mu''(x) < 0$ for all $x > 0$, one-shot resolution is strictly suboptimal for any T and any $K \geq 3$.*

Similar to the idea behind Proposition 12, the proof of Proposition 13 constructs an information structure to gradually deliver the good news that the state is the best one possible. By contrast, if μ is two-part linear with $\mu(x) = bx$ for $x \geq 0$, $\mu(x) = \lambda bx$ for some $b > 0, \lambda > 1$ (so that $\mu''(x) = 0$ for $x > 0$), then one-shot resolution is the uniquely optimal information structure [Kőszegi and Rabin, 2009].

Proposition 13 requires at least three distinct consumption levels, $K \geq 3$. In a binary-states world, the percentile-based news-utility function N only depends on the value of μ at two non-zero points. Thus every increasing μ is behaviorally indistinguishable from a two-part linear one, meaning the percentile-based model cannot capture diminishing sensitivity in a setting with $K = 2$.

As an analog to Corollary 2, we study a setting with percentile-based news utility and residual consumption uncertainty in Appendix B.2, finding that one-shot resolution is strictly suboptimal with *any* number of states for a power-function μ (Corollary 11).

2.3.3 Gradual Good News and Gradual Bad News

For the remainder of the paper, we focus on mean-based news-utility functions to study additional implications of diminishing sensitivity. Two classes of information structures will play important roles in the sequel. To define them, we write $v_t := \mathbb{E}_{\theta \sim \pi_t}[v(c_\theta)]$ for the expected future consumption utility based on the receiver's (random) belief at the end of period t. Partition states into two subsets, $\Theta = \Theta_B \cup \Theta_G$, where $v(c_\theta) < v_0$ for $\theta \in \Theta_B$ and $v(c_\theta) \geq v_0$ for $\theta \in \Theta_G$. Interpret Θ_B as the "bad" states and Θ_G as the "good" ones.

Definition 5. An information structure (M, σ) features *gradual good news, one-shot bad news* if

- $\mathbb{P}_{(M,\sigma)}[v_t \geq v_{t-1} \text{ for all } 1 \leq t \leq T \mid \theta \in \Theta_G] = 1$ and
- $\mathbb{P}_{(M,\sigma)}[v_t < v_{t-1} \text{ for no more than one } 1 \leq t \leq T \mid \theta \in \Theta_B] = 1$.

An information structure (M, σ) features *gradual bad news, one-shot good news* if

- $\mathbb{P}_{(M,\sigma)}[v_t \leq v_{t-1} \text{ for all } 1 \leq t \leq T \mid \theta \in \Theta_B] = 1$ and
- $\mathbb{P}_{(M,\sigma)}[v_t > v_{t-1} \text{ for no more than one } 1 \leq t \leq T \mid \theta \in \Theta_G] = 1$.

In the first class of information structures ("gradual good news, one-shot bad news"), the sender relays good news over time and gradually increases the receiver's expectation of future consumption. When the state is bad, the sender concentrates all the bad news in one period. The "one-shot bad news" terminology comes from noting that when $\theta \in \Theta_B$, the single period t where $v_t < v_{t-1}$ must satisfy $v_t = v(c_\theta)$ and $v_{t'} = v_t$ for all $t' > t$. The receiver gets negative information about his future consumption level for the first time in period t, and his expectation stays constant thereafter. On the other hand, we use the phrase "gradual bad news, one-shot good news" to refer to the "opposite" kind of information structure.

One-shot resolution falls into both of these classes. To rule out this triviality, we say that an information structure features *strictly gradual good news* if

$$\mathbb{P}_{(M,\sigma)}[v_t > v_{t-1} \text{ and } v_{t'} > v_{t'-1} \text{ for two distinct } 1 \leq t, t' \leq T \mid \theta \in \Theta_G] > 0.$$

That is, there is positive probability that the receiver's expectation strictly increases at least twice in periods 1 through T. Similarly define *strictly gradual bad news*.

We now prove that whenever μ satisfies diminishing sensitivity and (weak) loss aversion, information structures featuring strictly gradual bad news, one-shot good news are strictly worse than one-shot resolution. By contrast, under some additional restrictions, the optimal information structure falls into the strictly gradual good news, one-shot bad news class.

Proposition 14. *Suppose μ satisfies diminishing sensitivity and loss aversion. Any information structure featuring strictly gradual bad news, one-shot good news is strictly worse than one-shot resolution in expectation, and almost surely weakly worse ex-post.*

This result holds for arbitrary state space Θ, horizon T, and prior π_0.

For the rest of the paper, we specialize to the case of $K = 2$. The next result presents a necessary and sufficient condition for inconclusive bad news to be suboptimal when $T = 2$. We then verify the condition for quadratic news utility.

Proposition 15. *For $T = 2$, information structures with $\mathbb{P}_{(M,\sigma)}[\pi_1 < \pi_0$ and $\pi_1 \neq 0] > 0$ are strictly suboptimal if and only if there exists some $q \geq \pi_0$ so that the chord connecting $(0, U_1(0 \mid \pi_0))$ and $(q, U_1(q \mid \pi_0))$ lies strictly above $U_1(p \mid \pi_0)$ for all $p \in (0, \pi_0)$.*

Corollary 3. *Quadratic news utility satisfies the condition of Proposition 15.*

In particular, combining Corollaries 1 and 3, we infer that any optimal information structure for a receiver with quadratic news utility with $\alpha_n - \alpha_p \leq \beta_n + \beta_p$ with $T = 2$ must feature strictly gradual good news, one-shot bad news. Furthermore, since there exists an optimal information structure with binary messages by Proposition 11, in this environment there is an optimal information structure where the sender induces either belief 0 or belief $p_H > \pi_0$ in the only period of communication. The next subsection characterizes p_H as a function of the model parameters.

In summary, we have established a ranking between three kinds of information structures. For any time horizon and any state space, provided the condition in Proposition 12 holds and μ satisfies diminishing sensitivity and weak loss aversion, *some* information structure featuring gradual good news, one-shot bad news gives more news utility than one-shot resolution, which in turn gives more news utility than *any* information structure featuring strictly gradual bad news, one-shot good news. Further, under the additional restrictions in Proposition 15, a gradual good news, one-shot bad news information structure is optimal among all information structures.

2.3.4 Illustrative Example: Quadratic News Utility

We illustrate Proposition 11's concavification procedure by finding in closed-form the optimal information structure when the receiver has a quadratic news-utility function.

Suppose the parameters of μ satisfy $\alpha_n - \alpha_p \leq \beta_n + \beta_p$ in a $T = 2$ environment. From the arguments in Section 2.3.3, the optimal information structure induces either $\pi_1 = 0$ or $\pi_1 = p_H$ for some $p_H > \pi_0$. Proposition 11 implies $(\text{cav}U_1(\cdot \mid \pi_0))(x) > U_1(x \mid \pi_0)$ for all $x \in (0, p_H)$. The geometry of concavification shows the derivative of the value function at p_H, $\frac{\partial}{\partial x} U_1(x \mid \pi_0)(p_H)$, equals the slope of the chord from 0 to p_H on the function $U_1(\cdot \mid \pi_0)$. We use this equality to derive p_H as the solution to a cubic polynomial.

Proposition 16. *For $T = 2$ and quadratic news utility satisfying $\alpha_n - \alpha_p \leq \beta_n + \beta_p$, the optimal partial good news $p_H > \pi_0$ satisfies*

$$\pi_0(\alpha_n - \alpha_p) - (\beta_p + \beta_n)\pi_0^2 = p_H^2(\alpha_n - \alpha_p + \beta_n + \beta_p) - p_H^3(2\beta_p + 2\beta_n).$$

We have $\frac{dp_H}{d\pi_0} > 0$ for $\pi_0 < \frac{1}{2}\frac{\alpha_n - \alpha_p}{\beta_n + \beta_p}$ and $\frac{dp_H}{d\pi_0} < 0$ for $\pi_0 > \frac{1}{2}\frac{\alpha_n - \alpha_p}{\beta_n + \beta_p}$.

In other words, the optimal partial good news is in general non-monotonic in the prior belief. For low prior beliefs, p_H increases with prior. But for high prior beliefs, p_H decreases with prior. Figure 2.3.2 illustrates. In the case of $\alpha_n = \alpha_p$, and in particular when μ is symmetric around 0, $\frac{dp_H}{d\pi_0} > 0$ for any $\pi_0 \in (0, 1)$.

2.3.5 Endogenous Diversity of Information Preferences

Leaving aside the setting where the sender knows the state upfront and can choose any information structure, consider an environment where a sequence of exogenous signal realizations determine the state. We show that agents with

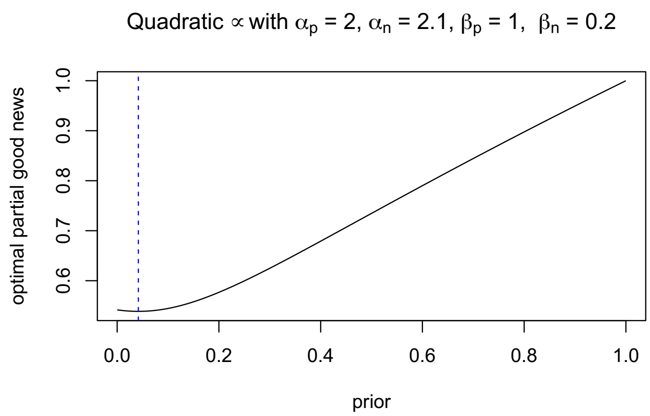

Figure 2.3.2: Optimal partial good news with quadratic news utility. $T = 2$, fixing parameters $\alpha_p = 2$, $\alpha_n = 2.1$, $\beta_p = 1$, $\beta_n = 0.2$ and considering different prior beliefs. The dashed blue line is at $\pi_0 = \frac{1}{2}\frac{\alpha_n - \alpha_p}{\beta_n + \beta_p} \approx 0.042$. The optimal partial good news is decreasing in the prior before this threshold, and increasing afterwards.

opposite consumption preferences over the two states can exhibit opposite preferences when choosing between observing the signals as they arrive (gradual resolution) or only learning the final state (one-shot resolution).

There are two states of the world, Alternative (A) and Baseline (B). In each period $t = 1, 2, \ldots, T$, a binary random variable X_t realizes, where $\mathbb{P}[X_t = 1] = q_t$ with $0 < q_t < 1$. If $X_t = 1$ for all t, then the state is A. Else, if $X_t = 0$ for at least one t, then the state is B. The agent's consumption utility in period T depends on the state, and is normalized without loss to be either 0 or 1.

At time 0, the agent chooses between observing the realizations of the random variables $(X_t)_{t=1}^{T}$ in real time, or only learning the state of the world at the end of period T. As an example, imagine a televised debate between two political candidates A and B where A loses as soon as she makes a "gaffe" during the debate.[3] If A does not make any gaffes, then A wins. An individual who strongly prefers one of the candidates to win must choose between watching the debate

[3] Augenblick and Rabin [2018] use a similar example of political gaffes to illustrate Bayesian belief movements.

live in the evening or only reading the outcome of the debate the following morning.

The agent forms Bayesian belief $\pi_t \in [0, 1]$ about the probability of state A at the end of each period t, starting with the correct Bayesian prior π_0. For notational convenience, we also write $\rho_t = 1 - \pi_t$ as the belief in state B at the end of t, with the prior $\rho_0 = 1 - \pi_0$. If the agent prefers state A, he gets news utility $\mu(\pi_t - \pi_{t-1})$ at the end of period t. If the agent prefers state B, then he gets news utility $\mu(\rho_t - \rho_{t-1})$. The function μ exhibits diminishing sensitivity, that is $\mu''(x) < 0$ and $\mu''(-x) > 0$ for $x > 0$. Also, to quantify the amount of loss aversion, we consider the parametric class of *λ-scaled news-utility functions*. We fix some $\tilde{\mu}_{pos} : [0, 1] \to \mathbb{R}_+$, strictly increasing and strictly concave with $\tilde{\mu}_{pos}(0) = 0$, and consider the family of μ's given by $\mu(x) = \tilde{\mu}_{pos}(x)$, $\mu(-x) = -\lambda\tilde{\mu}_{pos}(x)$ for $x > 0$ as we vary $\lambda \geq 1$.

Under diminishing sensitivity, someone rooting for state A wants to watch the events unfold in real time to celebrate the small victories, while someone hoping for state B prefers to only learn the final state to avoid piecemeal bad news. The next proposition formalizes this intuition.

Proposition 17. *Consider the class of λ-scaled news-utility functions. For any $\lambda \geq 1$, an agent who prefers state B will choose one-shot resolution of uncertainty over gradual resolution of uncertainty. There exists some $\bar{\lambda} > 1$ so that for any $1 \leq \lambda \leq \bar{\lambda}$, an agent who prefers state A will choose gradual resolution of uncertainty over one-shot resolution of uncertainty.*

This result suggests a possible mechanism for media competition: if the realization of some state A depends on a series of smaller events, then some news sources may cover these small events in detail as they happen, while other sources may choose to only report the final outcome. Viewers sort between these two kinds of news sources based on how they rank states A and B in terms of consumption. Opposite consumption preferences induce opposite informational preferences.

By contrast, other behavioral models do not predict a diversity of informational preferences in this environment.

Proposition 18. *The following models do not predict different informational preferences for agents with the opposite consumption rankings for the two states.*

A. *Two-part linear news-utility function μ.*

B. *Anticipatory utility where the agent gets either $u(\pi_t)$ or $u(\rho_t)$ in period t depending on his preference over states A and B, with u an increasing, weakly concave function.*

C. *Ely, Frankel, and Kamenica [2015]'s suspense and surprise utilities.*

Another application of Proposition 17 concerns the design of game shows. Consider a game show featuring a single contestant who will win either $100,000 or nothing depending on her performance across five rounds.[4] The audience, empathizing with the contestant, derives news utility $\mu(\pi_t - \pi_{t-1})$ at the end of round t, where π_t is the contestant's probability of winning the prize based on the first t rounds. One possible format ("sudden death") features five easy rounds each with $w = 0.5^{1/5} \approx 87\%$ winning probability, where the contestant wins $100,000 if she wins all five rounds. Another possible format ("repêchage") involves five hard rounds each with $1 - w$ winning probability, but the contestant wins $100,000 as soon as she wins any round. Both formats lead to the same distribution over final outcomes and generate the same amount of suspense and surprise utilities à la Ely, Frankel, and Kamenica [2015]. Proposition 17 shows the first format induces more news utility than one-shot resolution for audience members who are not too loss averse, while the second format is worse than one-shot resolution for all audience members. Consistent with our model, the vast majority of game shows resemble the first format more than the second format.

[4]This can be thought of as a stylized payout structure for game shows like *American Ninja Warrior* and *Who Wants to Be a Millionaire*.

2.4 The Credibility Problem of Gradual Good News

Section 2.3 studied the optimal disclosure of news when the sender has commitment power. We provided conditions for the optimal information structure to feature gradual good news, one-shot bad news. Information structures of this kind encounter a credibility problem when the commitment assumption is dropped. If the sender wishes to gradually reveal the good state to a Bayesian receiver over multiple periods, then she must also sometimes provide false hope in the bad state due to the cross-state constraints on beliefs. But without commitment, the benevolent sender may strictly prefer giving false hope over telling the truth in the bad state. This deviation improves the *total* news utility of a receiver with diminishing sensitivity, if the positive utility from today's good news outweighs the additional future disappointment from higher expectations. In fact, when news utility is symmetric and exhibits diminishing sensitivity, the above credibility problem is so severe that every equilibrium is payoff-equivalent to the babbling equilibrium. The same result also applies to asymmetric news-utility functions with $\mu(-x) = -\lambda\mu(x)$ for all $x > 0$, provided loss aversion $\lambda > 1$ is weak enough relative to μ's diminishing sensitivity in a way we formalize.

Sufficiently strong loss aversion can restore the equilibrium credibility of good-news messages. We show that the highest equilibrium payoff when the sender lacks commitment may be non-monotonic in the extent of loss aversion, in contrast to the conclusion that more loss-averse receivers are always strictly worse off when the sender has commitment power. We also completely characterize the class of equilibria that feature (a deterministic sequence of) gradual good news in the good state and study the equilibrium rate of learning. With the quadratic or the square-root news-utility function, equilibria within this class always release progressively larger pieces of good news over time, so the receiver's belief in the good state grows at an increasing rate.

2.4.1 Equilibrium Analysis When the Sender Lacks Commitment

We continue to maintain that state space $\Theta = \{G, B\}$ is binary. To study the case where the sender lacks commitment, we analyze the perfect-Bayesian equilibria of the cheap talk game between the two parties. Formally, the equilibrium concept is as follows.

Definition 6. Let a finite set of messages M be fixed. A *perfect-Bayesian equilibrium* consists of sender's strategy $\sigma^* = (\sigma_t^*)_{t=1}^{T-1}$ together with receiver's beliefs $p^* : \cup_{t=0}^{T-1} H^t \to [0, 1]$, where:

- For every $1 \leq t \leq T-1$, $h^{t-1} \in H^{t-1}$ and $\theta \in \{G, B\}$, σ^* maximizes the receiver's total expected news utility in periods $t, \ldots, T-1, T$ conditional on having reached the public history h^{t-1} in state θ at the start of period t.

- p^* is derived by applying the Bayes' rule to σ^* whenever possible.

We make two belief-refinement restrictions:

- If $t \leq T-1$, h^t is a continuation history of $h^{\underline{t}}$, and $p^*(h^{\underline{t}}) \in \{0, 1\}$, then $p^*(h^t) = p^*(h^{\underline{t}})$.

- The receiver's belief in period T when state is θ satisfies $\pi_T = 1_\theta$, regardless of the preceding history $h^{T-1} \in H^{T-1}$.

We will abbreviate a perfect-Bayesian equilibrium satisfying our belief refinements as an "equilibrium." Our definition requires that once the receiver updates his belief to 0 or 1, it stays constant through the end of period $T-1$. In period T, the receiver updates his belief to reflect full confidence in the true state of the world, regardless of his (possibly dogmatic) belief at the end of period $T-1$.. The receiver derives news utility in periods $1 \leq t \leq T$ based on changes in his belief, as in the model with commitment.

Let $\mathcal{V}_{\mu,M,T}(\pi_0) \subseteq \mathbb{R}$ denote the set of equilibrium payoffs with news-utility function μ, message space M, time horizon T, and prior π_0. Clearly, $\mathcal{V}_{\mu,M,T}(\pi_0)$ is non-empty. There is always the *babbling equilibrium*, where the sender mixes over all messages uniformly in both states and the receiver's belief never updates from the prior belief until period T. Denote the babbling equilibrium payoff by

$$V_{\mu}^{Bab}(\pi_0) := \pi_0\mu(1-\pi_0) + (1-\pi_0)\mu(-\pi_0)$$

and note it is independent of M or T.

We state two preliminary properties of the equilibrium payoffs set $\mathcal{V}_{\mu,M,T}(\pi_0)$.

Lemma 3. *We have:*

A. *For any finite M, $\mathcal{V}_{\mu,M,T}(\pi_0) \subseteq \mathcal{V}_{\mu,\{g,b\},T}(\pi_0)$*

B. *If $T \leq T'$, then $\mathcal{V}_{\mu,M,T}(\pi_0) \subseteq \mathcal{V}_{\mu,M,T'}(\pi_0)$.*

The first statement says any equilibrium payoff achievable with an arbitrary finite message space is also achievable with a binary message space. The second statement says the set of equilibrium payoffs weakly expands with the time horizon.

2.4.2 The Credibility Problem and Babbling

To understand the source of the credibility problem, let $N_B(x;\pi) := \mu(x-\pi) + \mu(-x)$ denote the total amount of news utility across two periods when the receiver updates his belief from π to $x > \pi$ today and updates it from p to 0 tomorrow. Suppose there exists a period $T-2$ public history $h^{T-2} \in H^{T-2}$ with $p^*(h^{T-2}) = \pi$ and some $x > \pi$ satisfying $N_B(x;\pi) > N_B(0;\pi)$. Then, the sender strictly prefers to induce belief x rather than belief 0 after arriving at the history h^{T-2} in the bad state. A good-news message m_x inducing belief x and a bad-news message m_0 inducing belief 0 cannot both be on-path following h^{T-2}, else the sender would strictly prefer to send m_x with probability 1 in the bad state.

Yet, the inequality $N_B(0;\pi) < N_B(x;\pi)$ automatically holds for any $x > \pi$, provided μ is strictly concave in the positive region and symmetric around 0.

Lemma 4. *If μ is symmetric around 0 and $\mu''(x) < 0$ for all $x > 0$, then for any $0 < \pi < x < 1$ it holds $N_B(0; \pi) < N_B(x; \pi)$.*

The intuition is that when the state is bad, the sender knows the receiver will inevitably get conclusive bad news in period T. Giving false hope in period $T - 1$ (i.e., inducing belief $x > \pi$ instead of 0) provides positive news utility at the cost of greater disappointment in the final period. Diminishing sensitivity limits the *incremental* cost of this additional disappointment.

The credibility problem implies that the babbling payoff is the unique equilibrium payoff.

Proposition 19. *Suppose μ is symmetric around 0 and $\mu''(x) < 0$ for all $x > 0$. For any M, T, π_0, $\mathcal{V}_{\mu,M,T}(\pi_0) = \{V_\mu^{Bab}(\pi_0)\}$.*

We now explore what happens when μ is asymmetric around 0 due to loss aversion. Say μ exhibits *greater sensitivity to losses* if $\mu'(x) \leq \mu'(-x)$ for all $x > 0$. We first establish a robustness check on Proposition 19 within this class of news-utility functions: when loss aversion is sufficiently weak relative to diminishing sensitivity in a $T = 2$ model, the babbling equilibrium remains unique up to payoffs.

Proposition 20. *Suppose μ exhibits greater sensitivity to losses. If $\min_{z\in[0,1-\pi_0]} \frac{\mu'(z)}{\mu'(-(\pi_0+z))} > 1$, then $\mathcal{V}_{\mu,M,2}(\pi_0) = \{V_\mu^{Bab}(\pi_0)\}$ for any M.*

When μ is symmetric and does not exhibit strict loss aversion, diminishing sensitivity implies $\mu'(-(\pi_0 + z)) = \mu'(\pi_0 + z) < \mu'(z)$ for every $z \in [0, 1 - \pi_0]$, so the inequality condition in Proposition 20 is always satisfied. This condition continues to hold if μ is slightly asymmetric due to a "small enough" amount of loss aversion relative to the size of the sensitivity gap $\mu'(z) - \mu'(\pi_0 + z)$. This interpretation is clearest for the λ-scaled news-utility functions, as formalized in the following corollary.

Corollary 4. *Suppose for some $\tilde{\mu}_{pos} : [0, 1] \to \mathbb{R}_+$ and $\lambda \geq 1$, the news-utility function μ satisfies $\mu(x) = \tilde{\mu}_{pos}(x)$, $\mu(-x) = -\lambda\tilde{\mu}_{pos}(x)$ for all $x \geq 0$. Provided $\lambda < \min_{z\in[0,1-\pi_0]} \frac{\tilde{\mu}'_{pos}(z)}{\tilde{\mu}'_{pos}(\pi_0+z)}$, $\mathcal{V}_{\mu,M,2}(\pi_0) = \{V_\mu^{Bab}(\pi_0)\}$ for any M.*

When μ is strictly concave in the positive region, Corollary 4 gives a non-degenerate interval of loss-aversion parameters for which the conclusion of Proposition 19 extends in a $T = 2$ setting. If $\tilde{\mu}_{pos}$ contains more curvature, then $\tilde{\mu}'_{pos}(z)/\tilde{\mu}'_{pos}(\pi_0 + z)$ becomes larger and the interval of permissible λ's expands.

What happens when loss aversion is high? The next proposition says a new equilibrium that payoff-dominates the babbling one exists for large λ, provided the marginal utility of an infinitesimally small piece of good news is infinite — as in the power-function specification.

Proposition 21. *Fix $\tilde{\mu}_{pos} : [0, 1] \to \mathbb{R}_+$ strictly increasing and concave, continuously differentiable at $x > 0$, $\tilde{\mu}_{pos}(0) = 0$, and $\lim_{x \to 0} \tilde{\mu}'_{pos}(x) = \infty$. Consider the family λ-indexed news-utility functions $\mu(x) = \tilde{\mu}_{pos}(x)$, $\mu(-x) = -\lambda \tilde{\mu}_{pos}(x)$ for $x \geq 0$. For each $\pi_0 \in (0, 1)$, there exists $\bar{\lambda} \geq 1$ so that whenever $\lambda \geq \bar{\lambda}$ and for any $T \geq 2$, $|M| \geq 2$, there exists $V \in \mathcal{V}_{\mu,M,T}(\pi_0)$ with $V > V_{\mu}^{Bab}(\pi_0)$.*

To help illustrate these results, suppose $\mu(x) = \sqrt{x}$ for $x \geq 0$, $\mu(x) = -\lambda\sqrt{-x}$ for $x < 0$, $T = 2$, and $\pi_0 = \frac{1}{2}$. Corollary 4 implies whenever $\lambda < \sqrt{2}$, the babbling equilibrium is unique up to payoffs. On the other hand, Proposition 21 says when λ is sufficiently high, there is another equilibrium with strictly higher payoffs. In fact, a non-babbling equilibrium first appears when $\lambda = 2.414$.

Figure 2.4.1 plots the highest equilibrium payoff for different values of λ. Receivers with higher λ may enjoy higher equilibrium payoffs. The reason for this non-monotonicity is that for low values of λ, the babbling equilibrium is unique and increasing λ decreases expected news utility linearly. When the new, non-babbling equilibrium emerges for large enough λ, the sender's behavior in the new equilibrium depends on λ. Higher loss aversion carries two countervailing effects: first, a *non-strategic effect* of hurting welfare when $\theta = B$, as the receiver must eventually hear the bad news; second, an *equilibrium effect* of changing the relative amounts of good news in different periods conditional on $\theta = G$. Receivers with an intermediate amount of loss aversion enjoy higher expected news utility than receivers with low loss aversion, as the equilibrium effect leads to better "consumption smoothing" of good news across time. But, the non-strategic effect eventually dominates and receivers with high loss aversion experience worse payoffs than receivers with low loss aversion.

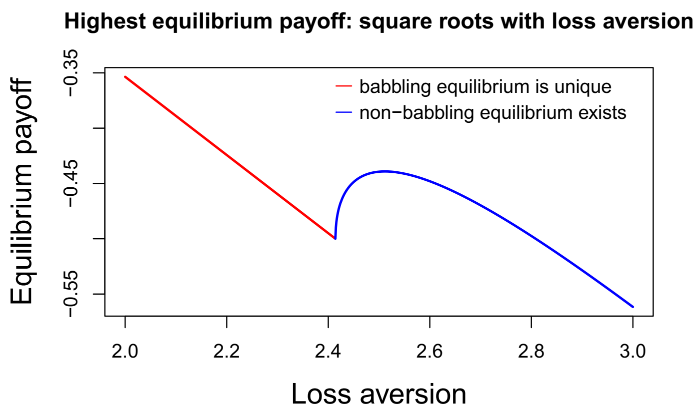

Figure 2.4.1: Highest equilibrium payoff with square roots and loss aversion. The babbling equilibrium is essentially unique for low values of λ, but there exists an equilibrium with gradual good news for $\lambda \geq 2.414$. Due to the role of loss aversion in sustaining credible partial news, a receiver with higher loss aversion may experience higher or lower expected news utility in equilibrium than a receiver with lower loss aversion.

2.4.3 Deterministic Gradual Good News Equilibria

An equilibrium (M, σ^*, p^*) features *deterministic*[5] *gradual good news* (GGN equilibrium) if there exist a sequence of constants $p_0 \leq p_1 \leq \ldots \leq p_{T-1} \leq p_T$ with $p_0 = \pi_0$, $p_T = 1$, and the receiver always has belief p_t in period t when the state is good. By Bayesian beliefs, in the bad state of any GGN equilibrium the

[5]This class of equilibria is slightly more restrictive than the gradual good news, one-shot bad news information structures from Definition 5, because the sender may not randomize between several increasing paths of beliefs in the good state.

sender must induce a belief of either 0 or p_t in period t, as any message not inducing belief p_t is a conclusive signal of the bad state.

The class of GGN equilibria is non-empty, for it contains the babbling equilibrium where $\pi_0 = p_0 = p_1 = \ldots = p_{T-1} < p_T$. The number of *intermediate beliefs* in a GGN equilibrium is the number of distinct beliefs in the open interval $(\pi_0, 1)$ along the sequence $p_0, p_1, \ldots, p_{T-1}$. The babbling equilibrium has zero intermediate beliefs.

The next proposition characterizes the set of all GGN equilibria with at least one intermediate belief.

Proposition 22. *Let $P^*(\pi) \subseteq (\pi, 1]$ be those beliefs x satisfying $N_B(x; \pi) = N_B(0; \pi)$. Suppose μ exhibits diminishing sensitivity and loss aversion. For $1 \leq J \leq T-1$, there exists a gradual good news equilibrium with the J intermediate beliefs $q^{(1)} < \ldots < q^{(J)}$ if and only if $q^{(j)} \in P^*(q^{(j-1)})$ for every $j = 1, \ldots, J$, where $q^{(0)} := \pi_0$.*

To interpret, $P^*(\pi)$ contains the set of beliefs $x > \pi$ such that the sender is indifferent between inducing the two belief paths $\pi \to x \to 0$ and $\pi \to 0$. Recall that when μ is symmetric, Lemma 4 implies this indifference condition is never satisfied, which is the source of the credibility problem for good-news messages. The same indifference condition pins down the relationship between successive intermediate beliefs in GGN equilibria.

We illustrate this result with the quadratic news utility.

Corollary 5. *1) With quadratic news utility,*
$P^*(\pi) = \left\{\pi \cdot \frac{\beta_p+\beta_n}{\beta_p-\beta_n} - \frac{\alpha_n-\alpha_p}{\beta_p-\beta_n}\right\} \cap (\pi, 1)$.

2a) If $\beta_n > \beta_p$, there cannot exist any gradual good news equilibrium with more than one intermediate belief.

2b) If $\beta_n < \beta_p$, there can exist gradual good news equilibria with more than one intermediate belief. For a given set of parameters of the quadratic news-utility function and prior π_0, there exists a uniform bound on the number of intermediate beliefs that can be sustained in equilibrium across all T.

3) In any GGN equilibrium with quadratic news utility, intermediate beliefs in the good state grow at an increasing rate.

Combined with Proposition 22, part 1) of this corollary says that in every GGN equilibrium, the successive intermediate beliefs are related by the linear map $x \mapsto x \cdot \frac{\beta_p+\beta_n}{\beta_p-\beta_n} - \frac{\alpha_n-\alpha_p}{\beta_p-\beta_n}$. When $\beta_n > \beta_p$, this map has a negative slope, so there cannot exist any GGN equilibrium with more than one intermediate belief. When $\beta_p > \beta_n$, this map has a slope strictly larger than 1. As a result, after eliminating periods where no informative signal is released, every GGN equilibrium releases progressively larger pieces of good news in the good state, $q^{(j+1)} - q^{(j)} > q^{(j)} - q^{(j-1)}$. Since equilibrium beliefs in the good state grow at an increasing rate, there exists some uniform bound $\bar{J}$ on the number of intermediate beliefs depending only on the prior belief π_0 and parameters of the news-utility function.

As an illustration, consider the quadratic news utility with $\alpha_p = 2$, $\alpha_n = 2.1$, $\beta_p = 1$, and $\beta_n = 0.2$. Starting at the prior belief of $\pi_0 = \frac{1}{3}$, Figure 2.4.2 shows the longest possible sequence of intermediate beliefs in any GGN equilibrium for arbitrarily large T. Since the P^* sets are either empty sets or singleton sets for the quadratic news utility, Figure 2.4.2 also contains all the possible beliefs in any state of any GGN equilibrium with these parameters.

The result that GGN equilibria release increasingly larger pieces of good news generalizes to other news-utility functions with diminishing sensitivity. The basic intuition is that if the sender is indifferent between providing d amount of false hope and truth-telling in the bad state when the receiver has prior belief π_L (i.e., $\pi_L + d \in P^*(\pi_L)$), then she strictly prefers providing the same amount of false hope over truth-telling at any higher prior belief $\pi_H > \pi_L$. The false hope generates the same positive news utility in both cases, but an extra d units of disappointment matters less when added a baseline disappointment level of π_H rather than π_L, thanks to diminishing sensitivity.

The next proposition formalizes this idea. It shows that when diminishing

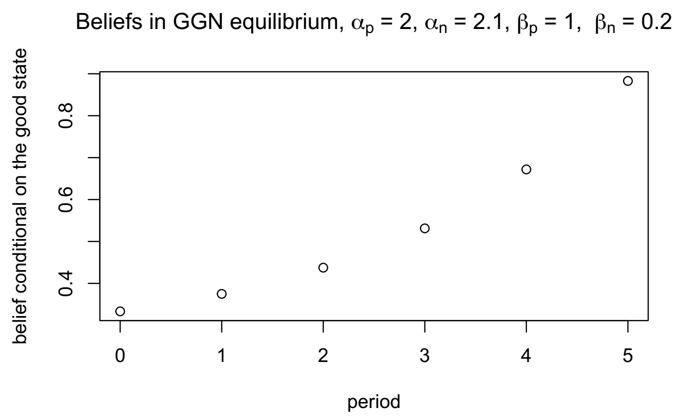

Figure 2.4.2: The longest possible sequence of GGN intermediate beliefs starting with prior $\pi_0 = \frac{1}{3}$.
For quadratic news utility, equilibrium GGN beliefs always increase at an increasing rate in the good state.

sensitivity is combined with a pair of regularity conditions, intermediate beliefs grow at an increasing rate in any GGN equilibrium.

Proposition 23. *Suppose μ exhibits diminishing sensitivity, $|P^*(\pi)| \leq 1$ and $\frac{\partial}{\partial x}N_B(x;\pi)|_{x=\pi} > 0$ for all $\pi \in (0,1)$. Then, in any GGN equilibrium with intermediate beliefs $q^{(1)} < \ldots < q^{(J)}$, we get $q^{(j)} - q^{(j-1)} < q^{(j+1)} - q^{(j)}$ for all $1 \leq j \leq J-1$.*

The first regularity condition requires that the sender is indifferent between the belief paths $\pi \to x \to 0$ and $\pi \to 0$ for at most one $x > \pi$. It is a technical assumption that lets us prove our result, but we suspect the conclusion also holds under some relaxed conditions. The second regularity condition says in the bad state, the total news utility associated with an ε amount of false hope is higher than truth-telling for small ε. These conditions are satisfied by the power-function news utility with $\alpha = \beta$, for example.

Corollary 6. *In any GGN equilibrium with power-function news utility with $\alpha = \beta$ and any $\lambda \geq 1$, intermediate beliefs in the good state grow at an increasing rate.*

2.5 A Random-Horizon Model

In this section, we study a version of our information design problem without a deterministic horizon. Each period, with probability $1 - \delta \in (0, 1]$, the true state of the world is exogenously revealed to the receiver and the game ends. Until then, the informed sender communicates with the receiver each period as in the model from Section 2.2. We verify that our results from the finite-horizon setting extend analogously into this random-horizon environment.

2.5.1 The Random-Horizon Model

Consider an environment where the consumption event takes place far in the future, but the sender is no longer the receiver's only source of information in the interim. Instead, a third party perfectly discloses the state to the receiver with some probability each period. For instance, the sender may be the chair of a central bank who has decided on the bank's monetary policy for next year and wishes to communicate this information over time, while the third party is an employee of the bank who also knows the planned policy. With some probability each period, the employee goes to the press and leaks the future policy decision.

Time is discrete with $t = 0, 1, 2, \ldots$ The sender commits to an information structure (M, σ) at time 0. The information structure consists of a finite message space M and a sequence of message strategies $(\sigma_t)_{t=1}^{\infty}$ where each $\sigma_t(\cdot \mid h^{t-1}, \theta) \in \Delta(M)$ specifies how the sender will mix over messages in period t as a function of the public history h^{t-1} so far and the true state θ.

The sender learns the state at the beginning of period 1 and sends a message according to σ_1. At the start of each period $t = 2, 3, 4, \ldots$, there is probability $(1 - \delta) \in (0, 1]$ that the receiver exogenously and perfectly learns the state θ. If so, the game effectively ends because no further communication from the sender can change the receiver's belief. If not, then the sender sends the next message according to σ_t. The randomization over exogenous learning is i.i.d. across periods, so the time of state revelation (i.e., the horizon of the game) is a geometric random variable.

2.5.2 The Value Function with Commitment

Let $V_\delta : [0, 1] \to \mathbb{R}$ be the value function of the problem with continuation probability δ — that is, $V_\delta(p)$ is the highest possible total expected news utility up to the period of state revelation, when the receiver holds belief p in the current period and state revelation does not happen this period. The value function satisfies the recursion $V_\delta(p) = \tilde{V}_\delta(p \mid p)$, where

$$\tilde{V}_\delta(\cdot \mid p) := \mathrm{cav}_q[\mu(q-p) + \delta V_\delta(q) + (1-\delta)(q \cdot \mu(1-q) + (1-q) \cdot \mu(-q))].$$

Ely [2017] studies an infinite-horizon information design problem whose value function also involves concavification. Unlike in Ely [2017], the current belief enters the objective function for our news-utility problem.

Our first result shows this recursion has a unique solution which increases in δ for any fixed $p \in [0, 1]$.

Proposition 24. *For every $\delta \in [0, 1)$, the value function V_δ exists and is unique. Furthermore, $V_\delta(p)$ is increasing in δ for every $p \in [0, 1]$.*

Figure 2.5.1 illustrates this result by plotting $V_\delta(p)$ for the quadratic news utility with $\alpha_p = 2$, $\alpha_n = 2.1$, $\beta_p = 1$, and $\beta_n = 0.2$ for three different values of δ : 0, 0.8, and 0.95. (In fact, the monotonicity of the value function in δ also holds when there are more than two states.)

The monotonicity of V_δ in δ says that when the sender is benevolent and has commitment power, third-party leaks are harmful for the receiver's expected welfare. This result can be explained intuitively as follows. Just as with increasing T in the finite-horizon model, increasing δ expands the set of implementable belief paths. The idea behind implementing a payoff from a shorter horizon / lower δ is that the sender switches to babbling forever after certain histories. This switching happens at a deterministic calendar time in the finite-horizon setting but at a random time in the random-horizon setup, mimicking the random arrival of the state revelation period.

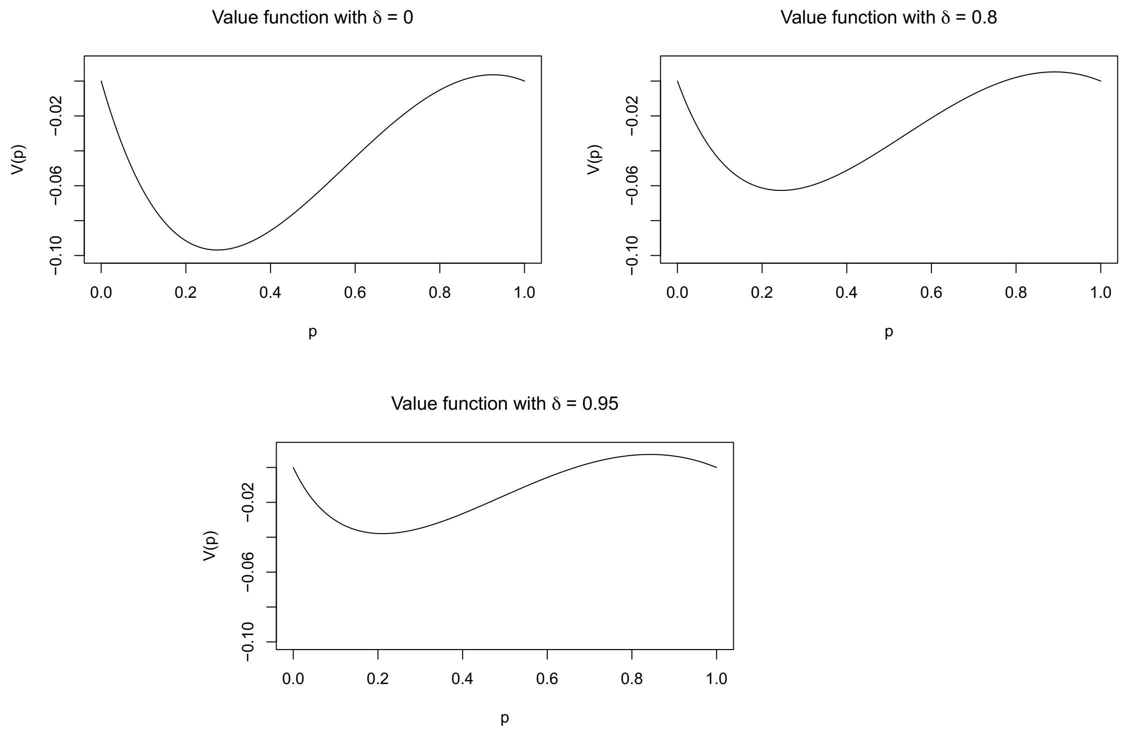

Figure 2.5.1: The value function for the random horizon model. $\delta = 0, 0.8, 0.95$. Consistent with Proposition 24, the value function is pointwise higher for higher δ.

2.5.3 Gradual Good News Equilibria Without Commitment

Now we turn to equilibria of the random-horizon cheap talk game when the sender lacks commitment power. Analogously to the case of finite horizon, a *strict gradual good news* equilibrium (strict GGN) features a deterministic sequence of increasing posteriors $q^{(0)} < q^{(1)} < \ldots$ such that $q^{(0)} = \pi_0$ is the receiver's prior before the game starts and $q^{(t)}$ is his belief in period t, provided state revelation has not occurred. An analog of Proposition 22 continues to hold.

Proposition 25. *Let $P^*(\pi) \subseteq (\pi, 1]$ be those beliefs p satisfying $N_B(p; \pi) = N_B(0; \pi)$. Suppose μ exhibits diminishing sensitivity and loss aversion. There exists a gradual good news equilibrium with a (possibly infinite) sequence of* *intermediate beliefs* $q^{(1)} < q^{(2)} < \ldots$ *if and only if* $q^{(j)} \in P^*(q^{(j-1)})$ *for every* $j = 1, 2, \ldots$, *where* $q^{(0)} := \pi_0$.

The P^* set is the same in the finite- and random-horizon environments. Corollary 6 then implies that even in the random-horizon environment where the game could continue for arbitrarily many periods, intermediate beliefs grow at an increasing rate in GGN equilibria for quadratic and square-roots μ, and there exists a finite bound on the number of periods of informative communication that applies for all $\delta \in [0, 1)$.

2.6 Other Models of Preference over Non-Instrumental Information

2.6.1 Diminishing Sensitivity over News

The literature on reference-dependent preferences and news utility has focused on two-part linear gain-loss utility functions, which violate diminishing sensitivity. If μ is two-part linear with loss aversion, then it follows from the martingale property of Bayesian beliefs that one-shot resolution is weakly optimal for the sender among all information structures. If there is strict loss aversion, then one-shot resolution does strictly better than any information structure that resolves uncertainty gradually. As our results have shown, more nuanced information structures emerge as optimal when the receiver exhibits diminishing sensitivity.

2.6.2 Anticipatory Utility

In our setup, a receiver who experiences anticipatory utility gets $A\left(\sum \pi_t(\theta) \cdot \nu(c_\theta)\right)$ if she ends period t with posterior belief $\pi_t \in \Delta(\Theta)$, where $A : \mathbb{R} \to \mathbb{R}$ is a strictly increasing anticipatory-utility function. When A is the identity function (as in Kőszegi [2006]), the solution to the sender's problem would be unchanged if we modified our model and let the receiver experience both anticipatory utility and news utility. This is because by the martingale property, the receiver's ex-ante expected anticipatory utility in a given period is the same across all information structures. So, the ranking of information structures entirely depends on the news utility they generate. For a general A, if the receiver only experiences anticipatory utility, not news utility, then the sender has an optimal information structure that only releases information in $t = 1$, followed by uninformative babbling in all subsequent periods. For instance, Schweizer and Szech [2018] show that the best non-instrumental medical test for a patient with a concave anticipatory-utility function A is fully uninformative. The above argument establishes that even if the doctor can give the patient a series of tests on different days and even if A is not concave, the optimal test design involves a possibly informative test on the first day, followed by uninformative tests on all subsequent days. The rich dynamics of the optimal information structure in our news-utility model are thus absent in an anticipatory-utility model.

2.6.3 State-Dependent Suspense or Surprise Utility

A key distinction of our model from Ely, Frankel, and Kamenica [2015] is that changes in beliefs may bring utility or disutility to the receiver, depending on the nature of the news. By contrast, agents with suspense or surprise utilities always derive greater utility from larger movements in beliefs, regardless of the directions of these movements.

Ely, Frankel, and Kamenica [2015] also discuss state-dependent versions of suspense and surprise utilities, but this extension does not embed our model either. Suppose there are two states, $\Theta = \{G, B\}$, and the agent has the suspense objective $\sum_{t=0}^{T-1} u\left(\mathbb{E}_t(\sum_\theta \alpha_\theta \cdot (\pi_{t+1}(\theta) - \pi_t(\theta))^2\right)$ or the surprise objective $\sum_{t=1}^{T} u\left(\sum_\theta \alpha_\theta \cdot (\pi_t(\theta) - \pi_{t-1}(\theta))^2\right)$, where $\alpha_G, \alpha_B > 0$ are state-dependent scaling weights. We must have $\pi_{t+1}(G) - \pi_t(G) = -(\pi_{t+1}(B) - \pi_t(B))$, so pathwise $(\pi_{t+1}(G) - \pi_t(G))^2 = (\pi_{t+1}(B) - \pi_t(B))^2$. This shows that the new objectives obtained by applying two possibly different scaling weights $\alpha_G \neq \alpha_B$ to states G and B are identical to the ones that would be obtained by applying the *same* scaling weight $\alpha = \frac{\alpha_G + \alpha_B}{2}$ to both states. Due to this symmetry in preference, the optimal information structure for entertaining an agent with state-dependent suspense or surprise utility does not treat the two states asymmetrically, in contrast to a central prediction of diminishing sensitivity in our model.

2.7 Concluding Discussion

In this work, we have studied how an informed sender optimally communicates with a receiver who derives diminishing gain-loss utility from changes in beliefs. If we think that diminishing sensitivity to the magnitude of news is psychologically realistic in this domain, then the stark predictions of the ubiquitous two-part linear models may be misleading. In the presence of diminishing sensitivity, richer information structures emerge as optimal for the committed sender. For example, the optimal information structure can feature asymmetric treatments of good and bad news. If the sender lacks commitment power, diminishing sensitivity leads to novel credibility problems that inhibit any meaningful communication when the receiver has no loss aversion.

Some of our predictions can empirically distinguish news utility with diminishing sensitivity from other models of belief-based preference over non-instrumental information, including the two-part linear news-utility model. Proposition 17, for example, suggests a laboratory experiment where a sequence of binary events determines whether a baseline state or an alternative state realizes, with the alternative state happening only if all of the binary events are "successful." Consider two treatments that have the same success probabilities for the binary events, but differ in terms of whether subjects get a higher consumption or a lower consumption in the alternative state compared with the baseline state. Diminishing sensitivity over news predicts that more subjects should prefer one-shot resolution when consumption is lower in the alternative state than when it is higher in the alternative state, a hypothesis we plan to test in future work.

3

Identification and Welfare Analysis in Sequential Sampling Models

This chapter is coauthored with Yi-Hsuan Lin. We are thankful to Larry Epstein, Drew Fudenberg and Tomasz Strzalecki for their continuous encouragement and support in this project. We also thank Jerry Green and Kevin He for insightful comments on this project. Any errors are ours.

3.1 Introduction

Many economic situations involve an agent conducting sequential learning about an unobservable state of the world before taking a decision. This paper shows how an outside analyst, who observes for choice menus the joint frequencies of choice out of the menu and of the decision time, can identify the agent's parameters, namely her taste, her prior as well as the costs of the information. We focus on two dynamic learning environments with costly information. In both environments information is only revealed gradually. In the first one the agent is impatient whereas in the second one the agent is repeatedly conducting the same experiment at a fixed cost unobservable to the analyst. We also show how, based only on random choice data about choice from menus and decision times, the analyst can conduct welfare analysis for agents in either of the two situations.

We have two illustrative examples in mind. First, consider a pharmaceutical company which is in the process of evaluating a new medical procedure before applying for a licence. Based on a verified procedure of experiments, she performs a first experiment on the viability of the medical procedure and depending on its outcome it may perform another experiment and so on, till enough information about the viability of the procedure is gathered to take a final decision on whether to apply for a licence. The sequence of experiments may exhibit history-dependence, in the sense that, what experiment becomes available and is appropriate at each stage depends on outcomes of previous experiments. This process takes time and the firm, which is impatient, is trading off the value of more information for decision making with the resulting delay in the realization of the benefits of the decision making. In fact, such trade-off considerations are not constrained to pharmaceutical research and are important for most R&D activities.

As a second example, consider a firm in the automotive industry which doesn't face impatience considerations and which is considering the introduction of a new production technology. She may seek the advice of an expert who can perform verifiable technical analysis of the new technology at a fee. The firm may seek the advice of the expert repeatedly so that information accumulates, before taking an informed decision on whether to introduce the new technology.

Both examples are unified by the fact that learning is costly, the technology for learning is inflexible and sequential. We say the agent is a *sequential sampler*. They differ in the type of costs the agent incurs. Whereas the costs in the case of a pharmaceutical firm come from impatience, the costs for the automotive firm come from acquiring advice repeatedly in the market. For ease in exposition we call in the following the case of the pharmaceutical firm the case of *costs from impatience* and the case of the automotive firm the case of *market costs* but, especially in the second case, different interpretations are possible.[1]

More precisely, in this paper we assume the analyst has access to random choice data over choice from a menu and decision time for the choice from same menu and allow variation in the menu of options the agent faces. We call this observable *Random Choice with Decision Times (RCDT)*. This observable allows for two interpretations. In the first one, the analyst is obtaining data from a single agent facing the same decision problem repeatedly, and randomness in choice and decision time comes from different learning outcomes. In the second one, the analyst faces a population of homogeneous agents who learn with the same technology and the randomness in choice and decision time comes from the heterogeneity in the experimental outcomes.[2]

The analyst in this paper is primarily interested in welfare analysis of different interventions but also in the identification of the information costs of the agent. For example, in the case of the pharmaceutical company the analyst might want to identify the welfare costs and the impact on learning of a lump sum income tax on the firm whereas in the case of the automotive firm the intervention to be evaluated could be subsidizing the market price for technical expertise. We show how both identification and welfare analysis are possible with RCDT without any knowledge of the sequential experiment the agent can access.

Our approach to the two problems of the analyst shows analogies to the classical consumer theory.[3] Just as in the classical theory, welfare analysis is

[1]One could also interpret the second case of costs as pertaining to a *patient* agent who performs a fixed experiment she herself has designed in a prior stage. The interpretation of market costs for the automotive firm is motivated by recent research on markets for information; see Bergemann and Bonatti [2018] for a survey.

[2]We note here that in the case of market costs 'decision time' is maybe best interpreted as the number of experiments the agent chooses to perform.

[3]See chapter 3 in Mas-Collel et al. [1995] and references therein.

closely tied to the identification of the parameters of the agent from the observables. As a first step in the welfare analysis we show how for each menu of options the analyst can recover the ex-ante valuation of the agent through random choice data augmented with decision time data. This is based on the analogue for our model of Roy's identity from classical consumer theory.[4] The formula in our set-up is 'simpler' than the classical one because our set up has the advantage that, since the agent follows subjective expected utility (SEU), random choice data contain cardinal information about agent's preferences. We establish several such 'Roy's identities' for different parameter changes in the decision environment of the agents. On a technical level, they all rely on the abstract envelope theorem arguments developed in Milgrom and Segal [2002].

The envelope theorem arguments lead to explicit formulas for the welfare of the agents in terms of their RCDT. One insight worthy of emphasis from the formulas is that in the case of market costs to recover the welfare of a single menu the analyst does not need decision time data. When the costs are from impatience the optimal decision time directly affects the value of a menu even for fixed decision time. As a consequence, to recover the welfare of a single menu for an agent who is impatient decision time data are indispensable. That decision time data may not be needed for the case of market costs remains overall a special case; e.g. they are indispensable in that same set up when evaluating welfare changes of a price change for experiments.

Another insight from the welfare calculations is, that subsidizing the duration of the experimentation by the agent, has different welfare implications based on whether the costs of learning are from impatience or are instead market costs. In the case of impatience, the subsidy needed to induce longer experimentation always costs more in expected discounted terms than the increase in welfare to the agent. This is because the optimal stopping decision of the agent already internalizes the impatience costs and the related trade-off between time and information. In the case of market costs, the supplier of experiments is a separate entity who does not internalize the negative effect that quoting a positive price

[4]See Proposition 3.G.4 in Mas-Collel et al. [1995] for the classical Roy's identity.

for experiments has on the optimal number of repetitions of experiments for the agent. Therefore, lowering the price of experiments and compensating the supplier of the experiments for the lost profits costs overall less than the resulting welfare increase for the agent.

We use the insights from welfare analysis to identify the costs of information. Before that step, the analyst identifies the prior and the (Bernoulli) taste of the agent by looking at choices from menus for which the agent decides not to learn and also at small perturbations thereof. For menus where the sequential experiment at the disposal of the agent has no additional decision value, the agent behaves as a SEU agent with a fixed taste and fixed information in the form of a prior. Therefore classical results of identification from ex-post choice apply. Once taste and prior are identified, we show how one can use the results from the welfare analysis to uniquely identify the costs of information from RCDT.

Identification of information costs is possible whenever there exists genuine learning in the data. This means that there exists a menu of options such that the agent has *strict* incentives to employ the sequential experiment when choosing from that menu. For the case of impatience costs, the analyst may identify the discount factor of the agent by looking at choices from two different perturbations of such a menu: 1) adding a safe option to the menu and 2) subsidizing choice from that menu in a lump-sum fashion unconditional on the ultimate choice of the agent. The two types of perturbations combined through one of the welfare results give an equation for the discount factor which has a unique solution, if the model of impatience costs is underlying the data.

In the case of market costs we use a two-step procedure for identification. First, we recover through perturbations of a menu for which the agent has strict incentives to learn, the ex-post random choice from *arbitrary* menus *under the condition that* the agent uses for *any* menu the optimal information acquisition strategy as for the original menu with strict incentives to learn. This 'fictional' random choice can be used to calculate the valuation of the original (unperturbed) menu with strict information incentives under the assumption that the agent is compensated *ex-post* for the incurred information costs. The difference between this fictional valuation for the original menu and the true one given by the RCDT as calculated through the welfare results, yields the average information costs. The average information costs are proportional to the expected decision time for the original menu. Because the decision time is also revealed from the RCDT, this leads to the identification of the constant marginal costs of experiments. In both cases of information costs, only perturbations of a single menu are needed to identify the costs of information. We calculate numerical examples for both cases of information costs. These illustrate that the identification procedure is numerically simple. We also give numerical examples related to welfare calculations illustrating the same degree of computational simplicity.[5]

A related question to identification is uncovering the set of conditions on RCDT data which ensure rationalizability through either of the sequential sampling models addressed in this paper. The analogue of this question in classical consumer theory is the *integrability* question: what conditions on a demand function imply that it arises from utility maximization.[6] In the last part of the paper we try to give a complete characterization of RCDT data coming from a sequential sampler. We first identify four axioms on RCDT which ensure that the identification procedures described above work. They have intuitive explanations. The first two axioms, *Sampling is costly* and *SEU behavior upon not sampling* ensure that taste and prior of the agent can be identified. The third axiom called *Taste stationarity* ensures that the taste of the agent does not change with time. On a technical level, once we pick a Bernoulli utility as identified from the first two axioms we can use the cardinality feature of these time-zero preferences to measure the exploration-exploitation trade-off the agent faces when deciding whether to continue or stop the sequential experiment. The quantification of this trade-off reveals the costs of information and is only meaningful if Taste stationarity is satisfied. Finally, the fourth axiom ensures that

[5]Nothing more than taking integrals of monotonic functions and simple algebraic calculations are needed from a computational perspective.

[6]See section 3.H in Mas-Collel et al. [1995].

the process of identifying the costs of information is well-defined in either case of impatience costs or market costs.

All of the first four axioms neither use nor try to identify the information flow of the sequential experiment. This implies that identification of the flow of information of the agent is unnecessary for the identification of the costs of information, or for welfare analysis in the two sequential sampling models we consider. Ultimately, rationalizability through the two sequential sampling models requires more conditions than the ones sufficient for identification. To deliver a full characterization we impose the additional assumption that the analyst knows the sequential experiment of the agent. This assumption is not unrealistic in some environments. Returning to the illustrative examples, testing of medical procedures may follow agreed-upon standards which are public knowledge. The same can be said for the technical analysis of new production technologies performed for a fee by an outside expert. Under this (admittedly in some cases strong) informational assumption, we introduce a final fifth axiom for the characterization, called *Data matching*. It implies consistency of the RCDT with an agent whose parameters are identified through the first four axioms. Data matching can be interpreted as a statistical test on the level of the sophistication of the agent with respect to her future behavior and the decision environment she is in.

Organization of the rest of the paper. In the next subsection we cover some decision theoretic literature related to this paper. The set up of the model is in section 3.2, the welfare results in section 3.3 and identification results in section 3.4. Section 3.5 offers numerical examples illustrating the identification and welfare results. Section 3.6 offers a behavioral characterization in terms of axioms. The last section concludes.

3.1.1 Related literature

There are by now many fields in economics in which sequential sampling models is used as a tool to describe agents in dynamic learning environments. We avoid here the insurmountable task of reviewing the applied theory literature related to sequential sampling models and instead focus on the theoretical one.

The agent's problem in a sequential sampling model is both *when* and *what* to choose (Wald [1947]; Arrow et al. [1949]; Fudenberg et al. [2018]). In more general models the agent is allowed to choose between experiments at each point in time (Che and Mierendorff [2018]; Hebert and Woodford [2018a]; Hebert and Woodford [2018b]; Zhong [2019]). All of these works assume parametric forms of the sequential experiment and focus on the resulting choices. In this paper, we take behavior of an agent or of a population of identical agents as given, and ask what we can learn from data about the parameters of the agent or of the population.

Our results highlight the usefulness of decision time data. In particular, we show that they are useful for identifying the discount factor/linear costs of an agent who performs sequential experiments and also for computing her welfare given a decision problem. Echenique and Saito [2017] also conduct a revealed preference analysis given data on choice and decision time. However, there is no uncertainty in their model and decision time is used to infer the difference between utilities of two alternatives.

If the agent faces the task of acquiring the same experiment repeatedly at a cost without delay and/or is patient, the cost of learning information is menu-independent. In that case, the model can be viewed as a special case of rational inattention models (De Oliveira et al. [2017]; Lin [2018]). As shown by Lin [2018], data in the form of random choice from menus are already sufficient to recover welfare and to identify the smallest information cost function consistent with the data in the case of rational inattention models. However, the identification strategy in Lin [2018] results in a complicated formula when applied to our set-up. Here we show how additional information about decision time allows for a simple procedure to identify the information costs.[7] The related papers Caplin and Dean [2011] and Denti [2018] show how information costs can be identified in the *static* rational inattention model when the analyst has

[7] Appendix C.5 contains the identification strategy without decision times based on Lin [2018].

access to *state-dependent* stochastic choice data and in addition knows the prior of the agent. Their approach allows to identify (a minimal version of) the information the agent acquires for each menu. This is then used to identify the cost function through duality arguments from convex analysis. In contrast, here we show in two dynamic settings of costly information how an analyst can identify from random choice data containing information about decision time, the costs of information without making the assumption that the analyst knows the prior of the agent. Our approach works without first identifying the information flow of the agent and analogous to classical works of consumer theory follows instead the route of envelope theorem arguments.

Time preferences and discounting of future payoffs have been studied in the literature under various frameworks, such as choice over consumption streams (Koopmans [1960]; Epstein [1983]), choice over dated rewards (Fishburn and Rubinstein [1982]), discrete choice models (Magnac and Thesmar [2002]). The novelty in our paper is that we consider characterizations of dynamic stochastic choice in the framework of sequential sampling. Here the agent does not face a decision problem of the type 'either take act f at time t or take act g at time s' as in the old literature but rather a decision of the type 'when and what to choose from a given, time-independent menu of options'.

In comparison to other recent work in dynamic stochastic choice (see Frick et al. [2018] or Duraj [2018a]) our data is limited in the sense that, the analyst does not observe data with recursive structure as those induced by temporal lotteries. Our agent does not face a dynamic decision problem with recursive structure as is the case in Frick et al. [2018] or Duraj [2018a].

Finally, in Duraj and Lin [2019] we cover the related case where the agent can only make use of a single experiment. This corresponds to switching off the optimal-stopping decision of the agent, which is central in this paper. We assume that the analyst observes random choice from menus, as well as whether the agent employs the experiment. We give a full axiomatic characterization result for the cases where the experiment is only available with a fixed amount of delay and the agent is impatient, or she acquires the experiment at a fixed, menu-independent cost. We show how this data is sufficient for the analyst to uncover not only the experiment the agent can employ, but also her information costs. Here we look at a set up which is genuinely dynamic so that the results from Duraj and Lin [2019] are not applicable (the experiment is sequential) and focus on the value of random choice and decision time data in welfare analysis. Moreover, our identification procedure for the learning costs of the agent works *without* identifying her information flow.

3.2 Set up

3.2.1 Notation and preliminaries.

Let $Z = \mathbb{R}_+$ be the prize space equipped with the topology of the Euclidean norm.[8]

Let S be a finite set of objective states, where the wording *objective* means that neither the analyst nor the agent can influence the realization of the state s. Let $\mathbb{F}$ be the set of Anscombe-Aumann *acts* with a typical element given by $f : S \to \Delta(Z)$ where $\Delta(Z)$ denotes the space of *simple* lotteries over prizes in Z.[9] Note that $\mathbb{F}$ has a cone structure, defined by scaling acts through multiplication of the prizes that can occur. Finally, denote by $\mathcal{A}$ the collection of finite, nonempty subsets of $\mathbb{F}$. A typical element in $\mathcal{A}$ is called a *menu* and denoted by a capital letter, e.g. $A, B \in \mathcal{A}$. $\mathcal{A}$ is a metric space equipped with the Hausdorff topology. In the following we identify constant acts with the state-independent lottery they offer. We also identify degenerate lotteries from $\Delta(Z)$ with the prize they offer with certainty. A menu A is called *constant* if it contains only constant acts. It is a *singleton* menu if it contains only one element.

Given a belief of the agent over S, i.e. an element π from $\Delta(S)$ and an Expected Utility function $u : \Delta(Z) \to \mathbb{R}$ to evaluate simple lotteries, we say the agent

[8] We take this prize space for ease in exposition. Results can be proven for general metric space Z.

[9] A lottery is called *simple* if only finitely many prizes can happen with positive probability. $\Delta(Z)$ is equipped with the topology of weak convergence of probability measures. The set of acts $\mathbb{F}$ is equipped with the product-topology over $\Delta(Z)^S$.

satisfies *Subjective Expected Utility (SEU) with beliefs q and Bernoulli taste u* if the utility of an act f is given by $\pi \cdot (u \circ f) := \sum_{s \in S} \pi(s) u(f(s))$.[10]

In the following we denote the value of an instantaneous decision problem/menu A for a SEU agent with taste u and belief π as

$$V(A, \pi) = \max_{f \in A} \pi \cdot (u \circ f).$$

Define also

$$M(A; u, \pi) = \{f \in A : \pi \cdot (u \circ f) \geq \pi \cdot (u \circ g), \forall g \in A\}.$$

This is the set of maximizers from A when the agent's belief about objective state of the world is π and her Bernoulli utility is u.

We equip the set $\mathcal{A}$ with a sum and multiplication operation based on the Minkowski sum. Thus in general, a scaled sum of menus $A, B \in \mathcal{A}$ with parameters $\lambda_1, \lambda_2 \in \mathbb{R}_+$ is given by

$$\lambda_1 A + \lambda_2 B = \{\lambda_1 f + \lambda_2 g : f \in A, g \in B\}.$$

In the case of Minkowski mixtures, $\lambda_2 = 1 - \lambda_1$ and $\lambda_1 \in (0, 1)$, one possible interpretation of $\lambda_1 A + (1 - \lambda_2) B$ is that the agent makes contingent plans: the uncertainty about the choice from the menu is resolved before the uncertainty about the state s and so she decides on what to pick from each of the two menus/contingencies that can happen. In detail, if the agent faces $\lambda_1 A + (1 - \lambda_1) B$ we are assuming under the contingencies interpretation that the analyst observes the contingent plan of choosing f if A gets realized and g if B, i.e. she observes that agent picks $\lambda_1 f + (1 - \lambda_1) g$. In terms of random choice this interpretation would presume the analyst observes the *likelihood* with which the agent picks $\lambda_1 f + (1 - \lambda_1) g$ from $\lambda_1 A + (1 - \lambda_1) B$.

[10]In the following we abuse notation for simplicity of exposition and often identify the EU-functional $u : \Delta(Z) \to \mathbb{R}$ with its Bernoulli utility from $\mathbb{R}^Z$.

3.2.2 The agent

Henceforth we denote $\mathcal{T}$ the set of possible decision times $\{0, \ldots, T\}$ for the agent with some finite $T \geq 2$. By time T the agent has to make a choice from the menu she faces.

Information. The agent possesses a prior $\pi_0 \in \Delta(S)$ with $\pi_0(s) > 0$ for every $s \in S$ and a (menu-independent) sequential experiment.

A sequential experiment is a random variable

$$\mathcal{E} : S \to \Delta(E_1 \times \cdots \times E_T), \tag{3.1}$$

where the sets E_i, $i = 0, \ldots, T$ are finite sets of experimental outcomes. We denote by $e^t := (e_1, \ldots, e_t)$ a history of length t of experimental outcomes. Every state of the world induces a probability distribution over outcomes in the sequential experiment. Denote by $\mathcal{E}_s(e_1, \ldots, e_t)$ the probability that the sequence of outcomes $e^t = (e_1, \ldots, e_t)$ is realized under the sequential experiment $\mathcal{E}$ if the state is s. For an element $e^T = (e_1, \ldots, e_T) \in E_1 \times \cdots \times E_T$ we write $\mathcal{E}_s(e^T)$ for its probability when the true state of the world is s. (3.1) allows for arbitrary correlation between the sequential outcomes of the experiment. In particular, it allows for the case that the experiment available to the agent at time $t \geq 3$ depends on the history of outcomes $(e_1, \ldots, e_{t-1})$ realized till that moment in time.

Given the outcomes of the experiments $(e_1, \ldots, e_t)$ till period t, if the Bayesian agent has prior $\pi_0 \in \Delta(S)$ over the hidden objective state s her posterior is given by

$$\pi(s|e_1, \ldots, e_t) = \frac{\mathcal{E}_s(e_1, \ldots, e_t)\pi_0(s)}{\sum_{s'} \mathcal{E}_{s'}(e_1, \ldots, e_t)\pi_0(s')}. \tag{3.2}$$

Suppose we have two sequences of experimental outcomes which under some (strictly positive) prior lead to the same sequence of posteriors. Formally, we have $e^T = (e_1, \ldots, e_T)$ and $e'^T = (e'_1, \ldots, e'_T)$ with $e^T \neq e'^T$ such that for some prior π_0 as above it holds

$$\pi(s|e^t) = \pi(s|e'^t), \quad 1 \leq t \leq T, s \in S. \tag{3.3}$$

Now this implies the following condition

$$\frac{\mathcal{E}_{s_1}(e^t)}{\mathcal{E}_{s_2}(e^t)} = \frac{\mathcal{E}_{s_1}(e'^t)}{\mathcal{E}_{s_2}(e'^t)}, \quad \forall s_1, s_2, 1 \leq t \leq T \qquad (ratio).$$

In fact, one can show that these are equivalent, namely, that for every distinct e^T, e'^T that satisfies $(ratio)$, (3.3) also holds true.[11] Thus, it follows that whenever (3.3) holds true for two distinct e^T, e'^T and some prior π_0 it will hold true for every prior. To avoid this redundancy we impose the following condition on $\mathcal{E}$ throughout the paper.

Non-redundant experiment. $\mathcal{E}$ is such that whenever $e^T \neq e'^T$ there exists $s, s' \in S$ and $1 \leq t \leq T$ such that

$$\frac{\mathcal{E}_s(e^t)}{\mathcal{E}_{s'}(e^t)} \neq \frac{\mathcal{E}_s(e'^t)}{\mathcal{E}_{s'}(e'^t)}.$$

Given prior $\pi_0 \in \Delta(S)$ Bayes rule implies a unique induced probability measure on sequential experimental outcomes. Denote this measure on $E_1 \times \cdots \times E_T$ by $\mu(\pi_0)$ and for a history $e^t = (e_1, \ldots, e_t) \in E_1 \times \cdots \times E_t$ denote $\pi(e^t)$ the belief from Bayes rule under $\mu(\pi_0)$ after history e^t of experimental outcomes. By standard results, $\mu(\pi_0)$ corresponds to a martingale belief process on $\Delta(S)^{T+1}$.

Information costs. We consider two special cases of information costs for the agent.

A. Information costs due to impatience. The agent determines through backwards induction the amount of information she gathers for a menu A using

[11] Pick $s_2 = \bar{s}$ and plug $(ratio)$ into (3.2).

the sequential experiment $\mathcal{E}$ at her disposal.[12]

If the agent is in next-to-last period $t = T - 1$ then with the definition

$$W_T(A, e^{T-1}) = \mathbb{E}_{e^T \sim \mu(\pi_0)(\cdot|e^{T-1})}[V(A, \pi(e^T))],$$

for the option value of continuing the sequential experiment till the last period the value function of the agent at time $T - 1$ is given by

$$\tilde{W}_{T-1}(A, e^{T-1}) = \max\{V(A, \pi(e^{T-1})), \delta \cdot W_T(A, e^{T-1})\}.$$

By backwards induction similar calculations hold for $1 \leq t \leq T - 1$. Define $W_{t+1}(A, e^t) = \mathbb{E}_{e^{t+1} \sim \mu(\pi_0)(\cdot|e^t)}[\tilde{W}_{t+1}(A, e^{t+1})]$ so that value function at time t is

$$\tilde{W}_t(A, e^t) = \max\{V(A, \pi(e^t)), \delta W_{t+1}(A, e^t)\}.$$

Finally, at $t = 0$ we arrive at

$$\tilde{W}_0(A, \pi_0) = \max\{V(A, \pi_0), \delta W_1(A)\},$$

for the value function of the agent at time $t = 0$ where $W_1(A) = \mathbb{E}_{e_1 \sim \mu(\pi_0)(e_1 \in \cdot)}[\tilde{W}_2(A, e_1)]$.

The value function at time 0 can also be written with help of optimal stopping strategies. Because the agent is dynamically consistent, she can be thought of as choosing a (randomized) optimal stopping time $\tau : E_1 \times \cdots \times E_T \to \Delta(\mathcal{T})$. This is a random variable satisfying the condition that the event $\{\tau = t\}$ depends only on $E_1 \times \cdots \times E_t$, i.e. only on experimental outcomes till time t.

For a stopping time τ, denote by $\mathcal{F}_\tau$ the induced sub-sigma algebra of the power set of $E_1 \times \ldots E_T$. Then we can write

$$\tilde{W}_0(A, \pi_0) = \sup_{\tau \text{ stopping time}} \mathbb{E}_{e^\tau \sim \mu(\pi_0)(\cdot|\mathcal{F}_\tau)}[\delta^\tau V(A, \pi(e^\tau))]. \tag{3.4}$$

[12]In the following $\mathbb{E}_{x \sim F}[G(x)]$ will generically denote the expectation of random variable $G(x)$ when x is distributed according to the probability measure F.

B. Additive costs of information This is similar to A. except for two important changes. We assume that $\mathcal{E}$ describes an i.i.d. sequence of experiments. That is, $E_t = E$ for all $t = 1, \ldots, T$ and the distribution $\mathcal{E}_s(e_{t+1}|e^t)$ is independent of t or history e^t. Our results hold for general sequential experiment as in (3.1) but the assumption of a constant flow cost of time is more sensible in the i.i.d. case. The calculation of the value functions proceeds as in A. with the substitution of the terms $\delta W_{t+1}(A, e^t)$ by $W_{t+1}(A, e^t) - c$ everywhere. The value function at time 0 now becomes

$$\tilde{W}_0(A, \pi_0) = \sup_{\tau \text{ stopping time}} \mathbb{E}_{e^\tau \sim \mu(\pi_0)(\cdot|\mathcal{F}_\tau)}[V(A, \pi(e^\tau)) - c\tau]. \tag{3.5}$$

Overall, an agent is prescribed by a tuple $\mathbb{A} = (u, \delta, \mathcal{E}, \pi_0)$ in case A. or $\mathbb{A} = (u, c, \mathcal{E}, \pi_0)$ in case B. Here, u, π_0 and $\mathcal{E}$ are respectively the Bernoulli taste, prior and the sequential experiment of the agent, whereas $\delta \in (0, 1)$ is the discount factor in case A. and $c > 0$ are the costs for the i.i.d. experiment in case B. In the following we call an agent $\mathbb{A} = (u, \delta, \mathcal{E}, \pi_0)$ as prescribed above a **SeSa-GD** and an agent $\mathbb{A} = (u, c, \mathcal{E}, \pi_0)$ as above a **SeSa-LC**.[13]

3.2.3 The analyst

Consider an analyst who observes data as we just described from an agent who can learn about a payoff-relevant objective state before picking from a menu by making use of a sequential experiment. An optimal stopping strategy for the agent induces a random decision time because different realizations of histories of experimental outcomes lead to different realizations of the stopping decision. Conditional on stopping at time t, the choice of act appears random to an outside analyst because the belief of the agent at that moment is random and unobservable. We assume that for every menu of acts A the analyst observes the joint probability of the choice of act from the menu as well as of the decision time $\tau \in \mathcal{T}$. We assume the analyst knows that the agent is Bayesian and that he

[13] GD stands for geometric discounting and LC for linear costs.

understands that if the additional information available to the agent were free, the agent would always use it. The analyst has in mind two possible cases of *costly* information described in the previous subsection: **SeSa-GD**, where the experimentation costs come from impatience and **SeSa-LC** where the agent can repeat an experiment for a fixed price. We assume w.l.o.g. that the analyst knows that $|\mathcal{T}| = T + 1$.[14]

Finally, we assume the analyst doesn't have information about the taste $u : \Delta(Z) \to \mathbb{R}$, discount factor δ/linear costs c.

For the welfare results we assume the analyst knows precisely which of the cases SeSa-GD or SeSa-LC is true as well as the discount factor of the agent/the additive costs. We justify this assumption in section 3.4.

The observable - formal definition. Formally we assume that for every menu $A \in \mathcal{A}$ the analyst possesses random choice data in the form $P_A \in \Delta(A \times \mathcal{T})$. We call the collection $\{P_A : A \in \mathcal{A}\}$ of such data *Random Choice with Decision Time*, in short *RCDT*.

Random choice as an observable may be interpreted in two distinct ways. In the *single-agent interpretation* the analyst observes the limiting frequency of choices as well as of decision times of a single agent coming from many, many repetitions of the same decision problem. In the *population interpretation* the analyst observes choices from menus as well as of decision times from a population of homogeneous individuals with the same taste, prior and costs of information having access to the same sequential experiment. Our results hold for both interpretations but we pick in the following the single-agent interpretation for ease of exposition.

When compared to other random choice data from other dynamic stochastic choice models (e.g. Frick et al. [2018], Duraj [2018a]) RCDT data are relatively

[14] This is w.l.o.g. because T is revealed from the data that we assume is available to the analyst. In the SeSa-GD case T is the last time the agent observes choices from *any* menu. In the SeSa-LC case, T is the overall maximal observed number of repetitions of the i.i.d. experiment. In the latter case this may be related to a time dimension or not. If it is, the implicit assumption in SeSa-LC is that the agent is patient.

scarce: for every menu A the analyst only observes when the agent makes a choice and what choice she makes. In particular, for every observation of the choices of a single agent from a menu A, the menu A does not change across time so that the recursivity of the observable as it appears in Frick et al. [2018] and Duraj [2018a] is lost. The main implication of this fact is that identifying the information flow of the agent (the sequential experiment $\mathcal{E}$) becomes a very difficult task, especially in the case of impatience costs.

Given an RCDT the analyst can determine the collection of menus where the agent has strict incentives not to learn and those where she has strict incentives to learn till period t with positive probability. Formally we define

$$\mathcal{A}_0 = \{A \in \mathcal{A} : \lim_{n\to\infty} P_{A_n}(\tau = 0) = 1 \text{ for every sequence } A_n \to A\},$$

as well as for $t \geq 1$

$$\mathcal{A}_t = \{A \in \mathcal{A} : \liminf_{n\to\infty} P_{A_n}(\tau = t) > 0 \text{ for every sequence } A_n \to A\}.$$

$\mathcal{A}_0$ is the collection of menus for which the agent has strict incentives not to start the learning process. If either of the models A. or B. are correct then this set is non-empty and contains all menus of constant acts as well as singleton menus.

We assume in the following that any RCDT satisfies the following degeneracy condition to varying degrees.

Condition N: Non-degenerate dynamic choice.

- **N1.** There exists $A \in \mathcal{A}$ such that $A \in \mathcal{A}_t$ for some $t \geq 1$,
- **N2.** There exists $A \in \mathcal{A}_t$ for some $t \geq 1$ and a prize k with $A + \{k\} \in \mathcal{A}_0$.

Condition N1. states that there exists a menu A where learning is strictly profitable. We always require this condition on any RCDT. Condition N2. is a strengthening of N1. and requires that if in addition a sufficiently high prize k is given for certain to the agent, unconditional on her choice from A, she has no incentives to acquire any information. N2. is only needed in the case of SeSa-GD and is a very weak requirement on the data. This is because, whenever SeSa-GD is true and non-trivial dynamic behavior is observed, N2. is automatically satisfied. Intuitively it says that, because of discounting, when k is large enough payoffs from $A + \{k\}$ are relatively insensitive to the realization of the objective state so that learning is not worthwhile.

N1. imposes joint restrictions on the taste of the agent, her sequential information as well as the discount factor or the additive costs of information. Namely, it is easy to construct examples where if the sequential experiment is overall very uninformative about the state of the world, the agent would never start any menu A.[15]

In section 3.4 we show that Condition N suffices for full identification of respectively discount factor of the agent (Case A.) or additive costs of information (Case B.).

One can show that Condition N is unnecessary, if the model includes prizes of arbitrarily negative utility.[16] Here we insist on limited liability of the agent, i.e. there is a lower bound for the utility of prizes and the 'stakes' the agent faces are limited from below. Lack of limited liability is especially untenable in experimental as well as most empirical settings the model of this paper aims to approximate.

In general there will be menus where the agent will have to break ties in either of the choices: when to stop and upon stopping, what to pick from the menu. This tie-breaking behavior will be incorporated in the observable P_A whenever A is such that ties occur with positive probability. We show in the appendix that data explained by either of SeSa-GD or SeSa-LC have the property that the set of menus where the agent needs to break ties is *meager* in the sense that any such menu can be approximated arbitrarily closely by a menu where there are no ties. Moreover, observing P_A for menus without ties is sufficient for the set of

[15] Examples are available upon request or the reader may try to construct them easily from the calculated examples presented in section 3.5.

[16] Proof is available upon request.

identification and characterization results we offer. Therefore we decide to remain agnostic about the precise tie-breaking behavior of the agent and focus only on menus without ties in the remainder of the main body of the paper. The appendix establishes the denseness result for menus without ties as well as describes its technical ramifications.

Formally, we show in the appendix that menus without ties are those that satisfy the following definition.

Definition 7. Say that a menu $A \in \mathcal{A}$ *is without ties* if there exists a neighborhood $\mathcal{N}_A$ of A in $\mathcal{A}$ such that for every sequence $A_n \in \mathcal{N}_A$ with $A_n \rightarrow A$ and $|A_n| = |A|$ for all n it holds uniformly for $t \in \{0, \dots, T\}$ that $P_{A_n}(f_n, t) \rightarrow P_A(f, t)$ if $f_n \in A_n$ and $f_n \rightarrow f \in A$.

For a menu without ties, the optimal stopping strategy of the agent in either (3.4) or (3.5) is unique and deterministic. Denote it in either case τ_A; whether we are in the case of impatience or market costs will be clear from context. Thus, for a menu without ties $\tau_A : E_1 \times \cdots \times E_T \rightarrow \mathcal{T}$ is a $\mathcal{T}$-valued random variable which is adapted to the natural filtration of the space of experimental outcomes.

The following definition gives the rationalizability concept we use in this paper.

Definition 8. 1) A RCDT $\{P_A : A \in \mathcal{A}\}$ corresponds to a SeSa-GD agent $\mathbb{A} = (u, \delta, \mathcal{E}, \pi_0)$ if for all menus $A \in \mathcal{A}$ without ties it holds

$$P_A(f, \tau = t) = \mu(\pi_0) \left(e^T \in E_1 \times \cdots \times E_T : \tau_A(e^T) = t, M(A; u, \pi(e^{\tau_A(e^T)})) = \{f\} \right),$$

where τ_A is the unique optimal stopping time in (3.4).

2) A RCDT $\{P_A : A \in \mathcal{A}\}$ corresponds to a SeSa-LC agent $\mathbb{A} = (u, c, \mathcal{E}, \pi_0)$ if for all menus $A \in \mathcal{A}$ without ties it holds

$$P_A(f, \tau = t) = \mu(\pi_0) \left(e^T \in E_1 \times \cdots \times E_T : \tau_A(e^T) = t, M(A; u, \pi(e^{\tau_A(e^T)})) = \{f\} \right),$$

where τ_A is the unique optimal stopping time in (3.5).

This definition says that data explained by a SeSa agent, be it in the case A. or B., need to fulfill the consistency condition that the probability of observing the agent pick f from A at time t corresponds to the probability that under the sequential experiment $\mathcal{E}$ the agent with prior π_0 and taste u observes histories of outcomes which rationalize stopping at time t for menu A and picking f upon stopping at that time t.

3.3 Welfare Analysis

In this section we show how welfare considerations are feasible in either SeSa-GD or SeSa-LC for an analyst who is oblivious of the agent's information. We assume that the analyst besides having access to a RCDT, also knows the discount factor $\delta \in (0, 1)$ in case of impatience costs as well as the flow costs of time $c > 0$ in case of market costs.[17] For ease of exposition only, we also assume the analyst knows the Bernoulli taste of the agent $u : Z \rightarrow \mathbb{R}_+$.

We show how such an analyst can gauge welfare by use of the RCDT data without any knowledge of the prior of the agent π_0 or the sequential experiment $\mathcal{E}$ in the possession of the agent.

In the next section we show how all of c or δ as well as u can be identified from RCDT alone. So that the main insight of our welfare analysis in this section, when coupled with the analysis in section 3.4, is that data in the form of RCDT are enough to conduct welfare analysis.

In the following we state the results for efficiency of exposition in terms of utility acts and menus of utility acts, which we now introduce.

We can use u to go over to *utility acts* by considering the prize space $u(Z) = \{u_0(z) : z \in Z\}$ instead of Z and defining acts to give lotteries over utilities as consequences. Thus, to an act $f \in \mathbb{F}$ corresponds the utility act $\tilde{f}$ defined as $\tilde{f}(s) = u \circ f(s), s \in S$. Note then that constant utility acts $a \in u(Z)$ are ranked in terms of the utility they provide: $a > b$ if and only if a corresponds to the utility of a prize $z' \in Z$ and b corresponds to $z \in Z$ and so that z' is strictly better than z for the agent. The same construction can be done for menus: one can consider menus of utility acts $\tilde{A} = \{u \circ f : f \in A\}$ corresponding to $A \in \mathcal{A}$.

[17] In particular, the analyst knows which case he is facing.

We also impose in the following the assumption that there is a worst prize and use the convention $u(w) = 0$ for the worst prize w of the agent.

Denote by $V_\delta(A)$ the welfare of the SeSa-GD agent when her discount factor is δ and she faces menu A. Similarly, let $V_c(A)$ the welfare of the SeSa-LC agent when her cost of repetition of the i.i.d. experiment is $c > 0$ and she faces menu A.

Theorem 4. *It holds for every menu A*

$$V_\delta(A) = \int_0^\infty \left(1 - \sum_{t=0}^{T} \delta^t P_{A\cup\{a\}}(a,t)\right) da,$$

and

$$V_c(A) = \int_0^\infty P_{A\cup\{a\}}(A) da.$$

Sketch of Proof. The proof is based on the envelope theorem in Milgrom and Segal [2002, Theorem 1 and 2]. Here, we provide a sketch of the proof for SeSa-GD model. The proof is then completed in the appendix. Recall from (3.4) that

$$V_\delta(A) = \sup_{\tau \text{ stopping time}} \mathbb{E}_{e_\tau \sim \mu(\pi_0)(\cdot|\mathcal{F}_\tau)}[\delta^\tau V(A, \pi(e^\tau))], \tag{3.6}$$

for all menus $A \in \mathcal{A}$ if the agent corresponds to $\mathbb{A} = (u, \delta, \mathcal{E}, \pi_0)$. One can show easily that the *sup* is actually a *max*, i.e. the supremum is attained for every menu A.

Fixing a menu A, we look at the menus of the type $A \cup \{r\}$ where r is a constant (utility) act. View $V_\delta(A \cup \{r\})$ as a function of r. Since A is finite, there exists $R > 0$ such that R dominates every act f in A. Note that $r \to V_\delta(A \cup \{r\})$ is weakly increasing, continuous and becomes $V_\delta(A \cup \{r\}) = r$ whenever r is higher than the highest utility prize an agent can get from A.[18]

First, one argues that $V_\delta(A \cup \{R\})$ can be expressed as the integral of its derivative:

[18] Continuity follows from the classical Berge maximum theorem.

$$V_\delta(A \cup \{R\}) = V_\delta(A \cup \{0\}) + \int_0^R \frac{d}{dr} V_\delta(A \cup \{r\}) dr.$$

One establishes then that the derivative of the value of menu $A \cup \{r\}$ with respect to r for an instantaneous decision from menu $A \cup \{r\}$ when the agent's belief π about the state of the world is distributed according to some $\bar{\pi} \in \Delta(\Delta(S))$ is given by

$$\frac{d}{dr} \mathbb{E}_{q \sim \bar{\pi}} \left[V(A \cup \{r\}, q) \right] = \bar{\pi} \left(q \in \Delta(S) : M(A \cup \{r\}; id, q) = \{r\} \right). \tag{3.7}$$

Here, $id : u(Z) \to u(Z), id(\tilde{z}) = \tilde{z}$ is the identity function on the utility space $u(Z)$. Thus this derivative is equal to the probability under $\bar{\pi}$ that a posterior belief results which rationalizes the choice of r from $A \cup \{r\}$.

Now note that we can decompose the value of using a specific stopping time τ for the agent as follows.

$$\mathbb{E}_{e^\tau \sim \mu(\pi_0)(\cdot|\mathcal{F}_\tau)} [\delta^\tau V(A, \pi(e^\tau))] = \tag{3.8}$$

$$\sum_{t \in \mathcal{T}} \mu(\pi_0)(\tau = t) \delta^t \mathbb{E}_{e^\tau \sim \mu(\pi_0)(\cdot|\tau = t)} [V(A, \pi(e^\tau))]. \tag{3.9}$$

Pick an optimal stopping time $\tau_{A \cup \{r\}}$ for menu A. The general envelope theorems from Milgrom and Segal [2002] are applicable and deliver, together with an interchange of derivatives and sums, that the derivative of $V_\delta(A \cup \{r\})$ with respect to r is given by

$$\frac{d}{dr} \mathbb{E}_{e^{\tau_{A \cup \{r\}}} \sim \mu(\pi_0)(\cdot|\mathcal{F}_{\tau_{A \cup \{r\}}})} \left[\delta^{\tau_{A \cup \{r\}}} V(A \cup \{r\}, \pi(e^{\tau_{A \cup \{r\}}})) \right] \tag{3.10}$$

$$= \sum_{t \in \mathcal{T}} \mu(\pi_0)(\tau_{A \cup \{r\}} = t) \delta^t \frac{d}{dr} \mathbb{E}_{e^{\tau_{A \cup \{r\}}} \sim \mu(\pi_0)(\cdot|\tau_{A \cup \{r\}} = t)} \left[V(A \cup \{r\}, \pi(e^{\tau_{A \cup \{r\}}})) \right].$$

That is, it is given by the derivative of the objective of the agent as described in (3.8), evaluated at the optimal stopping strategy $\tau_{A \cup \{r\}}$.

Now we can combine (3.10) and (3.7). We use for each $t \in \mathcal{T}$ as $\bar{\pi}$ the distribution of the posteriors of the agent under the optimal strategy $\tau_{A\cup\{r\}}$ under the condition that $\tau_{A\cup\{r\}} = t$, to arrive at

$$V_\delta(A \cup \{R\}) = V_\delta(A \cup \{0\}) + \int_0^R \sum_{t\in\mathcal{T}} \delta^t P_{A\cup\{r\}}(r,t)dr.$$

Pick R such that it dominates every utility act in A. It follows, $V_\delta(A \cup \{R\}) = R$. Since 0 is the worst utility act, $V_\delta(A \cup \{0\}) = V_\delta(A)$. Hence,

$$V_\delta(A) = R - \int_0^R \sum_{t\in\mathcal{T}} \delta^t P_{A\cup\{r\}}(r,t)dr = \int_0^R \left(1 - \sum_{t\in\mathcal{T}} \delta^t P_{A\cup\{r\}}(r,t)\right) dr.$$

We can take the upper bound of the integral to be infinity because the integrand becomes 0 for all $r > R$. □ The approach of recovering the indirect utility of the agent for a menu A or synonymously her ex-ante valuation for A through ex-post random choice first appears in Lu [2016]. In his model, private information is exogenous and menu-independent. Then the probability of r being chosen from $A \cup \{r\}$ turns out to be the probability of deriving utility at most r from A. Hence, by varying r, one obtains the distribution function for the utility from A. The integral is therefore the subjective expected utility from choosing from A. Thus, in the case of static, exogenous information the recoverability result can be proven without invoking technical envelope theorem arguments. However, this proof intuition fails in general when private information is menu-dependent as it is in our set up since it is the optimal stopping strategy of the agent for the menu A which determines the information she uses to pick from A.[19]

[19] When private information is optimally acquired as in a static rational inattention model (De Oliveira et al. [2017]), Lin [2018] shows that a similar envelope theorem argument applies and ex-ante valuation of a menu can be recovered in a similar way. That recoverability result in a static setting may be used to prove recoverability of welfare in SeSa-LC model, because the latter is equivalent to a static rational inattention model as in De Oliveira et al. [2017]. Here we apply a similar proof technique in the dynamic setting of SeSa-GD.

Theorem 4 shows that decision time data is not needed to recover welfare in the SeSa-LC model, but they are crucial in the SeSa-GD model. Intuitively, when changing the decision problem by increasing slightly a safe option the distribution of the optimal decision time does not change on the margin. In the SeSa-LC model decision time does not affect the value of the menu conditional on stopping at a fixed period with a fixed amount of information, but it does so in the SeSa-GD model. Therefore the envelope theorem argument gives in the SeSa-LC case a derivative w.r.t. the safe option which is independent of decision time, whereas the decision time still appears in the derivative in the case of SeSa-GD.[20]

In the following we consider several interesting cases of welfare analysis which illustrate the power of RCDT for an outside analyst interested in welfare analysis of agents who are performing sequential experiments.

We start with the case of a SeSa-GD who faces a menu of options A and of an outside party who has the option of offering the agent a one-time lump sum utility transfer $k > 0$, or of imposing a one-time lump sum tax $-k$. The tax or the subsidy occurs at the moment the agent picks from the menu, i.e. it occurs with delay if the agent decides to learn. Effectively, the outside party is changing the menu of the agent from A to $A \pm k$.[21] In the pharmaceutical R&D example this outside party may be thought of as a government or regulatory body.

Proposition 26. *Let $k > 0$ be a constant utility act. Then it holds*

$$V_\delta(A+k) = V_\delta(A) + \sum_{t=0}^{T} \delta^t \int_0^k P_{A+\lambda}(\tau = t)\,d\lambda. \tag{3.11}$$

Moreover, if k is weakly lower than the worst utility prize achievable in A it holds

$$V_\delta(A-k) = V_\delta(A) + \sum_{t=0}^{T} \delta^t \int_0^{-k} P_{A-\lambda}(\tau = t)\,d\lambda.$$

[20]Henceforth, *w.r.t.* denotes *with respect to*.

[21]We use the shortcut $A \pm k$ for $A \pm \{k\}$ in the following when A is a menu of utility acts.

To see the implications of Proposition 26 we first show the following Claim.

CLAIM 1. The function $\lambda \mapsto \mathbb{E}[\delta^{\tau_{A+\lambda}}] = \sum_{t=0}^{T} \delta^t P_{A+\lambda}(\tau = t)$ is weakly increasing in $\lambda > 0$ almost everywhere.[22]

Proof of Claim 1. Recall the valuation for an instantaneous decision from A under belief π given by $V(A, \pi) = \max_{f \in A} \pi \cdot f$. It satisfies the property

$$V(A + \lambda, \pi) = \lambda + V(A, \pi), \text{ for every } \lambda \geq 0.$$

In the following denote for simplicity of exposition μ_τ the random posterior induced by a stopping strategy τ for the sequential experiment the agent possesses. Then we have

$$V_\delta(A + \lambda) = \mathbb{E}[\delta^{\tau_{A+\lambda}} V(A, \mu_{\tau_{A+\lambda}})] + \lambda \mathbb{E}[\delta^{\tau_{A+\lambda}}].$$

Here $\tau_{A+\lambda}$ is an optimal stopping strategy for the menu $A + \lambda$ and expectations are w.r.t. the random realizations of $\tau_{A+\lambda}$ and $\mu_{\tau_{A+\lambda}}$. It holds by revealed preference

$$\mathbb{E}[\delta^{\tau_{A+\lambda}} V(A, \mu_{\tau_{A+\lambda}})] + \lambda \mathbb{E}[\delta^{\tau_{A+\lambda}}] \geq \mathbb{E}[\delta^{\tau_{A+\lambda'}} V(A, \mu_{\tau_{A+\lambda'}})] + \lambda \mathbb{E}[\delta^{\tau_{A+\lambda'}}]. \quad (3.12)$$

By combining the inequality (3.12) for the optimality in $A + \lambda$ with its analogue for the optimality in $A + \lambda'$ we get

[22]Note that in general there is no uniform monotonicity of the optimal stopping time $\tau_{A+\lambda}$ in λ. This is because stopping earlier after every history can lead to less information on average, even though the decision utility is discounted less. Therefore any of the Topkis's theorems (see subsection 2.8.1 in Topkis [1998]) is not applicable.

$$\frac{1}{\lambda}\left(\mathbb{E}[\delta^{\tau_{A+\lambda}}V(A,\mu_{\tau_{A+\lambda}})]-\mathbb{E}[\delta^{\tau_{A+\lambda'}}V(A,\mu_{\tau_{A+\lambda'}})]\right)$$
$$\geq \mathbb{E}[\delta^{\tau_{A+\lambda'}}]-\mathbb{E}[\delta^{\tau_{A+\lambda}}]$$
$$\geq \frac{1}{\lambda'}\left(\mathbb{E}[\delta^{\tau_{A+\lambda}}V(A,\mu_{\tau_{A+\lambda}})]-\mathbb{E}[\delta^{\tau_{A+\lambda'}}V(A,\mu_{\tau_{A+\lambda'}})]\right).$$

From here Claim 1 easily follows. □

One can use Proposition 26 and Claim 1. to gauge welfare change coming from a lump-sum subsidy of k utils. For example, one easily gets the bounds

$$k\mathbb{E}[\delta^{\tau_A}]\leq V_\delta(A+k)-V_\delta(A)\leq k\mathbb{E}[\delta^{\tau_{A+k}}].$$

In particular, if k is such that $P_{A+k}(\tau=0)<1$ and $P_A(\tau\geq 1)>0$ we have that ex-ante welfare increases by strictly less than k. The intuition for this is that because the agent chooses to learn when facing $A+k$, she gets the benefit from $A+k$ only with delay. This discounts also k for which learning has actually no benefit.

Alternatively, we can look at a lump-sum tax. Let's take a specific A so that every prize in it is strictly positive. In this case, we get for a tax $k>0$ such that every utility lottery in $A-k$ is non-negative the following relation.

$$k\mathbb{E}[\delta^{\tau_{A-k}}]\leq V_\delta(A)-V_\delta(A-k)\leq k\mathbb{E}[\delta^{\tau_A}]. \tag{3.13}$$

If A is such that $P_A(\tau=0)<1$ the fall in welfare is less than k. Intuitively, the agent can postpone the payment of the tax by starting to learn. This lowers revenues for the tax authority and mitigates some of the welfare effects of the tax.

Another intuition delivered by (3.13) is the following. Suppose the tax is paid at the moment of the transaction and that both agent and analyst have the same discount factor as well as same taste u. Then it holds per above that

$$V_\delta(A-k)+k\mathbb{E}[\delta^{\tau_{A-k}}]\leq V_\delta(A).$$

The term $k\mathbb{E}[\delta^{\tau_{A-k}}]$ are the expected tax-revenues which can be thought of as the welfare of the tax authority. The difference

$$V_\delta(A) - V_\delta(A-k) + k\mathbb{E}[\delta^{\tau_{A-k}}] \geq 0$$

is then the dead-weight loss from the lump-sum taxation. It is strictly positive whenever the function $s \mapsto \mathbb{E}[\delta^{\tau_{A-s}}]$ is strictly monotone for a range of s. This dead-weight loss comes from the effect of the taxation on learning. Namely the lump-sum taxation increases the sensitivity of the agent to the realization of the state of the world and therefore the agent starts to learn more often and for longer periods of time. Therefore the agent gets the benefit from the menu A *ceteris paribus* later. Random choice data in the form of RCDT give a way for the tax authority to calculate this welfare loss.

Finally, one may wonder about comparative statics of scaling the prizes up or down by a factor, i.e. changing the menu the agent faces from A to λA. As one can easily see, optimal learning behavior is unaffected by this type of change in the case of SeSa-GD so that the welfare analysis is trivial. It holds $V_\delta(\lambda A) = \lambda V_\delta(A)$.

We now turn to another type of intervention in the case of SeSa-GD: subsidizing the duration of the experimentation. In this case the analyst transfers an amount of $k > 0$ utils for every period that the agent continues with the sequential experiment. If the agent faces menu A and uses optimal stopping time $\tau_{A,k}$ the subsidy costs $k\mathbb{E}[1 + \delta + \cdots + \delta^{\tau_{A,k}}] = k\mathbb{E}\left[\frac{1-\delta^{\tau_{A,k}+1}}{1-\delta}\right]$. Denote by $V_\delta^k(A)$ the value function for the SeSa-GD agent when duration of the experimentation is subsidized by k utils per unit of time. Envelope theorem arguments lead to the following Proposition.

Proposition 27. *Let $k > 0$ be a constant utility act. Then it holds*

$$V_\delta^k(A) = V_\delta(A) + \frac{1}{1-\delta}\int_0^k \mathbb{E}\left[1 - \delta^{\tau_{A,\lambda}+1}\right] d\lambda.$$

Just as for the case of a lump-sum subsidy, the following claim comes from the usual revealed preference logic.

Claim 2. The function $\lambda \mapsto \mathbb{E}[1 - \delta^{\tau_{A,\lambda}}]$ is weakly increasing in $\lambda > 0$.

This Claim and Proposition 27 then imply that subsidizing the duration of the experimentation in the SeSa-GD case is not welfare-enhancing from an ex-ante perspective. Formally, it holds that

$$V_\delta^k(A) - V_\delta(A) - k\mathbb{E}\left[\frac{1 - \delta^{\tau_{A,k}+1}}{1-\delta}\right] \leq 0, \tag{3.14}$$

with strict inequality whenever $\lambda \mapsto \mathbb{E}[1 - \delta^{\tau_{A,\lambda}}]$ is strictly increasing for a range of λ.

We now turn to the SeSa-LC model. Another envelope theorem argument yields the following result.

Proposition 28. *Let $\{P_A^c : A \in \mathcal{A}\}$ be the stochastic choice rule of the agent when she has linear costs of experimenting equal to $c > 0$. Then it holds for the welfare of the agent*

$$\frac{d}{dc} V_c(A) = -\sum_{t=0}^{T} t P_A^c(\tau = t).$$

In particular, for $c > c'$ we have

$$V_{c'}(A) = V_c(A) + \int_{c'}^{c} \sum_{t=0}^{T} t P_A^s(\tau = t) ds. \tag{3.15}$$

Using the classical Theorem 2.8.2 in Topkis [1998] one can prove the following easy claim for the SeSa-LC. It also implies that the integrand in (3.15) is weakly decreasing in the costs s.

Claim 3. For $c > c'$ it follows $\tau_A^c \leq \tau_A^{c'}$ where τ_A^c is an optimal stopping time for the menu A when linear costs are c. In particular, $\sum_{t=0}^{T} t P_A^c(\tau = t) \leq \sum_{t=0}^{T} t P_A^{c'}(\tau = t)$.

From Claim 2 and (3.15) it follows that

$$V_{c'}(A) - V_c(A) - c\sum_{t=0}^{T} tP_A^c(\tau = t) + c'\sum_{t=0}^{T} tP_A^{c'}(\tau = t) \geq 0. \tag{3.16}$$

Here we have strict inequality whenever $s \mapsto \sum_{t=0}^{T} tP_A^s(\tau = t)$ is strictly decreasing in some range of (c', c). Consider again an outside party whose Bernoulli taste is the same as for the agent and assume he is interested in partly subsidizing the costs of the i.i.d. experiment for the SeSa-LC agent. If she lowers them from c to c' the expected cost of the intervention to the supplier of experiments is $c\sum_{t=0}^{T} tP_A^c(\tau = t) - c'\sum_{t=0}^{T} tP_A^{c'}(\tau = t)$. It follows that the cost subsidy is overall welfare-enhancing and Proposition 28 can be used to quantify the welfare gain.

Note that the conclusion above is different from the case of SeSa-GD where subsidizing the duration of the experimentation is not overall welfare-enhancing. This is because in the SeSa-GD case the costs of experimentation come from impatience and thus are a feature of the preferences: the agent is trading off time with more information and she cares about both dimensions. In the case of SeSa-LC the costs of the experimentation are *market* costs and the seller of the experiments does not internalize the negative effect of a positive price of the experiments on the optimal experimentation duration. Therefore in the case of SeSa-LC there is scope for an overall welfare-enhancing intervention in the market.

Finally, one may wonder about the comparative statics of the type of Proposition 26 in the case of SeSa-LC. As one can easily see, optimal learning behavior is unaffected by this type of change in the case of SeSa-LC so that the result is trivial: a lump-sum payment doesn't change optimal stopping behavior and just constitutes a lump-sum transfer from the analyst to the agent.

3.4 Identification

In this section we show how an analyst who possesses data in the form of RCDT can identify all the parameters of the agent, except her information, namely we show how she can identify the taste, the prior as well as the discount factor in the case of SeSa-GD and the additive costs of the i.i.d. experiment in the case of SeSa-LC. The identification techniques use extensively the insights from section 3.3 on welfare analysis but ultimately the only data used is RCDT.

We gather together all the identification results presented in this section in the following Theorem. This result formally justifies the statement made in section 3.3 that the welfare analysis in both of the SeSa models presented in that section does not need more than the observable assumed, the RCDT data. Note that the identification result does not need any information about the sequential experiment of the agent.

Theorem 5. *(i) Suppose that the RCDT P is rationalized by SeSa-GDs* (u, δ, π_0) *and* (u', δ', π'_0)*. Suppose that Condition N holds. Then u is an affine transformation of* u'*,* $\pi_0 = \pi'_0$ *and* $\delta = \delta'$*.*

(ii) Suppose that the RCDT P is rationalized by SeSa-LCs (u, c, π_0) *and* (u', c', π'_0)*. Suppose that Condition N1. holds. Then u is an affine transformation of* u'*,* $\pi_0 = \pi'_0$ *and* $c = c'$*.*

In the following subsections the identification result will be proved in text by explaining also the identification procedure of the analyst. The analyst will first identify the Bernoulli taste and the prior of the agent and use this information to identify the rest of the parameters (except for the sequential experiment).

3.4.1 Identification of prior and taste

The identification of prior and taste follows standard procedures modified to take into account the dynamic incentives of the agent. The analyst focuses on menus where the agent doesn't have incentives to learn to extract the taste and the prior.

Examples of these are singleton menus and menus of constant acts where learning about the state of the world does not have decision value and so agent chooses not to learn given that learning is costly. The taste can be uniquely determined by looking at menus of constant acts. In fact, looking only at binary menus suffices to identify the Bernoulli utility u. Given the knowledge of the taste u we focus in the following on utility acts w.r.t. u for the exposition.

Once the Bernoulli taste has been identified the analyst can extract the prior as follows.[23]

The analyst can perturb a menu with a single constant act with a bet on a particular state $s \in S$. The perturbation needs to be small enough so that the agent, given the costs of learning, still decides not to learn for the perturbed menu. Formally, we define as follows.

Consider the utility bet act on state s: $f_s := (0, \ldots, \underset{s-\text{th place}}{1}, \ldots, 0)$, i.e. 1 util is awarded in state s; otherwise agent gets worst prize which has utility zero. a denotes the constant act giving a utils in every state. Then for all $\lambda \in (0, 1)$ near enough to 1 the probability of state s under the prior is given as follows.

$$\pi_0(s) = \frac{1}{1-\lambda} \left(\int_0^\infty P_{\{(1-\lambda)f_s + \lambda\mathbf{1}, a\}}((1-\lambda)f_s + \lambda\mathbf{1} | \tau = 0) da - \lambda \right). \quad (3.17)$$

This is just a reformulation of the identity $V_{\delta/c}(\{(1-\lambda)f_s + \lambda\mathbf{1}\}) = (1-\lambda)\pi_0(s) + \lambda$ for all λ near enough to 1.

The qualification 'near enough to 1' for the perturbation factor $\lambda \in (0, 1)$ is needed to make sure the agent doesn't decide to learn for the menu $\{(1-\lambda)f_s + \lambda\mathbf{1}, a\}$ for all $a \geq 0$ and thus only uses her prior as information for the choice from $\{(1-\lambda)f_s + \lambda\mathbf{1}, a\}$. This qualification is possible because the learning is costly in both SeSa models: $\lambda > \delta$ in SeSa-GD or $\lambda > 1 - c$ in SeSa-LC is sufficient.

Because the data needed to identify the taste and the prior of the agent use information which is costless for the agent, the techniques for their identification

[23] This is by no means the only procedure, but the one we found most elegant for exposition.

are formally based on results from Lu [2016].

3.4.2 Identification of discount factor.

Now we come to the identification of the discount factor for the SeSa-GD. We assume the analyst has already identified the prior and the taste as in subsection 3.4.1.

Take a menu $A \in \mathcal{A}_t$ for some $t \geq 1$. Denote by $b(A) > 0$ the highest prize that can happen under any act from A. Note that if we take $k + V(\pi_0, A) \geq \delta(k + b(A))$, which is equivalent to $k \geq \frac{\delta}{1-\delta}(b(A) - V(\pi_0, A))$, the agent will not start for the menu $A + k$. This is because an upper bound for the benefit of learning for the menu $A + k$ is $\delta(k + b(A))$ and so when k is large enough this benefit is smaller than $V(A + k, \pi_0)$, the value of picking from $A + k$ without learning. Thus, we have a menu $A \in \mathcal{A}_t$ for some $t \geq 1$ and some $k > 0$ so that $k + A \in \mathcal{A}_0$. But the existence of such a menu is already postulated by Condition N.

One could think that from a practical perspective k as above may be too large, especially for patient agents where δ is near 1. But recall that the whole model is homogeneous of degree one in the space of utility acts.[24] So by scaling A down, we can also make the needed k small.

Now we use Theorem 4 and Proposition 26 to write the welfare result (3.11) for A and k as above as follows.

$$k+\max_{f\in A} \pi_0 \cdot (u \circ f) = \int_0^{\infty} \left(1 - \sum_{t=0}^{T} \delta^t P_{A\cup\{a\}}(a, t)\right) da + \sum_{t=0}^{T} \delta^t \int_0^k P_{A+a}(\tau = t) da.$$

The left hand side is the value of $A + k$ as calculated by the data, whereas the right hand side is just the right hand side of (3.11) where we have replaced the valuation for A through its formula from Theorem 4.

[24] In particular, $\mu b(A) = b(\mu A)$, $V(\mu A, \pi_0) = \mu V(A, \pi_0)$ for every $\mu > 0$ and this property propagates to all value functions after every history of experimental outcomes.

The right-hand side is strictly increasing and continuous as a function of $\delta \in [0, 1]$ whenever there exists some range of a-s and a $t > 0$ s.t.

$$P_{A\cup\{a\}}(\tau = t) + P_{A+a}(\tau = t) > 0.$$

But this is always the case for very small a since menu A was chosen to be in $\mathcal{A}_t$ for some $t \geq 1$, i.e. it offers strict incentives for learning.[25] By the mean value theorem for continuous functions there exists then a unique $\delta \in (0, 1)$ satisfying the equation.

The identification strategy above is the most general in the sense that it just needs a menu satisfying Condition N in section 3.2 and the latter condition in the SeSa-GD model is equivalent to observing some non-trivial dynamic choice in the respective RCDT. Other identification strategies based on Theorem 4 are possible when menu A has a specific form. The following (sufficient) condition describes such a case.

Condition 1. There exists a state $s \in S$, $a > 0$, and $t \geq 1$ such that $P_{\{f_s,a\}}(a, t) > 0$.

Thus the agent decides to learn some of the time when faced with the choice between a bet on a state and a sure option. Under Condition 1 one can use the following identity for identification.

$$\sum_{s\in S} \left(\int_0^1 \left(1 - \sum_{t=0}^{T} \delta^t P_{\{f_s,r\}}(r, t) \right) dr \right) = \sum_{s\in S} \pi_0(s) = 1. \tag{3.18}$$

To immediately see the relation is true note that every element of the sum on the left hand side of (3.18) is just the value of the bet f_s, which is equal to π_s. Condition 1 together with the mean value theorem for continuous functions then implies the uniqueness of the discount factor.

[25] That $V_\delta(A)$, which corresponds to the first integral on the right hand side, is increasing in δ follows from the model assumption. Formally, this follows immediately from (3.4) because utilities are assumed non-negative.

3.4.3 Identification of additive costs.

We now turn to the identification of additive costs in SeSa-LC.

Fix $A \in \mathcal{A}_t$ for some $t \geq 1$. Intuitively, the costs of learning for A are given as the difference between the valuation of A under the optimal learning strategy for A, if using that particular strategy would be costless for the agent, and the *actual* value of A for the agent as revealed through RCDT data in Theorem 4. We first describe how to tease out from RCDT the value of A under the assumption that its optimal learning strategy is costless.

For each menu $B \in \mathcal{A}$, define

$$\rho_B(f) = \lim_{\alpha \uparrow 1} P_{\alpha A+(1-\alpha)B}(\alpha A + (1-\alpha)\{f\}, \mathcal{T}), \forall f \in B. \tag{3.19}$$

Intuitively, when α approaches 1, the private information used to make choices from $\alpha A + (1-\alpha)B$ converges to the information acquired for A. Thus, ρ_B is interpreted as the random choice from B if the agent uses the information she optimally acquires for A. Note that the sequence $P_{\alpha A+(1-\alpha)B}(\alpha A + (1-\alpha)\{f\}, \mathcal{T})$ becomes constant for all α near enough to 1. This is because A is assumed to be without ties (see subsection 3.2.3) and it is common knowledge between analyst and agent that the experimentation technology has finite experimental outcomes.

Following Lu [2016], the expected utility gain from a menu B when the agent uses the fixed information structure induced by the optimal stopping strategy for A can be recovered from ρ and is equal to

$$\int_0^\infty \rho_{B\cup r}(A)dr.$$

The learning cost for A is the product of the flow cost c and the expected decision time. Then, by Theorem 4, the following relation holds:

$$\int_0^\infty \sum_{t\in\mathcal{T}} P_{A\cup\{r\}}(A, \tau = t)dr = \int_0^\infty \rho_{A\cup\{r\}}(A)dr - c\sum_{t\in\mathcal{T}} tP_A(\tau = t).$$

We conclude that the flow cost is uniquely pinned down by the formula

$$c = \frac{\int_0^\infty \left(\rho_{A\cup\{r\}}(A) - P_{A\cup\{r\}}(A)\right) dr}{\sum_{t=1}^{T} tP_A(\tau = t)}. \tag{3.20}$$

Finally, we mention as an aside that there is a way to identify the additive costs in SeSa-LC without using decision time data. This is because as mentioned in the literature review, SeSa-LC corresponds to a rational inattention model as considered in De Oliveira et al. [2017] and Lin [2018] and so identification results from ex-post random choice from menus from Lin [2018] are applicable. The details are in the appendix. The formula without decision time data is much more demanding than (3.20) though, because it requires making use of the full RCDT data, not just of the RCDT in the vicinity of a single menu A as (3.20) does.

3.5 Examples

In this section we offer examples which illustrate the power of RCDT for the identification of the agent's parameters and welfare analysis.

3.5.1 Example for SeSa-GD

Identification. Suppose that $S = \{s_1, s_2, s_3, s_4\}$ and $T = 2$. The sequential experiment consists of the following: At t=1, it is revealed whether event $\{s_1, s_2\}$ occurs or not. At $t = 2$, the true state is disclosed. The agent possesses a uniform prior belief over S. We assume the analyst only has access to a RCDT as prescribed in subsection 3.2.3. The analyst discovers through choices from menus of constant acts that the Bernoulli utility of the agent is the identity: $u : \mathbb{R}_+ \to \mathbb{R}_+, u(x) = x$.[26] Moreover, through menus of the type $\{\lambda \mathbf{1} + (1-\lambda)f_{s_i}, a\}$ for $i \in \{1, 2, 3, 4\}$ and λ near to one as well as sure prizes a,

[26]This is w.l.o.g. since we are passing to utility acts in any case.

the analyst uncovers the uniform prior of the agent through the formula (3.17).[27]

The stochastic choice observed by the analyst for the menu $\{f_s, r\}$ as a function of r is summarized in Table 1.[28] In Table 1 and all the following tables a pair (g, t) in the upper row with g an act and $t \in \mathcal{T}$ is the argument of the RCDT under consideration. Intuitively, information is valuable only when r is not too high or too low. Moreover, as the safe option r is increased the agent decides earlier and picks the safe option r more often.

$P_{\{f_s,r\}}$	$(f_s, 0)$	$(r, 0)$	$(f_s, 1)$	$(r, 1)$	$(f_s, 2)$	$(r, 2)$
$r \in [0, \frac{1}{8})$	1	0	0	0	0	0
$r \in (\frac{1}{8}, \frac{1}{4})$	0	0	$\frac{1}{2}$	$\frac{1}{2}$	0	0
$r \in (\frac{1}{4}, \frac{4}{11})$	0	0	0	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$
$r \in (\frac{4}{11}, 1]$	0	1	0	0	0	0

Table 3.5.1: $P_{\{f_s,r\}}$ for different r-s.

Condition 1 is fulfilled for the data in Table 1 and it follows from (3.18):

$$\sum_{s\in S}\left(\int_0^1\left(1-\sum_{t\in\mathcal{T}}\delta^t P_{\{f_s,r\}}(r,t)\right)dr\right) = 4\left(\frac{1}{8}+\frac{1}{8}\times(1-\frac{1}{2}\delta)+\frac{5}{44}(1-\frac{1}{2}\delta-\frac{1}{4}\delta^2)\right) = 1.$$

Solving this equation, the analyst identifies a unique discount factor $\delta = \frac{4}{5}$.

Thus, the analyst identifies the agent's parameter $(u = id, \pi_0 = uniform, \delta = \frac{4}{5})$ through RCDT data without having access to the sequential experiment of the agent.

Welfare example. Suppose that the agent whose parameters were just identified faces the menu $A = \{f_1, \frac{1}{3}\}$. An outside party which has the same taste and the same discount factor as the agent is contemplating to subsidize the menu

[27] This is obviously not the only way to identify the prior from RCDT data. One could look at choices from menus $\{\lambda\mathbf{1} + (1-\lambda)f_{s_i}, \lambda\mathbf{1} + (1-\lambda)f_{s_j}\}$ for $i, j \in \{1, 2, 3, 4\}$ and λ near to one.

[28] Due to symmetry, the Table 1 is the same for all $s \in S$.

to $A + 1$ by adding a lump-sum payment which is independent of the realization of the state of the world and of the choice of the agent. Note that the subsidy is paid when the agent actually chooses from the menu. Thus the actual subsidy costs to the outside party may be less than one in present value. The analyst observes the data below and wants to calculate whether there is any dead-weight welfare loss associated with the subsidy and if yes, its magnitude.

$P_{(A+1)\cup\{r\}}$	$(A+1,0)$	$(r,0)$	$(A+1,1)$	$(r,1)$	$(A+1,2)$	$(r,2)$
$r \in [0, \frac{4}{3})$	1	0	0	0	0	0
$r \in (\frac{4}{3}, \infty)$	0	1	0	0	0	0

Table 3.5.2: $P_{(A+1)\cup\{r\}}$ for different r-s.

$P_{A\cup\{r\}}$	$(A,0)$	$(r,0)$	$(A,1)$	$(r,1)$	$(A,2)$	$(r,2)$
$r \in [0, \frac{1}{3})$	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0
$r \in (\frac{1}{3}, \frac{4}{11})$	0	0	0	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$
$r \in (\frac{4}{11}, 1]$	0	1	0	0	0	0

Table 3.5.3: $P_{A\cup\{r\}}$ for different r-s.

Using Theorem 4 the analyst calculates

$$V_{\frac{4}{5}}(A) = \frac{26}{75} \text{ and } V_{\frac{4}{5}}(A+1) = \frac{4}{3}.$$

Finally, we see from the data that $P_{A+1}(\tau = 0) = 1$, i.e. the agent ceases to learn after the subsidy is awarded. Thus, the subsidy costs 1. Overall, the deadweight loss from the subsidy is calculated to be $\frac{4}{3} - \frac{26}{75} - 1 = -\frac{1}{75}$. Intuitively, even though the subsidy directly makes the agent better off, it also decreases her incentives to learn and this leads to a large subsidy which has to be paid already at time $t = 0$ (and is thus undiscounted), besides erasing any added decision value from learning.

3.5.2 Example for SeSa-LC

Identification. We assume two states of nature: $S = \{s_1, s_2\}$. Just as in the SeSa-GD example we assume the Bernoulli utility is the identity and the prior belief is uniform, that is, $\pi_0 = (\frac{1}{2}, \frac{1}{2})$. We assume an agent who has access to the same experiment which can be repeated two times. The signal space is $E = \{s_1, s_2\}$ so that the outcome of the experiment is a point estimate of the unobserved state. When the true state is s_i, the distribution of the signal is

$$Pr(s_j|s_i) = \begin{cases} a & \text{if } j = i, \\ 1 - a & \text{if } j \neq i. \end{cases}$$

Assume that $a > \frac{1}{2}$. Thus, the realization s_i favors the state s_i. a is the accuracy of the i.i.d. signal. The prior belief is uniform; that is, $\pi_0 = (\frac{1}{2}, \frac{1}{2})$. The analyst first uncovers the taste and the prior of the agent just as for the SeSa-GD case through choices from constant menus and small perturbations thereof.

Consider menu $A = \{f_1, \frac{7}{12}\}$. The constant act is better under the prior belief. The bet becomes better only if the agent accumulates enough positive signals. Consider the following observed choice from menus of the type $A \cup \{r\}$ for different r-s.

$P_{A\cup\{r\}}$	$(A, 0)$	$(r, 0)$	$(A, 1)$	$(r, 1)$	$(A, 2)$	$(r, 2)$
$r \in [0, \frac{7}{12})$	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0
$r \in (\frac{7}{12}, \frac{13}{20})$	0	0	0	$\frac{1}{2}$	$\frac{5}{18}$	$\frac{4}{18}$
$r \in (\frac{13}{20}, \infty)$	0	1	0	0	0	0

Table 3.5.4: $P_{A\cup\{r\}}$ for different r-s.

Then consider menu $\alpha A + (1 - \alpha)(A \cup \{r\})$. When α is close to 1 enough, the agent's choice from this mixed menu reveals how she would choose from $A \cup \{r\}$ if she follows the same stopping strategy as facing A. An analyst observes the following table.

$\rho_{A\cup\{r\}}$	$(A, 0)$	$(r, 0)$	$(A, 1)$	$(r, 1)$	$(A, 2)$	$(r, 2)$
$r \in [0, \frac{7}{12})$	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0
$r \in (\frac{7}{12}, \frac{4}{5})$	0	0	0	$\frac{1}{2}$	$\frac{5}{18}$	$\frac{4}{18}$
$r \in (\frac{4}{5}, \infty)$	0	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$

Table 3.5.5: $\rho_{A\cup\{r\}}$ for different r-s.

An analyst given above two tables can compute the ex-ante value of A:

$$\int_0^1 \sum_{t\in\mathbb{T}} P_{A\cup\{r\}}(A,t)dr = \frac{7}{12} \times 1 + (\frac{13}{20} - \frac{7}{12}) \times \frac{5}{18} = \frac{65}{108},$$

compute the expected utility gain from menu A:

$$\int_0^1 \rho_{A\cup\{r\}}(A)dr = \frac{7}{12} \times 1 + (\frac{4}{5} - \frac{7}{12}) \times \frac{5}{18} = \frac{139}{216},$$

and compute the average decision time when facing A:

$$\sum_{t\in\mathbb{T}} t \times P_A(A,t) = \frac{1}{2} \times 1 + \frac{1}{2} \times 2 = \frac{3}{2}.$$

Consequently, this analyst recovers the market price of an experiment:

$$c = \frac{\frac{139}{216} - \frac{65}{108}}{\frac{3}{2}} = \frac{9}{216} \times \frac{2}{3} = \frac{1}{36}.$$

Welfare. Consider an agent who faces menu A and would like to learn information about the state of the world. A researcher can run an experiment repeatedly at a fixed fee. The price (in utils) of the experiment is c. That is, if the agent wants to run the experiment $k \leq 2$ times, she has to pay $k \times c$ to the researcher.

Suppose that $c = \frac{1}{36}$. As the data above show the agent asks for the experiment to be run once with probability $\frac{1}{2}$ and to be run twice with probability $\frac{1}{2}$.

Now suppose that the government intends to tax the experiment for revenue purposes. Specifically, if the agent runs the experiment k times, she has to pay the government $k \times (\frac{1}{25} - \frac{1}{36})$. Thus, the effective cost of learning becomes $\frac{1}{25}$. The analyst observes that under the new conditions the agent samples the experiment only once.

Before tax, the welfare of the government is 0. The welfare of the researcher who performs the experiments is $\frac{1}{36} \times \frac{3}{2}$. The welfare of the agent (calculated through Theorem 4) is $V_{\frac{1}{36}}(A) = \frac{65}{108}$. Total welfare in the economy is the sum of all individual welfares and thus it's $\frac{139}{216}$.

After tax, the welfare of the government is $\frac{1}{25} - \frac{1}{36}$. The welfare of the scientist is $\frac{1}{36}$. The welfare of the agent is $V_{\frac{1}{25}}(A) = \frac{5}{8} - \frac{1}{25} = \frac{117}{200}$. The total welfare is $\frac{5}{8}$.

After tax, the agent is worse off, and the total welfare also decreases. Note that

$$V_{\frac{1}{36}}(A) - V_{\frac{1}{25}}(A) = \frac{91}{5400} > \frac{1}{36} \times \frac{3}{2} - \frac{1}{25} = \frac{1}{600}.$$

Therefore we have a dead-weight loss from taxation of

$$\frac{91}{5400} - \frac{1}{600} = \frac{41}{2700}.$$

Intuitively, this results because the tax is detrimental to the learning incentives of the agent, thus forcing her to take a decision earlier with less information in expectation than before.

3.6 Behavioral characterization

In this part of the paper we offer an axiomatic characterization result of RCDT data coming from an agent in either case of SeSa-GD or case of SeSa-LC. Our result is valid under the assumption that the analyst knows the sequential experiment $\mathcal{E}$ of the agent. This is not far-fetched in many practical situations. In many cases the licensing process, say for a new medical procedure, or financial due diligence procedures before entering a long-term financial partnership are public. In other situations, e.g. when a financial regulatory body is trying to detect insider trading, identifying the information of the agent becomes important and so our characterization result does not apply.

We split the axioms in two parts. The first four axioms don't make use of the knowledge of $\mathcal{E}$ and ensure that the taste, prior and the costs of the information of the agent are identified from a RCDT. The last axiom is a condition ensuring that the information and the parameters of the agent identified from the first four axioms match the RCDT for menus where the agent decides to learn.

Finally, a technical axiom of continuity ensures that the restriction to menus without ties which we stated in section 3.2 is without loss of generality. It says that menus where the agent needs to break ties can be approximated arbitrarily closely with menus where tie-breaking is not necessary because the agent faces strict incentives in both stopping decision and choice from menu decision. Since the discussion for that axiom is technical, it is relegated to the appendix. The crucial point is that menus without ties can be identified from data and that identification and welfare analysis rely only on them. This allows us to remain agnostic about the precise tie-breaking process of the agent.

3.6.1 Axioms that ensure identification.

The first two axioms are related to menus for which the agent does not have incentives to learn.

Axiom 1: Sampling is costly Let C be a constant menu or a singleton menu. Then for all menus $A \in \mathcal{A}$ it holds

$$P_{\lambda A+(1-\lambda)C}(\tau = 0) = 1,$$

whenever $\lambda \in (0, 1)$ is small enough.

In words this says that if one perturbs slightly a menu where there is no need to learn the state, the agent still doesn't want to learn about the state. This reflects the fact that costs of learning are not zero. This axiom ensures that $\mathcal{A}_0$ defined in section 3.2 is non-empty.

To give the second axiom we first define a preference relation.

Fix a constant menu C. Define a preference over acts $f, g : S \to \Delta(\mathbb{R}_+)$ by

$$f \succeq_C g \iff P_{\lambda C + (1-\lambda)\{f,g\}}(\lambda C + (1-\lambda)\{f\} | \tau = 0) \geq \frac{1}{2}$$

for all $\lambda \in (0, 1)$ near enough to 1.

Axiom 2: No-sampling is SEU with worst prize $\succeq_C$ is independent of the constant menu C and satisfies axioms of subjective expected utility with a worst prize $w \in \mathbb{R}_+$.

This axiom ensures the existence of the prior π_0 and taste u used at time zero.

Axiom 3: Taste Stationarity For every menu A and time $t > 0$ such that $P_A(\tau = t) > 0$ as well as constant acts $c, c' \neq w$ it holds

$$P_{\lambda A + (1-\lambda)\{c,c'\}}(\lambda A + (1-\lambda)\{c\} | \tau = t) = P_{\{c,c'\}}(\{c\} | \tau = 0),$$

for all $\lambda \in (0, 1)$ near enough to 1.

This axiom says that whenever perturbing a decision problem through constant acts, conditional on stopping, the marginal choice over the constant acts doesn't depend on the decision problem being perturbed or the respective decision times. It ensures that the risk preferences of the agent identified from Axiom 2 are time-independent. In particular, by classical results the agent may be assumed to use the same Bernoulli taste for all time periods $t \in \mathcal{T}$. This axiom ensures that the passage to utility acts in the following is meaningful and that we can use utility acts based on a Bernoulli taste identified in Axiom 2 to measure the exploitation-exploration trade-off of the agent. This in turn reveals the information costs of the agent as we showed in section 3.4.

In the following fix a taste u with $u(\text{worst prize}) = 0$ and go over to utility acts as introduced in section 3.3.

Axiom 4-δ: Wedge between the values of full information and no information. For any $A \in \mathcal{A}_t$, $t \geq 1$ and $k \geq 0$ s.t. $A + k \in \mathcal{A}_0$ we have

$$k + \max_{f \in A} \pi_0 \cdot (u \circ f) > \int_0^\infty P_{A \cup \{a\}}(A, 0)da + \int_0^k P_{A+a}(\tau = 0)da$$

and

$$k + \max_{f \in A} \pi_0 \cdot (u \circ f) < \int_0^\infty 1 - \sum_{t=0}^T P_{A \cup \{a\}}(a, t)da + \sum_{t=0}^T \int_0^k P_{A+a}(\tau = t)da.$$

This axiom ensures existence and uniqueness of a discount factor in the SeSa-GD model.

In both inequalities, the left-hand side is the value of choosing from $A + k$. In the first inequality, the right-hand side is the value of the menu $A + k$ when the agent is myopic ($\delta = 0$). Thus the first inequality says that the decision of the agent to not learn for the menu $A + k$ doesn't come from her being myopic and any tie-breaking considerations (given that $A \in \mathcal{A}_t$). On the other hand, the right-hand side of the second inequality is the value of $A + k$ if the agent receives all of the information potentially available to him within the first period, i.e. if the whole sequential experiment could be performed within the first period so that no costs due to impatience are incurred. The second inequality then says that the agent values getting all the information at once for menu $A + k$ more than choosing from $A + k$ without information. In particular, since her revealed choice is no information for $A + k$ this implies that information must be costly.[29]

[29]Nevertheless, this Axiom is clearly not implied by Axiom 1 and it uses the taste u identified

We now turn to the SeSa-LC model and the respective fourth axiom for the additive costs. Recall the definition in (3.19) of the random choice when fixing the information to the optimal one from a menu A.

Axiom 4-*c*: Additive, expected costs of time For every pair of menus A, A' such that $P_A(\tau \geq 1), P_{A'}(\tau \geq 1) > 0$ as well as every $k > 0$:

$$\int_0^\infty \rho_{A\cup\{a\}}(A)da - k\sum_{t=1}^{T} tP_A(\tau = t) \leq \int_0^\infty P_{A\cup\{a\}}(A)da,$$

if and only if

$$\int_0^\infty \rho'_{A'\cup\{a\}}(A')da - k\sum_{t=1}^{T} tP_{A'}(\tau = t) \leq \int_0^\infty P_{A'\cup\{a\}}(A')da,$$

where ρ, respectively ρ', are the induced random choice when agent uses the optimal information for A, respectively A'.

In words this says that the difference in value of a decision problem before and after information acquisition, which determines information costs, is proportionally dependent only on the first moment of the decision time. Moreover, it also ensures that the marginal costs of an additional experiment are not menu-dependent.

3.6.2 Data-matching and characterization result

Even though they ensure identification of Bernoulli taste, prior and costs of information Axioms 1-4 only use part of the information contained in a RCDT. E.g. they do not contain information about choices from menus of the type A so that:

i. A doesn't contain constant acts and nor can it be written as a mixture of two menus where one consists only of constant acts;

from Axioms 1 and 2.

ii. the agent strictly wants to learn for A.

In particular, Axioms 1-4 don't say anything related to optimality of choosing from A with respect to the taste u identified from Axioms 1 and 2 upon stopping at a particular time t. Nor do they imply that the agent's stopping behavior is optimal given $\mathcal{E}$.

Here we present a natural Data-matching condition which ensures that the parameters identified from Axioms 1-4 together with the sequential experiment $\mathcal{E}$ rationalize the choice data in the form of RCDT.[30]

Recall that an agent is a tuple $\mathbb{A} = (u, \delta, \mathcal{E}, \pi_0)$ or $\mathbb{A} = (u, c, \mathcal{E}, \pi_0)$. The prior identified as

$$\pi_0(s) = \frac{1}{1-\lambda}\left(\int_0^\infty P_{\{(1-\lambda)f_s+\lambda\mathbf{1},a\}}((1-\lambda)f_s + \lambda\mathbf{1}|\tau = 0)da - \lambda\right), \quad (3.21)$$

for all $\lambda \in (0, 1)$ near enough to 1.

Recall the measure on histories of experimental outcomes induced by $\mathcal{E}$ and π_0 through Bayes rule.

$$\mu(\pi_0) \text{ measure on } E_1 \times \cdots \times E_T.$$

Bayes rule implies the belief process $\pi(e^t)$, $e^t \in E_1 \times \cdots \times E_t$, $t \geq 1$ constitutes a Martingale.

Finally, recall the discount factor identified from the identity

$$k + \max_{f\in A'} \pi_0 \cdot (u \circ f) = \int_0^\infty \left(1 - \sum_{t=0}^{T} \delta^t P_{A'\cup\{a\}}(a, t)\right) da + \sum_{t=0}^{T} \delta^t \int_0^k P_{A'+a}(\tau = t)da, \quad (3.22)$$

[30] Recall that $\mathcal{E}$ is assumed known in this section of the paper.

for the case of SeSa-GD and the additive costs identified through

$$c = \frac{\int_0^\infty \left(\rho_{A' \cup \{r\}}(A') - P_{A' \cup \{r\}}(A') \right) dr}{\sum_{t=1}^{T} t P_{A'}(\tau = t)}, \tag{3.23}$$

for the case of SeSa-LC. Here A' is a menu satisfying Condition N from section 3.2.

Consider either of the SeSa-GD and SeSa-LC. For a menu A without ties the backwards induction procedure dictated by sophistication of the agent results in the recursive partition of the histories of experimental outcomes $E_1 \times \cdots \times E_t, 1 \leq t \leq T$ as follows:

$$S_t(A) = \{e^t \in E_1 \times \cdots \times E_t : V(A, \pi(e^t)) > \delta W_{t+1}(A, e^t)\},$$

$$C_{<t+1}(A) = \{e^t \in E_1 \times \cdots \times E_t : V(A, \pi(e^t)) < \delta W_{t+1}(A, e^t)\}.$$

$S_t(A)$ is the collection of histories of length t where the agent decides to stop experimenting, whereas $C_{<t+1}(A)$ the set of histories of depth t which give strict incentives for the agent to continue experimenting. Naturally, backwards induction results in the recursion

$$C_{<t}(A) = S_t(A) \cup C_{<t+1}(A), \quad 1 \leq t \leq T-1,$$

with a 'boundary' condition $S_T(A) = E_1 \times \cdots \times E_T$, which is dictated from the finite horizon assumption. Finally, we can refine the sets $S_t(A)$ by defining $S_t(A, f)$ for the subset of histories of $S_t(A)$ where it is optimal to pick $f \in A$. For an agent $\mathbb{A} = (u, \delta, \mathcal{E}, \pi_0)$ or $\mathbb{A} = (u, c, \mathcal{E}, \pi_0)$ we add then superscripts

$$S_t^{\mathbb{A}}(A), \quad C_{<t}^{\mathbb{A}}(A), \quad S_t^{\mathbb{A}}(A, f),$$

to denote that the sets $S_t^{\mathbb{A}}(A), C_{<t}^{\mathbb{A}}(A), S_t^{\mathbb{A}}(A, f)$ originate from an agent with parameters as in the tuple $\mathcal{A}$. Given the 'identification' Axioms 1-4, the following Axiom ensures rationalizability in the case of SeSa-GD.

Axiom 5-c: Data-matching For the u from Axioms 1-2, the π_0 in (3.21) and δ in (3.22) it holds

$$P_A(f, \tau = t) = \mu(\pi_0)(S_t^{u,\delta,\mathcal{E},\pi_0}(A, f)), \qquad (3.24)$$

whenever A has no ties and satisfies $\mathcal{A} \in \mathcal{A}_t$ for some $t \geq 1$.

For the case of SeSa-LC Data-matching looks as follows.

Axiom 5-c: Data-matching For the u from Axioms 1-2, the π_0 in (3.21) and c in (3.23) it holds

$$P_A(f, \tau = t) = \mu(\pi_0)(S_t^{u,c,\mathcal{E},\pi_0}(A, f)), \qquad (3.25)$$

whenever A has no ties and satisfies $\mathcal{A} \in \mathcal{A}_t$ for some $t \geq 1$.

Note that the Data-matching axioms are the first time we use the observability of the sequential experiment $\mathcal{E}$ in this paper. Our characterization Theorem then reads as follows.

Theorem 6. *Suppose that an analyst knows $\mathcal{E}$ at the disposal of the agent.*

1) A RCDT $\{P_A \in \Delta(A \times \mathcal{T}) : A \in \mathcal{A}\}$ which satisfies Condition N can be rationalized by a SeSa-GD if and only if it satisfies Axioms 1-3, 4-δ and 5-δ.

2) A RCDT $\{P_A \in \Delta(A \times \mathcal{T}) : A \in \mathcal{A}\}$ which satisfies Condition N can be rationalized by a SeSa-LC if and only if it satisfies Axioms 1-3, 4-c and 5-c.

This characterization result shows that the question of whether RCDT data for an agent with sequential experiment $\mathcal{E}$ is rationalized by a SeSa can be answered in two steps. In the first step, the analyst can identify the agent's preference and initial information as well as the costs of information without knowledge of the informational technology (sequential experiment) of the agent. That this step is meaningful is ensured by Axioms 1-4. The first step uses only a very limited collection of menus $A \in \mathcal{A}$ and is not sufficient to test rationalizability. Therefore, in the second step the analyst builds upon the identification result of the first step to test rationalizability for the rest of the menus.

In this sense, the Data-matching relations (3.24) and (3.25) can be understood as a statistical test the analyst can build to test the joint hypothesis that the agent $\mathbb{A} = (u, \delta, \mathcal{E}, \pi_0)$ or $\mathbb{A} = (u, c, \mathcal{E}, \pi_0)$ understands the environment she is in (i.e. is using indeed $\mathcal{E}$ to acquire information and, say, has additive costs of c in the case of SeSa-LC) and is sophisticated about her future behavior.

3.7 Concluding Discussion

We have shown how random choice data from menus augmented with stochastic decision time data allow the identification of the taste, prior and costs of information of the agent without any knowledge of the technology of sequential experiments. We consider in this paper two versions of costs of information. First, the agent may be impatient and discount future payoffs geometrically. Second, she may need to pay for every repetition of an i.i.d. experiment. We have also shown how random choice data on choice from menus and decision times enable welfare analysis of the agent, without any knowledge of the sequential experiment at her disposal.

Moreover, when the sequential experiment at her disposal is known to an outside analyst, we show how this analyst may test the joint hypothesis that the agent understands the choice environment she is in and is sophisticated about her future behavior.

In ongoing work we consider how the outside analyst can uncover the sequential experiment of the agent from her random ex-post choice from menus and decision times. This is a challenging problem due to the lack of recursivity in the observable.

A
Appendix to Chapter 1

A remark on notation in the appendix and online appendix:

For limit statements with respect to $\Delta \to 0$, I often use the Landau notation $o(\Delta)$, $O(\Delta)$.[1]

A.1 Proofs for section 1.2

A.1.1 Proofs of general properties of PBEs

Proof of Lemma 1 and its corollaries.

The proof of Lemma 1 relies heavily on similar arguments in the proofs of the Lemmas 1 and 2 in Fudenberg et al. [1985]. The arguments need to be adapted to account for Buyer valuation changing over time due to learning.

[1] See e.g. chapter 5, section 4 of Lang [1997] for formal definitions.

Before giving the proof of Lemma 1, I remark an important fact used in the proof of the Lemma.

Remark 1. *Suppose a Bayesian learner has a prior F over $\mathbb{R}$ for the variable θ and uses an unbiased experiment $\mathcal{E} : supp(F) \to \Delta(S)$, with S finite, to learn about θ where S is a signal space. Let $\mathbb{E}_F[\theta|s]$ be the posterior mean after signal s when prior for θ is F. If F' is another prior so that F' FOSD-dominates F then for every $s \in S$ it holds*

$$\mathbb{E}_{F'}[\theta|s] \geq \mathbb{E}_F[\theta|s],$$

i.e. the estimates increase pointwise when the prior increases in the FOSD sense.[2]

Proof of Remark 1. One can realize all random variables needed in one large enough probability space where θ, θ' are such that $\theta \sim F, \theta' \sim F'$ and $\theta' = \theta + y$ in distribution with $y \geq 0$ a random variable. This larger probability space has as sample space the collection of pairs (θ, y). Note here the signal space S as well as the experiment $\mathcal{E}$, which is a random variable from $\Theta \supseteq supp(F) \cup supp(F')$ to $\Delta(S)$ is being kept fixed. In this larger probability space it holds conditional on the realization of a signal s

$$\mathbb{E}[\theta'|s] = \mathbb{E}[\theta + y|s] = \mathbb{E}[\theta|s] + \mathbb{E}[y|s] \geq \mathbb{E}[\theta|s]. \tag{A.1}$$

Because $\mathbb{E}_{F'}[\theta|s], \mathbb{E}_F[\theta|s]$ depend only on the joint distributions of (θ, s) (in the case of F given by $F(\theta) \cdot \mathcal{E}(\theta)(s)$ and in the case of F' given by $F'(\theta) \cdot \mathcal{E}(\theta)(s)$), the conclusion follows from (A.1).

□

Proof of Lemma 1. First, I show some auxiliary claims. These imply that in any PBE any price $p < \underline{v}$ is accepted immediately by all Buyer types w. For the auxiliary claims let h be a public history which ends with a period t and the rejection of the price quoted at the end of history h by Seller. Let $V(w, h)$ be the *expected* equilibrium payoff of Buyer type w after public history h (starting from period $t + 1$).

[2] Note, it is in general not true that the *distribution* of posteriors increases in the FOSD-sense.

Claim 1. $w \mapsto V(w, h)$ is strictly increasing and Lipschitz with constant one.

Proof of Claim 1. Define $\tau(h, w)$ to be the stopping time that gives the agreement date in an equilibrium for Buyer of type w after history h.

The equilibrium payoff of type w after history h can be written as

$$V(w, h) = \sum_{u \geq 0} \delta^t a_{t+1+u}(w, h)(\mathbb{E}_w[v_{\tau(h,w)} | \tau(h, w) = t+1+u] - \mathbb{E}_w[p_{t+1+u}(w, h)]),$$

where $a_{t+1+u}(w, h)$ is the probability of agreement at time $t + 1 + u$ on path if w follows her strategy, $\mathbb{E}_w[v_{\tau(h,w)} | \tau(h, w) = t + 1 + u]$ is the *expected* value for Buyer of type w conditional on agreement at time $t + 1 + u$ and $\mathbb{E}_w[p_{t+1+u}]$ is the expected price Buyer pays conditional on agreement at time $t + 1 + u$.[3] Given an (adapted) learning strategy for Buyer, the expected valuation of the good for Buyer is a Martingale. Thus, it holds

$$\mathbb{E}_w[v_{\tau(h,w)}] = w. \tag{A.2}$$

Let $w' > w$ be another possible type after history h. If Buyer of type w instead uses the optimal stopping strategy of the higher type w', $\tau(h, w')$, it holds again by Martingale property.

$$\mathbb{E}_w[v_{\tau(h,w')}] = w. \tag{A.3}$$

Similar to the classical proof, I use a no-imitation argument. From the equilibrium property it follows

$$V(w, h) \geq \sum_{u \geq 0} \delta^t a_{t+1+u}(w', h)(\mathbb{E}_w[v_{\tau(h,w')} | \tau(h, w') = t+1+u] - \mathbb{E}_w[p_{t+1+u}(w', h)]).$$

Replacing the formula for $V(w', h)$ from above yields

$$V(w', h) - V(w, h) \leq \mathbb{E}_{w'}[\delta^{\tau(h,w)} v_{\tau(h,w)}] - \mathbb{E}_w[\delta^{\tau(h,w)} v_{\tau(h,w)}].$$

[3] Note that time has been shifted accordingly.

Next, I show the following relation.

$$\mathbb{E}_{w'}[\delta^{\tau(h,w)} v_{\tau(h,w)}] - \mathbb{E}_{w}[\delta^{\tau(h,w)} v_{\tau(h,w)}] \leq w' - w. \tag{A.4}$$

Proof of (A.4). Fix a stopping time τ and consider two bounded, non-negative stochastic processes $(v_t)_t$ and $(v'_t)_t$ such that for all t it holds $v_t \leq v'_t$. Then it follows

$$\mathbb{E}[v_\tau(1-\delta^\tau)] \geq \mathbb{E}[v'_\tau(1-\delta^\tau)] \quad \iff \quad \mathbb{E}[v_\tau]-\mathbb{E}[\delta^\tau v_\tau] \geq \mathbb{E}[v'_\tau]-\mathbb{E}[\delta^\tau v'_\tau]. \tag{A.5}$$

When Buyer starts with type $w' > w$ and follows the same strategy as if she started from w, after every learning opportunity she receives a pointwise weakly higher estimate of the value of the good than under w (see Remark 1). Since this holds pointwise, one can use the inequality (A.5) together with (A.2) and (A.3) to show (A.4).

End of proof of (A.4).

This establishes the Lipschitz continuity of $w \mapsto V(w,h)$. One writes

$$V(w,h) = \max_{\text{learning},\tau,h\text{-measurable}} \mathbb{E}[\delta^\tau v_\tau - p_\tau].$$

For every fixed learning and stopping strategy, it holds that v_τ increases pointwise with w. Therefore monotonicity of $w \mapsto V(w,h)$ follows from a simple optimization/envelope theorem argument.

End of proof of Claim 1.

As a next step, one establishes the following *skimming property*.

$$\text{Skimming:} \quad w-p \geq \delta V(w,h) \quad \Longrightarrow \quad w'-p > \delta V(w',h), \text{ whenever } w' > w.$$

To see this, one uses Lipschitz continuity of $w \mapsto V(w,h)$ as follows.

$$\begin{aligned} w - p \geq & \delta(V(w',h) + w - w') \quad \iff \quad w + \delta(w'-w) - p \geq \delta V(w',h) \\ & \Longrightarrow \quad w' - p > \delta V(w',h), \end{aligned}$$

where the last step used the fact that $w' > 0$.

Proof of part 1) of the Lemma. Suppose that after history h the highest possible on-path Buyer valuation is $\bar{w}$ and the lowest possible is $\underline{w}$. Note that any Buyer type can always ensure 0 by always rejecting after any history. Same holds for Seller; she can always ensure zero after any history by always offering prices which are too high (say prices above $\bar{v}$). This implies that after any history, the expected valuation of any Buyer type whenever there is agreement is below the highest valuation $\bar{w}$. Because the continuation payoff $V(w, h)$ is increasing in w it holds then that $V(\underline{w}, h) \leq \bar{w}$.

It follows through Lipschitz continuity that

$$V(w, h) \leq w + \bar{w} - \underline{w}.$$

In particular, it follows that all Buyer types accept any subsidy larger than $-\underline{w} + \bar{w}$ immediately. Given this, it is never optimal for Seller to charge prices below $\underline{w} - \bar{w}$. Knowing this, Buyer of type w accepts any p satisfying $w - p \geq \delta(w - (\underline{w} - \bar{w})) \Longleftrightarrow p \leq \underline{w} - \delta\bar{w}$. Iterating this argument just as in the classical proof, one finds that Seller requires prices strictly below $\underline{w} - \delta^n\bar{w}$. Send $n \to \infty$ to finish the proof.

Proof of part 2) of the Lemma. Suppose Seller after some public history h asks with positive probability for a price p smaller than the reservation price of a Buyer of type w. Denote the reservation price of this type by $r(w, h)$ in the PBE in question. It satisfies $w - r(w, h) = \delta V(w, h)$. It holds for any $w' > w$ that $w' - r(w, h) > \delta V(w', h)$. Because of the skimming property, increasing p to $p + \varepsilon$ with $p + \varepsilon < r(w, h)$ and keeping the same probability on it as previously on p leads to higher profits.

Next, assume towards a contradiction that Seller asks for a price above the highest reservation price that is accepted with positive probability after the history h is optimal. Let F_h be the distribution of types conditional on history h and set $\bar{w}$ for the highest type in the support of F_h. Denote by $V_S(h)$ the continuation payoff of Seller from the equilibrium strategy. Charging more than the highest reservation price $r(\bar{w}, h)$ results in a payoff $\delta V_S(h)$. Suppose that instead Seller charges with positive probability p also $\bar{w} - \varepsilon$ with some $\varepsilon > 0$ small. The payoff from this deviation is at least

$$p \cdot F_h(\bar{w})(\bar{w} - \varepsilon) + (1 - p) \cdot \delta V_S(h) > \delta V_S(h),$$

for all $\varepsilon > 0$ small enough. The inequality above follows from the fact that in the assumed strategy, game slides into next period with starting distribution F_h (Seller does not learn anything by charging a price which no Buyer type can afford). In particular it holds

$$V_S(h) \leq \sup\{r(w, h') : w \in supp(F_{h'}), h' \text{ a continuation history of } h\} \leq \bar{w}.$$

Here, the second inequality uses the fact $r(w, h) \leq w$ for all w and all h.

Proof of part 3) of the Lemma. If this probability were zero, then no equilibrium would exist where this price is quoted with positive probability because Seller, whenever the equilibrium would prescribe quoting that price with positive probability, would want to deviate to arbitrarily smaller price offers. In particular, there would not be a well-defined best response of Seller after such a history. This contradicts the existence of the equilibrium.

Proof of part 4) of the Lemma. Upon disclosure of current valuation Seller can calculate the valuation of Buyer and knows that this valuation will not change going forward. The proof of part 4) is finished by an argument that follows the logic of the uniqueness of payoffs for the usual Rubinstein-Stahl bargaining model. See Example 9.A.A.2 in Mas-Collel et al. [1995].[4]

[4] The details are as follows. Suppose for simplicity and normalization that valuation to Buyer after this history of disclosure is w. Because this is common knowledge after the verifiable disclosure, the subgame turns into a classical *share-the-pie* game. Let $\underline{v}_S$ be the lowest PBE payoff in the continuation game of Seller, $\bar{v}_S$ the highest PBE payoff in the continuation game of Seller and define in addition accordingly $\underline{v}_B, \bar{v}_B$ to be respectively the lowest and the highest PBE payoffs for Buyer. Note that, because of learning once, the subgame after a rejection is isomorphic to the game started after the history of disclosure. Suppose Seller offers more than $\delta\bar{v}_B$. Then this is accepted and so $\underline{v}_S \geq w - \delta\bar{v}_B$. Now note that $\bar{v}_B \leq w -$ PBE payoff of Seller $\leq w - \underline{v}_S$. Combining with the

□

The next Remark shows that Lemma 1 remains valid in the environment with costs.

Remark 2 (Lemma 1 in the case of costs)**.** *All the statements of Lemma 1 remain valid in the environments with costs from sections 1.3 and 1.4.1.*

Proof. To see this, note first that all arguments in the Claim 1 of the proof of part 1) of the Lemma remain true if one interprets the price $p_u(w, h)$ in the proof arguments as the expected costs incurred if agreement is reached at period u for Buyer of type w at history h. These costs contain the price costs of the purchase at time u as well as the information acquisition costs incurred in the period between u and history h.[5] Note also that the probability of agreement $a_u(w, h)$ used in the proof is history-dependent and thus does not depend on whether the arrival rate of information μ is history-independent. This, and the fact that the proof of part 1) of Lemma 1 involves only Buyer payoffs, yield that part 1) remains true in the case of costs.

It is trivial to see that the proof of parts 2) and 4) of Lemma 1 do not depend on whether information is costly or not, because they involve the pricing decision of Seller.

To see that part 3) of the Lemma remains true in the environment with costs just note that any information costs incurred in the past are sunk at the moment of Buyer's decision of whether to accept an offered price. □

Corollary 7. *After every private history h the reservation prices of Buyer types $r(w, h)$ are strictly increasing in w.*

Proof. From the definition of reservation prices $w - r(w, h) = \delta V(w, h)$ and the skimming property it holds that $w' > w$ implies $w' - r(w, h) > \delta V(w', h)$. It follows $r(w', h) > r(w, h)$. □

---

previous inequality, delivers overall that $\underline{v}_S \geq w$. This implies Seller leaves zero surplus to Buyer.

[5] The crucial assumption that allows this interpretation is the fact that information acquisition costs enter the overall payoff of Buyer linearly.

Corollary 8 (No quiet periods)**.** *There is no PBE where after a history in which Seller is called upon to play, the probability of trade is zero.*

Proof. This is very similar to the second part of the proof of part 2) of Lemma 1. □

Proof of Lemma 2. Disclosure never happens when learning good news and Buyer receives positive payoff in the PBE after receiving good news, since learning is only once and disclosure would lead to zero continuation payoff.

Delaying disclosure of bad news never increases payoffs. Since the learning happens only once, Buyer with bad news knows that she receives zero payoff in equilibrium, no matter disclosure decision. This is true after every private history. □

A.1.2 No sequential screening of valuations near the HFL

First, I define a refinement for a PBE.

Refinement for off-path beliefs: 'divinity in bargaining'. After an off-path history resulting from the rejection of a price, if the pool of Buyer types contains only one type who was indifferent between accepting and rejecting and all other types, for which Seller's belief had positive probability, had a **strict** incentive to accept the quoted price, then Seller starts new period with a belief that puts probability one on the type who was indifferent between accepting and rejecting in the last period.

This equilibrium selection is motivated by the 'divinity' criterion for signalling games, see Banks and Sobel [1987]. The strongly stationary equilibria constructed in this paper satisfy divinity in bargaining.

Proposition 29. *For all Δ small enough, there are no equilibria which satisfy the following properties.*

A. $\underline{v} > 0$ *and Buyer of type* $\underline{v}$ *discloses immediately,* *B. sequential screening of valuations.*

or

A. $\underline{v} = 0$ *and Buyer of type* $\underline{v}$ *discloses immediately*

B. stationarity,

C. sequential screening of valuations,

D. 'divinity in bargaining'.

Proof. I show that for all Δ small enough, generically, all PBEs do not satisfy all of the properties in the statement of the Proposition. The proof is split into several claims. First, focus on the case $\underline{v} > 0$.

Claim 1. a) Fix $\Delta > 0$. There is no PBE which satisfies A and has sequential screening dynamics of arbitrary length (i.e. $K = \infty$ is impossible for any Δ).

b) Under the requirement A the number of periods needed in any PBE with sequential screening dynamics does not grow faster than $\frac{1}{\Delta}$, i.e. $K(\Delta)\Delta = O(1)$.

To see a), suppose there is such a PBE. Then the price dynamics, given by the reservation price relation of type $\bar{v}$, satisfies

$$r_k = (1-\delta)\bar{v} + \delta r_{k+1}, \quad k \geq 1, \tag{A.6}$$

with $r_2 < \bar{v}$. Then one can solve for the dynamics to get

$$r_{k+1} = \frac{1}{\delta^{k-1}} r_2 + (1 - \frac{1}{\delta^{k-1}})\bar{v}.$$

This dynamics leads to $r_k \rightarrow -\infty$ due to $\delta < 1$ and this is a contradiction to results of Lemma 1.

To see b) consider a sequence of equilibria which satisfy A. and have sequential screening dynamics. The reservation price relation for the high type $\bar{v}$ given by $\bar{v} - r_i = \delta(\bar{v} - r_{i+1})$ delivers $r_i - r_{i+1} = (1-\delta)(\bar{v} - r_{i+1})$. A telescopic sum argument gives

$$r_1 - r_K = \sum_{j=1}^{K(\Delta)-1} (r_i - r_{i+1}) \geq (1-\delta)(K(\Delta)-1)(\bar{v} - r_1),$$

where the last step uses that prices $r_i, i \leq K(\Delta)$ are decreasing. Recall that $1-\delta = 1-e^{-r\Delta}$ and from iterating the reservation price relation for the type $\bar{v}$ one arrives at $\bar{v} - r_1 = \delta^{K(\Delta)-1}(\bar{v} - r_{K(\Delta)})$. Overall it follows

$$1 \geq \frac{r_1 - r_{K(\Delta)}}{\bar{v} - r_{K(\Delta)}} \geq (1-\delta)\delta^{K(\Delta)-1}(K(\Delta)-1).$$

This delivers through simple estimates

$$K(\Delta)\Delta \leq Ce^{-rK(\Delta)\Delta},$$

for some $C > 0$ which is independent of Δ. From here it is easy to see that $K(\Delta)\Delta$ remains bounded away from infinity.[6]

Claim 2. In the last period K on path, a price $0 < r_K \leq \hat{v}$ is quoted.

By the skimming property the sequence of decreasing prices r_l corresponds, up to the last period $K(\Delta) < \infty$, to the reservation prices of the high-type $\bar{v}$. In period $K(\Delta)$ $r_{K(\Delta)}$ corresponds to the reservation price of type $\hat{v}$. If it corresponds to the reservation price of type $\bar{v}$, then this implies that type $\hat{v}$ is never screened before period $K(\Delta)$ and thus, with positive probability, bargaining goes on into period $K(\Delta)$. This is a contradiction to the definition of $K(\Delta)$ and to Claim 1.

If $r_K > \hat{v}$, with positive probability the game does not end at or before K because Buyer type $\hat{v}$ who is present in period $K(\Delta)$ after non-disclosure with positive probability rejects $r_{K(\Delta)}$. Given that $r_{K(\Delta)}$ is the reservation price of $\hat{v}$, Lemma 1 delivers $r_{K(\Delta)} > 0$ (recall that type $\underline{v}$ discloses immediately and has zero payoff after every on-path history).

[6]The function $x \mapsto xe^{-rx}, \quad x \geq 0$ is strictly increasing, continuous and convex.

For ease of notation, in the following I sometimes suppress the dependence on Δ in $K(\Delta)$, whenever the argument does not rely on the precise magnitude of Δ.

Claim 3. If $\underline{v} > 0$ then $K \leq 2$. In particular, the game never continues past third period, whenever Δ is small enough.

To see this, assume that $K \geq 4$ in the equilibrium and look for a contradiction. The price at time $t = K - 1$, r_{K-1} necessarily satisfies

$$r_K \leq U(\gamma_{K-1})(r_{K-1} - \delta r_K) + \delta r_K,$$

which can be rewritten with the help of (A.6) from the proof of Claim 1 as $U(\gamma_{K-1}) \geq \frac{r_K}{\bar{v}}$. But if $K \geq 4$ then it holds $\gamma_{K-1}(\bar{v}) \leq U(0)$ if there are sequential screening dynamics, because it holds $\gamma_2(\bar{v}) \leq U(0)$ by virtue of the assumed updating of Seller in Definition 2. $\gamma \to U(\gamma)$ is strictly increasing and one calculates

$$U^{-1}\left(\frac{r_K}{\bar{v}}\right) = \frac{\frac{r_K}{\bar{v}}(1 - \mu(1 - \pi_0)) - \mu\pi_0}{1 - \mu\pi_0 - \frac{r_K}{\bar{v}}\mu(1 - \pi_0)}.$$

The following string of inequalities results.

$$\frac{\underline{v}}{\bar{v}} < \frac{r_K}{\bar{v}} \leq \frac{\pi_0\mu(2 - \mu)}{1 - (1 - \pi_0)\mu(1 - \mu)}. \tag{A.7}$$

Here the last inequality on the right follows from $U^{-1}\left(\frac{r_K}{\bar{v}}\right) \leq U(0)$. The right-hand side of (A.7) converges to zero as $\Delta \to 0$ and this establishes a contradiction if $\underline{v} > 0$.

Claim 4. There is no sequential screening of valuations with $\underline{v} > 0$, whenever Δ is small enough.

The case $K = 1$ is easy to exclude, because it would require that price be $\underline{v}$ in the first period and this is suboptimal as Δ vanishes, because with very high probability Buyer has not learned yet. Suppose therefore in the following that $\underline{v} > 0$ and focus on the case $K = 2$.

Optimality at $t = 1$ leads to

$$U(0)((1-\delta)\bar{v} + \delta r_2) + (1 - U(0))\delta r_2 \geq r_2.$$

This gives $r_2 \leq U(0)\bar{v}$. Since $U(0) \to 0$ as $\Delta \to 0$, a contradiction to the inequalities $r_2 > \underline{v} > 0$ results.

I now turn to the case $\underline{v} = 0$.

Claim 5. There is no PBE with sequential screening of valuations, whenever Δ is small if $\underline{v} = 0$, if in addition one requires stationarity and 'divinity in bargaining'.

Assume there exists such a PBE, of maximal length on path of $K \geq 1$ for a sequence of Δ small and show that for any such PBE and $\Delta > 0$ small enough, the price quoted in the last period $K(\Delta)$ converges to a unique positive number, uniformly and independently of the 'length' $K(\Delta)$, $\Delta \to 0$ of the equilibrium. Once this is established, it suffices to essentially repeat the argument in Claims 3 and 4 of the case $\underline{v} > 0$ to conclude the proof.

First note that due to 'divinity in bargaining' and stationarity, if the price r_K is rejected, then in the period $K+1$ Seller begins with belief of high type equal to zero and updates upon non-disclosure to the interim belief $U(0)$. Because r_K is the reservation of type $\hat{v}$ after a history of rejection of $r_1, \ldots, r_{K-1}$ it holds $\hat{v} - r_K = \delta V_B^{cont}(\hat{v})$ where $V_B^{cont}(\hat{v})$ is the continuation payoff according to the PBE in question. This subgame is isomorphic to the whole game and so because of the stationarity assumption it holds $\hat{v} - r_K \geq \delta V_B$. To ease notation, define for use in the following $p_k = r_{K-k}, k \leq K$. The reservation price relation for type $\bar{v}$ from the price dynamic leads to $\bar{v} - p_{k-l} = \delta^{k-l}(\bar{v} - p_0)$. The goal is to try to characterize p_0 explicitly.

Let V_k^{cont} be the payoff of Buyer at the beginning of period $K - k$ on path, if she has not learned yet. It holds

$$V_B = \mu\pi_0(\bar{v} - p_K) + 0 + (1-\mu)\delta V_{K-1}^{cont}. \tag{A.8}$$

Moreover, the following recursion holds

$$V_k^{cont} = \mu\pi_0(\bar{v} - p_k) + (1-\mu)\delta V_{k-1}^{cont}.$$

Using the recursion repeatedly leads to

$$V_{K-1}^{cont} = \mu\pi_0 \sum_{l=1}^{M} (\bar{v} - p_{K-l})((1-\mu)\delta)^{l-1} + ((1-\mu)\delta)^M V_{K-M-1}^{cont}.$$

Specialize to $M = K - 1$ to arrive at

$$V_{K-1}^{cont} = \mu\pi_0\delta^{K-1}(\bar{v} - p_0) \sum_{l=1}^{K-1} (1-\mu)^{l-1} + ((1-\mu)\delta)^{K-1} V_0^{cont}. \tag{A.9}$$

Here it holds $V_0^{cont} = \mu\pi_0(\bar{v} - p_0) + (1-\mu)(\hat{v} - p_0)$. Using this together with (A.8) and (A.9) leads after algebra to

$$V_B = \delta^K \pi_0(\bar{v} - p_0)(1 - (1-\mu)^{K+1}) + \delta^K(1-\mu)^{K+1}(\hat{v} - p_0).$$

Combining this with the definition of the reservation price p_0 leads to the following equation for p_0.

$$\hat{v} - p_0 = \delta^{K+1}\pi_0(\bar{v} - p_0)(1 - (1-\mu)^{K+1}) + \delta^{K+1}(1-\mu)^{K+1}(\hat{v} - p_0).$$

Here one has used the stationarity requirement. This can be solved for p_0 uniquely to give

$$p_0 = \frac{\pi_0 - \sigma}{1 - \sigma}\bar{v},$$

where σ is given by

$$\sigma(\Delta, K, \pi_0) = \frac{\delta^{K+1}(1 - (1-\mu)^{K+1})\pi_0}{1 - \delta^{K+1}(1-\mu)^{K+1}}.$$

Note in particular that $\sigma \in (0, \pi_0)$ for all Δ. p_0 remains bounded away from 0 as $\Delta \to 0$ as long as in the HFL σ remains bounded away from π_0, as $\Delta \to 0$.

It is easy to see that $\frac{\sigma(\Delta, K, \pi_0)}{\pi_0}$ remains bounded away from one, as $\Delta \to 0$.

Namely, this follows from

$$\frac{\sigma(\Delta, K, \pi_0)}{\pi_0} = \frac{e^{-r\Delta(K(\Delta)+1)}\left(1 - e^{-\lambda\Delta(K(\Delta)+1)}\right)}{1 - e^{-(r+\lambda)\Delta(K(\Delta)+1)}},$$

and the fact that $K(\Delta)\Delta$ remains bounded as $\Delta \to 0$. If $K(\Delta)\Delta$ remains bounded away from zero along a subsequence, then this is clear. Otherwise, one uses the elementary limit statement: $\frac{1-e^{-rt}}{1-e^{-(r+\lambda)t}} \to \frac{r}{r+\lambda}$ as $t \to 0$.

□

Proof of Proposition 5. Note that the proof of Claims 1,2,3 in the proof of Proposition 29 only uses the definition of sequential screening of valuations as well as the fact that $0 < \underline{v} < \bar{v}$. Therefore it can easily be adapted to give non-existence of sequential screening dynamics for the case of deterministic variable costs. This follows because with deterministic variable costs Buyer is never able to learn perfectly, because perfect learning is prohibitively costly. This results in Buyer types strictly above zero after every Seller-history, even in the case $\underline{v} = 0$.

By an analogous logic to the proof of the first part of Proposition 29, the result remains true for the case of stochastic fixed costs, provided that $\underline{v} > 0$. □

A.1.3 High-price stationary equilibria and Buyer payoff near HFL

First, I give a formal definition of a high-price equilibrium.

Definition 9. Say that a PBE is a high-price equilibrium if, on path, whenever it is Seller's turn to quote a price, he asks with probability one for the highest buyer type that has positive probability using public information at that moment in time.

Next I show that an equilibrium features zero Buyer payoff if and only if it is a high-price equilibrium.

Lemma 5. *An equilibrium with zero Buyer payoff must be a high-price equilibrium.*

Proof. To see this, suppose that Seller after an on-path history of some length t, asks with positive probability for a price lower than the highest valuation she deems feasible with positive probability after that history. Buyer with the highest valuation can then achieve, with positive probability, a strictly positive payoff in the continuation game after that history. Namely, she can wait until period t and accept any price strictly below her reservation price. Since in an equilibrium of the game Buyer can always ensure a non-negative payoff after each non-terminal history which is on path, an overall positive surplus would result for Buyer, whenever the equilibrium is not a high-price equilibrium. □

Proof of Proposition 3. The proof is constructive. I find parameter restrictions which ensure that the Seller-optimality conditions (1.4) and (1.5) are satisfied. The last part of the proof specifies off-path beliefs.

Note that $U(\gamma) > \gamma$ and it is strictly increasing in γ. Denote also $U(0) = \frac{\mu\pi_0}{1-\mu(1-\pi_0)}$. Note also that the map $\gamma \to B(\gamma, q) = \frac{\gamma(1-q)}{\gamma(1-q)+1-\gamma}$ has $B(\gamma, q) < \gamma$ and that it is decreasing in q (equivalently increasing in $(1-q)$) as well as increasing in γ. In particular, $\lim_{\gamma\to 0} B(\gamma, q) = 0$ uniformly in $q \in [0, 1]$.[7]

One can solve for the q in (1.3) to get

$$q(\gamma) = \frac{\mu\pi_0 + \gamma\mu(1-\pi_0)}{\gamma + (1-\gamma)\mu\pi_0} \in (0, 1]. \tag{A.10}$$

It holds $q(0) = 1$ and $q(1) = \mu$ and q is strictly decreasing in γ.

For future use, let us also note

$$U(\gamma)q(\gamma) = \frac{\mu\pi_0 + \gamma\mu(1-\pi_0)}{1-(1-\gamma)\mu(1-\pi_0)}. \tag{A.11}$$

$U(\gamma)q(\gamma)$ is increasing in μ and it is increasing in γ.[8] One has

[7]Just note that $B(\gamma, q) \leq \frac{\gamma}{1-\gamma}$, for all $q \in [0, 1]$.

[8]Namely it holds

$$\frac{d}{d\gamma}U(\gamma)q(\gamma) = \frac{\mu(1-\pi_0)(1-\mu)}{(1-(1-\gamma)\mu(1-\pi_0))^2}.$$

$U(\gamma)q(\gamma)_{|\gamma=0} = U(0)$ as well as $U(\gamma)q(\gamma)\to 0$ as $\mu\to 0$, uniformly in γ.

Moreover it holds

$$\lim_{\mu\to 1-} U(\gamma)q(\gamma) = \frac{\pi_0 + \gamma(1-\pi_0)}{1-(1-\gamma)(1-\pi_0)} = 1.$$

The payoff function $W(\gamma)$ of Seller from $t=2$ on satisfies the recursion

$$\begin{aligned} W(\gamma) &= (\gamma + (1-\gamma)\pi_0\mu)\bar{v}\cdot q(\gamma) \\ &\quad + (1-\gamma)\mu(1-\pi_0)\underline{v} \\ &\quad + \delta((\gamma+(1-\gamma)\mu\pi_0)(1-q(\gamma)) + (1-\gamma)(1-\mu))W(\gamma). \end{aligned}$$

After algebra this results in

$$W(\gamma) = \frac{(\gamma + (1-\gamma)\pi_0\mu)\bar{v}\cdot q(\gamma) + (1-\gamma)\mu(1-\pi_0)\underline{v}}{1-\delta((\gamma+(1-\gamma)\mu\pi_0)(1-q(\gamma)) + (1-\gamma)(1-\mu))}.$$

Because of the formula for $q(\gamma)$ this simplifies to

$$W(\gamma) = \frac{\mu}{1-\delta(1-\mu)}((\pi_0+\gamma(1-\pi_0))\bar{v}+(1-\gamma)(1-\pi_0)\underline{v}) = \frac{\mu}{1-\delta+\delta\mu}(\gamma\bar{v}+(1-\gamma)\hat{v}).$$

W is strictly increasing in γ with derivative

$$\frac{dW(\gamma)}{d\gamma} = \frac{\mu(1-\pi_0)(\bar{v}-\underline{v})}{1-\delta(1-\mu)}.[9]$$

Now I show that the binding constraint for Seller-optimality is that from $t=1$.

---

[9]In particular, it is also Lipschitz continuous in (μ,γ). We have $W(0) = \frac{\mu\hat{v}}{1-\delta(1-\mu)}$ and $W(1) = \frac{\mu\bar{v}}{1-\delta(1-\mu)}$.

But because γ is bounded above, we know that the highest value of W is actually

$$W(U(0)) = \frac{\mu}{1-\delta+\delta\mu}(U(0)\bar{v} + (1-U(0))\hat{v}).$$

Claim. Whenever q_1 is such that Seller-optimality condition for $p = \bar{v}$ at $t = 1$ holds, seller optimality holds also for $t \geq 2$.

Proof of the Claim. Note that, because $U(\gamma)q(\gamma)$ is increasing, it holds $U(0)q_1 \leq U(0) \leq U(\gamma)q(\gamma)$. Moreover, it holds that $\delta W(\gamma) < \bar{v}$, uniformly for all γ. Namely,

$$\delta W(\gamma) \leq \delta W(1) = \frac{\delta\mu\bar{v}}{1 - \delta + \delta\mu} < \bar{v}.$$

Given that $\gamma \mapsto U(\gamma)q(\gamma)$ is strictly monotonic, the result in the claim follows.

End of the proof of the Claim.

Note that for $q_1 = 0$ Seller-optimality condition at $t = 1$ is never satisfied, whereas for $q_1 = 1$ it is satisfied whenever

$$(C - high) \qquad U(0)\bar{v} + (1 - U(0))\delta W(U(0)) > \hat{v}.$$

By use of continuity this gives a sufficient condition for existence of the full-extraction PBE.

One can write $(C - high)$ as

$$U(0)(1 - \delta + \delta\mu + \delta\mu(1 - U(0)))\bar{v} > (1 - \delta\mu(1 - U(0))^2)\hat{v}.$$

As $\lambda \to \infty$ (equivalently $\mu \to 1$) it holds $U(0) \to 1$ and so $(C - high)$ in the limit $\mu \to 1$, and fixed other parameters, becomes the condition $\bar{v} > \hat{v}$. This is always true, by assumptions on the primitives.

It remains to specify off-path play. I focus on specifying play only after single deviations.[10] For any more complicated deviations of the players, general existence theorems show existence of some continuation PBE after such histories.

Seller off-path. Suppose that at the beginning of a period the current belief of Seller is $\hat{\gamma} \neq 0, \gamma$ (in particular play is at a period $t \geq 2$). Given belief about

[10] I specify off-path play only in this Proposition. For the sake of length, in other proofs in the following where stationary equilibria are constructed, I skip specifying off-path play whenever it follows an analogous logic to the one given here.

continuation play, and the fact that parameters are such that for given continuation play, Seller chooses optimally $p = \bar{v}$ at $\gamma = 0$, it holds that: incentives to charge $p = \bar{v}$ are even stricter when $\hat{\gamma} > \gamma$ and $t \geq 2$ as well as for all $\hat{\gamma} > 0$ when $t = 1$.

Let us consider now the other case: $\hat{\gamma} \in [0, \gamma)$ and $t \geq 2$. Because the beliefs need to be derived by equilibrium strategies of Buyer whenever possible, there are two cases to consider:

- Case 1: $\hat{\gamma}$ came about after a past rejection of a price above $\bar{v}$. In this case, in the very next period after the rejection the belief, from the specification of strategies, would be $U(\gamma)$. Given continuation play, this would mean that $\hat{\gamma} > \gamma$, a case already considered above.

- Case 2: $\hat{\gamma}$ came about after a past rejection of a price $p \in (\hat{v}, \bar{v})$. In this case, in the next period Seller should have started with belief 0 and in any future period the starting belief of Seller should be the stationary γ. Thus this case is covered by the specification of strategies on path.

Buyer off-path. - If Seller deviates to some price $p \neq \bar{v}$ every Buyer type responds according to her reservation price strategy: if the price is strictly lower than the reservation price of her type Buyer accepts immediately, if it is strictly higher she rejects.

- If Buyer of type $\underline{v}$ has not disclosed in the past, she still remains indifferent between disclosing and not disclosing. Prescribe disclosure in the current period after such a history.

- If Buyer has disclosed good news and Seller has not asked yet for all the surplus, Buyer accepts any price weakly lower than $\bar{v}$ and rejects any price strictly higher. This is optimal given the anticipation that the continuation payoff of Buyer is zero in the continuation game. □

Proof of Proposition 4. First I prove a couple of auxiliary claims.

Claim 1. An equilibrium with zero Buyer payoff necessarily has Seller quote a price $p = \bar{v}$ as long as bargaining goes on and there is no disclosure.

Proof of Claim 1. Suppose this is not the case and denote by t the first period in which, after no disclosure, the price quoted by Seller is lower than $\bar{v}$ with positive probability. W.l.o.g. assume that the prices up to and including period t upon non-disclosure are in $(\hat{v}, \bar{v}]$.[11] Suppose that Buyer uses the following strategy: unless there is informational arrival and $\theta = \underline{v}$, wait until period t and in period t accept the current price if and only if it is weakly above Buyer's period-t estimate of the value of the good. If information arrives before or at date t and $\theta = \underline{v}$, disclose immediately. For periods after t, disclose only if $\theta = \underline{v}$ (and immediately) and otherwise accept only if price is strictly below current estimate of value. Under this strategy Buyer has a strictly positive payoff with positive probability in period t, and otherwise non-negative payoff in all other histories. Overall, this leads to a contradiction to the assumption that the equilibrium payoff of Buyer is zero.

End of proof of Claim 1.

Claim 2. Under any disclosure equilibrium in which Seller quotes $p = \bar{v}$ on path, the overall payoff of Seller in the HFL is less than $\frac{\lambda}{\lambda+r}\hat{v}$.

Proof of Claim 2. Fix such an equilibrium as in the statement of the claim for a $\Delta > 0$. Denote by $A(\Delta)$ the random variable giving the agreement time and by $L(\Delta)$ the random variable giving the time at which Buyer learns. Note that $L(\Delta)$ is geometrically distributed. Because on path the price quoted upon non-disclosure is always $\bar{v}$ it holds for both of $\theta = \underline{v}, \bar{v}$ that $A(\Delta) \geq L(\Delta)$ almost surely. Moreover, Seller receives a payoff of either $\bar{v}$ or $\underline{v}$ at time $A(\Delta)$ under these strategies. One calculates

[11] Otherwise it is easy to show Buyer receives positive payoff with positive probability already in period one.

$$\begin{aligned}\text{Seller-payoff} &= \bar{v}\mathbb{E}[\delta^{A(\Delta)}, \theta = \bar{v}] + \underline{v}\mathbb{E}[\delta^{A(\Delta)}, \theta = \underline{v}] \\ &\leq \bar{v}\mathbb{E}[\delta^{L(\Delta)}, \theta = \bar{v}] + \underline{v}\mathbb{E}[\delta^{L(\Delta)}, \theta = \underline{v}] \\ &= \bar{v}\pi_0\mathbb{E}[\delta^{L(\Delta)}|\theta = \bar{v}] + \underline{v}(1-\pi_0)\mathbb{E}[\delta^{L(\Delta)}|\theta = \underline{v}] \\ &= \hat{v}\mathbb{E}[\delta^{L(\Delta)}].\end{aligned}$$

Here, the first inequality follows from $A(\Delta) \geq L(\Delta)$ which holds almost surely under the assumptions made, whereas the last equality follows from the fact that the arrival of the opportunity to learn is independent of θ. Now I show that in the HFL $\mathbb{E}[\delta^{L(\Delta)}]$ converges to $\frac{\lambda}{\lambda+r}$. Recall that $\delta = \delta(\Delta) = e^{-r\Delta}$ and that $L(\Delta)$ is geometrically distributed over $1, 2, \ldots$ with probability of success given by $\mu = 1 - e^{-\lambda\Delta}$. It follows

$$\mathbb{E}[\delta^{L(\Delta)}] = \frac{\mu}{1-\mu}\sum_{t=1}^{\infty} e^{-(r+\lambda)t} = \frac{\mu}{1-\mu}\frac{e^{-(r+\lambda)\Delta}}{1-e^{-(r+\lambda)\Delta}}.$$

One uses that $\frac{1-e^{-x\Delta}}{\Delta} \to x$ as $\Delta \to 0$ for all $x > 0$ to finish the proof of the claim.

End of proof of Claim 2.

Now I finish the proof of the Proposition. The condition of optimality of $p = \bar{v}$ in the first period under the assumption of a disclosure equilibrium is given by

$$U(0)(\bar{v})q_1\bar{v} + (1 - U(0)(\bar{v})q_1)\delta W \geq U(0)(\{\hat{v}\})\hat{v}, \tag{A.12}$$

where q_1 is the probability with which Buyer of type $\bar{v}$ accepts the price $\bar{v}$ in period 1 and W is the continuation payoff of Seller upon non-disclosure and rejection of price $\bar{v}$ in period 1. Here $U(0) \in \mathcal{P}(\{\underline{v}, \bar{v}, \hat{v}\})$ is the belief of Seller in $t = 1$ over Buyer types after non-disclosure. It is given explicitly by

$$U(0)(\bar{v}) = \frac{\mu\pi_0}{1 - q_d\mu(1-\pi_0)},$$
$$U(0)(\hat{v}) = \frac{1-\mu}{1 - q_d\mu(1-\pi_0)},$$
$$U(0)(\underline{v}) = \frac{(1-q_d)\mu(1-\pi_0)}{1 - q_d\mu(1-\pi_0)},$$

where $q_d \in [0,1]$ is the probability with which type $\underline{v}$ discloses if Buyer learns in $t = 1$ that $\theta = \underline{v}$.

In particular, it holds $U(0)(\{\underline{v}, \bar{v}\}), U(0)(\{\bar{v}\})q_1 \to 0$ as $\Delta \to 0$ uniformly in q_1, q_d. It holds

$$\delta W = \left(\frac{\text{Seller-payoff} - q_d\mu(1-\pi_0)\underline{v}}{1 - q_d\mu(1-\pi_0)} - U(0)(\bar{v})q_1 \right) \frac{1}{1 - U(0)(\bar{v})q_1}.$$

It follows that
$$\limsup_{\Delta \to 0} \delta W \leq \frac{\lambda}{r+\lambda}\hat{v} < \hat{v}.$$

Since the left-hand side of (A.12) in the limit converges to $\hat{v}$ this shows that the condition of optimality of $p = \bar{v}$ upon non-disclosure at $t = 1$ cannot be satisfied for $\Delta \to 0$ (deviating to $\hat{v} - \varepsilon$ for small and positive ε is strictly better for Seller near the HFL).

This finishes the proof of the Proposition. □

A.1.4 Proof of Theorem 1

I prove Theorem 1 through a series of Lemmas and Propositions. I focus first on the case of strongly stationary equilibria with mixed pricing before analysing pure pricing. This is because their analysis is more involved and several proof steps for the case of pure pricing are similar to the case of mixed pricing. The same organizational principle holds true for all results characterizing HFL of strongly stationary equilibria in the paper.

The following Lemma is helpful for all existence results of strongly stationary equilibria in the paper.

Lemma 6. *The map* $[0,1] \ni q \mapsto \frac{q}{1-\delta(1-\mu)(1-(1-p)q)}$ *is strictly increasing and strictly concave for any* $p \in [0,1)$.

Proof. One calculates that

$$\frac{\partial}{\partial q}\left(\frac{q}{1-\delta(1-\mu)(1-(1-p)q)}\right) = \frac{1-\delta(1-\mu)}{(1-\delta(1-\mu)(1-(1-p)q))^2},$$

which shows that this function is strictly increasing and strictly concave for $q \in [0,1]$. □

The case of mixed pricing

An auxiliary remark follows which is helpful in the proof of Theorem 1.

Remark 3. *The function* $g : (0,1) \to [0,1], p \to \frac{\delta(1-p)}{1-\delta p}$ *is decreasing in* p *and converges uniformly to 1, as* $\Delta \to 0$.

Proof. It suffices to note that $0 \leq 1 - \frac{\delta(1-p)}{1-\delta p} \leq 1-\delta$. □

The prices p_H, p_L are required to satisfy

$$\underline{v} < p_L < p_H < \bar{v}, \text{ and } p_L < \hat{v}. \tag{A.13}$$

Suppose Seller mixes among the prices p_H, p_L with probability $(p, 1-p)$. The triple (p_H, p_L, p) necessarily satisfies

$$\begin{aligned}\hat{v} - p_L = \delta(&\mu\pi_0(p(\bar{v}-p_H) + (1-p)(\bar{v}-p_L)) \\ &+ (1-\mu)(1-p)(\hat{v}-p_L) + (1-\mu)p\delta V_L),\end{aligned}$$

where V_L is the continuation payoff from starting a period in the stationary phase with type $\hat{v}$.

It holds due to the reservation pricing feature of PBEs (part 2 of Lemma 1) that $\delta V_L = \hat{v} - p_L$. This leads to

$$\hat{v} - p_L = \frac{\mu\pi_0(\bar{v} - p_H) + \delta(1-\mu)(1-p)(\hat{v} - p_L)}{1 - \delta p(1-\mu)}.$$

Some algebra leads to

$$\hat{v} - p_L = \frac{\mu\pi_0(\bar{v} - p_H)}{1 - \delta(1-\mu)},$$

so that in total the two reservation-pricing relations about the prices p_H, p_L are given by

$$\frac{\hat{v} - p_L}{\bar{v} - p_H} = \frac{\mu\pi_0}{1 - \delta(1-\mu)}, \quad \frac{\bar{v} - p_L}{\bar{v} - p_H} = \frac{1 - \delta p}{\delta(1-p)}.$$

Lemma 7. *1) In any strongly stationary equilibrium with mixed pricing, p_H, p_L as a function of p are given by*

$$p_L(p) = \hat{v} - \frac{\mu\pi_0 \frac{\delta(1-p)}{1-\delta p}}{1 - \delta(1-\mu) - \mu\pi_0 \frac{\delta(1-p)}{1-\delta p}}(\bar{v} - \hat{v}), \tag{A.14}$$

$$p_H(p) = \bar{v} - \frac{(1 - \delta(1-\mu))\frac{\delta(1-p)}{1-\delta p}}{1 - \delta(1-\mu) - \mu\pi_0 \frac{\delta(1-p)}{1-\delta p}}(\bar{v} - \hat{v}). \tag{A.15}$$

2) p_H, p_L are strictly increasing in $p \in [0,1]$.

3) The price spread $ps(p) = p_H(p) - p_L(p)$ is given by

$$ps(p) = \left[1 - \frac{(1 - \delta(1-\mu))\frac{\delta(1-p)}{1-\delta p} - \mu\pi_0 \frac{\delta(1-p)}{1-\delta p}}{1 - \delta(1-\mu) - \mu\pi_0 \frac{\delta(1-p)}{1-\delta p}}\right](\bar{v} - \hat{v}). \tag{A.16}$$

Proof. This is straightforward algebra. Solving the two reservation price relations of the two Buyer types as a function of the mixing probability p leads to (A.14) and (A.15). The rest is also straightforward calculations. □

One checks easily by taking first derivatives, that the price spread $p_H(p) - p_L(p)$ is strictly increasing in p. Moreover, one checks easily that for every $p \in (0,1)$ it holds

$$\underline{v} < p_L(p) < \hat{v} < p_H(p) < \bar{v}.$$

The boundary values of p_H, p_L are given by (p_H, p_L are continuous in $p \in [0,1]$).

$$p_L(p=0) = \hat{v} - \frac{\mu\pi_0\delta}{1-\delta(1-\mu)-\mu\pi_0\delta}(\bar{v}-\hat{v})$$

$$p_H(p=0) = \bar{v} - \frac{\delta(1-\delta(1-\mu))}{1-\delta(1-\mu)-\mu\pi_0\delta}(\bar{v}-\hat{v})$$

$$p_L(p=1) = \hat{v}, \qquad p_H(p=1) = \bar{v}. \tag{A.17}$$

As a final boundary value I note down the price spread at $p = 0$.

$$p_H(0) - p_L(0) = (\bar{v} - \underline{v})(1-\pi_0)\frac{(1-\delta(1-\mu))(1-\delta)}{1-\delta(1-\mu)-\mu\pi_0\delta}.$$

Rewriting Seller's reservation price relation (1.7) leads to the definition of a function $f(p,q)$ which satisfies

$$\frac{f(p,q)}{\mu(1-U(0))} = \frac{\pi_0}{1-\mu}ps(p) - \frac{q}{1-\delta(1-\mu)(1-(1-p)q)}\delta\left(\frac{1-\delta(1-\mu)}{\mu}p_L(p) - \pi_0\hat{p} - (1-\pi_0)\underline{v}\right).$$

Existence of a two-price equilibrium with mixing probability p for Seller is tantamount to finding a root in q of $\frac{f(p,q)}{\mu(1-U(0))}$.

First let us show the following auxiliary lemma.

Lemma 8. *It holds as* $\Delta \to 0$ *uniformly in* $p \in [0,1]$ *that*

A.

$$p_L(p) \quad \to \quad \psi := \frac{r\pi_0\bar{v} + (r+\lambda)(1-\pi_0)\underline{v}}{r+(1-\pi_0)\lambda},$$

B.

$$p_H(p) - p_L(p) \to 0.$$

Proof. A. Note that

$$p_L(p) = \hat{v} - \frac{\mu\pi_0 \frac{\delta(1-p)}{1-\delta p}}{1-\delta(1-\mu) - \mu\pi_0 \frac{\delta(1-p)}{1-\delta p}}(\bar{v} - \hat{v}) = \hat{v} - \frac{\frac{\mu}{\Delta}\pi_0(1+O(\Delta))}{\frac{1-\delta(1-\mu)}{\Delta} - \frac{\mu}{\Delta}\pi_0(1+O(\Delta))}(\bar{v} - \hat{v})$$
$$\rightarrow \hat{v} - \frac{\lambda\pi_0}{r+(1-\pi_0)\lambda}(\bar{v} - \hat{v}) = \psi.$$

Here the last step follows from algebra and the uniformity in p comes from Remark 3.

B. Use (A.16) from Lemma 7 to estimate that

$$\frac{(1-\delta(1-\mu))\frac{\delta(1-p)}{1-\delta p} - \mu\pi_0\frac{\delta(1-p)}{1-\delta p}}{1-\delta(1-\mu) - \mu\pi_0\frac{\delta(1-p)}{1-\delta p}} = (1+O(\Delta))\frac{\frac{1-\delta(1-\mu))}{\Delta} - \frac{\mu}{\Delta}\pi_0}{\frac{1-\delta(1-\mu)}{\Delta} - \frac{\mu}{\Delta}\pi_0(1+O(\Delta))}$$
$$\rightarrow 1.$$

Here the uniformity of the convergence of $\frac{\delta(1-p)}{1-\delta p}$ as well as the fact that $\frac{1-\delta}{\Delta}\rightarrow r$, $\frac{\mu}{\Delta}\rightarrow\lambda$ as $\Delta\rightarrow 0$ have been used. This and (A.16) establishes the result. □

I note a sharper result for B. which is also used later. In particular, this Remark also proves that the price spread disappears in the HFL.

Remark 4. *It holds* $\frac{ps(\Delta)}{\Delta}\rightarrow\frac{r+\lambda}{r+(1-\pi_0)\lambda}\frac{r}{1-p}\pi_0(1-\pi_0)(\bar{v}-\underline{v})$.

Proof of Remark 4. Note that after some algebra

$$\frac{ps(p,\Delta)}{\Delta} = \pi_0(\bar{v} - \hat{v})\frac{1}{1-\delta p}\frac{1-\delta}{\Delta}\frac{\frac{1-\delta(1-\mu)}{\Delta}}{\frac{1-\delta(1-\mu)}{\Delta} - \frac{\mu}{\Delta}\pi_0},$$

from which the result follows immediately, because $\bar{v} - \hat{v} = (1-\pi_0)(\bar{v} - \underline{v})$. □

Now I establish existence of mixed stationary equilibria for any $p \in (0,1)$, whenever Δ small enough.

Proposition 30. *For any $p \in (0,1)$ there exist strongly stationary equilibria for all Δ small enough. Moreover, for fixed Δ and fixed p such that existence is ensured, the prices p_L, p_H and the mixing probability q of the type $\hat{v}$ when she faces p_L are unique.*

Proof. Fix $p \in (0,1)$.

Existence. Note that $\frac{f(p,0)}{\mu(1-U(0))} = \frac{\pi_0}{1-\mu}(p_H(p) - p_L(p)) > 0$ by construction. Thus, it suffices to show that $\frac{f(p,1)}{\mu(1-U(0))} < 0$ for all Δ small enough. First note that

$$\frac{1-\delta(1-\mu)}{\delta\mu} p_L(p,\Delta) - \pi_0 \hat{p} - (1-\pi_0)\underline{v} \quad \to \quad \left(\frac{r+\lambda}{\lambda} - \pi_0\right)\psi - (1-\pi_0)\underline{v} > 0. \tag{A.18}$$

The last inequality follows from the fact that $\psi > \underline{v}$, as one can easily check by looking at the statement in part A. of Lemma 8. (A.18) and Lemma 6 show that $[0,1] \ni q \mapsto \frac{f(p,q)}{\mu(1-U(0))}$ is strictly decreasing and strictly convex in q. From the results in Lemma 8 (namely part B. there) one arrives at

$$\lim_{\Delta \to 0} \frac{f(p,q)}{\mu(1-U(0))} = -\frac{1}{1-p}\left(\frac{r}{\lambda}\psi + (1-\pi_0)(\psi - \underline{v})\right) < 0. \tag{A.19}$$

Here, one uses that $\delta(1-\mu) = 1 + O(\Delta)$. Fixing some $\bar{\Delta}(p)$ where $\frac{f(p,1)}{\mu(1-U(0))} < 0$ with $\Delta < \bar{\Delta}(p)$ it follows that there exists a unique zero for $\frac{f(p,q)}{\mu(1-U(0))}$ whenever $\Delta < \bar{\Delta}(p)$, denoted $q(\Delta,p)$. Moreover, $q(\Delta,p)$ is in $(0,1)$.

Uniqueness. The arguments above show uniqueness of $q(\Delta,p)$. Uniqueness of the prices $p_L(\Delta,p)$ and $p_H(\Delta,p)$ follows from (A.15) and (A.14) in Lemma 7. □

Next I characterize the HFL of mixed pricing equilibria. This involves several steps. First, one calculates explicitly $q(\Delta,p)$ for fixed p and all Δ small enough. From the condition $\frac{f(p,q)}{\mu(1-U(0))} = 0$ it follows that $q(\Delta,p)$ satisfies the relation

$$\frac{\pi_0}{1-\mu} ps(p) = \frac{q(\Delta,p)}{1-\delta(1-\mu)(1-(1-p)q(\Delta,p))} h(p,\Delta,\pi_0,\underline{v}),$$

The function $h(p,\Delta,\pi_0,\underline{v})$ here is strictly positively valued and it becomes the constant function $h^* := \frac{r}{\lambda}\psi + (1-\pi_0)(\psi - \underline{v})$ in the HFL. This follows from straightforward algebra. Solving explicitly for $q(\Delta, p)$ one arrives at

$$q(\Delta, p) = \frac{(1 - \delta(1 - \mu))G(p, \Delta))}{1 - \delta(1 - \mu)(1 - p)G(p, \Delta)},$$

where

$$G(\Delta, p) = \frac{\pi_0}{1 - \mu} \frac{ps(p)}{h(p, \Delta, \pi_0, \underline{v})}.$$

Using Remark 4 one calculates that

$$\frac{G(\Delta, p)}{\Delta} \rightarrow G^*(p) := \frac{r}{1 - p} \frac{\pi_0}{h^*} \frac{r + \lambda}{r + (1 - \pi_0)\lambda} \pi_0(1 - \pi_0)(\bar{v} - \underline{v}).$$

It follows that

$$\frac{q(\Delta, p)}{\Delta^2} = \frac{\frac{1 - \delta(1 - \mu)}{\Delta} \frac{G(\Delta, p)}{\Delta}}{1 - \delta(1 - \mu)(1 - p)G(p, \Delta)},$$

which leads to the HFL statement

$$\frac{q(\Delta, p)}{\Delta^2} \rightarrow (r + \lambda)G^*(p), \quad \Delta \rightarrow 0. \tag{A.20}$$

The date of agreement is a geometric random variable with success probability $1 - (1 - \mu)(1 - (1 - p)q(\Delta, p))$. One calculates using (A.20)

$$\frac{1 - (1 - \mu)(1 - (1 - p)q(\Delta, p))}{\Delta} = \frac{\mu}{\Delta} + (1 - p)\frac{q(\Delta, p)}{\Delta^2} \cdot \Delta \rightarrow \lambda.$$

In all, the expected delay in real time converges to $\frac{1}{\lambda}$, irrespective of p.

Finally, one calculates the stationary payoffs of Buyer and Seller. Recall the relation

$$V_\Delta(q, p) = \frac{\mu\pi_0(p_H p + p_L(1 - p)) + (1 - \mu)(1 - p)qp_L + \mu(1 - \pi_0)\underline{v}}{1 - \delta(1 - \mu)(1 - (1 - p)q)},$$

for Seller's stationary payoff. Using Lemma 8 and (A.20) one arrives easily at the limit

$$V_S = \frac{\lambda}{r+\lambda}(\pi_0\psi + (1-\pi_0)\underline{v}), \tag{A.21}$$

for the payoff of Seller in the HFL. The payoff of Buyer in any two-price stationary equilibria satisfies

$$V_{\Delta,B}(q,p) = \frac{\mu\pi_0(\bar{v}-\hat{p}) + (1-\mu)(1-p)q(\hat{v}-p_L)}{1-\delta(1-\mu)(1-(1-p)q)}.$$

By similar steps as the case of Seller's payoff this converges in the HFL to

$$V_B = \frac{\lambda}{r+\lambda}\pi_0(\bar{v}-\psi). \tag{A.22}$$

Note that V_B, V_S are independent of $p \in (0,1)$. Straightforward algebra leads to the measure of inefficiency in the HFL given by

$$\hat{v} - (V_B + V_S) = \frac{r}{r+\lambda}\hat{v}.$$

THE CASE OF PURE PRICING

For simplicity of exposition define in the following the function $\zeta : [0,1] \times (0\infty) \rightarrow \mathbb{R}_+$ by

$$\zeta(q,\Delta) = \frac{q}{1-\delta(1-\mu)(1-q)}.$$

After algebra one can rewrite the Seller optimality condition in (1.6) as

$$\frac{1-U(0)}{U(0)}\zeta(q,\Delta)\frac{p_L-\delta V_\Delta(q,0)}{1-\delta}(1-\delta(1-\mu)(1-q)) \geq \bar{v}-p_L. \tag{A.23}$$

Lemma 9. *1) Fix any $\Delta > 0$. The reservation price relation of type $\hat{v}$ is solvable for a unique $p_L(\Delta)$. It holds $p_L(\Delta) \rightarrow \psi$, as $\Delta \rightarrow 0$.*

2) Let $q(\Delta)$ satisfy $\frac{q(\Delta)}{\Delta} \to \varkappa$ for some $\varkappa \in [0, \infty]$. It holds

$$\zeta(q(\Delta), \Delta) \to \frac{\varkappa}{\varkappa + r + \lambda}, \quad \Delta \to 0,$$

where $\frac{\varkappa}{\varkappa+r+\lambda}$ is to be understood as equal to 1, if $\varkappa = \infty$.

3) For any sequence $q(\Delta)$, $\Delta \to 0$ it holds

$$\frac{p_L(\Delta) - \delta V_\Delta(q(\Delta), 0)}{1 - \delta}(1-\delta(1-\mu)(1-q(\Delta))) \to \psi + \frac{\lambda}{r}(1-\pi_0)(\psi - \underline{v}) > 0, \quad \Delta \to 0.$$

Proof. 1) One solves explicitly the reservation pricing relation

$$\hat{v} - p_L = \frac{\delta\mu}{1 - \delta + \delta\mu}\pi_0(\bar{v} - p_L),$$

to get

$$p_L(\Delta) = \frac{\hat{v} - \frac{\delta\mu}{1-\delta+\delta\mu}\pi_0\bar{v}}{1 - \frac{\delta\mu}{\delta\mu+(1-\delta)}\pi_0}.$$

Limit algebra leads to the conclusion that $p_L(\Delta) \to \psi$.

2) It holds for every $\Delta > 0$ that

$$\zeta(q(\Delta), \Delta) = \frac{\frac{q(\Delta)}{\Delta}}{\frac{1-\delta(1-\mu)}{\Delta} + \frac{q(\Delta)}{\Delta}\delta(1 - \mu)}.$$

The rest is simple limit algebra.

3) Simple algebra leads to

$$\begin{aligned} &\frac{p_L(\Delta) - \delta V_\Delta(q(\Delta), 0)}{1 - \delta}(1 - \delta(1 - \mu)(1 - q(\Delta))) \\ &= p_L\frac{1 - \delta(1 - \mu)}{1 - \delta} - (1 - \mu)q(\Delta)p_L - \frac{\mu}{1 - \delta}(\pi_0 p_L + (1 - \pi_0)\underline{v}). \end{aligned}$$

The rest is simple limit algebra. □

Lemma 9 leads to the following result about solvability of (A.23).

Lemma 10. *Fix any $\underline{v}, \bar{v}, \pi_0, r, \lambda$ as in the setup of the model in section 1.2. Furthermore, fix any $q \in (0,1)$.*

Then (1.6) is solvable whenever Δ is small enough.

Proof. This is an easy consequence of the fact that $\frac{1-U(0)}{U(0)} = \frac{1-\mu}{\mu\pi_0} \to \infty$ as $\Delta \to 0$ and of results in Lemma 9. □

Lemma 10 and the preceding calculations show existence of strongly stationary equilibria with pure pricing for any baseline game parameters in section 1.2. In particular, because $\frac{1-U(0)}{U(0)} \to \infty$, one can always pick a sequence of positive probabilities $q(\Delta) > 0$, $\Delta \to 0$ such that $\frac{q(\Delta)}{\Delta} \to \varkappa$ as $\Delta \to 0$, for any $\varkappa \in [0, \infty]$. This shows that for any $\varkappa \in [0, \infty]$ there exists HFL of strongly stationary equilibria with pure pricing corresponding to $\varkappa$.[12]

Uniqueness up to the acceptance probability q near the HFL for strongly stationary equilibria with mixed pricing is immediate from the arguments above.

Next I characterize HFL payoffs depending on $\varkappa$. It holds for Seller payoff in strongly stationary equilibrium with pure pricing

$$V_\Delta(q(\Delta), 0) = \frac{(\mu\pi_0 + (1-\mu)q(\Delta))p_L + \mu(1-\pi_0)\underline{v}}{1 - \delta(1-\mu)(1-q(\Delta))}.$$

Limit algebra shows that in the HFL this converges to

$$V_S(\varkappa) = \begin{cases} \psi, & \text{if } \varkappa = \infty, \\ \frac{(\lambda\pi_0 + \varkappa)\psi + \lambda(1-\pi_0)\underline{v}}{r+\lambda+\varkappa}, & \text{if } \varkappa \in [0, \infty). \end{cases}$$

For Buyer payoff one calculates

$$V_{B,\Delta}(q(\Delta), 0) = \frac{\mu\pi_0(\bar{v} - p_L) + (1-\mu)q(\Delta)(\hat{v} - p_L)}{1 - \delta(1-\mu)(1-q(\Delta))}.$$

[12]For $\varkappa = 0$ one can also use the strongly stationary equilibria with mixed pricing of the previous subsection, whereas for $\varkappa = \infty$ one can pick $q(\Delta) = q_0 \in (0, 1]$ with the understanding that if $q_0 = 1$ the equilibrium satisfies 'divinity in bargaining'.

Limit algebra shows that in the HFL this converges to

$$V_B(\varkappa) = \begin{cases} \hat{v} - \psi, & \text{if } \varkappa = \infty, \\ \frac{\lambda\pi_0\bar{v} + \varkappa\hat{v} - (\lambda\pi_0 + \varkappa)\psi}{r + \lambda + \varkappa}, & \text{if } \varkappa \in [0, \infty). \end{cases}$$

The sum of Buyer and Seller payoffs in the HFL is

$$\begin{cases} \hat{v}, & \text{if } \varkappa = \infty, \\ \frac{\varkappa + \lambda}{\varkappa + \lambda + r}\hat{v}, & \text{if } \varkappa \in [0, \infty). \end{cases}$$

Note that the sum of Buyer and Seller payoffs in the HFL is an increasing function of ψ and converges to its corresponding value $\hat{v}$ as $\varkappa$ moves along a finite sequence which converges to ∞.

Finally, it remains to calculate the expected delay in the HFL for a sequence of strongly stationary equilibria with pure pricing corresponding to some $\varkappa \in (0, \infty]$. Equilibrium construction shows that the agreement date is a geometric random variable with success probability $1 - (1 - \mu)(1 - q(\Delta))$. One calculates thus for the expected delay

$$\frac{\Delta}{1 - (1 - \mu)(1 - q(\Delta))} = \frac{1}{\frac{\mu}{\Delta} + \frac{q(\Delta)}{\Delta}(1 - \mu)} \to \frac{1}{\lambda + \varkappa}, \quad \Delta \to 0,$$

where the limit value is to be understood as zero if $\varkappa = \infty$.

This finishes the proof of Theorem 1.

A.2 Proofs for section 1.3

A.2.1 *'Positive Buyer payoff in every equilibrium with costs'* and its corollaries

I show a slightly more general statement than that of Theorem 2. The generalization consists in

- I also look at the case of endogenous choice of intensity (see section 1.4.1)

- I assume that Buyer can pick *any* two-dimensional experiments.

The two-parametric experiments are modeled as follows.

Remark 5. *A general experiment is given by* $\mathcal{E} : \{\underline{v}, \bar{v}\} \to \Delta(\{H, L\})$. $s \in \{H, L\}$ *is the signal Buyer sees after performing the experiment. An experiment* $\mathcal{E}$ *is fully identified with the two accuracy parameters* $a_H = \mathbb{P}(s = H | \theta = \bar{v})$ *and* $a_L = \mathbb{P}(s = L | \theta = \underline{v})$. *The experiment is uninformative if and only if* $\frac{a_H}{1-a_L} = 1 \Longleftrightarrow \frac{a_L}{1-a_H} = 1$.

Because otherwise one can always relabel signals, one can assume w.l.o.g. that the region of possible accuracy parameters is given by

$$\nabla = \{(a_H, a_L) : a_H, a_L \in [0, 1], a_H + a_L \geq 1\}.$$

I restrict in the rest of this subsection of the appendix to experiments parametrized by pairs (a_H, a_L) *in* ∇.

Define the function $L : int(\nabla) \to \{(l_1, l_2) \in (1, \infty)^2, l_1 + l_2 < l_1 l_2 + 1\}$ *given by* $L(a_H, a_L) = \left(\frac{a_H}{1-a_L}, \frac{a_L}{1-a_H}\right)$. *This map is a diffeomorphism.*[13] *Note that the two variables* (l_1, l_2) *are in* $(1, \infty)$ *and independent of each other, as long as they are different from* 1 *and satisfy the condition* $l_1 + l_2 > l_1 l_2 + 1$.

They correspond to the informativeness of the two signals H, L *w.r.t. the two states. To see this, note e.g. that*

$$l_1 = \frac{a_H}{1 - a_L} = \frac{P(s = H | \theta = \bar{v})}{P(s = H | \theta = \underline{v})}.$$

Let $v(H)$ *be the valuation of the good, starting from prior* π_0, *if* $s = H$ *and analogously* $v(L)$ *the valuation if* $s = L$. *It holds*

$$v(H) = \frac{a_H \pi_0 \bar{v} + (1 - a_L)(1 - \pi_0)\underline{v}}{a_H \pi_0 + (1 - a_L)(1 - \pi_0)}, \quad v(L) = \frac{(1 - a_H)\pi_0 \bar{v} + a_L(1 - \pi_0)\underline{v}}{(1 - a_H)\pi_0 + a_L(1 - \pi_0)},$$

for the Bayesian estimates of the value of good, upon observing a high (H) *or low* (L)

[13] See section on (non-)concavity of value of information in the online appendix for a proof.

signal.

One calculates

$$v(H) - \hat{v} = \pi_0(1-\pi_0)(\bar{v} - \underline{v})\frac{1}{\pi_0 + \frac{1}{\frac{a_H}{1-a_L} - 1}}$$

and

$$\hat{v} - v(L) = \pi_0(1-\pi_0)(\bar{v} - \underline{v})\frac{1}{\frac{1}{1-\frac{1-a_H}{a_L}} - \pi_0}.$$

As straightforward algebra shows, it holds that $v(H) - \hat{v}$ *is increasing and concave in* $\frac{a_H}{1-a_L}$ *and* $\hat{v} - v(L)$ *is increasing and concave in* $\frac{a_L}{1-a_H}$.

It is therefore economically meaningful to put the costs on the pair $(\frac{a_H}{1-a_L}, \frac{a_L}{1-a_H})$. *Namely, one can interpret the differences* $v(H) - \hat{v}$ *and* $\hat{v} - v(L)$ *as 'outputs' from a production process of information in which the 'inputs' are precisely* l_1, l_2.

Say $c : [1, \infty)^2 \to \mathbb{R}$ *is a* cost function on informativeness *if it satisfies*

A. $c(1, 1) = 0$,

B. *c is strictly convex and increasing in each argument,*

C. $\lim_{t\to\infty} c'(t, b) = \lim_{t\to\infty} c'(a, t) = +\infty$ *for every* $a, b \in [1, \infty)$.

Finally, say that $C : [\frac{1}{2}, 1)^2 \to \mathbb{R}_+$ *is a* cost function *if it holds* $C(a_H, a_L) = c(L(a_H, a_L))$ *for all pairs* $(a_H, a_L) \in [\frac{1}{2}, 1)^2$ *and a c which is a cost function on informativeness.*

In the model with deterministic variable costs assume for this section of the appendix only, that Buyer possesses a cost function as defined in Remark 5.

Proposition 31. *There is no high-price equilibrium in either of the following cases:*
- deterministic variable accuracy costs with $\partial_a C(\frac{1}{2}, \frac{1}{2}) = 0$ *for* $a = a_H$ *or* $a = a_L$
- stochastic fixed accuracy costs.

Proof. The proof is by contradiction. Pick a Seller-history *h* on path. Because it is an on-path history and because of part 4) of Lemma 1 (learning is possible only once), one can focus on h such that there has not been any disclosure until that point in time and in which Buyer has rejected all prices up to that point in time.

Assume first, h is the shortest on-path history that has Seller put positive probability on Buyer having received some news.

Let $\bar{w}(h)$ be the highest type feasible from the perspective of Seller after history h. Thus, it corresponds to an agent who has learned. According to the conjecture, the reservation price $\bar{p}(h)$ of Buyer who has received good news, is quoted with probability one. This is because reservation prices move co-monotonically with the type of Buyer (see Corollary 7), and because learning happens once. Because of learning once, the type of Buyer won't change over time so going forward $\bar{p}(h)$ is the only price being quoted as long as no agreement is reached. Thus, equilibrium payoff for this Buyer type going forward is also $\bar{w}(h) - \bar{p}(h)$. From the reservation pricing relation of Buyer with good news, $\bar{w}(h) - \bar{p}(h) = \delta(\bar{w}(h) - \bar{p}(h))$, one sees that this type has zero continuation payoff, i.e. $\bar{w}(h) = \bar{p}(h)$.

Now look at the history h' which precedes h by one period and has Buyer get a chance to learn with positive probability. Assume that h' exists.

Seller after no-disclosure at h' thinks that Buyer has not learned yet (this is because of the definition of h). Thus, Seller charges at the end of that period the reservation price of $\hat{v}$, which is based on the specified continuation play in the equilibrium.

Consider first the model with stochastic fixed costs of accuracy. I show that Buyer has incentives to learn some of the time. This is because with positive probability the opportunity to learn will arrive and the costs will be small enough to justify learning θ conclusively. In the case of good news, Buyer makes a profit of at least $\bar{v} - \hat{v}$, whereas in the case of bad news the payoff is zero.[14] Thus, when c is so low that

$$\pi_0(\bar{v} - \hat{v}) - c > 0, \tag{A.24}$$

Buyer has strict incentives to learn. Note that (A.12) happens with positive

[14]Here, one uses that reservation prices of a type are weakly below their valuation.

probability under the assumptions on F in section 1.3.

Consider next the model with deterministic variable costs on accuracy, with two-parametric experiments and so that the assumption in the statement is satisfied.

I show that for some pair (a_H, a_L) it holds
$\pi_0 a_H(\bar{v} - \hat{v}) + (1 - \pi_0)(1 - a_L)(\underline{v} - \hat{v}) - C(a_H, a_L) > 0$.

To see this, suppose the costs satisfy $\partial_a C(\frac{1}{2}, \frac{1}{2}) = 0$ for either $a = a_H, a_L$. Assume it for $a = a_H$, the other case being analogous. Then for a_H, a_L near $\frac{1}{2}$ given by $a_H = \frac{1}{2} + \varepsilon, a_L = \frac{1}{2}$ with some small $\varepsilon > 0$[15], one has

$$\pi_0 a_H(\bar{v} - \hat{v}) + (1 - \pi_0)(1 - a_L)(\underline{v} - \hat{v}) - C(a_H, a_L) = \varepsilon\pi_0(1 - \pi_0)(\bar{v} - \underline{v}) - C(\frac{1}{2} + \varepsilon, \frac{1}{2}) \tag{A.25}$$

$$= \varepsilon\pi_0(1 - \pi_0)(\bar{v} - \underline{v}) - O(\varepsilon^2).$$

Here the last equality uses Taylor formula. Thus, there are incentives to learn at least just very little, just before the last period in h.

Thus, it has to be that h corresponds to a history started in period one (i.e. h' does not exist). In particular, the equilibrium must prescribe that Buyer chooses to learn with positive probability in period one, if she gets the chance to learn already in that period.

Take first the model where intensity is exogenous. Then Buyer has zero benefit from learning because of the conjectured high-price structure of Seller's continuation strategy. Since learning is costly, Buyer does not learn so the conjecture of Seller about Buyer learning with positive probability already in the first period is wrong. This is a contradiction.

Consider next the model in which intensity is endogenously chosen at a cost (see assumptions in section 1.4.1). The same argument leads to a contradiction, because, given that the benefit of learning is zero in the continuation play, and picking a positive intensity is costly, Buyer does not pick a positive intensity at all.

[15] I.e. signal structure is informative only in the case of good news but by very little.

Overall, the assumption of a high-price equilibrium leads to a contradiction under the assumptions on costs made in the statement of the Proposition.

□

Proof of Theorem 2. Recall Lemma 5. It suffices to show the following claim.

Claim. For any $\Delta > 0$ there are no high-price equilibria, if in the case of deterministic variable costs the set of available experiments to Buyer is constrained to the one-parametric one in section 1.3.

To see this, one adapts the proof of Proposition 31 to show that there are no high-equilibria with costly learning in the set up of restricted experiments from section 1.3. The only change necessary is in the case of deterministic variable costs. Assume that instead of general experiments, Buyer only has access to the one-parametric ones from section 1.3. One looks at a deviation to $a = \frac{1}{2} + \varepsilon$ and one replaces (A.25) with

$$\pi_0(\frac{1}{2} + \varepsilon)(\bar{v} - \hat{v}) + (1 - \pi_0)(\frac{1}{2} - \varepsilon)(\underline{v} - \hat{v}) - c\left(\frac{\frac{1}{2} + \varepsilon}{\frac{1}{2} - \varepsilon}\right) = \varepsilon\pi_0(1 - \pi_0)(\bar{v} - \underline{v}) - O(\varepsilon^2).$$

Here one uses that $\frac{\frac{1}{2}+\varepsilon}{\frac{1}{2}-\varepsilon} = 1 + \frac{4\varepsilon}{1-2\varepsilon} = 1 + O(\varepsilon)$, as $\varepsilon \to 0$. The remaining formal arguments to conclude the proof of the Claim are verbatim the same as in the proof of Proposition 31. The combination of Lemma 5 and of the Claim finishes the proof of Theorem 2. □

Theorem 2 and Proposition 31 have several important implications. Call an equilibrium a *no-learning* equilibrium if on path, Buyer learns with probability zero.

Corollary 9. *There is no no-learning equilibrium under the conditions of Proposition 31.*

Proof. This follows from the arguments in the proof of Proposition 31. In an equilibrium in which Buyer never learns, the price quoted by Seller in every period, as long as the barganing goes on, is given by $\hat{v}$. But under this requirement, it occurs with positive probability that Buyer has strict incentives to learn whenever she gets the chance. □

Corollary 10. *In the presence of costs there is no stationary high-price equilibrium.*

Proof. This follows immediately from Proposition 31 when specializing to stationary equilibria. □ In particular, none of the stationary high-price equilibria from Proposition 3 survives the introduction of learning costs.

A.2.2 Proof of Proposition 6

The proof is split according to the assumptions on learning costs (deterministic variable or stochastic fixed) and makes use of a series of auxiliary results stated in the following in the form of Lemmas.

As in the case of costless learning I exhibit first the proof for the case of mixed pricing and then add details for the case of pure pricing.

The case of mixed pricing

The case of deterministic variable costs. To save on notation introduce the shortcut $C(a) = c(I(a))$ for $a \in [\frac{1}{2}, 1)$.

Lemma 11. *There exists $\varepsilon = \varepsilon(c, \pi_0, \hat{v}) > 0$ such that for all $\hat{p}$ and Δ with $|\hat{v} - \hat{p}| < \varepsilon, \Delta < \varepsilon$ the reservation pricing relations of Buyer with good news and Buyer of type $\hat{v}$ are uniquely solvable.*

Proof. Corollary 9 implies that any PBE with the parametric assumptions made on costs in section 1.3 involves some amount of learning by Buyer.

Focus in the following only on $\hat{p}$ such that
$BL(\hat{p}) := V_A(a(\hat{p}), \hat{p}) - C(a(\hat{p})) > 0$. Because $BL(\hat{v}) > 0$ and continuity, the requirement $BL(\hat{p}) > 0$ is fulfilled in a small enough open neighborhood of $\hat{v}$.

Fix an average price $\hat{p}$ such that $\underline{v} < \hat{p} < \bar{v}$. With the assumption on costs it follows that $a(\hat{p}) > \frac{1}{2}$. Therefore, the possible Buyer valuations $\bar{w}, \underline{w}$ are also functions of $\hat{p}$, i.e. one writes $\bar{w}(\hat{p}), \underline{w}(\hat{p})$. From the reservation price relation for type $\bar{w}$ one can write $p_H(\hat{p}) = (1 - \delta)\bar{w}(\hat{p}) + \delta\hat{p}$, whereas from the reservation price relation of type $\hat{v}$ one can write

$$p_L(\hat{p}) = \hat{v} - \frac{\delta\mu}{1-\delta+\delta\mu}\left(V(a(\hat{p}),\hat{p}) - C(a(\hat{p}))\right).$$

Note that the reservation price relation for $\hat{v}$ does not deliver $p_L < \hat{p}$. It only delivers $p_L < \hat{v}$. To ensure that for given $\hat{p}$ it holds for $p_L(\hat{p})$ that $p_L(\hat{p}) < \hat{p}$, ask for $\hat{p}$ near enough to $\hat{v}$. It holds that $V_A(a(\hat{v}),\hat{v}) - C(a(\hat{v})) > V_A(\frac{1}{2},\hat{v}) - C(\frac{1}{2}) = \frac{1}{2}(\hat{v}-\hat{v}) = 0$.[16] Therefore in the following restrict to a neighborhood $\mathcal{N}$ of $\hat{v}$ such that

$$\hat{v} - \hat{p} < \frac{1}{2}\frac{\lambda}{\lambda+r}(V_A(a(\hat{p}),\hat{p}) - C(a(\hat{p}))). \tag{A.26}$$

In addition $\mathcal{N}$ is required to satisfy $BL(\hat{p}) > 0$ for all $\hat{p} \in \mathcal{N}$.

Fulfillment of (A.26) for all $\hat{p}$ near $\hat{v}$ is ensured because the inequality is true for $\hat{p} = \hat{v}$ and the involved functions are continuous.[17]

Now let $\bar{\Delta} > 0$ be small enough and pick $\mathcal{N}_1$, a compact non-empty subinterval of $\mathcal{N}$ such that for all $\hat{p} \in \mathcal{N}_1$ and $\Delta \leq \bar{\Delta}$ it holds

$$\hat{v} - \hat{p} < \frac{\delta\mu}{1-\delta+\delta\mu}(V_A(a(\hat{p}),\hat{p}) - C(a(\hat{p}))).$$

This is again possible due to continuity and the fact that $\frac{\delta\mu}{1-\delta+\delta\mu} \to \frac{\lambda}{\lambda+r}$, as $\Delta \to 0$.

This, together with the trivial bound $\frac{\delta\mu}{1-\delta+\delta\mu} \leq 1$, leads to $p_L(\hat{v}) < \min\{\hat{p},\hat{v}\}$, as solved from the reservation price relation $\hat{v}$. This finishes the proof of the Lemma.

□

Focus in the following on Δ small and $\hat{p}$ with $|\hat{v} - \hat{p}| < \varepsilon$, as needed in statement of Lemma 11.

Once $p_L(\hat{p})$ is determined with the property $p_L(\hat{p}) < \hat{p}$, the stationary mixing probability of Seller upon non-disclosure is uniquely determined the formula

---

[16] The inequality is strict because of the fact that C is strictly convex and Corollary 9.

[17] Continuity of $[\underline{v},\bar{v}) \ni \hat{p} \mapsto V(a(\hat{p}),\hat{p}) - C(a(\hat{v}))$ follows from Berge's maximum theorem.

$p(\hat{p}) = \frac{\hat{p}-p_L(\hat{p})}{p_H(\hat{p})-p_L(\hat{p})}$. This proves that the sufficient statistic for the construction of the strongly stationary equilibria is $\hat{p}$.

Lemma 12. *It holds $\underline{w}(\hat{p}) < \hat{p}$ in a suitable open neighborhood of $\hat{v}$.*

Proof. To see this, note that $\underline{w}(\hat{p})$ is a continuous function of $\hat{p}$ and that $\underline{w}(\hat{v}) < \hat{v}$, because valuation $\underline{w}$ is induced by bad news. □

In the following take $\mathcal{N}$, an open neighborhood around $\hat{v}$ and $\Delta \leq \bar{\Delta}$ with $\bar{\Delta} > 0$ such that

– $BL(\hat{p}) > 0, \hat{p} \in \mathcal{N}$,

– Lemmas 11 and 12 are true for the neighborhood $\mathcal{N}$ of $\hat{v}$ and $\Delta \leq \bar{\Delta}$.

It holds automatically that $V(a, \hat{p}) < \hat{v}$ whenever $\hat{p} > 0$. To see this, use the definition to get $V_A(a, \hat{p}) < \pi_0 \bar{v} + (1 - \pi_0)\underline{v} - (\pi_0 a + (1 - \pi_0)(1 - a))\hat{p} < \hat{v}$. This is a uniform bound which does not use the parametric form of the costs nor the value of λ, Δ. One sharpens this estimate by noticing that $\pi_0 a + (1 - \pi_0)(1 - a) \geq c(\pi_0) := \min\{\pi_0, 1 - \pi_0\}$.[18] This leads to the inequality

$$V_A(a, \hat{p}) \leq \hat{v} - c(\pi_0)\hat{p}. \tag{A.27}$$

This inequality is uniform in the specification of costs, λ, Δ. It also does not depend on the reservation price relation for type $\hat{v}$. Therefore, by using the reservation pricing relation for the type $\hat{v}$ one arrives at the uniform estimate

$$p_L(\hat{p}) \geq \frac{1-\delta}{1-\delta+\delta\mu}\hat{v} + \frac{\delta\mu}{1-\delta+\delta\mu}c(\pi_0)\hat{p}, \tag{A.28}$$

whenever $\hat{p} \in \mathcal{N}$.

From (A.28) and the reservation pricing relation for Buyer with good news $\bar{w}$ one arrives at the estimate

$$p_H(\hat{p}) - p_L(\hat{p}) \leq (1-\delta)\bar{w}(\hat{p}) + \delta\left(1 - \frac{\mu}{1-\delta+\delta\mu}c(\pi_0)\right)\hat{p} - \frac{1-\delta}{1-\delta+\delta\mu}\hat{v}. \tag{A.29}$$

Seller's indifference condition reduces to the study of zeros of the function

[18] Note that $c(\pi_0) > 0$.

$$f(\hat{p}, q) = U(0)(p_H(\hat{p}) - p_L(\hat{p})) + (1 - U(0))q(\delta V_\Delta(q, \hat{p}) - p_L(\hat{p})).$$

Here $V_\Delta(q, \hat{p})$ is Seller's payoff which is given in equilibrium by

$$V_\Delta(q, \hat{p}) = \frac{\mu GN(a(\hat{p}))\hat{p} + (1-\mu)(1-p)qp_L(\hat{p}) + \mu(1 - GN(a(\hat{p}))\underline{w}(\hat{p})}{1 - \delta(1-\mu)(1-(1-p)q)}, \tag{A.30}$$

where the shortcut $GN(a)$ denotes the stationary probability of good news given by $GN(a) = a\pi_0 + (1-a)(1-\pi_0)$. To see (A.30), note that $V_\Delta(q, \hat{p})$ satisfies the recursion

$$V_\Delta(q, \hat{p}) = \mu GN(a(\hat{p}))\hat{p} + (1-\mu)(1-p)qp_L + \mu(1 - GN(a(\hat{p})))\underline{w}(\hat{p}) + \delta(1-\mu)(1-(1-p)q)V_\Delta(q, \hat{p}).$$

One calculates

$$\delta V_\Delta(q, \hat{p}) - p_L(\hat{p}) = \frac{\delta\mu GN(a(\hat{p})\hat{p} + (\delta(1-\mu)-1)p_L(\hat{p}) + \delta\mu(1 - GN(a(\hat{p})))\underline{w}(\hat{p})}{1 - \delta(1-\mu)(1-(1-p(\hat{p}))q)} \tag{A.31}$$

$$\leq \frac{\delta\mu\hat{p} + (\delta(1-\mu)-1)p_L(\hat{p})}{1 - \delta(1-\mu)(1-(1-p(\hat{p}))q)} \leq \frac{\delta\mu(1-c(\pi_0))\hat{p} - (1-\delta)\hat{v}}{1 - \delta(1-\mu)(1-(1-p(\hat{p}))q)}.$$

Here the first inequality uses that $\underline{w}(\hat{v}) < \hat{v}$ and that $\hat{p}$ is close to $\hat{v}$ (recall the restrictions on $\mathcal{N}$). The second inequality in (A.31) uses (A.28). Impose now the following assumption on δ, μ:

$$\text{agents are not too patient:} \quad \frac{\delta\mu}{1-\delta}(1 - c(\pi_0)) < 1. \tag{A.32}$$

Note that in the high-frequency limit this assumption corresponds to $\lambda(1 - c(\pi_0)) < r$. A sufficient condition for satisfying (A.32) irrespective of the prior is $r > \lambda$, i.e. the discount rate is higher than the arrival rate of opportunities to learn.

Combine (A.32) with (A.31) to get that uniformly for all $\hat{p} \in \mathcal{N}$, $\Delta \leq \bar{\Delta}$ and all δ, μ satisfying (A.32) that

$$\begin{aligned} q(\delta V(q,\hat{p}) - p_L(\hat{p})) &\leq \frac{q}{1-\delta(1-\mu)(1-(1-p(\hat{p}))q)}(\delta\mu(1-c(\pi_0))\hat{p} - (1-\delta)\hat{v}) \\ &\leq q(\delta\mu(1-c(\pi_0))\hat{p} - (1-\delta)\hat{v}). \end{aligned}$$

Look now at $\frac{f(\hat{p},q)}{\mu(1-U(0))} = \frac{\pi_0}{1-\mu}(p_H(\hat{p}) - p_L(\hat{p})) + \frac{q}{\mu}(\delta V(q,\hat{p}) - p_L(\hat{p}))$. Overall the following estimate results for all $\hat{p} \in \mathcal{N}$, $\Delta \leq \bar{\Delta}$ and all δ, μ satisfying (A.32)

$$\begin{aligned} \frac{f(\hat{p},q)}{\mu(1-U(0))} \leq \frac{\pi_0}{1-\mu}\left((1-\delta)\bar{w}(\hat{p}) + \delta\left(1 - \frac{\mu}{1-\delta+\delta\mu}c(\pi_0)\right)\hat{p} - \frac{1-\delta}{1-\delta+\delta\mu}\hat{v}\right) \\ + q\left(\delta(1-c(\pi_0))\hat{p} - \frac{(1-\delta)}{\mu}\hat{v}\right). \end{aligned} \tag{A.33}$$

I make further assumptions on r, λ and redefine $\mathcal{N}$, $\bar{\Delta}$ appropriately so that the right-hand side of (A.33) becomes negative in near the HFL. For this, one sets first $q = 1$ and looks at $\Delta \to 0$. In the HFL the right-hand side of (A.33) becomes

$$\pi_0(1 - \frac{\lambda}{r+\lambda}c(\pi_0))\hat{p} - \frac{r}{r+\lambda}\pi_0\hat{v} + (1-c(\pi_0))\hat{p} - \frac{r}{\lambda}\hat{v}. \tag{A.34}$$

The coefficient in front of $\hat{p}$ is positive so that if one replaces $\hat{p} = \hat{v}$ it is sufficient in the limit to require the restriction

$$\pi_0(1 - \frac{\lambda}{r+\lambda}c(\pi_0)) + (1-c(\pi_0)) < \frac{r}{r+\lambda}\pi_0 + \frac{r}{\lambda}, \tag{A.35}$$

and in addition, that $\hat{p}$ is near enough to $\hat{v}$ so that (A.34) remains valid for these $\hat{p}$ close to $\hat{v}$. This additional restriction is on top of the other previous restrictions set above on $\mathcal{N}$.

Suppose first that $c(\pi_0) = \pi_0$ which is equivalent to $\pi_0 \leq \frac{1}{2}$. Then the left-hand side of (A.35) is strictly lower than the right-hand side due to (A.32).

To see this, replace $c(\pi_0)$ in (A.35) to arrive at the sufficient condition $1 - \frac{r}{\lambda} < \frac{r}{r+\lambda}\pi_0 + \frac{\lambda}{r+\lambda}\pi_0^2$. To get the result uniformly on $\pi_0 \leq \frac{1}{2}$ require $r \geq \lambda$ for the case $\pi_0 \leq \frac{1}{2}$.

Consider now $c(\pi_0) = 1 - \pi_0$ which is equivalent to $\pi_0 \geq \frac{1}{2}$. Plugging in $c(\pi_0) = 1 - \pi_0$ in (A.35), this inequality becomes $\frac{\lambda}{\lambda+r}\pi_0^2 + \pi_0 < \frac{r}{\lambda}$. Here the left-hand side $\frac{\lambda}{\lambda+r}\pi_0^2 + \pi_0$ is decreasing in $\frac{r}{\lambda}$, whereas the right-hand side is increasing in $\frac{r}{\lambda}$ (it being the identity map on $\frac{r}{\lambda}$). To get a condition which is uniform for all $\pi_0 > \frac{1}{2}$ one needs the condition $\frac{\lambda}{r+\lambda} + 1 < \frac{r}{\lambda}$. Algebraic manipulation shows that this is equivalent to $\frac{r}{\lambda} > \sqrt{2}$.

Require thus overall, that $\Delta \leq \hat{\Delta} \leq \bar{\Delta}$ for some $\hat{\Delta} > 0$ appropriate such that

$$\frac{\delta\mu}{1-\delta} < 1 \text{ if } \pi_0 \leq \frac{1}{2}, \quad \frac{\delta\mu}{1-\delta} < \frac{1}{\sqrt{2}} \text{ if } \pi_0 > \frac{1}{2}. \tag{A.36}$$

The corresponding HFL assumption for (A.36) is the one required in the statement of Proposition 6.[19]

With regards to the continuous function $f : (0,1) \times (\underline{v}, \bar{v}) \to \mathbb{R}$ the above analysis has shown the following two facts.

- A. There exists an open neighborhood of $\hat{v}$ such that for all $\hat{p}$ in that neighborhood it holds

$$\limsup_{\Delta \to 0} f(1, \hat{p}) < 0.$$

- B. Fix any $\hat{p} \in (\underline{v}, \bar{v})$ such that solvability of p_L, p_H is guaranteed (Lemma 11). Then for any $\mu \in (0,1)$ and $\delta < 1$ it holds

$$\lim_{q \to 0} f(q, \hat{p}) = U(0)(p_H(\hat{p}) - p_L(\hat{p})) > 0.$$

Pick first some neighborhood $\mathcal{N}$ of $\hat{v}$, $\bar{\Delta}$ so that Lemmas 11 and 12 are ensured for $\hat{p} \in \mathcal{N}$, $\bar{\Delta} > \Delta > 0$. Redefine $\bar{\Delta} > 0$ small such that for all $\Delta < \bar{\Delta}$ (A.36)

[19] To see this for the case $\pi_0 > \frac{1}{2}$, note that the function $g : (0, \infty) \to \mathbb{R}_+$ given by $g(t) = \frac{t}{1+t} + t$ has derivative $g'(t) = 1 + \frac{1}{(1+t)^2}$ and so is strictly decreasing and that $g(\sqrt{2}) = 2$.

holds true. This implies then that $f(1, \hat{p}) < 0$ for all $\hat{p} \in \mathcal{N}$.[20] Using B. and the intermediate-value theorem for continuous functions, one finds the required zero $q(\Delta, \hat{p}) \in (0, 1)$. This establishes that equilibria exist for the sufficient statistic $\hat{p} \in \mathcal{N}$ and $\Delta \leq \bar{\Delta}$.

Lemma 6 can be used to show easily that for all $\Delta > 0$ small, the function $\frac{f(\hat{p},q)}{\mu(1-U(0))}$ is strictly decreasing and convex in q. This shows uniqueness of $q(\Delta, \hat{p})$ for fixed $\hat{p}$, Δ as above. This finishes the proof of 1) and 2) from Proposition 6 for the case of deterministic variable costs.

The case of stochastic fixed costs. Note that the two reservation pricing relations for Buyer imply

$$\hat{v} - p_L \leq \frac{\delta\mu}{1 - \delta + \delta\mu}(\pi_0(\bar{v} - \hat{p})) < \pi_0(\bar{v} - \hat{p}).$$

In particular, $\mu > 0$ for any pair of prices (p_H, p_L) satisfying the reservation pricing relations of types $\bar{v}$, $\hat{v}$.

Remark 6. *The function $x \to \mathbb{E}[c|c \leq x]$ is weakly increasing in x for all F that are 'well-behaved', i.e. are limits in total variation norm of distributions having a positive and smooth density w.r.t. Lebesgue measure over $\mathbb{R}_+$.*[21]

Proof of Remark 6. *Step 1.* First, verify the claim for an F that has a smooth density f over $(0, \infty)$. It holds

$$\frac{d}{dx}\mathbb{E}[c|c \leq x] = \frac{d}{dx}\left(\frac{\int_0^x cf(c)dc}{F(x)}\right) = \frac{xf(x)F(x) - f(x)\int_0^x cf(c)dc}{F(x)^2} = \frac{f(x)}{F(x)}\left(x - \mathbb{E}[c|c \leq x]\right).$$

This derivative is strictly positive for all $x > 0$ because of the assumption that F puts positive probability on arbitrarily low costs (from section 1.3).

[20] This is because the function $f(1, \hat{p})$ is continuous also in Δ (via its continuity on μ and r) which are in turn continuous functions of Δ.

[21] 'smooth' in the following means infinitely often continuously differentiable.

Step 2. Let F be continuous with finite first moment and support on $[0, \infty)$ and assume that $F(0) = 0$.

CLAIM. There exists a sequence of distributions F_n which have smooth densities over $(0, \infty)$ as in Step 1 and so that F_n converges to F uniformly.

To see the Claim, note the following facts.

A. the set of distribution functions with finite support is dense in the topology of weak convergence of probability measures,[22]

B. if F_n over $\mathbb{R}$ converges weakly to a continuous distribution F, then the convergence is uniform,[23]

C. if F_n converges weakly to F and F is continuous then $F_n(\cdot|\cdot \leq x)$ converges weakly to $F(\cdot|\cdot \leq x)$, for any $x > 0$ (the convergence is even uniform),

   [*Proof.* To see this, recall that $F(x) > 0$ for any $x > 0$ by the assumptions in section 1.3. In particular, $(0, \infty) \ni c \mapsto F(c|c \leq x)$ is a continuous distribution function. Pick any $y < x$. Then obviously $F_n(x) \to F(x)$ and $F_n(y) \to F(y)$ as $n \to \infty$. This leads to the result stated in C.]

D. any distribution G which is a step function can be approximated from above and below pointwise by a sequence of distribution functions which are infinitely differentiable functions;

   [*Proof.* To see this, first approximate the step function through a continuous function by interpolating 'near' the (finitely many) discontinuities of G, to get two continuous distribution functions G', G'' which differ from G only in intervals of size less than ε around the discontinuities, are continuous and satisfy $G''(x) \leq G(x) \leq G'(x)$ for all $x \geq 0$. Then one performs a 'mollification' procedure on G', G'' around

[22] This is easy to see and uses the generally known fact of undergraduate analysis: any bounded, measurable function can be approximated from below by step functions with finitely many discontinuities.

[23] See e.g. exercise 3.2.9. in Durrett [2010].

the finitely many points of their non-differentiabilities to get smooth distribution functions H'', H' which satisfy $H''(x) \leq G''(x) \leq G(x) \leq G'(x) \leq H'(x)$ for all $x \geq 0$. See chapter 5 and in particular section 5.3 in Evans [2010] for more on mollification arguments.]

Now I finish the proof of Step 2. above and thus also of the Remark.

From Step 1 one knows that $\mathbb{E}_H[c|c \leq x]$ are increasing in x (even strictly) for H of the type H', H'' as given in the proof of D above. One can use the argument in D. above to show that monotonicity of $x \mapsto \mathbb{E}_F[c|c \leq x]$ holds true for F a step function. To see this, suppose there is $y > x > 0$ and F a step distribution function so that $\mathbb{E}_F[c|c \leq y] < \mathbb{E}_F[c|c \leq x]$. It holds in general with the Fubini Theorem that

$$\mathbb{E}_F[c, c \leq y] = \mathbb{E}_F\left[\int_0^c dt, c \leq y\right] = \int_0^y F(y) - F(t)dt. \tag{A.37}$$

Because of monotonicity of the approximation in point D. above and (A.37) one sees easily that:

$$\text{for } z = x, y \text{ and step function } F, \quad \mathbb{E}_F[c, |c \leq z] = \lim_{n \to \infty} \mathbb{E}_{H_n}[c|c \leq z], \tag{A.38}$$

for a sequence $\{H_n : n \geq 1\}$ of distribution functions which are smooth.

Ultimately, (A.38) and the assumption $\mathbb{E}_F[c|c \leq y] < \mathbb{E}_F[c|c \leq x]$ lead to contradiction of the monotonicity in Step 1.

It follows that $x \mapsto \mathbb{E}_F[c|c \leq x]$ is increasing for F a step distribution function.

Finally, by points A., B. and C. above, one establishes monotonicity of $\mathbb{E}_F[c|c \leq x]$ for general F by an approximation argument.

□

Denote in the following $BL(\hat{p}, p_L) = \pi_0(\bar{v} - \hat{p}) - \mathbb{E}[c|c \leq \pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L)]$. Note that this definition is independent of Δ.

Because F puts positive probability on arbitrarily small costs, $BL(\hat{p}, p_L)$ is strictly positive for any $p_L < \hat{v}, \hat{p} < \bar{v}$. Note that $BL(\hat{p}, p_L)$ is decreasing in p_L and the direction of the monotonicity in $\hat{p}$ is ambiguous. Finally, note that

$$BL(\hat{p}, p_L) \leq \pi_0(\bar{v} - \hat{p}) \leq \hat{v} - c(\pi_0)\hat{p}. \tag{A.39}$$

This helps in proving again an estimate precisely as in (A.28), whenever one has solvability of the reservation price relation for type $\hat{v}$.

Moreover, because of the assumption that F puts positive probability on a neighborhood of zero it follows that

$$BL(\hat{p}, p_L) - (\hat{v} - p_L) > 0, \text{ for every } \hat{p} < \bar{v}, p_L < \hat{v}.^{24} \tag{A.40}$$

The two reservation pricing relations give a system of equations to be solved for the pair (p_H, p_L).

For $\hat{p} \in (\underline{v}, \bar{v})$ and $p_L \in (\underline{v}, \hat{p})$ define for ease in notation $\bar{\mu}(\hat{p}, p_L) = F(\pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L))$. Denote $p_L^{min}(\hat{p}) = \hat{v} - \pi_0(\bar{v} - \hat{p}) = (1 - \pi_0)\underline{v} + \pi_0\hat{p} < \hat{p}$.

It holds $\bar{\mu}(\hat{p}, p_L^{min}(\hat{p})) = 0$ and $\bar{\mu}(\hat{p}, \hat{p}) = F((1 - \pi_0)(\hat{p} - \underline{v})) > 0$. Note that $\bar{\mu}(\hat{p}, p_L)$ is falling in $\hat{p}$ and increasing in p_L.

Lemma 13. *1) Suppose $\mu_0(\Delta) = 1 - e^{-\Delta\lambda}$. The reservation pricing equation for $\hat{v}$ is solvable for p_L whenever $\hat{p}$ is in a suitable open neighborhood $\mathcal{N}$ of $\hat{v}$ and Δ is smaller than some $\bar{\Delta} > 0$.*

Moreover, whenever solvable, the solution $p_L(\hat{p})$ is unique for each pair $(\Delta, \hat{p})$ and weakly increasing in $\hat{p}$ for fixed Δ.

2) Let $\hat{p} \in \mathcal{N}$, where $\mathcal{N}$ comes from 1). As $\Delta \to 0$, $p_L(\hat{p})$ converges to a $\bar{p}_L(\hat{p})$ that satisfies

$$\hat{v} - \bar{p}_L(\hat{p}) = \frac{\lambda\bar{\mu}(\hat{p}, \bar{p}_L(\hat{p}))}{r + \lambda\bar{\mu}(\hat{p}, \bar{p}_L(\hat{p}))} BL(\hat{p}, \bar{p}_L(\hat{p})). \tag{A.41}$$

[24] Formally, this follows because $\mathbb{E}[c|c \leq x] < x$ for $x > 0$ under the assumption that F puts a positive weight on a neighborhood of zero (recall assumptions in section 1.3).

Proof. 1) **Existence.**

One needs to solve for p_L in

$$\hat{v} - p_L = \frac{\delta\mu}{1 - \delta + \delta\mu} BL(\hat{p}, p_L), \tag{A.42}$$

for given $\hat{p}$ and $\Delta > 0$. Moreover $p_L \leq \hat{v}$. I relax this second requirement in the first part of the proof of existence and make sure it is satisfied in the end.

Recall $p_L^{min}(\hat{p}) = \hat{v} - \pi_0(\bar{v} - \hat{p}) = (1 - \pi_0)\underline{v} + \pi_0\hat{p} < \hat{p}$. Note that automatically from the definition of $p_L^{min}(\hat{p})$ it follows $p_L^{min}(\hat{p}) < \hat{v}$, because $\hat{p} < \bar{v}$. It follows that the right-hand side of the reservation price relation for type $\hat{v}$ is zero if one plugs $p_L = p_L^{min}(\hat{p})$, whereas the left-hand side is positive. On the other hand, when setting $p_L = \hat{p}$ the left hand side becomes $\hat{v} - \hat{p}$ and the right hand side becomes $\frac{\delta\mu_0\bar{\mu}(\hat{p},\hat{p})}{1-\delta+\delta\mu_0\bar{\mu}(\hat{p},\hat{p})} BL(\hat{p}, \hat{p})$. Note that both $\hat{v} - \hat{p}$ and $\frac{\delta\mu_0\bar{\mu}(\hat{p},\hat{p})}{1-\delta+\delta\mu_0\bar{\mu}(\hat{p},\hat{p})} BL(\hat{p}, \hat{p})$ are continuous in the parameters $\hat{p} \in (\underline{v}, \bar{v})$, $\Delta \in [0, \infty)$. Evaluated at $\hat{p} = \hat{v}$ and $\Delta = 0$ the expression $\hat{v} - \hat{p}$ is zero, whereas $\frac{\delta\mu_0\bar{\mu}(\hat{p},\hat{p})}{1-\delta+\delta\mu_0\bar{\mu}(\hat{p},\hat{p})} BL(\hat{p}, \hat{p})$ is $\frac{\lambda\bar{\mu}(\hat{v},\hat{v})}{r+\lambda\bar{\mu}(\hat{v},\hat{v})} > 0$. Thus, there exists an open neighborhood of $\hat{v}$, denoted by $\mathcal{N}$ and a $\bar{\Delta} > 0$ such that for all $\hat{p} \in \mathcal{N}$ and $0 < \Delta < \bar{\Delta}$ it holds true that $\hat{v} - \hat{p} < \frac{\lambda\bar{\mu}(\hat{p},\hat{p})}{r+\lambda\bar{\mu}(\hat{p},\hat{p})} BL(\hat{p}, \hat{p})$. Overall the intermediate-value theorem for continuous functions gives existence of $p_L(\Delta, \hat{p}) \in (p_L^{min}(\hat{p}), \hat{p})$ that satisfies (A.42) for $\hat{p} \in \mathcal{N}$, $\Delta \leq \bar{\Delta}$.

Finally, for existence one needs to ensure that $p_L(\hat{p}) < \hat{v}$ for the $\hat{p}$ from a small, suitable neighborhood of $\hat{v}$. Note that $BL(\hat{p}, p_L)$ is decreasing in p_L and that $BL(\hat{v}, \hat{v}) = \pi_0(\bar{v} - \hat{v}) - \mathbb{E}[c | c \leq \pi_0(\bar{v} - \hat{v})] > 0$. Therefore, there exists a small neighborhood of $\hat{v}$ such that $BL(\hat{p}, p_L) > BL(\hat{p}, \hat{p}) > 0$ for the $\hat{p}$ in this neighborhood. Note that this argument is independent of $\Delta > 0$ because the function $(\hat{p}, p_L) \mapsto BL(\hat{p}, p_L)$ does not depend on Δ. Restrict in the following $\mathcal{N}$ to a strict, non-empty subset of itself so that in addition it is satisfied: $BL(\hat{p}, \hat{p}) > 0$ for all $\hat{p} \in \mathcal{N}$. It follows from the reservation price relation for type $\hat{v}$ that $p_L(\Delta, \hat{p}) < \hat{v}$ for such $\hat{p}$, independently of Δ.

To see uniqueness, note that the reservation price relation for type $\hat{v}$ can be written as

$$\frac{1-\delta+\delta\mu_0\bar{\mu}(\hat{p},p_L)}{\delta\mu_0\bar{\mu}(\hat{p},p_L)}(\hat{v}-p_L)+\mathbb{E}[c|c\leq\pi_0(\bar{v}-\hat{p})-(\hat{v}-p_L)]=\pi_0(\bar{v}-\hat{p}).$$

This can be transformed into

$$(1-\delta)(\hat{v}-p_L)=\delta\mu_0\mathbb{E}[(\pi_0(\bar{v}-\hat{p})-(\hat{v}-p_L)-c),c\leq\pi_0(\bar{v}-\hat{p})-(\hat{v}-p_L)]. \tag{A.43}$$

Note that the right-hand side is strictly increasing in p_L, whereas the left-hand side is strictly decreasing in p_L. This gives uniqueness of $p_L(\hat{p})$.[25]

Monotonicity of $p_L(\hat{p})$ is shown in two steps. First, rewrite the reservation price relation of type $\hat{v}$ as

$$\frac{1-\delta}{\delta\mu_0}(\hat{v}-p_L)=\mathbb{E}_F[p_L-\pi_0\hat{p}-(1-\pi_0)\underline{v}-c,c\leq p_L-\pi_0\hat{p}-(1-\pi_0)\underline{v}]. \tag{A.44}$$

Step 1. Suppose first that F has smooth density. The left-hand side of (A.44) is not directly dependent of $\hat{p}$ and is strictly increasing in p_L, for any fixed $\hat{p}$. One can use the implicit function theorem locally, because of the differentiability of the functions involved. The right-hand side is strictly increasing in $p_L-\pi_0\hat{p}$ and the derivative of $p_L(\hat{p})-\pi_0\hat{p}$ with respect to $\hat{p}$ is given locally by $\frac{dp_L}{d\hat{p}}(\hat{p})-\pi_0$. Thus, the implicit function theorem together with the chain rule for differentiation delivers a relation of the type $\frac{d(LHS)}{dp_L}\frac{dp_L}{d\hat{p}}=\frac{d(RHS)}{dp_L}\left(\frac{dp_L}{d\hat{p}}(\hat{p})-\pi_0\right)$.[26] Because $\frac{d(RHS)}{dp_L}>0$, $\frac{d(LHS)}{dp_L}<0$, this delivers a positive derivative $\frac{dp_L}{d\hat{p}}(\hat{p})$.

Step 2. Suppose now that F does not have a smooth density.

Claim. Let F_n converge weakly to F, with $F_n, n\geq 0$, F satisfying the conditions about stochastic fixed costs in the main body of the paper. Then $G_n:=F_n(c\in\cdot,c\leq x_n)$ converges to $F(c\in\cdot,c\leq x)$, whenever $x_n\to x$, $n\to\infty$, $x>0$.

[25]Here, for ease of notation, I suppress the dependence on Δ without loss of meaning.
[26]Here, LHS and RHS denote respectively the left-hand side and right-hand side of (A.44).

Proof of Claim. Recall from the proof of Remark 6 that for any distribution F it holds

$$\mathbb{E}_F[c, c \le y] = \mathbb{E}_F\left[\int_0^c dt, c \le y\right] = \int_{\mathbb{R}_+} (F(y) - F(t))\mathbf{1}_{\{t \le y\}} dt. \qquad \text{(A.45)}$$

Now, since $x_n \to x$, the sequence of functions $t \mapsto (F(x_n) - F(t))\mathbf{1}_{\{t \le x_n\}}$ is bounded and has support contained on a compact set K of $\mathbb{R}_+$. Moreover, due to weak convergence, it follows that $(F(x_n) - F(t))\mathbf{1}_{\{t \le x_n\}} \to (F(x) - F(t))\mathbf{1}_{\{t \le x\}}$ for all $t \in K$. One applies finally Lebesgue dominated convergence to get the result. □

The Claim shows that the reservation price relation of type $\hat{v}$ is 'stable' with respect to weak limits of F. Recall the proof of Remark 6. Arguments there imply, that for any F satisfying the conditions for stochastic fixed costs in section 1.3, one can pick a sequence F_n with the same properties, converging weakly to F, and satisfying the additional requirement that F_n have smooth densities w.r.t. Lebesgue measure.[27]

Take then such a sequence F_n and apply Step 1 to each F_n. The uniqueness of the solution $p_L(\Delta, \hat{p})$ for fixed Δ and any F, delivers that $p_L^n(\Delta, \hat{p}) \to p_L(\Delta, \hat{p}), n \to \infty$. It follows that $\hat{p} \mapsto p_L(\Delta, \hat{p})$ is weakly increasing.

2) To see that all the limit points of $p_L(\Delta, \hat{p})$ are the same, divide (A.43) by Δ and take $\Delta \to 0$ to arrive at

$$r(\hat{v} - p_L) = \lambda \mathbb{E}[(\pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L) - c), c \le \pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L)].$$

This is a relation that has to be satisfied for all limit points of $p_L(\Delta, \hat{p})$ and thus the limit $\bar{p}_L(\hat{p})$ exists, because the right-hand side is strictly increasing in p_L,

[27] To see this, recall that from A. and B. in that proof, F can be approximated uniformly by distribution functions that are step functions, and then use C. there to get a pointwise approximation of F through a sequence of F_n which are smooth. This uses the fact that the approximation of step functions through step functions is done monotonically and the differences between the step function and the approximands are in 'small' sets around a finite number of discontinuities.

whereas the left-hand side is strictly decreasing in p_L.

$\square$

Let $V_\Delta(q,\hat{p})$ be stationary payoff of Seller.[28] By the same logic as in the case of determinstic costs it satisfies the recursion

$$\delta V_\Delta(q,\hat{p}) = \mu\pi_0\hat{p} + (1-\mu)(1-p)qp_L + \mu(1-\pi_0)\underline{v} + \delta(1-\mu)(1-(1-p)q)V_\Delta(q,\hat{p}),$$

which leads to

$$\delta V_\Delta(q,\hat{p}) - p_L(\hat{p}) = \frac{\delta\mu\pi_0\hat{p} + (\delta(1-\mu)-1))p_L(\hat{p}) + \delta\mu(1-\pi_0)\underline{v}}{1-\delta(1-\mu)(1-(1-p(\hat{p})q))}.$$

Analogous to the cases of costless information and the case of deterministic variable costs, one considers the function $g(q,\hat{p}) = \frac{f(q,\hat{p})}{\delta\mu(1-U(0))} = \frac{\pi_0}{\delta(1-\mu)}(p_H(\hat{p}) - p_L(\hat{p})) + q\frac{\delta V(q,\hat{p}) - p_L(\hat{p})}{\delta\mu}$. It holds that $g(0,\hat{p}) > 0$ for all $\Delta > 0$ and $\hat{p}$ from a neighborhood $\mathcal{N}$ as required in Lemma 13.

In order to find sufficient conditions for $g(1,\hat{p}) < 0$ one uses (A.42) from the proof of Lemma 13 to get through the use of (A.39) an analogous estimate as (A.29) from the case of deterministic variable costs. This leads to the estimate for the price spread

$$p_H(\hat{p}) - p_L(\hat{p}) \leq (1-\delta)\bar{v} + \delta\left(1 - \frac{\mu}{1-\delta+\delta\mu}c(\pi_0)\right)\hat{p} - \frac{1-\delta}{1-\delta+\delta\mu}\hat{v}. \quad (A.46)$$

All ingredients are present to apply the same procedure for existence of strongly stationary equilibria as for the case of deterministic variable costs. One uses a similar chain of estimates as for the case of deterministic variable costs to show that $g(1,\hat{p}) < 0$, whenever $\hat{p}$ is near enough to $\hat{v}$ and Δ is small enough. The

[28] The recursion leading to Seller's payoff, is exactly the same as for the case of no costs, or deterministic variable costs.

replacements needed are $\underline{w}(\hat{p}) \rightsquigarrow \underline{v}$, $\bar{w}(\hat{p}) \rightsquigarrow \bar{v}$. Moreover, one has to replace $\frac{\mu(\Delta)}{\Delta}$ from the deterministic variable model with $\frac{\mu_0 \bar{\mu}(\hat{p}, p_L(\hat{p},\Delta))}{\Delta}$ throughout. Near the HFL one gets sufficient conditions ensuring existence just as for the case of deterministic variable costs of the type '$\frac{r}{\lambda\bar{\mu}(\hat{p},p_L(\hat{p}))}$ is high enough'. These are implied by the same conditions of the type '$\frac{r}{\lambda}$ is high enough' that appear in the case of deterministic variable costs.

This finishes the proof of Proposition 6 for the case of stochastic fixed costs of accuracy.

A.2.3 Proof of Theorem 3

The case of deterministic variable costs. In the following whenever it is said *uniformly in* $\hat{p}$, it is meant that the statement holds uniformly for all $\hat{p} \in \mathcal{N}$ where the neighborhood $\mathcal{N}$ comes from Proposition 6.

Note first the following easy-to- prove facts:

- $p_H(\hat{p})$ converges in HFL uniformly to the identity function $id(\hat{p}) = \hat{p}$. The difference $p_H(\hat{p}) - \hat{p}$ is $O(\Delta)$, uniformly in $\hat{p}$.

- Because information choice does not depend on μ, δ and $V(a(\hat{p}), \hat{p}) - C(a(\hat{p}))$ is bounded uniformly in $\hat{p}$, it holds

  $$p_L(\hat{p}) \rightarrow \bar{p}_L(\hat{p}) := \hat{v} - \frac{\lambda}{\lambda + r}(V(a(\hat{p}), \hat{p}) - C(a(\hat{p}))), \text{ in HFL uniformly in } \hat{p}.$$

  Note that $\bar{p}_L(\hat{p})$ is decreasing in $\hat{p}$. For future use I express the HFL of $p_L(\hat{p})$ in a more helpful form. Recall that $V(a(\hat{p}), \hat{p}) = GN(a(\hat{p}))(\bar{w}(\hat{p}) - \hat{p})$. One can use this to express the HFL of $p_L(\hat{p})$ as

  $$\bar{p}_L(\hat{p}) = \hat{v} - \frac{\lambda}{\lambda + r}(GN(a(\hat{p}))(\bar{w}(\hat{p}) - \hat{p}) - C(a(\hat{p}))).$$

- It follows from the first two claims that

$$p_H(\hat{p}) - p_L(\hat{p}) \to \frac{\lambda}{\lambda + r}(V_A(a(\hat{p}), \hat{p}) - C(a(\hat{p}))) - (\hat{v} - \hat{p}), \text{ in HFL uniformly in } \hat{p}. \tag{A.47}$$

  Denote this limit by $ps(\hat{p})$ where 'ps' stands for price spread.

  Note that in HFL, uniformly in $\hat{p}$ it holds $\hat{v} - \hat{p} < \hat{v} - \bar{p}_L(\hat{p}) = \frac{\lambda}{\lambda+r}(V_A(a(\hat{p}), \hat{p}) - C(a(\hat{p})))$, where the equality follows from the reservation pricing relation for type $\hat{v}$. This, and the proof of Proposition 6 (especially the proof of Lemma 11) for the case of deterministic variable costs ensures that the price spread remains bounded away from zero for all $\hat{p}$ from the neighborhood $\mathcal{N}$.

  Another simple, but important implication of (A.47) is that it shows that the HFL of a sequence of mixed pricing equilibria cannot correspond to a pure pricing equilibrium, whenever the limit average price $\hat{p}$ satisfies $\hat{p} \geq \hat{v}$.

- From the formula $p(\hat{p}) = \frac{\hat{p} - p_L(\hat{p})}{p_H(\hat{p}) - p_L(\hat{p})}$ for the mixing probability of Seller one sees convergence to 1 of $p(\hat{p})$. This convergence is uniform in $\hat{p}$, as one can see from the calculation

$$1 - p(\hat{p}) = \frac{p_H(\hat{p}) - \hat{p}}{p_H(\hat{p}) - p_L(\hat{p})} = (1 - \delta)\frac{\bar{w}(\hat{p}) - \hat{p}}{p_H(\hat{p}) - p_L(\hat{p})}.$$

Now I calculate explicitly $q(\Delta, \hat{p})$, the probability with which the $\hat{v}$-type ends the game when facing price $p_L(\hat{p})$. Recall that q is determined through the equation $\frac{f(\hat{p},q)}{\mu(1-U(0))} = 0$. Define the quantities $A(\Delta, \hat{p}) = \frac{\pi_0}{1-\mu}(p_H(\Delta, \hat{p}) - p_L(\Delta, \hat{p}))$ and $B(\Delta, \hat{p}) = \delta GN(a(\hat{p}))\hat{p} - \frac{1-\delta(1-\mu)}{\mu}p_L(\hat{p}, \Delta) + \delta(1 - GN(a(\hat{p})))\underline{w}(\hat{p})$. By using the HFL of p_L one calculates the HFL of $B(\Delta, \hat{p})$ to be

$$B(\hat{p}) := -\frac{r}{\lambda}\hat{v} - C(a(\hat{p})),$$

where the fact that $GN(a(\hat{p}))\bar{w}(\hat{p}) + (1 - GN(a(\hat{p})))\underline{w}(\hat{p}) = \hat{v}$ (Martingale property of beliefs) has been used.

Note that $B(\hat{p})$ is uniformly bounded away from zero and negative. One easily calculates the HFL of A to be

$$A(\hat{p}) := \pi_0 ps(\hat{p}).$$

It follows that

$$\frac{f(\hat{p}, q)}{\mu(1 - U(0))} = A(\Delta, \hat{p}) + \frac{q}{1 - \delta(1-\mu) + \delta(1-\mu)(1 - p(\Delta, \hat{p}))q} B(\Delta, \hat{p}).$$

One can define $C(\Delta, \hat{p}) = -\frac{A(\Delta,\hat{p})}{B(\Delta,\hat{p})}$ and see that $q(\Delta, \hat{p})$ can be solved in close form as

$$q(\Delta, \hat{p}) = \frac{(1 - \delta(1-\mu))C(\Delta, \hat{p})}{1 - \delta(1-\mu)(1 - p(\Delta, \hat{p})C(\Delta, \hat{p})}.$$

The HFL of C is easily calculated to be

$$C(\hat{p}) = \frac{\pi_0 \lambda \cdot ps(\hat{p})}{r\hat{v} + \lambda C(a(\hat{p}))}.$$

In particular, $C(\hat{p})$ is bounded away from zero for all $\hat{p} \in \mathcal{N}$. Now recall that $1 - p(\Delta, \hat{p}) = O(\Delta)$ and that $\frac{1-\delta(1-\mu)}{\Delta} \to r + \lambda$ to get that in the HFL

$$\frac{q(\Delta)}{\Delta} \to (r + \lambda)C(\hat{p}).$$

Next, I calculate Seller's payoff in the HFL. Recalling (A.30) one looks first at the denominator and concludes that

$$\begin{aligned}
&\frac{1 - \delta(1-\mu)(1 - (1 - p(\Delta, \hat{p}))q(\Delta, \hat{p}))}{\Delta} \\
&= \frac{1 - \delta(1-\mu)}{\Delta} + \delta(1-\mu)\frac{q(\Delta)}{\Delta}(1 - p(\Delta, \hat{p})) \to r + \lambda.
\end{aligned}$$

Looking at the numerator of Seller's payoff one notes that

$$\frac{\mu GN(a(\hat{p}))\hat{p} + (1-\mu)(1-p(\Delta,\hat{p}))q(\Delta,\hat{p})p_L(\hat{p},\Delta) + \mu(1-GN(a(\hat{p})))\underline{w}(\hat{p})}{\Delta}$$
$$\to \lambda GN(a(\hat{p}))\hat{p} + \lambda(1-GN(a(\hat{p})))\underline{w}(\hat{p}).$$

Overall this delivers for Seller's payoff in the HFL

$$V_S(\hat{p}) = \frac{\lambda}{r+\lambda}\left(GN(a(\hat{p}))\hat{p} + (1-GN(a(\hat{p})))\underline{w}(\hat{p})\right). \tag{A.48}$$

Turning to Buyer's payoff in the HFL: $V_B(\Delta,\hat{p})$ satisfies the recursion

$$\begin{aligned} V_B(\Delta,\hat{p}) = {} & \mu(V_A(a(\hat{p}),\hat{p}) - C(a(\hat{p}))) + (1-\mu)(1-p(\Delta,\hat{p}))q(\Delta,\hat{p})(\hat{v} - p_L(\Delta,\hat{p})) \\ & + \delta(1-\mu)(1-(1-p(\Delta,\hat{p}))q(\Delta,\hat{p}))V_B(\Delta,\hat{p}), \end{aligned}$$

which can be solved for

$$V_B(\Delta,\hat{p}) = \frac{\mu(V_A(a(\hat{p}),\hat{p}) - C(a(\hat{p}))) + (1-\mu)(1-p(\Delta,\hat{p}))q(\Delta,\hat{p})(\hat{v} - p_L(\Delta,\hat{p}))}{1-\delta(1-\mu)(1-(1-p(\Delta,\hat{p}))q(\Delta,\hat{p}))}.$$

The same (limit-)algebra as in the case of Seller delivers the HFL

$$V_B(\hat{p}) = \frac{\lambda}{r+\lambda}\left(GN(a(\hat{p}))(\bar{w}(\hat{p}) - \hat{p}) - C(a(\hat{p}))\right). \tag{A.49}$$

The sum of the payoffs in the HFL is

$$V_B(\hat{p}) + V_S(\hat{p}) = \frac{\lambda}{\lambda+r}\left(\hat{v} - C(a(\hat{p}))\right), \tag{A.50}$$

where again the Martingale property of beliefs has been used. Finally, I turn to the expected delay in the HFL.

Note that the date of agreement is a geometric random variable with success probability $1-(1-\mu)(1-(1-p(\Delta,\hat{p}))q(\Delta,\hat{p}))$. One calculates

$$\begin{aligned}\frac{1-(1-\mu)(1-(1-p(\Delta,\hat{p}))q(\Delta,\hat{p}))}{\Delta} &= \frac{\mu}{\Delta}+(1-\mu)(1-p(\Delta,\hat{p}))\frac{q(\Delta,\hat{p})}{\Delta}\\ &\to\lambda,\end{aligned}$$

given the HFL behavior of $q(\Delta,\hat{p})$ and $p(\Delta,\hat{p})$.

This finishes the proof of Theorem 3 for the case of deterministic variable costs.

The case of stochastic fixed costs. Pick a $\hat{p}$ as needed in Proposition 6 for existence near the HFL. By exactly the same steps as for the case of deterministic variable costs one arrives at similar results, but for the only changes that $\underline{w}(\hat{p}) \rightsquigarrow \underline{v}$, $\lambda \rightsquigarrow \lambda\bar{\mu}$ and $GN(a(\hat{p})) \rightsquigarrow \pi_0$.

One arrives at

$$V_B(\hat{p})=\frac{\lambda\bar{\mu}(\hat{p})}{r+\lambda\bar{\mu}(\hat{p})}\left(\pi_0(\bar{v}-\hat{p})-\mathbb{E}_F[c|c\le\pi_0(\bar{v}-\hat{p})-(\hat{v}-\bar{p}_L(\hat{p}))\right),$$

and

$$V_S(\hat{p})=\frac{\lambda\bar{\mu}(\hat{p})}{r+\lambda\bar{\mu}(\hat{p})}\left(\pi_0\hat{p}+(1-\pi_0)\underline{v}\right).$$

On multiplicity of equilibria with mixed pricing in the HFL of costly learning

In contrast to the case of costless learning the equilibrium multiplicity of strongly stationary equilibria with mixed pricing survives the HFL. This multiplicity remains when accuracy costs become vanishingly small. To see this, fix π_0, r, λ satisfying the conditions in the statement of Proposition 6 and of Theorem 3. Suppose that the respective costs c become arbitrarily small or F approaches the zero-cost distribution. The proof of Proposition 6 shows that in this situation, one can pick the respective existence neighborhoods $\mathcal{N}$ of $\hat{v}$ independently of the information costs.

The intuition for the multiplicity in the HFL is as follows. If Seller decides to quote a lower $\hat{p}$, this leads *ceteris paribus* to a higher option value from learning for Buyer of type $\hat{v}$. Because learning is costly, this lowers the reservation price of the type $\hat{v}$. In addition, the lower $\hat{p}$ makes Buyer who has received good news more willing to accept. Moreover, due to *higher incentives for learning*, a lower price spread is needed to incentivize learning in the limit. Given lower reservation prices for Buyer who has not learned or Buyer who has learned good news, Seller is then indeed only able to extract a lower average price $\hat{p}$.[29]

The following Lemma contains a formal result that underlies the intuition of this equilibrium multiplicity.

Lemma 14. *Look a the HLF of the strongly stationary equilibria for $\hat{p} \in \mathcal{N}$, the existence neighborhood of $\hat{v}$ from Proposition 6.*

In both the case of deterministic variable and stochastic fixed costs the limit as $\Delta \to 0$ of $p_L(\hat{p})$ given by $\bar{p}_L(\hat{p})$ is strictly increasing in $\hat{p}$. The same holds for the limit of the spread $ps(\hat{p})$. Clearly, the limit of $p_H(\hat{p})$ is increasing, being equal to $\hat{p}$.

Proof. I give only the proof for the case of deteriministic costs. The proof of monotonicity of $\bar{p}_L(\hat{p})$ in the case of stochastic fixed costs follows from Lemma 13, whereas the proof of the monotonicity of the price spread for the case of stochastic fixed costs is analogous to the proof of the case of deterministic variable costs.

Recall that the price spread in the HLF is given by

$$\frac{\lambda}{\lambda + r}(V_A(a(\hat{p}), \hat{p}) - C(a(\hat{p}))) - (\hat{v} - \hat{p}).$$

Taking a derivative of this expression w.r.t. $\hat{p}$ and using the envelope theorem results in the expression

$$\frac{\partial}{\partial \hat{p}} ps(\hat{p}) = 1 - \frac{\lambda}{\lambda + r} GN(a(\hat{p})) > 0.$$

[29] In the case of costless information there is no reason for Seller to keep a price spread in the limit. Buyer always waits for the arrival of costless information and thus there is no gain in the HFL from screening between Buyer who has not learned and the one who has learned good news.

Moreover, taking a derivative of the formula for $\bar{p}_L(\hat{p})$, given by $p_L(\hat{p}) = \hat{v} - \frac{\lambda}{\lambda+r}(V(a(\hat{p}), \hat{p}) - C(a(\hat{p})))$ delivers that $\bar{p}_L(\hat{p})$ increasing in $\hat{p}$. $\square$

Thus, a higher $\hat{p}$ *co-moves* with a higher $\bar{p}_L(\hat{p})$, higher limit of $p_H(\hat{p})$ as well as a higher price spread.

Finally, I give the arguments showing that multiplicity persists as accuracy costs become vanishingly small.

Remark 7. ***Deterministic variable costs.***

Note that Lemmas 11 and 12 require the neighborhood $\mathcal{N}$ to be dependent on the costs c. It is easy to see from the proofs, that whenever passing from some costs c_1 to some c_2 with $c_1 > c_2$ (for $a \neq \frac{1}{2}$), the neighborhood $\mathcal{N}$ can be chosen to be strictly smaller (in the sense of set-inclusion) for the costs c_2 than c_1.

Stochastic fixed costs. *This is similar to the case of deterministic variable costs, except that one replaces the inequality $c_1 \geq c_2$ with $F_1 >_{FOSD} F_2$.*

Proof of Proposition 6 and Theorem 3 in the case of pure pricing

Given the analysis of the case of mixed pricing and because of the analogy in the proof between the two cases accuracy costs, I shorten exposition of the case of pure pricing by only spelling out the proof for the case of deterministic variable costs. In the case of stochastic fixed costs, the only changes are the replacements $\bar{w}(\hat{p}) \rightsquigarrow \bar{v}$, $\underline{w}(\hat{p}) \rightsquigarrow \underline{v}$, $GN(a(\hat{p})) \rightsquigarrow \pi_0$, $\mu \rightsquigarrow \mu_0 \bar{\mu}(\hat{p}, \Delta)$, $\lambda \rightsquigarrow \lambda\bar{\mu}(\hat{p})$ and $c(I(a(\hat{p}))) \rightsquigarrow \mathbb{E}[c|c \leq (1-\pi_0)(\hat{p} - \underline{v})]$.

As a first step, one writes after algebra the Seller optimality condition as follows.

$$\frac{1-U(0)}{U(0)}\zeta(q,\Delta)\frac{p_L - \delta V_\Delta(q,0)}{1-\delta}(1-\delta(1-\mu)(1-q)) \geq \bar{w}(\hat{p},\Delta) - p_L. \quad \text{(A.51)}$$

Lemma 15. *1) Fix any $\Delta > 0$. The reservation price relation of type $\hat{v}$ is solvable for some $p_L(\Delta) = \hat{p}(\Delta)$.*

2) It holds for any limit point $\bar{p}$ of $p_L(\Delta)$ as $\Delta \to 0$ that $\bar{p} > \underline{w}(\bar{p})$.

3) For any sequence $q(\Delta), \Delta \to 0$ with limit point $\bar{q} \in [0,1]$ and any limit point $\bar{p}$ of $p_L(\Delta)$ from part 1) and 2) it holds

$$\frac{p_L(\Delta) - \delta V_\Delta(q(\Delta), 0)}{1-\delta}(1 - \delta(1-\mu)(1-q(\Delta)))$$
$$\to \left(1 + \frac{\lambda}{r} - \bar{q}\right)\bar{p} - \frac{\lambda}{r}\left(GN(a(\bar{p}))\bar{p} + (1 - GN(a(\bar{p}))\underline{w}(\bar{p}))\right) > 0, \quad \Delta \to 0.$$

Proof. 1) One needs to show existence of $p_L \in (\underline{v}, \hat{v})$ that satisfies

$$\hat{v} - p_L = \frac{\delta\mu}{1-\delta+\delta\mu}(GN(a(p_L))(\bar{w}(p_L) - p_L) - c(I(a(p_L)))). \qquad (A.52)$$

If $p_L = \hat{v}$ then the left-hand side of (A.52) is zero, whereas the right-hand side is strictly positive, as shown in the proof of Proposition 6 for the case of mixed pricing. If $p_L > 0$ but very close to 0, then the inequality is reversed because of the inequality (A.27) shown in the case of mixed pricing.[30]

This, and the intermediate-value theorem for continuous functions shows existence.

2) The HFL of (A.52) is given by

$$\hat{v} - \bar{p} = \frac{\lambda}{\lambda + r}(GN(a(\bar{p}))(\bar{w}(\bar{p}) - \bar{p}) - c(I(a(\bar{p})))).$$

This can be transformed through simple algebra into

$$\bar{p} = \frac{r\hat{v} + \lambda(1 - GN(a(\bar{p})))\underline{w}(\bar{p}) + \lambda c(I(a(\bar{p})))}{r + \lambda(1 - GN(a(\bar{p})))}. \qquad (A.53)$$

The right-hand side of (A.53) can be estimated from below as follows.

[30]Note that (A.27) only depends on the average price quoted upon non-disclosure and not the price distribution.

$$\frac{r\hat{v} + \lambda(1 - GN(a(\bar{p})))\underline{w}(\bar{p}) + \lambda c(I(a(\bar{p})))}{r + \lambda(1 - GN(a(\bar{p})))} > \frac{r\hat{v} + \lambda(1 - GN(a(\bar{p})))\underline{w}(\bar{p})}{r + \lambda(1 - GN(a(\bar{p})))} > \underline{w}(\bar{p}).$$

3) Simple algebra delivers

$$\begin{aligned}
&\frac{p_L(\Delta) - \delta V_\Delta(q(\Delta), 0)}{1 - \delta}(1 - \delta(1 - \mu)(1 - q(\Delta))) \\
&= p_L(\Delta)\frac{1 - \delta(1 - \mu)}{1 - \delta} - (1 - \mu)q(\Delta)p_L(\Delta) \\
&- \frac{\mu}{1 - \delta}(GN(a(p_L(\Delta)))p_L(\Delta) + (1 - GN(a(p_L(\Delta))))\underline{w}(p_L(\Delta))).
\end{aligned}$$

The limit statement follows from simple limit algebra. The limit is strictly positive because it can be rewritten as

$$\begin{aligned}
&\left(1 + \frac{\lambda}{r} - \bar{q}\right)\bar{p} - \frac{\lambda}{r}\left(GN(a(\bar{p}))\bar{p} + (1 - GN(a(\bar{p}))\underline{w}(\bar{p}))\right) \\
&= (1 - \bar{q})\bar{p} + \frac{\lambda}{r}(1 - GN(a(\bar{p}))(\bar{p} - \underline{w}(\bar{p})).
\end{aligned}$$

This finishes the proof of the Lemma. □

From here, the proof of existence and HFL characterization of strongly stationary equilibria with pure pricing follows the same steps as the corresponding case for costless learning (see proof arguments in subsection A.1.4). I note here down Buyer and Seller payoffs in the HFL with price upon non-disclosure $\bar{p}$ as well as their sum.

$$V_B(\bar{p}, \varkappa) = \frac{\lambda\left(GN(a(\bar{p})))(\bar{w}(\bar{p}) - \bar{p}) - c(I(a(\bar{p})))\right) + \varkappa(\bar{v} - \bar{p})}{r + \lambda + \varkappa},$$

for Buyer payoff in any HFL with $\varkappa$ and price upon non-disclosure equal to $\bar{p}$.

$$V_S(\bar{p},\varkappa) = \frac{(\lambda GN(a(\bar{p})) + \varkappa)\,\bar{p} + \lambda(1 - GN(a(\bar{p})))\underline{w}(\bar{p})}{r+\lambda+\varkappa},$$

for Seller payoff in any HFL with $\varkappa$ and price upon non-disclosure equal to $\bar{p}$. The sum of payoffs in a HFL with $\varkappa$ and price upon non-disclosure equal to $\bar{p}$ is calculated to be

$$\hat{v} - \frac{\lambda c(I(a(\bar{p}))) + r\hat{v}}{r+\lambda+\varkappa}.$$

A.3 Proofs for sections 1.4.1 and 1.4.2

A.3.1 Results with costly choice of intensity

The case of mixed pricing

Proposition 7 contains only some of the results included in the Propositions of this section of the appendix. The results presented here have the same structure as the results for the case of exogenous intensity: first existence results near the HFL and then the analysis in the HFL.

Throughout the proofs of this section I focus first on the special case $\mathcal{C}(\Delta,\mu) = \Delta \cdot f(\mu)$ for ease of exposition and then comment on the changes needed for the general case of costs $\mathcal{C}(\Delta,\mu)$ which satisfy the conditions stated in section 1.4.1.

Deterministic variable costs on accuracy with endogenous intensity

To save on notation, in the following I suppress the dependence of $\mathcal{C}$ on Δ and I recall it only when looking at arguments near the HFL or in cases where the dependence on Δ is important for the argument. Thus, with some abuse of notation, when I write $\mathcal{C}'(\mu)$ I actually mean $\frac{\partial}{\partial\mu}\mathcal{C}(\Delta,\mu)$.

At the beginning of the period of a strongly stationary equilibrium Buyer who has not learned yet solves the following maximization problem

$$\max_{\mu} \mu(V_A(a(\hat{p}),\hat{p}) - C(a(\hat{p}))) + (1-\mu)(\hat{v} - p_L) - \mathcal{C}(\mu).$$

On the equilibrium path, after the information acquisition, it is true that $V(a(\hat{p})) - C(a(\hat{p})) > \hat{v} - p_L$. Therefore, the FOC condition delivers a unique μ characterized implicitly by the FOC condition

$$(V_A(a(\hat{p}),\hat{p}) - C(a(\hat{p}))) - (\hat{v} - p_L) = \mathcal{C}'(\mu). \tag{A.54}$$

This gives a unique $\mu(\hat{p}, p_L) \in (0,1)$ and therefore defines a map $(\hat{p}, p_L) \mapsto \mu(\hat{p}, p_L)$ as an intensity-reaction function of Buyer. One solves explicitly $\mu(\hat{p}, p_L) = \mathcal{C}'^{-1}((V_A(a(\hat{p})) - C(a(\hat{p}))) - (\hat{v} - p_L))$ and it is trivial to see that the reaction function of Buyer is smooth in $(\hat{p}, p_L)$.

Envelope theorem for the stage of accuracy choice in a strongly stationary equilibrium with deterministic variable costs delivers $\frac{\partial}{\partial \hat{p}}\{V_A(a(\hat{p})) - C(a(\hat{p}))\} < 0$. (A.54) implies for the partial derivatives of μ:

$$\frac{\partial}{\partial \hat{p}}\mu(\hat{p}, p_L) < 0, \quad \frac{\partial}{\partial p_L}\mu(\hat{p}, p_L) > 0.$$

The reservation pricing relations for $\bar{w}$, $\hat{v}$ remain the same as in the model with exogenous intensity, except for the fact that μ is determined at the preceding stage of intensity choice and is therefore a function of $\hat{p}$, besides of Δ.

Just as in the case of deterministic variable costs with exogenous λ, I restrict in the following the analysis to $\hat{p} \in (\underline{v}, \bar{v})$ such that $BL(\hat{p}) = V_A(a(\hat{p}),\hat{p}) - C(a(\hat{p})) > 0$.

Denote $p_L^{min}(\hat{p}) = \hat{v} - BL(\hat{p})$, where $a(\hat{p})$, as always, is the accuracy-reaction function of Buyer once she gets an opportunity to learn. Recall the estimate $BL(\hat{p}) < \hat{v} - c(\pi_0)\hat{p}$ shown in the case of exogenous λ. This estimate remains true in this more general set up as well. This is a consequence of stationary play on path.

Note that $\mu(\hat{p}, p_L^{min}(\hat{p})) = 0$ for all $\hat{p}$. Therefore, $\mu(\hat{p}, p_L) > 0$ if and only if $p_L > p_L^{min}$.

Lemma 16. *Let* $\Gamma = \{\hat{p} \in (\underline{v}, \hat{v}) : \hat{v} - \hat{p} < BL(\hat{p})\}$.[31] *For every compact and non-empty interval I contained in* Γ*, there exists a constant* $\varkappa(I) > 0$ *such that for all* $\hat{p} \in (\underline{v}, \hat{v}]$ *it holds* $\mu_\Delta(\hat{p}, \hat{p}) \geq \varkappa(I)\Delta$ *for all* $0 < \Delta < \bar{\Delta}$*, where* $\bar{\Delta} > 0$ *is uniformly on* $\hat{p} \in I$*.*

Proof. Recall that $\mu_\Delta(\hat{p}, p_L)$ satisfies

$$\frac{1}{1-\mu} f' \left(-\frac{\log(1-\mu)}{\Delta} \right) = BL(\hat{p}) - (\hat{v} - p_L).$$

Because the right-hand side of the FOC condition is uniformly bounded across all $\hat{p} \leq \bar{v}$ and $p_L \leq \hat{v}$ it holds for the unique solution $\mu_\Delta(\hat{p}, p_L)$ (for any $\hat{p} \in (\underline{v}, \bar{v}), p_L \leq \hat{v}$) that all limit points of $\mu_\Delta(\hat{p}, p_L)$ with respect to Δ are zero, and so overall $\mu_\Delta(\hat{p}, p_L) \to 0$, as $\Delta \to 0$. Now pick a $\bar{\Delta}$ such that $\mu_\Delta(\underline{v}, \hat{v}) \leq \frac{1}{2}$ for all $\Delta < \bar{\Delta}$. One uses the uniform upper bound $\mu_\Delta(\hat{p}, p_L) \leq \mu_\Delta(\underline{v}, \hat{v})$, so that the convergence of $\mu_\Delta(\hat{p}, p_L)$ to zero is uniform. This delivers the uniform estimates

$$f'\left(2\frac{\mu_\Delta(\hat{p}, \hat{p})}{\Delta}\right) \geq f'\left(-\frac{log(1-\mu_\Delta(\hat{p}, \hat{p}))}{\Delta}\right) \geq (1-\mu_\Delta(\underline{v}, \hat{v}))(BL(\hat{p}) - (\hat{v} - \hat{p})).$$

Here the first estimate uses the elementary inequality $-\log(1-\mu_\Delta(\hat{p}, \hat{p})) \leq 2\mu_\Delta(\hat{p}, \hat{p})$ for all $\hat{p}$, which holds whenever $\Delta < \bar{\Delta}$ because of the uniform estimate $\mu_\Delta(\hat{p}, p_L) \leq \mu_\Delta(\underline{v}, \hat{v})$ and the Taylor series of the logarithm.

The function $\hat{p} \mapsto BL(\hat{p}) - (\hat{v} - \hat{p})$ is continuous and strictly positive on Γ and so has a positive minimum m on I. It thus follows

$$f'\left(2\frac{\mu_\Delta(\hat{p}, \hat{p})}{\Delta}\right) \geq (1-\mu_\Delta(\underline{v}, \hat{v}))m \geq \frac{1}{2}m, \text{ for all } \hat{p} \in I, \Delta < \bar{\Delta}.$$

This delivers the result, recalling that f' is strictly positive for positive arguments and also strictly increasing. □

[31] Note that the interior of Γ is non-empty as it contains $\hat{v}$.

Remark 8. *In the more general case of intensity costs satisfying* (1.18) *one uses instead that relation to prove the same statement as in Lemma 16. The only fact used in the new proof is the uniformity of the limit in* (1.18). *This is used twice. Once to prove that* $\frac{\mu_\Delta(\hat{p},p_L)}{\Delta}$ *remains bounded as* $\Delta \to 0$*, for any* $\hat{p}, p_L$ *and the second time to show that the same expression is also bounded from below. The details are very similar to the ones in the proof of Lemma 16 and so are skipped.*

Lemma 17. *1-a) Fix any* $\Delta > 0$. *The reservation pricing relation for type* $\hat{v}$ *is solvable for* p_L *for any* $\hat{p} \geq \hat{v}$.

1-b) Let I as in Lemma 16 and let $\varkappa(I)$ *be the corresponding positive constant delivered by Lemma 16. Then there is a* $\bar{\Delta}(I) > 0$ *such that when* $0 < \Delta < \bar{\Delta}(I)$*, the reservation pricing relation for type* $\hat{v}$ *is solvable for all* $\hat{p}$ *in I that satisfy*

$$\hat{v} - \hat{p} < \frac{1}{2}\frac{\varkappa(I)}{r + \varkappa(I)} BL(\hat{p}).^{32} \tag{A.55}$$

Moreover, whenever the reservation pricing relation for type $\hat{v}$ *is solvable for some* $\hat{p}$*, the solution is unique for fixed* Δ *and* $\hat{p}$. $p_L(\Delta, \hat{p})$ *is increasing in* $\hat{p}$ *for fixed* Δ.

2) There exists an open neighborhood of the form $(\underline{p}, \bar{v})$ *of* $\hat{v}$ *and a* $\bar{\Delta} > 0$ *such that the reservation pricing relation for type* $\hat{v}$ *is solvable for* p_L *as long as* $\hat{v} \geq \hat{p} > \underline{p}$*, whenever* $\Delta < \bar{\Delta}$ *and for all* $\Delta > 0$ *as long as* $\hat{p} \geq \hat{v}$*. Moreover, the solution* $p_L(\Delta, \hat{p})$ *is unique and continuous in the parameters* $\Delta, \hat{p}$ *as long as* $\Delta < \bar{\Delta}$ *and* $\hat{p} \in (\underline{p}, \bar{v})$.

Proof. Fix a $\hat{p} \in (\underline{v}, \bar{v})$. One needs to solve for p_L in

$$\hat{v} - p_L = \frac{\delta\mu(\hat{p}, p_L)}{1 - \delta + \delta\mu(\hat{p}, p_L)} BL(\hat{p}). \tag{A.56}$$

Here necessarily $p_L \leq \min\{\hat{p}, \hat{v}\}$. Suppose $\hat{v} \leq \hat{p}$. Then, if one sets $p_L = \hat{v}$ on the left hand side it follows that the right-hand side of (A.56) is strictly larger

[32] There is nothing special about the constant $\frac{1}{2}$. The proof would go through when replacing (A.55) with an estimate of the type

$$\hat{v} - \hat{p} < \tilde{a}\frac{\varkappa(I)}{r + \varkappa(I)} BL(\hat{p}),$$

as long as $\tilde{a} \in (0, 1)$.

because it is positive. If one sets $p_L = p_L^{min}(\hat{p})$ then right-hand side becomes zero, whereas left-hand side is strictly positive.[33] It follows that there exists a $p_L \in (p_L^{min}, \hat{v})$ such that (A.56) is satisfied.

Suppose now that $\hat{p} < \hat{v}$. The condition (A.55) implies that $\hat{v} - \hat{p} < BL(\hat{p})$, which in turn implies that $p_L^{min}(\hat{p}) < \hat{p}$. It holds again that if one sets $p_L = p_L^{min}$ the right-hand side in (A.56) is zero, and the left-hand side is strictly positive. Now note that the function $\delta \mapsto \frac{\delta\mu(\hat{p},\hat{p})}{1-\delta+\delta\mu(\hat{p},\hat{p})}$ is increasing in $\mu(\hat{p},\hat{p})$. Pick a compact interval I in the interior of Γ where Γ is defined in the statement of Lemma 16. Using Lemma 16 one picks a $\bar{\Delta}(I)$ such that for all $\mu_\Delta(\hat{p},\hat{p}) \geq \varkappa(I)\Delta$, whenever $\hat{p} \in I$, $0 < \Delta < \bar{\Delta}(I)$. The condition (A.55) together with the fact that $B(\hat{p})$ is bounded away from zero and from above for $\hat{p} \in I$, imply immediately

$$\hat{v} - \hat{p} < \frac{\delta\Delta\varkappa(I)}{1 - \delta + \delta\Delta\varkappa(I)} BL(\hat{p}), \tag{A.57}$$

if one picks $0 < \bar{\Delta}_1(I) \leq \bar{\Delta}(I)$ small enough. $\Delta_1(I)$ has to ensure that one can go over from (A.55) to (A.57) for all $\hat{p} \in I$. (A.57) together with the estimate $\mu_\Delta(\hat{p},\hat{p}) \geq \varkappa(I)\Delta$ delivers that for $\hat{p} \in I$ and $\Delta < \bar{\Delta}_1(I)$ it holds

$$\hat{v} - \hat{p} < \frac{\delta\mu_\Delta(\hat{p},\hat{p})}{1 - \delta + \delta\mu_\Delta(\hat{p},\hat{p})} BL(\hat{p}).$$

Now the intermediate-value theorem for continuous functions delivers existence.

To see uniqueness for fixed $\Delta, \hat{p}$ note that one can easily rewrite (A.56) as

$$\frac{1 - \delta + \delta\mu_\Delta(\hat{p}, p_L)}{\delta\mu_\Delta(\hat{p}, p_L)} (\hat{v} - p_L) = BL(\hat{p}). \tag{A.58}$$

The left-hand side of (A.58) is strictly decreasing in p_L and so uniqueness follows immediately. Since both sides of (A.58) are strictly decreasing in respectively $\hat{p}$ and p_L it follows that $p_L(\hat{p})$ is strictly increasing in $\hat{p}$.

2) This is an easy consequence of 1) and its proof arguments. For the case $\hat{p} < \hat{v}$, pick an I as in Lemma 16 and for the corresponding $\varkappa(I)$ delivered from

[33] Note also that $p_L^{min}(\hat{p}) < \hat{p}$ automatically in this case.

that Lemma, pick $\underline{p}$ such that (A.55) in the proof of 1-b) above is ensured for any $\hat{p} \in (\underline{p}, \hat{v})$. □

To continue with the proof of existence near HFL I establish first the following auxiliary result.

Lemma 18. *For any fixed $\hat{p}$ such that $BL(\hat{p}) > 0$ and $p_L < \min\{\hat{p}, \hat{v}\}$ such that $\mu_\Delta(\hat{p}, p_L)$ satisfies* (A.54) *and p_L solves the reservation price relation for the type $\hat{v}$ for $\hat{p}$ and Δ, the sequence $\frac{\mu_\Delta(\hat{p}, p_L)}{\Delta}$ converges to a strictly positive limit which is a function of $\hat{p}$.*

Proof. For ease of notation I drop the arguments in this proof μ depends on.

Replacing the reservation price relation of type $\hat{v}$ into (A.54) leads to the equation

$$BL(\hat{p})\frac{1-\delta}{1-\delta+\delta\mu} = \frac{1}{1-\mu}f\left(-\frac{\log(1-\mu)}{\Delta}\right). \tag{A.59}$$

From the estimate $f\left(-\frac{\log(1-\mu)}{\Delta}\right) \leq BL(\hat{p})$ and $\mu \leq -\log(1-\mu)$ one sees that $\frac{\mu}{\Delta}$ remains bounded from above. On the other hand, $\frac{\mu}{\Delta}$ is clearly bounded from below so that overall one has that the sequence is bounded. Let now $\lambda(\hat{p})$ be a limit point of $\frac{\mu}{\Delta}$ as $\Delta \to 0$. It follows from (A.59) that $\lambda(\hat{p})$ satisfies

$$f(\lambda)\frac{r+\lambda}{r} = BL(\hat{p}). \tag{A.60}$$

Here one sees that λ only depends on $\hat{p}$ and it is unique because the left-hand side of (A.60) is strictly increasing in λ. It easily follows, that $\lambda(\hat{p})$ is strictly decreasing in $\hat{p}$. Positivity of the limit is straightforward from the assumptions on f. □

Remark 9. *In the case of general costs of intensity satisfying* (1.18)*, the proof of Lemma 18 follows the same steps, except that one replaces in* (A.59) *the right-hand side with $\frac{\partial}{\partial\mu}\mathcal{C}(\Delta, \Delta \cdot \frac{\mu}{\Delta})$.*

In the following, for a neighborhood $\mathcal{N}$ around $\hat{v}$ so that the reservation pricing of type $\hat{v}$ is solvable for $p_L < \min\{\hat{p}, \hat{v}\}$ for all small Δ, I write $\lambda(\hat{p})$ for the unique solution λ to the equation (A.60).

A careful analysis of the proof of existence near HFL of the strongly stationary equilibria shows that the steps of the proof carry through by just replacing $\lambda \rightsquigarrow \lambda(\hat{p})$ when the following conditions for $\hat{p}$ from a suitable open neighborhood of $\hat{v}$ are required.

$$\lambda(\hat{p})\hat{p}(1 - c(\pi_0)) < r\hat{v}, \tag{A.61}$$

and in addition, depending on whether $\pi_0 \leq \frac{1}{2}$ or $\pi_0 > \frac{1}{2}$ require[34]

$$\text{if } \pi_0 \leq \frac{1}{2}, \text{ then } r > \lambda(\hat{p}), \quad \text{if } \pi_0 > \frac{1}{2}, \text{ then } r > \sqrt{2}\lambda(\hat{p}). \tag{A.62}$$

If $\hat{p}$ is chosen close enough to $\hat{v}$, the conditions in (A.62) are sufficient to ensure (A.61). Another sufficient condition can be given by requiring (A.62) for $\hat{p} = \hat{v}$ and then requiring $\hat{p}$ in a neighborhood of $\hat{v}$ such that (A.61) and (A.62) are then satisfied in that neighborhood. In this way, using (A.60) and the properties of f one arrives at the following sufficient conditions for (A.62) evaluated at $\hat{p} = \hat{v}$.

$$\text{if } \pi_0 \leq \frac{1}{2}, \text{ then } f'(r) > \frac{1}{2}BL(\hat{v}), \quad \text{if } \pi_0 > \frac{1}{2}, \text{ then } f'\left(\frac{1}{\sqrt{2}}r\right) > \frac{\sqrt{2}}{\sqrt{2}+1}BL(\hat{v}). \tag{A.63}$$

Recalling the uniform estimate $BL(\hat{p}) < \pi_0\bar{v}$ one can strengthen these even more by requiring a sufficient condition on f' alone.

$$\text{if } \pi_0 \leq \frac{1}{2}, \text{ then } f'(r) > \frac{1}{2}\pi_0\bar{v}, \quad \text{if } \pi_0 > \frac{1}{2}, \text{ then } f'\left(\frac{1}{\sqrt{2}}r\right) > \frac{\sqrt{2}}{\sqrt{2}+1}\pi_0\bar{v}. \tag{A.64}$$

Remark 10. *One can always ensure that* (A.64) *and other weaker variants of it are satisfied* for some *f, because the class of functions f yielding C through the requirement*

[34]Note that (A.62) can be weakened even more by steps similar to the ones in the proof of Proposition 6. I skip these for the sake of length.

(1.18) *is a cone within the space of convex, strictly increasing, differentiable functions with a zero at zero and derivative zero at zero. Thus, if a* f *as needed in* (1.18) *does not work for* (A.64)*, one can always go over to some rescaling of* f *by a factor* $\alpha \in (0, 1)$.

The following Propositions result.

Proposition 32. *[Existence of mixed SSE with deterministic variable accuracy costs and endogenous intensity of learning] Pick any* $r, \pi_0, \underline{v}, \bar{v}$. *Assume that* (A.64) *holds true.*

Then there is a neighborhood $\mathcal{N}$ *of* $\hat{v}$ *and an* $\varepsilon > 0$ *such that for all* $\Delta < \varepsilon$ *and* $\hat{p} \in \mathcal{N}$ *there exist stationary two-price disclosure equilibria with average price* $\hat{p}$. *Moreover, for any fixed average price* $\hat{p}$ *the quantities* $\mu(\Delta, \hat{p}), a(\hat{p}), q(\hat{p}, \Delta), p_L(\hat{p}, \Delta), p_H(\hat{p}, \Delta), p(\hat{p}, \Delta)$ *are uniquely determined.*

Next follows the analysis in the HFL.

Proposition 33. *For any* $\hat{p}$ *in the neighborhood* $\mathcal{N}$ *of* $\hat{v}$ *coming from Proposition 32 the following hold true in HFL.*

A. *Expected delay in real time is equal to* $\frac{1}{\lambda(\hat{p})}$ *and is increasing in* $\hat{p}$.

B. *The price spread* $ps(\hat{p})$ *is bounded away from zero but the low price is charged with vanishingly small probability.*

C. *Denoting by* $a(\hat{p})$ *the reaction function of Buyer and* $GN(a(\hat{p}))$ *the resulting probability of good news, Buyer's and Seller's payoffs are given by*

$$V_B(\hat{p}) = \frac{\lambda(\hat{p})}{r + \lambda(\hat{p})} \left(GN(a(\hat{p}))(\bar{w}(\hat{p}) - \hat{p}) - C(a(\hat{p}))\right) - \frac{f(\lambda(\hat{p}))}{r + \lambda(\hat{p})}. \tag{A.65}$$

and

$$V_S(\hat{p}) = \frac{\lambda(\hat{p})}{r + \lambda(\hat{p})} \left(GN(a(\hat{p}))\hat{p} + (1 - GN(a(\hat{p})))\underline{w}(\hat{p})\right). \tag{A.66}$$

D. *The shortfall in efficiency (i.e. the difference between $\hat{v}$ and sum of payoffs) is given by*

$$\frac{r}{r+\lambda(\hat{p})}\hat{v} + \frac{\lambda(\hat{p})}{r+\lambda(\hat{p})}C(a(\hat{p})) + \frac{f(\lambda(\hat{p}))}{r+\lambda(\hat{p})},$$

and is always strictly positive.

Details on proofs of Proposition 32 and 33. One finds first a neighborhood $\mathcal{N}$ of $\hat{v}$ and $\bar{\Delta} > 0$ such that $\underline{w}(\hat{p}) < \hat{p}$ for $\hat{p} \in \mathcal{N}$ and so that the reservation pricing for type $\hat{v}$ is solvable for p_L as a function of $\hat{p} \in \mathcal{N}$ whenever $\Delta < \bar{\Delta}$. The estimates (A.28) and (A.29) remain true in this setting, when one also allows the additional dependence of μ on $\hat{p}$. The estimate (A.31) remains the same with the added dependency on $\hat{p}$ for μ. From here the proof follows the same steps as for the case of exogenous intensity.

One has to ensure that in the HFL the equivalent of (A.32) in the proof of the case of exogenous λ remains true. One also has to ensure that the relation (A.34) for the $\hat{p}$ near enough to $\hat{v}$ is valid. The arguments are the same as in the case of exogenous λ, except that now $\lambda \rightsquigarrow \lambda(\hat{p})$ and $\lambda(\hat{p})$ satisfies (A.60).

Once one picks a neighborhood of $\hat{v}$ $\mathcal{N}$ and $\bar{\Delta} > 0$ (where $\bar{\Delta} > 0$ is small enough to work for all $\hat{p} \in \mathcal{N}$) so that the corresponding $\hat{p}$-dependent version of (A.32) and the near-HFL version of the corresponding $\hat{p}$-dependent version of (A.34) is ensured for $\hat{p} \in \mathcal{N}$ and $\Delta < \bar{\Delta}$, the proof of existence follows precisely the same steps as in Proposition 6 (case of deterministic variable costs).

For the HFL limit, one just follows the same steps as in the proof of Theorem 3 (case of deterministic variable costs) and uses the fact that $\frac{\mu_\Delta(\hat{p}, p_L(\hat{p}))}{\Delta} \to \lambda(\hat{p})$, as $\Delta \to 0$ instead of its exogenous version $\frac{\mu(\Delta)}{\Delta} \to \lambda$, as $\Delta \to 0$. □

A.3.2 Stochastic fixed costs of information with endogenous intensity

First note, that $\bar{\mu}(\hat{p}, p_L)(BL(\hat{p}, p_L) - (\hat{v} - p_L))$ for $\hat{p} \in (\underline{v}, \bar{v})$ and $p_L < \min\{\hat{p}, \hat{v}\}$ is again increasing in p_L and decreasing in $\hat{p}$. To see this, just note that

$$\bar{\mu}(\hat{p}, p_L)(BL(\hat{p}, p_L) - (\hat{v} - p_L)) = \mathbb{E}\left[\pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L) - c, c \leq \pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L)\right].$$

The FOC for the optimal choice of intensity $\mu_0(\Delta, \hat{p})$ at the beginning of a period satisfies

$$\bar{\mu}(\hat{p}, p_L)(BL(\hat{p}, p_L) - (\hat{v} - p_L)) = \frac{\partial}{\partial \mu_0} C(\Delta, \mu_0).^{35} \tag{A.67}$$

To ease on notation in the following, I suppress again the dependence of C on Δ when this is not explicitly needed for the argument.

Throughout I use the fact that the function

$$\mathbb{E}_{c \sim F}[x - c, c \leq x], \text{ is strictly increasing in } x \in supp(F),$$

where $supp(F)$ denotes the support of F. This implies that in terms of the best response of Buyer, it holds $\mu_0(\Delta, \hat{p}, p_L) = \rho(\Delta, p_L - \pi_0 \hat{p})$ for some function ρ which is strictly increasing in both of its arguments. In particular, for fixed $\hat{p}, p_L$ the sequence $\frac{\mu_0(\Delta, \hat{p}, p_L)}{\Delta}$ has a limit since it is bounded. Thus, (A.67) delivers a unique $\mu_0(\Delta, \hat{p}, p_L) \in (0, 1)$ for every $\Delta > 0, \hat{p} \in (\underline{v}, \bar{v}), p_L^{min}(\hat{p}) < p_L \leq \min\{\hat{v}, \hat{p}\}$. $\mu_0(\Delta, \hat{p}, p_L) \in (0, 1)$ is strictly decreasing in $\hat{p}$ and strictly increasing in p_L.

Moreover, note that for fixed $\hat{p}, p_L$ the sequence $\frac{\rho(\Delta, p_L - \pi_0 \hat{p})}{\Delta}$ converges to some $\bar{\rho}(p_L - \pi_0 \hat{p})$ with $\bar{\rho}(\cdot)$ increasing.

In all, this delivers a function $\mu_\Delta(\hat{p}, p_L) = \mu_0(\Delta, \hat{p}, p_L)\bar{\mu}(\hat{p}, p_L)$ which is continuous in its arguments and strictly increasing in p_L as well as strictly decreasing in $\hat{p}$. Furthermore it holds again $\mu_\Delta(\hat{p}, p_L^{min}(\hat{p})) = 0$ with $p_L^{min}(\hat{p}) = \hat{v} - \pi_0(\bar{v} - \hat{p})$.

Finally, note that the left-hand side of (A.67) is independent of Δ and therefore, with the same arguments as in the case of deterministic variable costs, the sequence $\frac{\mu_0(\Delta, \hat{p}, p_L)}{\Delta}$ is bounded away from zero, and converges to a function which is denoted henceforth by $\lambda(\hat{p}, p_L)$.

Lemma 19. *1-a) Fix any $\Delta > 0$. The reservation pricing relation for type $\hat{v}$ is always solvable for p_L for any $\hat{p} \geq \hat{v}$.*

[35] Note here that $\pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L) = p_L - \pi_0 \hat{p} - (1 - \pi_0)\underline{v}$ which is strictly positive whenever $p_L > p_L^{min} = (1 - \pi_0)\underline{v} + \pi_0 \hat{p}$.

1-b) There exists a $\underline{p}$ such that the reservation pricing relation for type $\hat{v}$ is always solvable for $\hat{p} \in (\underline{p}, \hat{v})$.

Moreover, whenever the reservation pricing relation for type $\hat{v}$ is solvable for some $\hat{p}$, the solution is unique for fixed Δ and $\hat{p}$. The solution $p_L(\Delta, \hat{p})$ is increasing in $\hat{p}$ for fixed Δ.

2) There exists an open neighborhood of the form $(\underline{p}, \bar{v})$ of $\hat{v}$ and a $\bar{\Delta} > 0$ such that the reservation pricing relation for type $\hat{v}$ is solvable for p_L as long as $\hat{v} \geq \hat{p} > \underline{p}$, whenever $\Delta < \bar{\Delta}$ and for all $\Delta > 0$ as long as $\hat{p} \geq \hat{v}$. Moreover, the solution $p_L(\Delta, \hat{p})$ is unique and continuous in the parameters $\Delta, \hat{p}$ as long as $\Delta < \bar{\Delta}$ and $\hat{p} \in (\underline{p}, \bar{v})$. Finally, $p_L(\hat{p}, \Delta)$ is strictly increasing in $\hat{p}$.

3) As $\Delta \to 0$ the solution $p_L(\Delta, \hat{p})$ converges to $\bar{p}_L(\hat{p})$ which satisfies the relation

$$\hat{v} - \bar{p}_L(\hat{p}) = \frac{\lambda(\hat{p})\bar{\mu}(\hat{p}, \bar{p}_L(\hat{p}))}{r + \lambda(\hat{p})\bar{\mu}(\hat{p}, \bar{p}_L(\hat{p}))} BL(\hat{p}, \bar{p}_L(\hat{p})).$$

Proof. 1-a) is proven just as in Lemma 17.

1-b). Look at the case $\hat{p} < \hat{v}$. I just need to establish one direction of the inequality to justify the use of the intermediate-value theorem. Note that
$BL(\hat{p}, \hat{p}) = \pi_0(\bar{v} - \hat{p}) - \mathbb{E}[c | c \leq (1 - \pi_0)(\hat{p} - \underline{v})]$,
$\mu(\Delta, \hat{p}, \hat{p}) = \mu_0(\Delta, \hat{p}, \hat{p}) F((1 - \pi_0)(\hat{p} - \underline{v}))$, where μ_0 satisfies

$$\mathbb{E}[(1 - \pi_0)(\hat{p} - \underline{v}) - c, c \leq (1 - \pi_0)(\hat{p} - \underline{v})] = C'(\mu_0(\hat{p}, \hat{p})).$$

In particular, $\mu(\Delta, \hat{p}, \hat{p}) > 0$ for all $\hat{p} \in (\underline{v}, \hat{v})$. Clearly $\bar{\mu}(\hat{p}, \hat{p})$ and $\mu_0(\Delta, \hat{p}, \hat{p})$ are increasing and continuous in $\hat{p}$. Take some $\underline{p} < \hat{v}$. Then $\lambda(\hat{p}, \hat{p}) \geq \lambda(\underline{p}, \underline{p})$, $\bar{\mu}(\hat{p}, \hat{p}) \geq \bar{\mu}(\underline{p}, \underline{p})$ for all $\hat{p} \in (\underline{p}, \hat{v})$. First pick $p_0 \in (\underline{p}, \hat{v})$ such that

$$\hat{v} - \hat{p} < \frac{1}{2} \frac{\lambda(\underline{p}, \underline{p})\bar{\mu}(\underline{p}, \underline{p})}{r + \lambda(\underline{p}, \underline{p})\bar{\mu}(\underline{p}, \underline{p})} BL(\underline{p}, \underline{p}),$$

for $\hat{v} > \hat{p} > p_0$. Now pick $\bar{\Delta} > 0$ such that for $\Delta < \bar{\Delta}$ it holds true

$$\frac{1}{2} \frac{\lambda(\underline{p}, \underline{p})\bar{\mu}(\underline{p}, \underline{p})}{r + \lambda(\underline{p}, \underline{p})\bar{\mu}(\underline{p}, \underline{p})} < \frac{\delta \mu_0(\Delta, \underline{p}, \underline{p})\bar{\mu}(\underline{p}, \underline{p})}{1 - \delta + \delta \mu_0(\Delta, \underline{p}, \underline{p})\bar{\mu}(\underline{p}, \underline{p})}$$

Combining and using monotonicity delivers for $\hat{p} \in (p_0, \hat{v}), \Delta < \bar{\Delta}$

$$\hat{v} - \hat{p} < \frac{\delta \mu_0(\Delta, \hat{p}, \hat{p})\bar{\mu}(\hat{p}, \hat{p})}{1 - \delta + \delta \mu_0(\Delta, \hat{p}, \hat{p})\bar{\mu}(\hat{p}, \hat{p})} BL(\hat{p}, \hat{p}).$$

Thus, for all $\hat{p} \in (p_0, \hat{v}), \Delta < \bar{\Delta}$ solvability follows for some $p_L(\Delta, \hat{p}) \in (p_L^{min}(\hat{p}), \hat{p})$.

Uniqueness and part 2), except for monotonicity, follow exactly as in Lemma 17. Monotonicity follows the same steps as in the proof of Lemma 13 and is thus skipped.

3) is an easy consequence of the continuity of the functions involved and of the fact that $\frac{\mu_0(\Delta, \hat{p}, p_L)}{\Delta}$ is continuous in all of the arguments and converges as $\Delta \to 0$. □

Combining (A.67) and the solvability of the reservation price relation of type $\hat{v}$ from Lemma 19 one proves again that there exists a neighborhood of $\hat{v}$ such that $\frac{\mu_0(\hat{p}, p_L(\hat{p}))}{\Delta} \to \lambda(\hat{p})$ (proof is analogous to the one of Lemma 18).

Define

$$\bar{BL}(\hat{p}) = \mathbb{E}\left[\pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L(\hat{p})) - c, c \leq \pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L(\hat{p}))\right]. \tag{A.68}$$

One combines (A.67) with solvability of the reservation price relation of the type $\hat{v}$ to write the relation

$$\bar{BL}(\hat{p}) = f'(\lambda(\hat{p})),$$

as

$$r(\hat{v} - p_L(\hat{p})) = f'(\lambda(\hat{p}))\lambda(\hat{p}). \tag{A.69}$$

Lemma 19 delivers that $\hat{p} \mapsto \lambda(\hat{p})$ is strictly decreasing in $\hat{p}$ when the reservation price relation of the type $\hat{v}$ is satisfied. Thus, a higher average price has a detrimental effect on the intensity chosen to learn. I note also for future use that it follows $p_L(\hat{p}) - \pi_0\hat{p}$ is decreasing in $\hat{p}$. Since this is a sufficient statistic for all of $\bar{\mu}$, μ, $\bar{BL}$, it also follows that all of these are decreasing in $\hat{p}$.

From here on the proof remains very similar to the case of deterministic variable costs. One does the replacements $\lambda(\hat{p}) \rightsquigarrow \lambda(\hat{p})\bar{\mu}(\hat{p})$, where the shortcut $\bar{\mu}(\hat{p}) := \bar{\mu}(\hat{p}, p_L(\hat{p}))$ has been used. To find conditions which are not minimal, but are not directly dependent on the distribution of costs F one can just estimate $\bar{\mu}(\hat{p}) \leq 1$ to get, depending on whether $\pi_0 \leq \frac{1}{2}$ or $\pi_0 > \frac{1}{2}$, a condition just as (A.36) in the proof of Proposition 6 for the case of deterministic variable costs. Finally, one uses (A.69) to estimate $r(\hat{v} - p_L(\hat{p})) < \hat{v} - p_L^{\min}(\hat{p}) = \pi_0(\bar{v} - \hat{p}) < \pi_0\bar{v}$. This delivers the following sufficient condition for existence in the neighborhood of $\hat{v}$.

$$\text{if } \pi_0 \leq \frac{1}{2}, \text{ then } f'(r) > \pi_0\bar{v}, \quad \text{if } \pi_0 > \frac{1}{2}, \text{ then } f'\left(\frac{r}{\sqrt{2}}\right) > \sqrt{2}\pi_0\bar{v}. \quad \text{(A.70)}$$

Note that (A.70) is independent of the specification of F.[36]

The analysis delivers the following Propositions.

Proposition 34. *[Existence with stochastic fixed accuracy costs and endogenous intensity of learning] Pick any $r, \pi_0, \underline{v}, \bar{v}$. Assume that* (A.70) *holds.*

Then there is a neighborhood $\mathcal{N}$ of $\hat{v}$ and an $\varepsilon > 0$ such that for all $\Delta < \varepsilon$ and $\hat{p} \in \mathcal{N}$ there exist strongly stationary equilibria with average price $\hat{p}$. Moreover, for any fixed average price $\hat{p}$ the quantities $\mu(\Delta, \hat{p}), \mu_0(\Delta, \hat{p}), q(\hat{p}, \Delta), p_L(\hat{p}, \Delta), p_H(\hat{p}, \Delta), p(\hat{p}, \Delta)$ are uniquely determined.

And now to the HFL of this case. The proof of the following Proposition is a simple adaptation of the proof from the case of exogenous intensity, with the added dependence of λ on $\hat{p}$.

Proposition 35. *Pick any $r, \pi_0, \underline{v}, \bar{v}$. Assume that* (A.70) *holds. Denote $\bar{\mu}(\hat{p}) = F(\pi_0(\bar{v} - \hat{p}) - (\hat{v} - p_L(\hat{p})))$.*

For any $\hat{p}$ as in Proposition it holds in the HFL

A. *Expected delay in real time is equal to $\frac{1}{\lambda(\hat{p})\bar{\mu}(\hat{p})}$ and it is increasing in $\hat{p}$.*

[36] Again, the stated conditions are not minimal conditions on the parameters and their relaxation follows in a similar way to the relaxation of the existence conditions in the case of deterministic variable costs on accuracy.

B. The price spread $ps(\hat{p})$ is bounded away from zero but the low price is charged with vanishingly small probability

C. Buyer's and Seller's payoffs are given by

$$V_B(\hat{p}) = \frac{\lambda(\hat{p})\bar{\mu}(\hat{p})}{r + \lambda(\hat{p})\bar{\mu}(\hat{p})} \left[\pi_0(\bar{v} - \hat{p} - \mathbb{E}[c|c \leq (\bar{v} - \hat{p}) - (\hat{v} - p_L(\hat{p}))]]\right) - \frac{f(\lambda(\hat{p}))}{r + \lambda(\hat{p})}. \quad (A.71)$$

and

$$V_S(\hat{p}) = \frac{\lambda(\hat{p})\bar{\mu}(\hat{p})}{r + \lambda(\hat{p})\bar{\mu}(\hat{p})} \left(\pi_0 \hat{p} + (1 - \pi_0)\underline{v}\right). \quad (A.72)$$

D. The shortfall in efficiency (i.e. the difference between $\hat{v}$ and sum of payoffs) is given by

$$\frac{r}{r + \lambda(\hat{p})\bar{\mu}(\hat{p})}\hat{v} + \frac{\lambda(\hat{p})\bar{\mu}(\bar{p})}{r + \lambda(\hat{p})\bar{\mu}(\bar{p})}\pi_0 \mathbb{E}[c|c \leq \pi_0(\bar{v} - \hat{p}) - (\hat{v} - \bar{p}_L(\hat{p}))] + \frac{f(\lambda(\hat{p}))}{r + \lambda(\hat{p})},$$

and is positive.

The case of pure pricing

The case of strongly stationary equilibria with pure pricing is a straightforward combination of arguments from the case of mixed pricing and of the case of costless intensity. For brevity's sake I only give a sketch of the arguments for the solvability of the $\hat{v}-$indifference condition.

In the case of deterministic variable costs intensity choice satisfies (recall that $\mu = 1 - e^{-\lambda\Delta}$)

$$BL(\hat{p}) - (\hat{v} - \hat{p}) = C'(\mu).$$

In the case of stochastic fixed costs it satisfies

$$\bar{\mu}(\hat{p}, \hat{p})\left(BL(\hat{p}, \hat{p}) - (\hat{v} - \hat{p})\right) = C'(\mu).$$

In both cases this results in an equilibrium value for μ which I denote by $\mu(\hat{p}, \hat{p})$

using the notation of subsection A.3.1.

The $\hat{v}$-indifference relation in the case of pure pricing is given by

$$\hat{v} - \hat{p} = \frac{\delta \mu(\hat{p}, \hat{p})}{1 - \delta + \delta \mu(\hat{p}, \hat{p})} BL(\hat{p}). \tag{A.73}$$

Solvability of the indifference relation for the case of deterministic variable costs for some $\hat{p} \in (\underline{v}, \hat{v})$ follows very closely that of subsection A.2.3. Therefore I skip it. In the case of stochastic fixed costs, note that the right-hand side of (A.73) is strictly positive if $\hat{p} = \hat{v}$, while the left-hand side is zero. In case $\hat{p} = \underline{v}$ the left-hand side is strictly positive, whereas the right-hand side is zero, because $\bar{\mu}(\underline{v}, \underline{v}) = F((1 - \pi_0)(\underline{v} - \underline{v})) = 0$. From here, the proof follows very closely a combination of the cases of pure pricing for exogenous intensity and the case of mixed pricing with endogenous intensity.

Comparative statics for information acquisition

In this subsection I don't comment separately on the cases of pure and mixed pricing, because the proofs are verbatim the same for both cases.

The results stated in Proposition 8 are part of the statements in the following two Propositions.

Proposition 36. *1) Suppose there are two strongly stationary equilibria in the HFL with the same average price $\hat{p}$ and all parameters the same except for $r_1 > r_2$. Then the equilibrium intensity is higher for r_1 than r_2. Equilibrium accuracy is the same in both cases.*

2) Suppose there are two strongly stationary equilibria with the same average price $\hat{p}$ and all parameters the same except for $\pi_0^1 > \pi_0^2$. Equilibrium accuracy is higher for π_0^1 if $\frac{\bar{v}+\underline{v}}{2} > \hat{p}$, whereas it is higher for π_0^2 if $\frac{\bar{v}+\underline{v}}{2} < \hat{p}$. Equilibrium intensity is always higher for π_0^1.

Proof. Recall that in this case, the benefit of learning in the HFL, $BL(\hat{p})$ is dependent only on $\underline{v}, \bar{v}, \pi_0$. It follows from the relation determining $\lambda(\hat{p})$, given in (A.60), that for fixed $\hat{p}$, λ is increasing in r and also in π_0. Next, look at accuracy

$a(\hat{p})$. Recall that it does not depend on any other parameters, except $\hat{p}$, π_0, $\underline{v}$, $\bar{v}$. Using the first order condition related to the incentive constraint (OL-intensive) one arrives easily at the required result. □

Next the comparative statics for the case of stochastic fixed costs of accuracy.

Proposition 37. *1) Suppose there are two strongly stationary equilibria in the HFL with the same average price $\hat{p}$ and all parameters the same, except for $r_1 > r_2$. Then the equilibrium intensity is higher for r_1 than r_2.*

2) Suppose there are two strongly stationary equilibria with the same average price $\hat{p}$ and all parameters the same, except for $\pi_0^1 > \pi_0^2$. Then the equilibrium intensity is higher for π_0^1.

3) Suppose in the case of stochastic fixed costs there are two strongly stationary equilibria with the same average price $\hat{p}$ and all parameters the same, except for $F_1 >_{FOSD} F_2$. Then λ_1 is lower than λ_2.

Proof. By using the definition (A.68) one writes the relation in part 3) of the statement of Lemma 19 as follows.

$$\frac{r}{\lambda}(\hat{v} - p_L) = \bar{BL}(\hat{p}, p_L). \tag{A.74}$$

Recall that the choice of λ at the beginning of a period satisfies the FOC condition in the HFL.

$$\bar{BL}(\hat{p}, p_L) = f'(\lambda) \tag{A.75}$$

1) Suppose that the environment changes so that there is a higher r, but the other parameters stay the same. There are two choice variables that can adjust to keep the reservation pricing relation for type $\hat{v}$ satisfied, i.e. relation (A.74) intact. Either λ can increase to reinstate the balance in (A.74), or if λ weakly falls, then p_L necessarily goes up. But, other else equal, this leads to a higher level of $\bar{BL}(\hat{p}, p_L)$, which, via (A.75) also leads to an increase of λ. Overall it follows the equilibrium intensity λ must increase with r in the HFL.

2) Suppose that π_0 increases but other parameters are kept the same. Because $\bar{BL}(\hat{p}, p_L)$ is strictly increasing in π_0, all else kept equal, λ and/or p_L necessarily change to reinstate (A.74). Suppose p_L falls enough so that $\bar{BL}$ remains the same in both situations. Then λ necessarily increases. Suppose next, that p_L falls but by little enough so that $\bar{BL}$ increases overall in the situation with the higher prior. Via (A.75) the equilibrium value of λ necessarily increases in this case. Suppose for the remaining case, that p_L remains the same after the change in π_0, so that the (A.74) is reinstated through a fall in λ. But in this case, $\bar{BL}$ necessarily increases as well, so that via (A.75), λ has to increase as well. Thus, this last case cannot arise because it leads to a contradiction.

Overall, it follows that the equilibrium intensity λ in the HFL necessarily increases with the prior.

3) Suppose that F increases in the FOSD-sense. Then for fixed $\hat{p}, p_L$, $\bar{BL}$ decreases. If p_L adjusts upwards so that $\bar{BL}$ overall does not change, (A.74) implies that the equilibrium value λ necessarily decreases. If p_L adjusts upwards so that overall $\bar{BL}$ still decreases, then equilibrium λ falls because of (A.75). Suppose instead that p_L adjusts downwards to a shift of F in the FOSD-sense. Then overall, $\bar{BL}$ decreases. This implies via (A.74) that λ necessarily increases, because $\hat{v} - p_L$ increases in this case. But (A.75) implies that λ has to decrease. This is a contradiction.

Overall, the equilibrium intensity λ in the HFL necessarily decreases whenever F increases in the FOSD-sense. □

A.3.3 Proofs for subsections 1.4.2 and 1.4.3

Proofs for the extension to pre-learning negotations

Proof of Proposition 9. I show first that there is no PBE in which game continues past $t = 0$. Suppose for the sake of contradiction there is such a PBE. In particular, it has to prescribe an offer p on path by Seller with positive probability, which is subsequently rejected with positive probability by Buyer. Let $V_B(p)$ be the payoff of Buyer in the continuation bargaining game after she has rejected p. It holds $\hat{v} - p \leq \delta V_B(p)$. Moreover, the payoff from offering p of Seller is then at most $\delta V_S(p)$ where $V_S(p)$ is Seller-payoff in the continuation bargaining game played after Buyer has rejected p. It holds $\delta V_S(p) \leq \delta(\hat{v} - V_B(p)) \leq \delta\hat{v} + p - \hat{v} = p - (1-\delta)\hat{v}$. Suppose Seller deviates instead to $p' = \hat{v} - \delta V_B(p) - \varepsilon$ for some $\varepsilon > 0$ very small. Then if Buyer acccepts she receives more than $\delta V_B(p)$ and Seller receives $p' > \delta(\hat{v} - V_B(p))$, i.e. strictly more than $\delta V_S(p)$. To arrive at the desired contradiction, assume now in addition that p is so that

– *it is the price on path offered with positive probability and rejected with positive probability, with the highest $V_B(p)$ upon rejection.*

Whenever the set of prices offered on path and rejected with positive probability is compact, this assumption is w.l.o.g. in that, a price p can always be chosen to satisfy it. Otherwise, one can use an approximation of the supremum and pick $\varepsilon(p_n) > 0$ very small, for $p_n, n \geq 1$ a sequence of prices such that they are offered and rejected on path with positive probability and so that $V_B(p_n), n \geq 1$ converge to the supremum in question. In that case one arrives at a contradiction similarly.

Efficiency of equilibria results trivially from the fact that agreement is always at $t = 0$.

□

Discussion of the case of negative values

First, I discuss any necessary adaptations of the proof arguments for existence and characterization of strongly stationary equilibria from the case $\underline{v} \geq 0$ to the case $\underline{v} < 0$.[37]

When both disclosure and walking away are possible, then Buyer has zero continuation payoff whenever bad news arrives. To see this, suppose that in equilibrium the accuracy chosen leads to a bad-news valuation which is

[37]The general results for all PBE for any $\Delta > 0$ hold true verbatim in this variation of the model. Details are available upon request.

non-negative. Then Buyer can disclose bad news immediately and receive a zero continuation payoff. If Buyer instead picks an accuracy such that it leads to a negative valuation after bad news then Buyer can walk away. This again results in zero continuation payoff.

Overall, Buyer's information acquisition decision is analogous to the case of $\underline{v} \geq 0$, because payoff upon exercising the strategic outside option has the same structure as in the baseline model. Seller's payoff in equilibria is different, depending on whether Buyer's type can become negative or not. To include both cases, one needs to just do the following replacement $\underline{w} \rightsquigarrow \max\{\underline{w}, 0\}$ in Seller's payoffs in this case.[38] Otherwise the analysis remains the same.

Note that one does not need to put additional restrictions on parameters to ensure existence, besides the one that $BL(\hat{p}) > 0$, i.e. option value from learning in the HFL is positive. This is always possible for $\hat{p} > 0$ near $\hat{v}$, because any amount of learning leads to a valuation after good news $\bar{w}$ which is strictly positive.[39]

Proof of Proposition 10. 1) is trivial so I focus on parts 2) and 3).

2) The sum of Buyer and Seller payoffs in the HFL is given by

$$V_S(\hat{p}) + V_B(\hat{p}) = \frac{\lambda}{\lambda + r}\left(GN(a(\hat{p}))\bar{w}(\hat{p}) + (1 - GN(a(\hat{p})))\underline{w}^+(\hat{p}) - C(a(\hat{p}))\right).$$

The change in welfare due to the possibility to learn is given by

$$\frac{\lambda}{r+\lambda}C(a(\hat{p})) + \hat{v} - \frac{\lambda}{r+\lambda}\max\{\hat{v}, GN(a(\hat{p}))\bar{w}(\hat{p})\}.$$

To show that learning can be beneficial from a welfare point of view, assume $\pi_0 > \frac{1}{2}$, choose $r \approx \sqrt{2}\lambda$ and let the exploitation costs c be very small where needed, i.e. $c \approx 0$ in the topology of pointwise convergence. The welfare change

[38] Here $\underline{w}$ stands for $\underline{w}(\hat{p})$ or $\underline{v}$ depending on type of accuracy costs.

[39] In the case of pure pricing equilibria $\hat{p}$ is positive and of course strictly below $\hat{v}$.

with respect to the case that learning is impossible is bounded from above by

$$\frac{\lambda}{r+\lambda}o(1) + \hat{v} - \frac{1}{\sqrt{2}+1}\pi_0\bar{v}.$$

Note that the condition $\hat{v} < \frac{1}{\sqrt{2}+1}\pi_0\bar{v}$ is equivalent to

$$\frac{\sqrt{2}}{\sqrt{2}+1}\frac{\pi_0}{1-\pi_0} < -\frac{\underline{v}}{\bar{v}},$$

whereas the condition that $\hat{v} > 0$ is equivalent to

$$\frac{\pi_0}{1-\pi_0} > -\frac{\underline{v}}{\bar{v}}.$$

Overall, it follows that under the condition that

$$\frac{\sqrt{2}}{\sqrt{2}+1}\frac{\pi_0}{1-\pi_0} < -\frac{\underline{v}}{\bar{v}} < \frac{\pi_0}{1-\pi_0},$$

whenever exploitation costs are small enough and $r \approx \sqrt{2}\lambda$, learning is beneficial from a welfare perspective.

3) The sum of Buyer and Seller payoffs in the HFL is given by

$$V_S(\hat{p}) + V_B(\hat{p}) = \frac{\lambda\bar{\mu}(\hat{p})}{r+\lambda\bar{\mu}(\hat{p})}\left(\pi_0\bar{v} - \mathbb{E}[c|c \leq \pi_0(\bar{v}-\hat{p}) - (\hat{v} - \bar{p}_L(\hat{p}))]\right).$$

The efficiency loss with respect to the case where learning is impossible (i.e. $\hat{v} - (V_S + V_B)$) is given by

$$\begin{aligned}&\frac{r}{r+\lambda\bar{\mu}(\hat{p})}\pi_0\bar{v} + \frac{\lambda\bar{\mu}(\hat{p})}{r+\lambda\bar{\mu}(\hat{p})}\mathbb{E}[c|c \leq \pi_0(\bar{v}-\hat{p}) - (\hat{v} - p_L(\hat{p}))] + (1-\pi_0)\underline{v} \\ &= \frac{r}{r+\lambda\bar{\mu}(\hat{p})}\hat{v} + \frac{\lambda\bar{\mu}(\hat{p})}{r+\lambda\bar{\mu}(\hat{p})}\left(\mathbb{E}[c|c \leq \pi_0(\bar{v}-\hat{p}) - (\hat{v} - p_L(\hat{p}))] + (1-\pi_0)\underline{v}\right).\end{aligned}$$

Fix $\hat{p} = \hat{v}$ and consider a sequence of F_n that are smooth and converge uniformly to the distribution of the Dirac measure on zero, i.e. I let the accuracy costs converge to zero approach zero uniformly. Then $\bar{\mu}(\hat{v})$ has a limit point which is strictly above zero. In fact, it is 1, because there should not be any divergence between the learning and exploration rate in the limit $n \to \infty$ as costs of accuracy become vanishingly small. So that near the limit one gets approximately an inefficiency given by

$$\frac{r}{r+\lambda}\pi_0\bar{v} + (1-\pi_0)\underline{v}.$$

Note that this is strictly negative whenever

$$\frac{r}{r+\lambda}\pi_0\bar{v} < -(1-\pi_0)\underline{v} < \pi_0\bar{v},$$

which is exactly the condition given in the statement of Proposition 10.

This finishes the proof of the Proposition.

□

A.4 Other results

This last subsection of Appendix A contains results from the online appendix of my job market paper *Bargaining with endogenous learning*.

This section is organized as follows. The first subsection gives the proof of Proposition 1. The second subsection generalizes the result about *no sequential screening of valuations near the HFL* to its real-time counterpart. Subsection 3 comments on strongly stationary equilibria, whereas subsection 4 comments on the multiplicity of equilibria in the costless case. Subsection 5 analyzes the case of costless choice of learning intensity. Finally, subsection 6 comments on the non-concavity of the value of information in the model in which Buyer can pick general, two-parametric experiments.

A.5 Proof of Proposition 1

Because of the requirement that payoff of Buyer after choice $a < 1$ is positive, it follows that Buyer never decides to disclose *irrespective of the signal she receives*. Distinguish therefore two cases.

Case 1: Suppose Buyer discloses good news on path with positive probability. Suppose the history is such that after $a < 1$ she learns good news. If she reveals it she gets zero continuation payoff. On-path Seller puts positive probability on all three types $\{\underline{w}, \hat{v}, \bar{w}\}$ upon non-disclosure (with $\underline{w} < \hat{v} < \bar{w}$). Therefore, upon non-disclosure she (potentially) mixes between the reservation price of either of the three types. Suppose she asks with positive probability for the reservation price of the types $\{\underline{w}, \hat{v}\}$. Then, for Buyer who has just learned after $a < 1$ that $s = H$ it is strictly better to not disclose, defeating the assumption at the start. Therefore, it must be the case that Seller asks after any history of non-disclosure with probability one for the reservation price of the type $\bar{w}$ in this PBE. But then all types would have zero payoff in this PBE after history h. This contradicts one of the requirements in the statement of the Proposition.

Case 2: Suppose Buyer discloses (if at all) only bad news with positive probability. Consider the case of a PBE where Buyer learns less than perfectly, i.e. with $a \in [\frac{1}{2}, 1)$. Then, just as in Case 1, after seeing signal $s = L$ she has implied valuation $\underline{w}$ and after seeing signal $s = H$ she has implied valuation $\bar{w}$ and it holds

$$\underline{v} < \underline{w} \leq \hat{v} \leq \bar{w} < \bar{v}.$$

For this purported PBE let $\beta^a(w, t)$ be the probability of the game ending at time t (restart the time at zero after h to save on notation) if Buyer valuation after learning is $w \in \{\underline{w}, \bar{w}\}$. Let similarly $\beta^a(w, t, \underline{v})$, $\beta^a(w, t, \bar{v})$ be the probability that type w (after learning with accuracy $a < 1$), learns ex-post that the value of the good is respectively $\underline{v}$ and $\bar{v}$ (after the game ends and payoffs are realized). It holds for either $w = \underline{w}, \bar{w}$:

$$\beta^a(w,t,\underline{v})\underline{v} + \beta^a(w,t,\bar{v})\bar{v} = \beta^a(w,t)w. \tag{A.76}$$

The payoff of Buyer after h in the PBE is given by

$$\sum_{t\geq 0} \delta^t \sum_{w\in\{\underline{w},\bar{w}\}} \beta^a(w,t)(w - p(t,w)),$$

where $p(t,w)$ is the expected payment to Seller upon stopping at time t. Note that because Buyer payoff from the strategy is positive (requirement from statement of Proposition) it cannot be that $\beta^a(w,t)w = 0$ for all t and all w. Using (A.76) one writes the payoff at any fixed period t and any $w \in \{\underline{w}, \bar{w}\}$ as

$$\beta^a(w,t)(w-p(t,w)) = \beta^a(w,t,\underline{v})(\underline{v}-p(t,w)) + \beta^a(w,t,\bar{v})(\bar{v}-p(t,w)). \tag{A.77}$$

Part 1) of Lemma 1 in the main paper implies that $\underline{v} - p(t,w) < \underline{w} - p(t,w) \leq 0$. Thus, under the condition that $a < 1$ Buyer makes *ex-post* a loss whenever she learns that the good has low value. Look at the following deviation for Buyer.

- Deviate to $a = 1$
- If learn $\theta = \bar{v}$ never disclose and imitate strategy as if using $a < 1$ from PBE and having valuation $\bar{w}$
- If learn $\theta = \underline{v}$ immediately disclose.

The deviation leads to higher payoff for Buyer because it leads to payoff from any fixed period t (leaving out discounting) and any $w \in \{\underline{w}, \bar{w}\}$ given by

$$\beta^a(w,t,\bar{v})(\bar{v}-p(t,w)) > \beta^a(w,t)(w-p(t,w)), \text{ whenever } \beta^a(w,t) > 0 \text{ and zero otherwise.} \tag{A.78}$$

(A.78) holds for all $t \geq 0$ and all w. Therefore, integrating and summing up the payoffs across w and time t establishes a contradiction to the equilibrium property. This finishes the proof.

A.5.1 A result about 'No sequential screening of valuations near the HFL' in real-time

I look again at the case of disclosure equilibria, i.e. Buyer with bad news discloses immediately.

Moreover, the same definition of 'divinity in bargaining' as in the appendix of the main paper is used here as well. I also use extensively the notation from the subsection A.2 in the appendix of the main paper. Introduce the following modification of the definition of sequential screening of valuations.

Definition 10 (Sequential screening of valuations.)**.** Say that a PBE features sequential screening of valuations (SSV) if on-path

- Seller quotes a strictly decreasing sequence of deterministic prices $\{r_l, l \leq K\}$ $(K \leq \infty)$ upon non-disclosure,
- the sequence of beliefs of the high type $\gamma_l(\bar{v})$ at the beginning of every period $2 \leq l \leq T+1$ is strictly decreasing over time

Say that a sequence of equilibria with SSV with respective final period on-path $T(\Delta)$, exhibits *SSV in the high-frequency limit* if $T(\Delta)\Delta \rightarrow 0$ as Δ goes to zero along the respective subsequence.

For future use note the uniform inequality

$$U(\gamma) - \gamma = \frac{(1-\gamma)\mu}{1-(1-\gamma)\mu(1-\pi_0)} \leq 2\mu, \quad \forall \gamma \in [0,1], \text{ whenever } \mu < \frac{1}{2}. \tag{A.79}$$

(A.79) implies that if on-path the belief at the start of period T is γ_T then the following estimate is valid[40]

$$\gamma_T = \sum_{t=1}^{T} \gamma_t - \gamma_{t-1} \leq \sum_{t=1}^{T} U(\gamma_{t-1}) - \gamma_{t-1} \leq 2T\mu, \text{ whenever } \mu < \frac{1}{2}. \quad (A.80)$$

Here the last inequality uses (A.79) and the first inequality follows from the skimming property and the fact that the game goes on into a new period only after a rejection of the price quoted in the last period.[41]

Proposition 38. *In both the costless case as well as the case of stochastic fixed costs on accuracy the following holds true. For all Δ small enough, there are no PBE which satisfy the following properties.*

A. The type $\underline{v}$ discloses immediately and $\underline{v} > 0$,

B. SSV in the HFL.

The same is true for the case of deterministic variable costs on accuracy for all $\underline{v} \geq 0$.

Proof. The proof is a modification of the proof of the 'no SSV near the HFL' from the appendix of the main paper. I focus on the costless case and comment on the extension needed for the other cases near the end of the proof. Claims 1. and 2. remain the same, including their proofs. In Claims 3. and 4. one replaces everywhere $U(0)$ with $U(\gamma_T)$, where $T = T(\Delta)$ is the history-length of the pre-SSV part of the on-path play. The calculations otherwise follow the same steps. To give a few more details for Claim 3.: the string of inequalities in its proof is now replaced by

$$\frac{\underline{v}}{\overline{v}} < \frac{r_K}{\overline{v}} \leq \frac{\mu\pi_0 + (1 - \mu\pi_0)U(\gamma_T)}{1 - (1 - \pi_0)\mu(1 - U(\gamma_T))}. \quad (A.81)$$

[40]Recall that $\gamma_0 = 0$ as there is no initial private information.

[41]In particular, the negative selection effect lowers belief of type $\overline{v}$ after the rejection of a price.

To establish the equivalent of Claim 3. for this setting, it is enough to show that $U(\gamma_T) \to 0$ as $\Delta \to 0$. Using (A.79) it is enough to show that $\gamma_T \to 0$ and this is true due to the definition of SSV in the HFL and (A.80).

It is easy to adapt the proof for the costless case to the case of deterministic variable costs or stochastic fixed costs with $\underline{v} > 0$. Just as in the proof of the 'no SSV near the HFL' for these two cases in the main body of the paper, this is because the proof above is based only on arguments of Seller-pricing. □

A.5.2 On strongly stationary equilibria

The results of this section hold for stationary equilibria of all model versions in the paper, but for ease of exposition I only use the notation from the costless case.

Recall the definition of strongly stationary equilibria in section 2 of the main paper: A stationary equilibrium is called *strongly stationary*, if as long as bargaining goes on Seller starts each period with belief concentrated on Buyer type $\hat{v}$.

Let for a sequence of period-lengths $\{\Delta, \Delta \to 0\}$ be $\mathcal{E} = \{\mathcal{E}(\Delta), \Delta\}$ a sequence of equilibria.

Proposition 39. *Suppose that $\mathcal{E} = \{\mathcal{E}(\Delta), \Delta\}$ is a sequence of disclosure equilibria. Let $\mathcal{T}(\Delta)$ be the random variable of the terminal date of on-path play for $\mathcal{E}(\Delta)$ valued on $\{\Delta, 2\Delta, \dots\}$. Suppose that the sequence of random variables $\mathcal{T}(\Delta)$ converges weakly as $\Delta \to 0$. Then the beliefs at the beginning of every period on-path converge to the degenerate distribution on $\hat{v}$.*

Proof. Let $\gamma_t(\bar{v}, \Delta), t \geq 2$ be the sequence of beliefs at the beginning of period t in $\mathcal{E}(\Delta)$ when in the previous period the reservation price of Buyer with good news was rejected. Here it is assumed that $t\Delta$ is in the support of $\mathcal{T}(\Delta)$ for all Δ small enough.

Because of the non-disclosure property, one can repeat the arguments in (A.79) and (A.80) to show that whenever Δ is small enough so that

$\mu(\Delta) = 1 - e^{-\lambda\Delta} < \frac{1}{2}$, it holds

$$\gamma_t(\bar{v}, \Delta)(\bar{v}) \leq 2t\mu(\Delta). \tag{A.82}$$

Moreover, let $\gamma_t(\hat{v}, \Delta), t \geq 2$ be the sequence of beliefs on-path at the beginning of period t in $\mathcal{E}(\Delta)$ when in the previous period the reservation price of Buyer of type $\hat{v}$ was rejected. The assumption of disclosure equilibrium implies $\gamma_t(\hat{v}, \Delta)(\hat{v}) = 1$. This is because Lemma 1 from the main paper implies that after rejection of the reservation price of type $\hat{v}$ Seller puts zero probability on type $\bar{v}$, whereas the assumption on the disclosure choice of type $\underline{v}$ implies that type $\underline{v}$ has zero probability after non-disclosure. (A.82) together with the fact that $\mu(\Delta) = 1 - e^{-\lambda\Delta} \to 0, \Delta \to 0$ and the assumption of disclosure equilibria establishes the result.

□

Note that a sequence of strongly stationary equilibria which are also disclosure equilibria automatically satisfies $\gamma_t(\bar{v}, \Delta)(\bar{v}) = 0$. In particular, Buyer with good news never rejects on-path her reservation price.

A.5.3 On multiplicity of equilibria in the costless case

As a preliminary step in the analysis I note down a remark related to the condition $(C - high)$ from Proposition 3 in the main paper.

Remark 11. *For $\delta \to 1$, $(C - high)$ in the limit becomes*

$$\frac{\bar{v}}{\hat{v}} > \frac{1 - (1 - U(0))^2 \mu}{U(0)(1 - \mu U(0))}.$$

This can be expressed equivalently as

$$\frac{\bar{v} - \hat{v}}{\hat{v}} > \frac{1 - \mu + 2\mu U(0) - U(0)}{U(0)(1 - \mu U(0))}.$$

I now look for stationary mixed equilibria in which μ can be intermediate or high. Specifically, look for stationary mixed equilibria in which in the stationary phase Seller mixes between a high price p_H and a low price p_L and so that $\theta = \bar{v}$ accepts with probability one, both p_H and p_L whereas $\theta = \hat{v}$ accepts p_L with probability $q \in (0, 1)$ and rejects p_H. In case there is a PBE with $q = 1$ one can use the refinement "divinity in bargaining": after a rejection of p_L in any period, next period's belief is again $\gamma = 0$.

Proposition 40. *Suppose that the following condition on the parameters is satisfied.*

$$(C - general) \quad \frac{\mu\pi_0(1-\pi_0)}{1-\mu+\mu\pi_0}(\bar{v} - \underline{v}) < \frac{1-\delta}{1-\delta+\delta\mu}\hat{v}.$$

Then there exists a stationary mixed equilibrium with stationary belief $\gamma = 0$ such that, upon non-disclosure, the low price is accepted with probability $q \in (0, 1]$ from the low type and with probability one from the high type, whereas the high price is accepted only by the $\theta = \bar{v}$ buyer (with probability one).

The condition $C - general$ is always satisfied for fixed parameters $\hat{v}, \delta, \mu$ whenever the spread $\bar{v} - \underline{v}$ is small enough.

Proof. Given that the triple (p_H, p_L, p) will be chosen to make both buyer types $\{\bar{v}, \hat{v}\}$ indifferent between accepting and rejecting in the stationary phase the formulas $p_H(p)$ and $p_L(p)$ from the proof of Theorem 1 from the main paper apply in this setting as well. Just as in the main paper, the payoff function of Seller is calculated to be

$$V_S(q,p) = \frac{\mu\pi_0(p_H p + p_L(1-p)) + (1-\mu)(1-p)qp_L + \mu(1-\pi_0)\underline{v}}{1-\delta(1-\mu)(1-(1-p)q)}.$$

Define the continuous function $f : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$

$$f(p, q) = U(0)(p_H(p) - p_L(p)) + q(1 - U(0))(\delta V_S(q, p) - p_L(p)).$$

For every $p \in [0,1]$ it holds $f(p,0) > 0$. One calculates that

$$f(1,1) = U(0)(\bar{v} - \underline{v})(1 - \pi_0) - \frac{1-\delta}{1-\delta+\delta\mu}\hat{v}.$$

Condition $C - \textit{general}$ implies that $f(1,1) < 0$ and the intermediate-value theorem for continuous functions delivers existence.

Seller and buyer off-path. This is virtually the same, with the obvious replacements, as in Proposition 3 in the main paper. In the case of an equilibrium with $q = 1$ the "divinity in bargaining" refinement delivers a unique off-path belief compatible with equilibrium play.

□

Remark 12. *Propositions 3 in the main paper and Proposition 40 here, together with Remark 11 establish, that there is multiplicity of equilibria, when μ is high enough. To see this formally, note that passing in the limit $\mu \rightarrow 1$ and fixing the other parameters of the game, the condition $C - high$ becomes*

$$\bar{v} > \hat{v}, \tag{A.83}$$

whereas the condition $C - general$ becomes

$$\frac{\bar{v} - \hat{v}}{\hat{v}} < 1 - \delta. \tag{A.84}$$

(A.83) *and* (A.84) *are satisfied whenever δ is small enough or $\bar{v} - \hat{v}$ is very small, e.g. when $\underline{v} = 0$ and π_0 is near enough to one. Thus, there is multiplicity of equilibria for all μ near enough to 1, for a range of other parameters of the game form which has 'positive measure'.*

In contrast to the classical uniqueness results Gul et al. [1986] and Fudenberg et al. [1985], one sees that the multiplicity of equilibria arises due to the two channels of seller-learning: from disclosure and from rejection of offers.

A.5.4 Costless choice of the learning intensity

I assume for this section that both μ, a are *costless* choice parameters of Buyer in every period.

Proposition 41. *Suppose that in the model with costless information one expands the action set of Buyer in one of the following ways.*

1) Buyer can choose in the first period once and for all the intensity of learning μ privately,

or

2) Buyer can choose in every period the intensity of learning.

Then picking $\mu = 1$ at every feasible choice moment is a weakly optimal action for Buyer, irrespective of her type.

Proof. Suppose that there is a PBE where after some history in which Buyer can choose intensity, she choses an intensity λ less than infinite. One knows from Proposition 1 in the main body of the paper that Buyer learns conclusively, whenever there is a chance to learn, unless the PBE features zero payoff for Buyer. Suppose it is the case that Buyer has positive payoff in the PBE. Consider the (undetected) deviation of picking infinite intensity to learn and otherwise following the PBE strategy as prescribed in the original equilibrium. There are two cases to consider.

Case 1. Suppose the original PBE is such that Buyer never agrees to a price when still uncertain about the value of the good.

In this case, the imitation of the strategy of the original PBE, except for the choice of intensity, leads to weakly higher payoffs.

Case 2. Suppose the original PBE is such that after some private history h, Buyer agrees with positive probability to a quoted price when still uncertain about the value of the good. Because of learning once, the current valuation can only be $\hat{v} \in (\underline{v}, \overline{v})$.

Let $\tilde{h}$ be the same history as h for Buyer, but with choice of $\lambda = \infty$ at the beginning of the game. Suppose the previous equilibrium had specified agreement with positive probability for a price of $p = p(h) \leq \hat{v}$. Let also $\beta(h, p)$ be the probability that history h occurs and Seller quotes p right after h. By imitating play in h except for the choice of $\lambda = \infty$ in the first period as well as disclosing immediately the value $\underline{v}$ at $t = 1$, if that happens to be the case, Buyer ensures a payoff of at least $\pi_0 \beta(h, p)(\bar{v} - p) > \beta(h, p)(\hat{v} - p)$ by agreeing to p after h, if the game has still not ended at $t = 1$ and zero otherwise. She would have received a negative payoff under the previous strategy if she had continued with $\theta = \underline{v}$ at $t = 1$. Thus, in the Case 2, the deviation to $\lambda = \infty$ at $t = 1$ leads to a strict improvement for Buyer.

Suppose for the remaining case that Buyer has zero payoff in the PBE and picks an intensity less than ∞. She has zero payoff in the PBE if and only if Seller quotes with probability one the highest valuation of Buyer she deems feasible after every public history. Suppose this buyer valuation is $\bar{w}(h)$ after every public history h on-path. If $\bar{w}(h) < \bar{v}$, Buyer can ensure positive payoff by picking $\lambda = \infty$ in the first period, disclosing immediately if value is $\underline{v}$ and otherwise waiting for the history h to arrive (by rejecting prices and not disclosing till h). This contradicts the assumption of zero payoff for Buyer in the purported PBE. It remains to look at the case that $\bar{w}(h) = \bar{v}$ after *every* public history in which there has been no disclosure. But then, picking $\lambda = \infty$, $a = 1$ and disclosing the value immediately is also optimal, given Seller strategy. □

The next result shows that the generic uniqueness result from the classical works Fudenberg et al. [1985] and Gul et al. [1986] depends crucially *on the assumption of commitment from Buyer's side to have private initial information before she approaches Seller*. Namely, when this assumption is relaxed, the positive-buyer-payoff PBEs which feature Coasian dynamics, co-exist with high-price equilibria in which Buyer receives zero payoff.

Proposition 42. *If both accuracy and intensity of learning are endogenous there is a Coasian equilibrium where Buyer has positive payoff. This equilibrium is not unique as it co-exists with the stationary high equilibria from Proposition 3 in the main paper.*

Proof. In the first step, suppose that Buyer picks $\lambda = \infty$ in the first period. Then the game turns into a classical bargaining game as in Fudenberg et al. [1985] and Gul et al. [1986], with initial private information described by a two-point support distribution. Thus, results from Fudenberg et al. [1985] and Gul et al. [1986] apply.[42]

In the second step, look at the following equilibrium. In each period (or at the start of the game), Buyer picks some $\lambda < \infty$ so that $(C - high)$ from Proposition 3 in the main paper is satisfied. Morever, whenever she gets the opportunity to learn she picks $a = 1$ and discloses $\underline{v}$ immediately. Otherwise, she follows the same price acceptance strategy as in the construction of stationary high-price equilibria from Proposition 3 in the main paper. Seller asks for $\bar{v}$ on-path, just as in the construction for Proposition 3 in the main paper. This constitutes again an equilibrium in the expanded game defined in the statement of this Proposition and Buyer has zero payoff in this equilibrium. □

A.5.5 On (non-)concavity of the value of information

Recall from the main paper the two-parameter experiments. A general experiment is given by $\mathcal{E} : \{\underline{v}, \bar{v}\} \to \Delta(\{H, L\})$ and is fully identified with the two accuracy parameters $a_H = \mathbb{P}(s = H|\theta = \bar{v})$ and $a_L = \mathbb{P}(s = L|\theta = \underline{v})$. Because one can always relabel signals, I focus in the following w.l.o.g. only on the case:

$$\nabla = \{(a_H, a_L) : a_H, a_L \in [0, 1], a_H + a_L \geq 1\}.$$

I show that generally the value of information, given by the option value from learning, is not concave in this model if one allows for general experiments. This warrants the restriction to one-dimensional experiments in the main body of the paper. The possibility that the value of information may be non-concave near uninformative experiments is a well-known phenomenon since the work Radner and Stiglitz [1984]. Chade and Schlee [2002] deliver a generalisation and careful analysis of this phenomenon. In these works the set of experiments that the agent

[42] Note also the similar isomorphism to the classical model in Fudenberg et al. [1985] in the public outside option model of Hwang and Li [2017] (see Appendix A. there).

can choose is parametrized through a one-dimensional variable. In the main paper I consider a class of one-parameter experiments which deliver concave value of information. This section shows that in the model of the main paper the non-concavity of the value of information reappears for two-dimensional experiments.

Define the function $L : int(\nabla) \to \{(l_1, l_2) \in (1, \infty)^2, l_1 + l_2 < l_1 l_2 + 1\}$ given by $L(a_H, a_L) = \left(\frac{a_H}{1-a_L}, \frac{a_L}{1-a_H}\right)$.

Lemma 20. *The map L defined above is a C^∞−diffeomorphism.*

Proof. I give explicitly the diffeomorphism. Given $(l_1, l_2) \in (1, \infty)^2$ one needs to solve for $a_H, a_L \in \nabla$ s.t.

$$l_1 = \frac{a_H}{1 - a_L}, \quad l_2 = \frac{a_L}{1 - a_H}.$$

Algebra shows that necessarily

$$a_H = \frac{l_1 l_2 - l_1}{l_1 l_2 - 1}, \quad a_L = \frac{l_1 l_2 - l_2}{l_1 l_2 - 1}.$$

The condition that $a_H + a_L > 1$ is equivalent to $l_1 + l_2 < l_1 l_2 + 1$ which is precisely the condition that appears in the definition of L. This shows surjectivity. Now I show the map is also injective.

Suppose there are two pairs $(a, b), (a', b') \in \nabla$ such that $L(a, b) = L(a', b')$. Thus it holds that

$$\frac{a}{1 - b} = \frac{a'}{1 - b'}, \quad \frac{b}{1 - a} = \frac{b'}{1 - a'}.$$

Algebra manipulations lead to the identities

$$b' = \frac{1 - a'}{1 - a} b, \quad b' = 1 - \frac{a'}{a}(1 - b).$$

Equating and manipulating one arrives at the conclusion $b = 1 - a$, which contradicts the definition of L. Thus, L is injective. Now it is easy to see that L^{-1}

is continuously differentiable because each component of it is a rational function which is well-defined for every pair in $(1, \infty)^2$. $\square$

While it is not a problem to make the costs convex in l_1, l_2 by just defining them to be so, the benefit of information in this model is not globally concave. This can be illustrated by showing that in general the function

$$f(l_1, l_2) = A\frac{l_1 l_2 - l_1}{l_1 l_2 - 1} + B\frac{1 - l_2}{l_1 l_2 - 1}, \quad A, B \geq 0.$$

In the model with deterministic variable costs on accuracy, the values of A, B correspond respectively to $A = \pi_0 \bar{v}$ and $B = (1 - \pi_0)\underline{v}$.

It is enough to show that the functions $f(l_1, l_2) = \frac{l_1 l_2 - l_1}{l_1 l_2 - 1}$ and $g(l_1, l_2) = \frac{1 - l_2}{l_1 l_2 - 1}$ are not concave. One calculates

$$\frac{\partial^2}{\partial^2 l_1} f(l_1, l_2) = \frac{(l_2 - 1) l_2}{(l_1 l_2 - 1)^4} > 0, \quad \frac{\partial^2}{\partial^2 l_2} g(l_1, l_2) = -2\frac{(l_1 - 1) l_1}{(l_1 l_2 - 1)^3} < 0,$$
$$\frac{\partial^2}{\partial^2 l_1} g(l_1, l_2) = -2\frac{(l_2 - 1) l_2^2}{(l_1 l_2 - 1)^3}, \quad \frac{\partial^2}{\partial l_1 \partial l_2} g(l_1, l_2) = \frac{2 l_2 - 1 - l_1 l_2}{(l_1 l_2 - 1)^3}.$$

One sees immediately that f cannot be concave in (l_1, l_2) because the second derivative of f w.r.t. l_1 is strictly positive, whereas the determinant of the Hessian of g is calculated to be

$$|Hess_g(l_1, l_2)| = 4\frac{(l_1 - 1) l_2^2 l_1 (l_2 - 1)}{(l_1 l_2 - 1)^6} - \frac{(l_1 l_2 - 2 l_2 + 1)^2}{(l_1 l_2 - 1)^6}.$$

It is easily seen that for l_1 near 1 and l_2 large, the determinant of the Hessian becomes negative, so that g cannot be globally concave.

B
Appendix to Chapter 2

B.1 Proofs

In the proofs, we will often use the following fact about news-utility functions with diminishing sensitivity. We omit its simple proof.

Fact. *Let* $d_1, d_2 > 0$ *and suppose* $\mu(0) = 0$.

- *(sub-additivity in gains) If* $\mu''(x) < 0$ *for all* $x > 0$, *then* $\mu(d_1 + d_2) < \mu(d_1) + \mu(d_2)$.
- *(super-additivity in losses) If* $\mu''(x) > 0$ *for all* $x < 0$, *then* $\mu(-d_1 - d_2) > \mu(-d_1) + \mu(-d_2)$

B.1.1 Proof of Proposition 11

Proof. We first justify by backwards induction that the value function is indeed given by

$$U_t^*(x) = (\text{cav}U_t(\cdot \mid x))(x),$$

for all $x \in \Delta(\Theta)$ and all $t \leq T-1$, and that it is continuous in x.

If the receiver enters period $t = T-1$ with the belief $x \in \Delta(\Theta)$, the sender faces the following maximization problem.

$$(P_{T-1}) \qquad \max_{\mu \in \Delta(\Delta(\Theta)), \mathbb{E}[\mu]=x} \int_{\Delta(\Theta)} U_{T-1}(p \mid x) d\mu(p).$$

This is because any sender strategy σ_{T-1} induces a Bayes plausible distribution of posterior beliefs, μ with $\mathbb{E}[\mu] = x$, and conversely every such distribution can be generated by some sender strategy, as in Kamenica and Gentzkow [2011]. It is well-known that the value of problem P_{T-1} is $(\text{cav}U_{T-1}(\cdot \mid x))(x)$, justifying $U_{T-1}^*(x)$ as the value function for any $x \in \Delta(\Theta)$. The objective in P_{T-1} is continuous in p (by assumption on N) and hence in μ, and furthermore the constraint set $\{\mu \in \Delta(\Delta(\Theta)) : \mathbb{E}[\mu] = x\}$ is continuous in x. Therefore, $x \mapsto U_{T-1}^*(x)$ is continuous by Berge's Maximum Theorem.

Assume that we have shown that value function is continuous and given by $U_t^*(x)$ for all $t \geq S$. If the receiver enters period $t = S-1$ with belief x, then the sender's value must be:

$$(P_t) \qquad \max_{\mu \in \Delta(\Delta(\Theta)), \mathbb{E}[\mu]=x} \int_{\Delta(\Theta)} N(p \mid x) + U_{t+1}^*(p) d\mu(p)$$

using the inductive hypothesis that $U_{t+1}^*(p)$ is the period $t+1$ value function. But $N(p \mid x) + U_{t+1}^*(p) = U_t(p \mid x)$ by definition, and it is continuous by the inductive hypothesis. So by the same arguments as in the base case, $U_{S-1}^*(x)$ is the time-$(S-1)$ value function and it is continuous, completing the inductive step.

In the first period, by Carathéodory's theorem, there exist weights $w^1, \ldots, w^K \geq 0$, beliefs $q^1, \ldots, q^K \in \Delta(\Theta)$, with $\sum_{k=1}^K w^k = 1$, $\sum_{k=1}^K w^k q^k = x$, such that $U_1^*(\pi_0) = \sum_{k=1}^K w^k U_1(q^k \mid \pi_0)$. Having now shown U_2^* is the period-2 value function, there must exist an optimal information structure where $\sigma_1(\cdot \mid \theta)$ induces beliefs q^k with probability w^k. This information structure induces one of the beliefs $q^1, \ldots, q^K$ in the second period. Repeating the same procedure for subsequent periods establishes the proposition. □

B.1.2 Proof of Proposition 12

Proof. Suppose $T = 2$. Consider the following family of information structures, indexed by $\varepsilon > 0$. Order the states based on $\mathbb{E}_{c \sim F_\theta}[v(c)]$ and label them $\theta_L, \theta_2, \ldots, \theta_{K-1}, \theta_H$. Let $M = \{m_L, m_2, \ldots, m_{K-1}, m_H\}$. Let $\sigma_t(\theta_k)(m_k) = 1$ for $2 \leq k \leq K-1$, $\sigma_t(\theta_H)(m_H) = 1$, and $\sigma_t(\theta_L)(m_L) = x$, $\sigma_t(\theta_L)(m_H) = 1 - x$ for some $x \in (0, 1)$ so that the posterior belief after observing m_H is $(1-\varepsilon)1_H \oplus \varepsilon 1_L$.

For every $\varepsilon > 0$, the information structure just described leads to one-shot resolution of states $\theta \notin \{\theta_L, \theta_H\}$. The difference between its expected news utility and that of one-shot resolution is $W(\varepsilon)$, given by

$$\pi_0(\theta_H) \cdot [N((1-\varepsilon)1_H \oplus \varepsilon 1_L \mid \pi_0) + N(1_H \mid (1-\varepsilon)1_H \oplus \varepsilon 1_L) - N(1_H \mid \pi_0)]$$
$$+ \frac{\varepsilon}{1-\varepsilon}\pi_0(\theta_H) \cdot [N((1-\varepsilon)1_H \oplus \varepsilon 1_L \mid \pi_0) + N(1_L \mid (1-\varepsilon)1_H \oplus \varepsilon 1_L) - N(1_L \mid \pi_0)].$$

W is continuously differentiable away from 0 and $W(0) = 0$. To show that $W(\varepsilon) > 0$ for some $\varepsilon > 0$, it suffices that $\lim_{\varepsilon \to 0^+} W'(\varepsilon) > 0$. Using the continuous differentiability of N except when its two arguments are identical, this limit is

$$\lim_{\varepsilon \to 0^+} \frac{N((1-\varepsilon)1_H \oplus \varepsilon 1_L \mid \pi_0) - N(1_H \mid \pi_0)}{\varepsilon} + \lim_{\varepsilon \to 0^+} \frac{N(1_H \mid (1-\varepsilon)1_H \oplus \varepsilon 1_L)}{\varepsilon}$$
$$+ N(1_H \mid \pi_0) + N(1_L \mid 1_H) - N(1_L \mid \pi_0).$$

Simple rearrangement gives the expression from Proposition 12. The expression for the case of mean-based μ follows by algebra, noting that $N((1-x)1_H \oplus x 1_L \mid \pi_0) = \mu((1-x) - v_0)$ for $x \in [0, 1]$.

If $T > 2$, then note the sender's T-period problem starting with prior π_0 has a value at least as large as the 2-period problem with the same prior. On the other hand, one-shot resolution brings the same total expected news utility regardless of T. □

B.1.3 Proof of Corollary 1

Proof. We verify Proposition 12's condition

$$\mu'(0^+) + \mu(1-\pi_0) - \mu(-\pi_0) > -\mu(-1) + \mu'(1-\pi_0).$$

We have that

$$LHS = \alpha_p + \alpha_p(1-\pi_0) - \beta_p(1-\pi_0)^2 - [\beta_n \pi_0^2 - \alpha_n \pi_0]$$
$$RHS = [-\beta_n + \alpha_n] + [\alpha_p - 2\beta_p(1-\pi_0)]$$

By algebra,

$$LHS - RHS = (1-\pi_0)(\alpha_p - \alpha_n) + (1-\pi_0^2)(\beta_p + \beta_n).$$

Given that $(\alpha_n - \alpha_p) \leq (\beta_p + \beta_n)$ and $1 - \pi_0^2 > 1 - \pi_0$ for $0 < \pi_0 < 1$,

$$LHS - RHS > -(1-\pi_0^2)(\beta_p + \beta_n) + (1-\pi_0^2)(\beta_p + \beta_n) = 0.$$

□

B.1.4 Proof of Corollary 2

Proof. This follows from Proposition 12 because $\mu'(0^+) = \infty$ for the power function. □

B.1.5 Proof of Proposition 13

Proof. Suppose $\Theta = \{\theta_1, \ldots, \theta_K\}$ and assume without loss the states are associated with consumption levels $c_1 < \ldots < c_K$.

Let the message space be $M = \{m_1, \ldots, m_K, m_*\}$. In the first period,

- $\sigma_1(m_k \mid \theta_k) = 1$ for $1 \leq k \leq K-2$,
- $\sigma_1(m_* \mid \theta_{K-1}) = 1$,
- $\sigma_1(m_* \mid \theta_K) = \frac{\pi_0(\theta_{K-1})}{1-\pi_0(\theta_K)}$,
- $\sigma_1(m_K \mid \theta_K) = 1 - \sigma_1(m_* \mid \theta_K)$.

So, message m_k perfectly reveals state θ_k, whereas m_* is a "muddled" message that implies the state is either θ_{K-1} or θ_K. By simple algebra, the probability that the receiver assigns to state θ_K after m_* is the same as the prior belief,

$$\mathbb{P}[\theta_K \mid m_*] = \frac{\pi_0(\theta_K) \cdot \sigma_1(m_* \mid \theta_K)}{\pi_0(\theta_K) \cdot \sigma_1(m_* \mid \theta_K) + \pi_0(\theta_{K-1}) \cdot 1} = \pi_0(\theta_K).$$

In the second period, the information structure perfectly reveals the true state regardless of the last message, $\sigma_2(m_k \mid \theta_k) = 1$ for all $1 \leq k \leq K$.

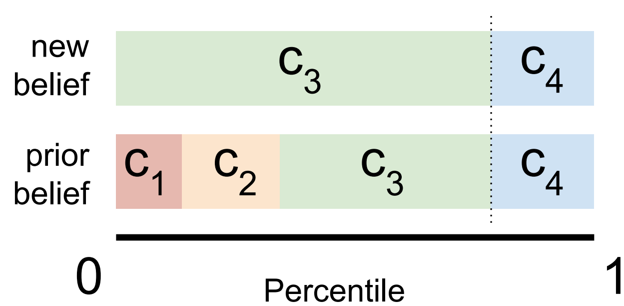

Figure B.1.1: New belief about consumption after the muddled message m_*. Environment with 4 states; new belief compared with the old belief given by the prior π_0.

To compute the news utility of the muddled message m_*, note that at percentiles $p \in [0, \pi_0(\theta_1))$, the change in p-percentile consumption utility is

$v(c_{K-1}) - v(c_1)$. Similarly, for $2 \le k \le K-2$, the change in consumption utility at percentile $p \in \left[\sum_{j=1}^{k-1} \pi_0(\theta_j), \pi_0(\theta_k) + \sum_{j=1}^{k-1} \pi_0(\theta_j)\right)$ is $v(c_{K-1}) - v(c_k)$. There are no changes at percentiles above $\sum_{j=1}^{K-2} \pi_0(\theta_j)$.

If $\theta = \theta_{K-1}$, total news utility from receiving m_* then m_{K-1} is

$$\underbrace{\left[\sum_{k=1}^{K-2} \pi_0(\theta_k) \cdot \mu(v(c_{K-1}) - v(c_k))\right]}_{\text{from } m_* \text{ in period 1}} + \underbrace{\pi_0(\theta_K) \cdot \mu(v(c_{K-1}) - v(c_K))}_{\text{from } m_{K-1} \text{ in period 2}}.$$

This is identical to the news utility from one-shot resolution in state θ_{K-1}. Similarly, the information structure just constructed gives the same news utility as one-shot resolution when the state is θ_k for $1 \le k \le K-2$, and when the state is θ_K and the receiver gets m_K in period 1.

When the receiver sees m_* in period 1 and m_K in period 2 in state θ_K, an event that happens with strictly positive probability since $\pi_0(\theta_{K-1}) < 1 - \pi_0(\theta_K)$ as $K \ge 3$, he gets strictly more news utility than from one-shot resolution.

If $\theta = \theta_K$, total news utility from receiving m_* then m_K is

$$\underbrace{\left[\sum_{k=1}^{K-2} \pi_0(\theta_k) \cdot \mu(v(c_{K-1}) - v(c_k))\right]}_{\text{from } m_* \text{ in period 1}} + \underbrace{\left[\sum_{k=1}^{K-1} \pi_0(\theta_k) \cdot \mu(v(c_K) - v(c_{K-1}))\right]}_{\text{from } m_K \text{ in period 2}},$$

while one-shot resolution gives

$$\sum_{k=1}^{K-1} \pi_0(\theta_k) \cdot \mu(v(c_K) - v(c_k)).$$

For each $1 \le k \le K-2$ (non-empty since $K \ge 3$),

$$\mu(v(c_K) - v(c_{K-1})) + \mu(v(c_{K-1}) - v(c_k)) > \mu(v(c_K) - v(c_k))$$

by sub-additivity in gains. This shows the constructed information structure gives strictly more news utility. □

B.1.6 Proof of Proposition 15

Proof. Let (!) be the following geometric condition: the concavification of $U_1(p|\pi_0)$ involves a linear segment starting at the pair $p = 0$, $U_1(0|\pi_0)$ which is strictly above $U_1(\pi_0|\pi_0)$ when evaluated at $p = \pi_0$. We need to show that (!) holds true if and only if partial bad news are suboptimal. It is clear that whenever the geometric condition (!) is satisfied, partial bad news are suboptimal as the posterior induced in the bad state must be equal to 0 with probability one. On the other hand, knowing that the posterior induced in the bad state is 0 with probability 1 implies two possibilities: (i) either perfect revelation of the state is optimal, or (ii) the optimal information structure involves partial good news and perfect revelation of the bad state. In either case, the only posterior induced in the bad state is that of 0, i.e. the concavification has to include the point $(0, U_1(0|\pi_0))$. From the definition of concavification and the fact that it is supported on two points of the graph of $q \to U_1(q|\pi_0)$, it follows that the concavification has to include a linear segment starting at $(0, U_1(0|\pi_0))$, thus (!) should hold true.

Because of the two-point support feature of the concavification and the fact that the average of the posteriors needs to be equal to the prior $\pi_0 \in (0, 1)$, this implies that there is a $q > \pi_0$, which is the second point of support for the concavification and the linear segment. In case (i) above, it holds $q = 1$ whereas in case (ii) it holds $q < 1$. □

B.1.7 Proof of Corollary 3

We first prove a sufficient condition for the sub-optimality of information structures with partial bad news with $T = 2$. Consider the chord connecting $(0, U_1(0 \mid \pi_0))$ and $(\pi_0, U_1(\pi_0 \mid \pi_0))$ and let $\ell(x)$ be its height at $x \in [0, \pi_0]$. Let $D(x) := \ell(x) - U_1(x \mid \pi_0)$.

Lemma 21. *For this chord to lie strictly above $U_1(p \mid \pi_0)$ for all $p \in (0, \pi_0)$, it suffices that $D'(0) > 0$, $D'(\pi_0) < 0$, and $D''(p) = 0$ for at most one $p \in (0, \pi_0)$.*

Proof. We need $D > 0$ in the region $(0, \pi_0)$. We know that $D(0) = D(\pi_0) = 0$. Given the conditions in the statement and the twice-differentiability of D in $(0, \pi_0)$ it follows that D'' changes sign only once. Moreover, it also follows that $D > 0$ in a right-neighborhood of $x = 0$ and a left-neighborhood of $x = \pi_0$. Suppose D has an interior minimum at $x_0 \in (0, \pi_0)$. Then it holds $D''(x_0) \geq 0$.

Suppose $D''(x) > 0$ for all small x. Then it follows $x_0 \leq p$, where we set $p = \pi_0$ if p doesn't exist. Because $D''(x) \geq 0$ for all $x \leq p$ we have that $D'(x) > 0$ for all $x \leq p$. In particular also $D(x) > 0$ for all such x due to the Fundamental Theorem of Calculus. Thus, the interior minimum is positive and so the claim about D in $(0, \pi)$ is proven in this case.

Suppose instead that $D''(x) < 0$ for all x near enough to 0. Then it follows that $x_0 \geq p$. In particular, for all $x > p$ we have $D''(x) > 0$. Since the derivative is strictly increasing for all $x \in (x_0, \pi_0)$ and $D'(\pi_0) < 0$ we have that $D'(x) < 0$ for all $x \in (x_0, \pi_0)$. In particular, from the Fundamental Theorem of Calculus, $D(\pi_0)$ is strictly below $D(x_0)$. Since $D(\pi_0) = 0$ we have again that $D(x_0) > 0$.

Given the boundary values of D and the signs of the derivatives at $0, \pi_0$ and that any interior minimum of D is strictly positive, we have covered all cases and so shown that $D > 0$ in $(0, \pi_0)$. $\square$ Now we verify that the condition in Lemma 21 holds for the quadratic news utility, which in turn verifies the condition of Proposition 15 for $q = \pi_0$ and shows partial bad news information structures to be strictly suboptimal.

Proof. Clearly, $D(p)$ is a third-order polynomial, so $D^{''}(p)$ has at most one root.

For $p < \pi_0$, we have the derivative

$$\begin{aligned}\frac{d}{dp}U(p \mid \pi_0) =& 2\beta_n(p - \pi_0) + \alpha_n + \alpha_p(1 - p) - \beta_p(1 - p)^2 \\ &+ p(-\alpha_p + 2\beta_p(1 - p)) - (\beta_n p^2 - \alpha_n p) + (1 - p)(2\beta_n p - \alpha_n)\end{aligned}$$

The slope of the chord between 0 and π_0 is:
$\alpha_p - \beta_p + (2\beta_p - \alpha_p + \alpha_n)\pi_0 - (\beta_p + \beta_n)\pi_0^2$. So, after straightforward algebra, $D^{'}(0) = (2(\beta_p + \beta_n) - (\alpha_p - \alpha_n))\pi_0 - (\beta_p + \beta_n)\pi_0^2$. Applying weak loss

aversion with $z = 1$, $\alpha_p - \alpha_n \leq \beta_p - \beta_n$. This shows

$$\begin{aligned} D^{'}(0) &\geq (2(\beta_p + \beta_n) - (\beta_p - \beta_n))\pi_0 - (\beta_p + \beta_n)\pi_0^2 \\ &= (\beta_p + \beta_n)\pi_0(1 - \pi_0) + 2\beta_n\pi_0 > 0 \end{aligned}$$

for $0 < \pi_0 < 1$.

We also derive $D^{'}(\pi_0) = (\alpha_p - 2\beta_p - 2\beta_n - \alpha_n)\pi_0 + (2\beta_p + 2\beta_n)\pi_0^2$. Note that this is a convex parabola in π_0, with a root at 0. Also, the parabola evaluated at 1 is equal to $\alpha_p - \alpha_n \leq 0$, where the inequality comes from the weak loss aversion with $z = 0$. This implies $D^{'}(\pi_0) < 0$ for $0 < \pi_0 < 1$. □

B.1.8 Proof of Proposition 14

Proof. We show that one-shot resolution gives weakly higher news utility conditional on each state, and strictly higher news utility conditional on at least one $\theta \in \Theta_B$.

When $\theta \in \Theta_B$, $\mathbb{P}_{(M,\sigma)}$-almost surely the expectations in different periods form a decreasing sequence $v_0 \geq v_1 \geq ... \geq v_T = v(c_\theta)$. By super-additivity in losses, $\sum_{t=1}^{T} \mu(v_t - v_{t-1}) \leq \mu(v_T - v_0) = \mu(v(c_\theta) - v_0)$. This shows $\mathbb{P}_{(M,\sigma)}$-almost surely the ex-post news utility in state θ is no larger than $\mu(v(c_\theta) - v_0)$, the news utility from one-shot resolution.

Let E be the event where the receiver's expectation strictly decreases two or more times. From the definition of strict gradual bad news, there exists some $\theta^* \in \Theta_B$ so that $\mathbb{P}_{(M,\sigma)}[E \mid \theta^*] > 0$. On $E \cap \{\theta^*\}$, $\sum_{t=1}^{T} \mu(v_t - v_{t-1}) < \mu(v(c_{\theta^*}) - v_0)$ from super-additivity in losses, which means the expected news utility conditional on $E \cap \{\theta^*\}$ is strictly lower than that of one-shot resolution. Combined with the fact that the ex-post news utility in state θ^* is always weakly lower than $\mu(v(c_{\theta^*}) - v_0)$, this shows expected news utility in state θ^* is strictly lower than that of one-shot resolution.

Conditional on any state $\theta \in \Theta_G$, there is some random period $t^* \in \{0, ..., T-1\}$ so that v_t is weakly decreasing up to $t = t^*$ and $v_t = v(c_\theta)$ for

$t > t^*$. If $t^* = 0$, then this belief path yields the same news utility as one-shot resolution. If $t^* \geq 1$, then the total news utility is

$$\sum_{t=1}^{t^*} \mu(v_t - v_{t-1}) + \mu(v(c_\theta) - v_{t^*}).$$

By sub-additivity in gains,

$$\sum_{t=1}^{t^*} \mu(v_t - v_{t-1}) \leq \mu(v_{t^*} - v_0),$$

and for the same reason,

$$\mu(v(c_\theta) - v_{t^*}) \leq \mu(v_0 - v_{t^*}) + \mu(v(c_\theta) - v_0)$$

as we must have $v_{t^*} \leq v_0$. Total news utility is therefore bounded above by

$$\mu(v_{t^*} - v_0) + \mu(v_0 - v_{t^*}) + \mu(v(c_\theta) - v_0).$$

By weak loss aversion, $\mu(v_{t^*} - v_0) + \mu(v_0 - v_{t^*}) \leq 0$, therefore total news utility is no larger than that of one-shot resolution, $\mu(v(c_\theta) - \pi_0)$. □

B.1.9 Proof of Proposition 16

Proof. We have

$$\frac{d}{dp}U(p \mid \pi_0) = 2\alpha_p - \alpha_n - \beta_p + 2\beta_p\pi_0 + p(-2\alpha_p + 2\beta_p + 2\alpha_n + 2\beta_n) + p^2(-3\beta_p - 3\beta_n)$$

Further, p times slope of chord is:

$$\begin{aligned} U(p \mid \pi_0) - U(0 \mid \pi_0) =& U(p \mid \pi_0) - (\beta_n\pi_0^2 - \alpha_n\pi_0) \\ =& \pi_0(-\alpha_p + \alpha_n) + \pi_0^2(-\beta_p - \beta_n) + p(2\alpha_p - \alpha_n - \beta_p) \\ &+ p^2(-\alpha_p + \beta_p + \alpha_n + \beta_n) + p^3(-\beta_p - \beta_n) + p\pi_0(2\beta_p) \end{aligned}$$

Equating $p \cdot \frac{d}{dp} U(p \mid \pi_0) = U(p \mid \pi_0) - U(0 \mid \pi_0)$, we get

$$\pi_0(\alpha_n - \alpha_p) - (\beta_p + \beta_n)\pi_0^2 = p^2(\alpha_n - \alpha_p + \beta_n + \beta_p) - p^3(2\beta_p + 2\beta_n).$$

Define $c = \frac{\alpha_n - \alpha_p}{\beta_n + \beta_p}$. Then, we can write the implicit function as

$$\pi_0 c - \pi_0^2 = p^2(1 + c) - 2p^3.$$

That for every $0 \leq c \leq 1$ and $\pi_0 \in (0, 1)$ this has a solution we know it from the fact that the chord condition is always satisfied for quadratic specification. We want to take derivatives though and maybe try and solve explicitly for the function $p(\pi_0, c)$.

The condition to apply implicit function theorem: define the function $f(\pi_0, p, c) = p^2(1 + c) - 2p^3 - \pi_0 c + \pi_0^2$ with domain $(0, 1)^3$; then we need $\partial_p f(\pi_0, p, c) \neq 0$. If this is true, then we can solve locally for $p(\pi_0, c)$ and then also calculate the local derivative/comparative statics we do below locally. We note that $\partial_p f(\pi_0, p, c) = 2p(1 + c) - 6p^2$. Thus, the only constellation where this would be zero is if $p = \frac{1+c}{3} =: \hat{p} \in [\frac{1}{3}, \frac{2}{3}]$, given the sufficient condition $c \in (0, 1)$ that we are imposing (see Corollary 1) but we leave out the boundary values for a second). Now, let us see for a fixed c, what π_0 would give $\hat{p}$ (because we only focus on region where implicit function gives out a solution). This would mean solving for π_0 the quadratic equation

$$\pi_0^2 - \pi_0 c + \frac{1}{27}(1 + c)^3 = 0. \tag{B.1}$$

Let's calculate the discriminant as a function of c. It is given as $D(c) = c^2 - \frac{4}{27}(1 + c)^3$. Note that $D'(c) = \frac{2}{9}(2 - c)(2c - 1)$. In particular, D is falling from $c = 0$ till $c = \frac{1}{2}$ and increasing from then on till $c = 1$. We note also that $D(0) < 0, D(1) < 0$ so that overall it follows that $D(c) < 0$ for all $c \in [0, 1]$. In particular, it holds that Equation (B.1) is never solvable! This means that $\partial_p f$ never changes sign in $(0, 1)^3 \cap \{(p, \pi_0, c) : \pi_0 c - \pi_0^2 = p^2(1 + c) - 2p^3\}$ (f is a smooth function on its domain). Thus, implicit function theorem is applicable for all $(\pi_0, c) \in (0,1)^2$.

Totally differentiating, we get:

$$d\pi_0 \cdot (\alpha_n - \alpha_p) - (\beta_p + \beta_n)2\pi_0 \cdot d\pi_0 = 2p \cdot dp \cdot (\alpha_n - \alpha_p + \beta_n + \beta_p) - 3p^2 \cdot dp \cdot (2\beta_p + 2\beta_n),$$

which can be rearranged to $\frac{dp}{d\pi_0}\frac{1}{p} = \frac{c-2\pi_0}{2p^2(1+c)-6p^3}$. We note that we showed above the denominator of this expression never changes sign. Given that we know it's negative at $c = 0$ we conclude that it's always negative for all c and all $\pi_0 \in (0,1)$. It follows that unless $c = 0$, $p(\pi_0)$ is falling till some prior and increasing afterwards. For $c = 0$ it is strictly increasing all the way. Note that an implication of the shape for the case of $c > 0$ is that $p(0, c) = \frac{1+c}{2}$ (because the other root which is zero would lead to a contradiction of the shape, given that $p \in [0,1]$)). Thus, the amount of partial good news in the good state remains bounded away from zero as the prior indicates more and more that overall the state is bad with high probability. □

B.1.10 Proof of Proposition 17

Proof. Consider an agent who prefers B over A. In state A, he gets $\mu(-\rho_0)$ with one-shot information, but $\sum_{t=1}^{T} \mu(\rho_t - \rho_{t-1})$ with gradual information. For each t, $\rho_t - \rho_{t-1} < 0$, and furthermore $\sum_{t=1}^{T} \rho_t - \rho_{t-1} = -\rho_0$ by telescoping and using the fact that $\rho_T = 0$. Due to super-additivity in losses, we get that $\mu(-\rho_0) > \sum_{t=1}^{T} \mu(\rho_t - \rho_{t-1})$. In state B, he gets $\mu(1 - \rho_0)$ with one-shot information. With gradual information, let $\hat{T} \leq T$ be the first period where the coin toss comes up tails. His news utility is $\left[\sum_{t=1}^{\hat{T}-1} \mu(\rho_t - \rho_{t-1})\right] + \mu(1 - \rho_{\hat{T}-1})$ where each $\rho_t - \rho_{t-1} < 0$ for $1 \leq t \leq \hat{T} - 1$. Again by super-additivity in losses, $\sum_{t=1}^{\hat{T}-1} \mu(\rho_t - \rho_{t-1}) < \mu(\rho_{\hat{T}-1} - \rho_0)$. By sub-additivity in gains, $\mu(1 - \rho_{\hat{T}-1}) < \mu(\rho_0 - \rho_{\hat{T}-1}) + \mu(1 - \rho_0) \leq -\mu(\rho_{\hat{T}-1} - \rho_0) + \mu(1 - \rho_0)$, where the weak inequality follows since $\lambda \geq 1$. Putting these pieces together,

$$\left[\sum_{t=1}^{\hat{T}-1} \mu(\rho_t - \rho_{t-1})\right] + \mu(1 - \rho_{\hat{T}-1}) < \mu(\rho_{\hat{T}-1} - \rho_0) - \mu(\rho_{\hat{T}-1} - \rho_0) + \mu(1 - \rho_0)$$

$$= \mu(1 - \rho_0)$$

as desired.

Now consider an agent who prefers A over B. We show that when $\lambda = 1$, the agent *strictly* prefers gradual information to one-shot information. By continuity of news utility in λ, the same strict preference must also hold for λ in an open neighborhood around 1.

In state A, the agent gets $\mu(1 - \pi_0)$ with one-shot information, but $\sum_{t=1}^{T} \mu(\pi_t - \pi_{t-1})$ with gradual information. For each t, $\pi_t - \pi_{t-1} > 0$, and furthermore $\sum_{t=1}^{T} \pi_t - \pi_{t-1} = 1 - \pi_0$ by telescoping and using the fact that $\pi_T = 1$. Due to sub-additivity in gains, we get that $\sum_{t=1}^{T} \mu(\pi_t - \pi_{t-1}) > \mu(1 - \pi_0)$. In state B, he gets $\mu(-\pi_0)$ with one-shot information. With gradual information, let $\hat{T} \leq T$ be the first period where the $X_{\hat{T}} = 0$. His news utility is $\left[\sum_{t=1}^{\hat{T}-1} \mu(\pi_t - \pi_{t-1})\right] + \mu(-\pi_{\hat{T}-1})$ where each $\pi_t - \pi_{t-1} > 0$ for $1 \leq t \leq \hat{T} - 1$. Again by sub-additivity in gains, $\sum_{t=1}^{\hat{T}-1} \mu(\pi_t - \pi_{t-1}) > \mu(\pi_{\hat{T}-1} - \pi_0)$. By super-additivity in losses, $\mu(-\pi_{\hat{T}-1}) > \mu(-(\pi_{\hat{T}-1} - \pi_0)) + \mu(-\pi_0) = -\mu(\pi_{\hat{T}-1} - \pi_0) + \mu(-\pi_0)$, where the equality comes from the fact that $\lambda = 1$ so μ is symmetric about 0. Putting these pieces together,

$$\left[\sum_{t=1}^{\hat{T}-1} \mu(\pi_t - \pi_{t-1})\right] + \mu(-\pi_{\hat{T}-1}) > \mu(\pi_{\hat{T}-1} - \pi_0) - \mu(\pi_{\hat{T}-1} - \pi_0) + \mu(-\pi_0)$$

$$= \mu(-\pi_0)$$

as desired. $\square$

B.1.11 Proof of Proposition 18

Proof. (1) Suppose μ is two-part linear with $\mu(x) = kx$ for $x \geq 0$, $\mu(x) = \lambda kx$ for $x < 0$, where $k > 0, \lambda \geq 1$. Then, agents preferring either state will strictly prefer one-shot information over gradual information. Indeed, since news utility is proportional to negative of expected movement in beliefs, and since $\mathbb{E}[\sum_{t=1}^{T} |\pi_t - \pi_{t-1}|] = \mathbb{E}[\sum_{t=1}^{T} |(1-\rho_t) - (1-\rho_{t-1})|] = \mathbb{E}[\sum_{t=1}^{T} |\rho_t - \rho_{t-1}|]$, agents preferring state A and state B also derive the same amount of news utility from each informational structure and hence have the same intensity of preference for one-shot information.

If $\lambda = 1$, agents do not exhibit strict preference for either information structure.

(2) Anticipatory utility. If u is linear, then agents are indifferent between gradual and one-shot information, so (up to tie-breaking) the agents preferring states A and B have the same preference over information structure. If u is strictly concave, then for $1 \leq t \leq T-1$, $\mathbb{E}[u(\pi_t)] < u(\pi_0)$ and $\mathbb{E}[u(\rho_t)] < u(\rho_0)$ by combining the martingale property and Jensen's inequality. So all agents strictly prefer to keep their prior beliefs until the last period and will therefore all choose one-shot information.

(3) Suspense and surprise. Ely, Frankel, and Kamenica [2015] mention a "state-dependent" specification of their surprise and suspense utility functions. With two states, A and B, their specification uses weights $\alpha_A, \alpha_B > 0$ to differentially re-scale belief-based utilities for movements in the two different directions. Specifically, their re-scaled suspense utility is

$$\sum_{t=0}^{T-1} u\left(\mathbb{E}_t\left[\alpha_A \cdot (\pi_{t+1} - \pi_t)^2 + \alpha_B \cdot (\rho_{t+1} - \rho_t)^2\right]\right)$$

and their re-scaled surprise utility is

$$\mathbb{E}\left[\sum_{t=1}^{T} u\left(\alpha_A \cdot (\pi_{t+1} - \pi_t)^2 + \alpha_B \cdot (\rho_{t+1} - \rho_t)^2\right)\right].$$

We may consider agents with opposite preferences over states A and B as agents with different pairs of scaling weights (α_A, α_B). Specifically, say there are $\alpha^{\text{High}} > \alpha^{\text{Low}} > 0$. For an agent preferring A, $\alpha_A = \alpha^{\text{High}}$, $\alpha_B = \alpha^{\text{Low}}$. For an agent preferring B, $\alpha_A = \alpha^{\text{Low}}$, $\alpha_B = \alpha^{\text{High}}$. But note that we always have $\pi_{t+1} - \pi_t = -(\rho_{t+1} - \rho_t)$, so along every realized path of beliefs, $(\pi_{t+1} - \pi_t)^2 = (\rho_{t+1} - \rho_t)^2$. This means these two agents with the opposite scaling weights actually have identical objectives and therefore will have the same preference over gradual or one-shot information. □

B.1.12 Proof of Lemma 3

Proof. **Part 1.** Fix a prior π_0 and a pair $(\bar{M}, \bar{\sigma})$ which induces an equilibrium as in Definition 6. We focus on the case that $|\bar{M}| > 2$ as the other cases are trivial.

Let $M = \{g, b\}$ and we will inductively define the sender's strategy σ_t on t so that (M, σ) is another equilibrium which delivers the same expected utility as $(\bar{M}, \bar{\sigma})$. In doing so we will successively define a sequence of subsets of histories, $H^t_{int} \subseteq M^t$ and $\bar{H}^t_{int} \subseteq \bar{M}^t$, which are length t histories associated with interior equilibrium beliefs about the state in the new and old equilibria, as well as a map φ that associates new histories to old ones.

Let $H^0_{int} = \bar{H}^0_{int} := \{\varnothing\}$, $\varphi(\varnothing) = \varnothing$.

Once we have defined σ_{t-1}, H^{t-1}_{int}, $\bar{H}^{t-1}_{int}$ and $\varphi : H^{t-1}_{int} \to \bar{H}^{t-1}_{int}$, we then define σ_t. If $h^{t-1} \notin H^{t-1}_{int}$, then simply let $\sigma_t(h^{t-1}, \theta)(g) = 0.5$ for both $\theta \in \{G, B\}$. For each $h^{t-1} \in H^{t-1}_{int}$, by the definition of $\bar{H}^{t-1}_{int}$, the equilibrium belief π_{t-1} associated with $\varphi(h^{t-1})$ in the old equilibrium satisfies $0 < \pi_{t-1} < 1$. Let $\Phi_G(h^{t-1})$ and $\Phi_B(h^{t-1})$ represent the sets of posterior beliefs that the sender induces with positive probability in the good and bad states following public history $\varphi(h^{t-1}) \in \bar{H}^{t-1}_{int}$ in $(\bar{M}, \bar{\sigma})$.

We must have $\Phi_G(h^{t-1}) \backslash \Phi_B(h^{t-1}) \subseteq \{1\}$ and $\Phi_B(h^{t-1}) \backslash \Phi_G(h^{t-1}) \subseteq \{0\}$, since any message unique to either state is conclusive news of the state. We construct $\sigma_t(h^{t-1}, \theta)$ based on the following four cases.

Case 1: $1 \in \Phi_G(h^{t-1})$ and $0 \in \Phi_B(h^{t-1})$. Let $\sigma_t(h^{t-1}, G)$ assign probability 1 to

g and let $\sigma_t(h^{t-1}, B)$ assign probability 1 to b.

Case 2: $1 \in \Phi_G(h^{t-1})$ but $0 \notin \Phi_B(h^{t-1})$. By Bayesian plausibility, there exists some smallest $q^* \in (0, \pi_{t-1})$ with $q^* \in \Phi_G(h^{t-1}) \cap \Phi_B(h^{t-1})$, induced by some message $\bar{m}_b \in \bar{M}$ sent with positive probabilities in both states. Also, some message $\bar{m}_g \in \bar{M}$ sent with positive probability in state G induces belief 1. Let $\sigma_t(h^{t-1}, B)(b) = 1$ and let $\sigma_t(\varnothing, G)(b) = x$ where $x \in (0, 1)$ solves $\frac{\pi_{t-1}x}{\pi_{t-1}x + (1-\pi_{t-1})} = q^*$.

Case 3: $1 \notin \Phi_G(h^{t-1})$ but $0 \in \Phi_B(h^{t-1})$. By Bayesian plausibility, there exists some largest $q^* \in (\pi_{t-1}, 1)$ with $q^* \in \Phi_G(h^{t-1}) \cap \Phi_B(h^{t-1})$. Let $\sigma_t(h^{t-1}, G)(g) = 1$ and let $\sigma_t(h^{t-1}, B)(g) = x$ where $x \in (0, 1)$ solves $\frac{\pi_{t-1}}{\pi_{t-1} + (1-\pi_{t-1})x} = q^*$.

Case 4: $1 \notin \Phi_G(h^{t-1})$ and $0 \notin \Phi_B(h^{t-1})$. By Bayesian plausibility, $\Phi_G(h^{t-1}) = \Phi_B(h^{t-1})$, and there exist some largest $q_L \leq \pi_{t-1}$ and smallest $q_H \geq \pi_{t-1}$ in this common set of posterior beliefs, and further there exist $x, y \in (0, 1)$ so that $\frac{\pi_{t-1}x}{\pi_{t-1}x + (1-\pi_{t-1})y} = q_H$ and $\frac{\pi_{t-1}(1-x)}{\pi_{t-1}(1-x) + (1-\pi_{t-1})(1-y)} = q_L$. Let $\sigma(h^{t-1}, G)(g) = x$ and $\sigma(h^{t-1}, B)(g) = y$.

Having constructed σ_t, let H^t_{int} be those on-path period t histories with interior equilibrium beliefs, that is $h^t = (h^{t-1}, m) \in H^t_{int}$ if and only if $h^{t-1} \in H^{t-1}_{int}$ and $\sigma(h^{t-1}, \theta)(m) > 0$ for both $\theta \in \{G, B\}$. A property of the construction of σ_t is that if $h^{t-1} \in H^{t-1}_{int}$, then both (h^{t-1}, g) and (h^{t-1}, b) are on-path. That is, off-path histories can only be continuations of histories with degenerate beliefs in $\{0, 1\}$.

Let $\bar{H}^t_{int}$ be on-path period t histories with interior equilibrium beliefs in $(\bar{M}, \bar{\sigma})$. By the definition of σ_t, there exists $\bar{m} \in \bar{M}$ so that h^t induces the same equilibrium belief in the new equilibrium as the history $(\varphi(h^{t-1}), \bar{m}) \in \bar{H}^t_{int}$ in the old equilibrium, and we define $\varphi(h^t) := (\varphi(h^{t-1}), \bar{m})$.

The receiver's expected payoff in both the B and G states are the same as in the old equilibrium. To see this, note that by our construction, the receiver's expected payoff in state B is the same as if we took a deterministic selection of messages $m_1, m_2, \ldots$ in the old equilibrium with the property that $\sigma_1(\varnothing, B)(m_1) > 0$ and, for $t \geq 2$, $\sigma_t(m_1, \ldots, m_{t-1}, \theta)(m_t) > 0$. Then, we had the sender play message m_t in period t. Since this sequence of messages is played with positive probability in state B of the old equilibrium, it must yield the expected payoff under B — if it yields higher or lower payoffs, then we can construct a deviation that improves the receiver's ex-ante expected payoffs in the old equilibrium. A similar argument holds for state G.

It remains to check that (M, σ) is an equilibrium by ruling out one-shot deviations. We argued before that all off-path histories must follow an on-path history with equilibrium belief in 0 or 1. There are no profitable deviations at off-path histories or at on-path histories with degenerate beliefs, because the receiver does not update beliefs after such histories regardless of the sender's play.

So consider an on-path history with a non-degenerate belief, i.e. a member $h^t \in H^t_{int}$. A one-shot deviation following h^t corresponds to a deviation following $\varphi(h^t)$ in $(\bar{M}, \bar{\sigma})$, and must not be strictly profitable.

Part 2. We now turn to the second claim. If $T \leq T^{'}$, then for any equilibrium with horizon T, we may construct an equilibrium of horizon $T^{'}$ which sends messages in the same way in periods $1, \ldots, T-1$, but babbles starting in period T. This equilibrium has the same expected payoff as the old one. □ Note that the first claim of Lemma 3 also holds for the infinite horizon model of subsection 2.5.3. Nothing in the argument relies on T being finite. This is because the proof argument relies on the one-shot deviation property which holds for equilibria in both finite and infinite horizon models. Thus, in particular, in the proof of Proposition 25 we can also focus on a binary signal space.

B.1.13 Proof of Lemma 4

Proof. Due to sub-additivity,

$$\mu(p) < \mu(p-\pi) + \mu(\pi). \tag{B.2}$$

Note that symmetry implies $\mu(-p) = -\mu(p)$ and that $\mu(-\pi) = -\mu(\pi)$. Rearranged (B.2) is precisely $N(0;\pi) < N(p;\pi)$. □

B.1.14 Proof of Proposition 19

We begin by giving some additional definition and notation.

For $p, \pi \in [0, 1]$, let $N_G(p; \pi) := \mu(p - \pi) + \mu(1 - p)$.

We state and prove a preliminary lemma about N_G and N_B.

Lemma 22. *Suppose μ exhibits diminishing sensitivity and greater sensitivity to losses. Then, $p \mapsto N_G(p; \pi)$ is strictly increasing on $[0, \pi]$ and symmetric on the interval $[\pi, 1]$. For each $p_1 \in [\pi, 1]$, there exists exactly one point $p_2 \in [\pi, 1]$ so that $N_G(p_1; \pi) = N_G(p_2; \pi)$. For every $p_L < \pi$ and $p_H \geq \pi$, $N_G(p_L; \pi) < N_G(p_H; \pi)$. Also, $N_B(p; \pi)$ is symmetric on the interval $[0, \pi]$. For each $p_1 \in [0, \pi]$, there exists exactly one point $p_2 \in [0, \pi]$ so that $N_B(p_1; \pi) = N_B(p_2; \pi)$.*

Proof. We have $\frac{\partial N_G(p;\pi)}{\partial p} = \mu^{'}(p - \pi) - \mu^{'}(1 - p)$. For $0 \leq p < \pi$ and under greater sensitivity to losses, $\mu^{'}(p - \pi) \geq \mu^{'}(\pi - p)$. Since $\mu^{''}(x) < 0$ for $x > 0$, $\mu^{'}(\pi - p) > \mu^{'}(1 - p)$. This shows $\frac{\partial N_G(p;\pi)}{\partial p} > 0$ for $p \in [0, \pi)$.

The symmetry results follow from simple algebra and do not require any assumptions.

Note that $\frac{\partial^2 N_G(p;\pi)}{\partial p^2} = \mu^{''}(p - \pi) + \mu^{''}(1 - p) < 0$ for any $p \in [\pi, 1]$, due to diminishing sensitivity. Combined with the required symmetry, this means $\frac{\partial N_G(p;\pi)}{\partial p}$ crosses 0 at most once on $[\pi, 1]$, so for each $p_1 \in [\pi, 1]$, we can find at most one p_2 so that $N_G(p_1; \pi) = N_G(p_2; \pi)$. In particular, this implies at every intermediate $p_1 \in (\pi, 1)$, we get $N_G(p_1; \pi) > N_G(\pi; \pi)$ since we already have $N_G(1; \pi) = N_G(\pi; \pi)$. This shows $N_G(\cdot; \pi)$ is strictly larger on $[\pi, 1]$ than on $[0, \pi)$.

A similar argument, using $\mu^{''}(x) > 0$ for $x < 0$, establishes that for each $p_1 \in [0, \pi]$, we can find at most one p_2 so that $N_B(p_1; \pi) = N_B(p_2; \pi)$. □

Consider any period $T - 2$ history h_{T-2} in any equilibrium (M, σ^*, p^*) where $p^*(h_{T-2}) = \pi \in (0, 1)$. Let P_G and P_B represent the sets of posterior beliefs induced at the end of $T - 1$ with positive probability, in the good and bad states. The next lemma gives an exhaustive enumeration of all possible P_G, P_B.

Lemma 23. *The sets P_G, P_B belong to one of the following cases.*

A. $P_G = P_B = \{\pi\}$

B. $P_G = \{1\}, P_B = \{0\}$

C. $P_G = \{p_1\}$ for some $p_1 \in (\pi, 1)$ and $P_B = \{0, p_1\}$

D. $P_G = \{\pi, 1\}$ and $P_B = \{0, \pi\}$

E. $P_G = \{p_1, p_2\}$ for some $p_1 \in (\pi, \frac{1+\pi}{2}), p_2 = 1 - p_1 + \pi, P_B = \{0, p_1, p_2\}$.

Proof. Suppose $|P_G| = 1$.

If $P_G = \{\pi\}$, then any equilibrium message not inducing π must induce 0. By the Bayes' rule, the sender cannot induce belief 0 with positive probability in the bad state, so $P_B = \{\pi\}$ as well.

If $P_G = \{1\}$, then any equilibrium message not inducing 1 must induce 0. Furthermore, the sender cannot send equilibrium messages inducing belief 1 with positive probability in the bad state, else the equilibrium belief associated with these messages should be strictly less than 1. Thus $P_B = \{0\}$.

If $P_G = \{p_1\}$ for some $0 \leq p_1 < \pi$, then any equilibrium message not inducing p_1 must induce 0. This is a contradiction since the posterior beliefs do not average out to π.

This leaves the case of $P_G = \{p_1\}$ for some $\pi < p_1 < 1$. Any equilibrium message not inducing p_1 must induce 0. Furthermore, the sender must induce the belief p_1 in the bad state with positive probability, else we would have $p_1 = 1$. At the same time, the sender must also induce belief 0 with positive probability in the bad state, else we violate Bayes' rule. So $P_B = \{0, p_1\}$.

Now suppose $|P_G| = 2$.

In the good state, the sender must be indifferent between two beliefs p_1, p_2 both induced with positive probability. By Lemma 22, $N_G(p; \pi)$ is strictly increasing on $[0, \pi]$ and strictly higher on $[\pi, 1]$ than on $[0, \pi)$, while for each $p_1 \in [\pi, 1]$, there exists exactly one point $p_2 \in [\pi, 1]$ so that $N_G(p_1; \pi) = N_G(p_2; \pi)$. This means we must have $p_1 \in [\pi, \frac{1+\pi}{2}], p_2 = 1 - p_1 + \pi$.

If $P_G = \{\pi, 1\}$, any equilibrium message not inducing π or 1 must induce 0. Also, $1 \notin P_B$, because any message sent with positive probability in the bad state cannot induce belief 1. We cannot have $P_B = \{0\}$, because then the message inducing belief π actually induces 1. We cannot have $P_B = \{\pi\}$ for then we violate Bayes' rule. This leaves only $P_B = \{0, \pi\}$.

If $P_G = \{p_1, p_2\}$ for some $p_1 \in (\pi, \frac{1+\pi}{2})$, then any equilibrium message not inducing p_1 or p_2 must induce 0. Also, $p_1, p_2 \in P_B$, else messages inducing these beliefs give conclusive evidence of the good state. By Bayes' rule, we must have $P_B = \{0, p_1, p_2\}$.

It is impossible that $|P_G| \geq 3$, since, by Lemma 22, $N_G(p; \pi)$ is strictly increasing on $[0, \pi]$ and strictly higher on $[\pi, 1]$ than on $[0, \pi)$, while for each $p_1 \in [\pi, 1]$, there exists exactly one point $p_2 \in [\pi, 1]$ so that $N_G(p_1; \pi) = N_G(p_2; \pi)$. So the sender cannot be indifferent between 3 or more different posterior beliefs of the receiver in the good state. $\square$

We now give the proof of Proposition 19.

Proof. Consider any period $T-2$ history h^{T-2} with $p^*(h^{T-2}) \in (0, 1)$. By Lemma 4, $N_B(p; p^*(h^{T-2})) > N_B(0; p^*(h^{T-2}))$ for all $p \in (p^*(h^{T-2}), 1]$. Therefore, cases 3 and 5 are ruled out from the conclusion of Lemma 23. This shows that after having reached history h^{T-2}, the receiver will get total news utility of $\mu(1 - p^*(h^{T-2}))$ in the good state and $\mu(-p^*(h^{T-2}))$ in the bad state. This conclusion applies to all period $T-2$ histories (including those with equilibrium beliefs 0 or 1). So, the sender gets the same utility as if the state is perfectly revealed in period $T-1$ rather than T, and the equilibrium up to period $T-1$ form an equilibrium of the cheap talk game with horizon $T-1$. By backwards induction, we see that along the equilibrium path, whenever the receiver's belief updates, it is updated to the dogmatic belief in θ. $\square$

B.1.15 Proof of Proposition 20

Proof. The conclusions of Lemmas 22 and 23 continue to hold, since these only depend on μ exhibiting greater sensitivity to losses. As in the proof of Proposition

19, we only need to establish $N_B(p;\pi_0) > N_B(0;\pi_0)$ for all $p \in (\pi_0, 1]$ to rule out cases 3 and 5 from Lemma 23 and hence establish our result.

For $p = \pi_0 + z$ where $z \in (0, 1-\pi_0]$,

$$N_B(p;\pi_0) - N_B(0;\pi_0) = \mu(z) + \mu(-(\pi_0 + z)) - \mu(-\pi_0).$$

Consider the RHS as a function $D(z)$ of z. Clearly $D(0) = 0$, and $D^{'}(z) = \mu^{'}(z) - \mu^{'}(-(\pi_0 + z))$. Since $\min_{z\in[0,1-\pi_0]} \frac{\mu^{'}(z)}{\mu^{'}(-(\pi_0+z))} > 1$, we get $D^{'}(z) > 0$ for all $z \in [0, 1-\pi_0]$, thus $D(z) > 0$ on the same range. □

B.1.16 Proof of Corollary 4

Proof. First, μ exhibits greater sensitivity to losses, because $\mu(-x) = -\lambda\mu(x)$ for all $x > 0$ and we have $\lambda \geq 1$.

To apply Proposition 20, we only need to verify that $\min_{z\in[0,1-\pi_0]} \frac{\mu^{'}(z)}{\mu^{'}(-(\pi_0+z))} > 1$. For the λ-scaled μ, $\min_{z\in[0,1-\pi_0]} \frac{\mu^{'}(z)}{\mu^{'}(-(\pi_0+z))} = \frac{1}{\lambda} \cdot \min_{z\in[0,1-\pi_0]} \frac{\tilde{\mu}^{'}_{pos}(z)}{\tilde{\mu}^{'}_{pos}(\pi_0+z)}$. The assumption that $\min_{z\in[0,1-\pi_0]} \frac{\tilde{\mu}^{'}_{pos}(z)}{\tilde{\mu}^{'}_{pos}(\pi_0+z)} > \lambda$ gives the desired conclusion. □

B.1.17 Proof of Proposition 21

Proof. By the proof of Proposition 22, which does not depend on this result, there is a GGN equilibrium with one intermediate belief $p \in (\pi_0, 1)$ whenever $N_B(p;\pi_0) = N_B(0;\pi_0)$. In this equilibrium, the sender induces a belief of either p or 0 by the end of period 1, then babbles in all remaining periods of communication. Since the sender is indifferent between inducing belief p or 0 in the bad state, this equilibrium gives the same payoff as the babbling one in the bad state. But, since $\mu(p-\pi_0) + \mu(1-p) > \mu(1-\pi_0)$ due to strict concavity of $\tilde{\mu}_{pos}$, the receiver gets strictly higher news utility in the good state.

To find $\bar{\lambda}$ that guarantees the existence of a p solving $N_B(p;\pi_0) = N_B(0;\pi_0)$, let $D(p) := N_B(p;\pi_0) - N_B(0;\pi_0)$. We have $D(\pi_0) = 0$ and

$\lim_{p \to \pi_0^+} D'(p) = \lim_{x \to 0^+} \tilde{\mu}'_{pos}(x) - \mu'(-\pi_0) = \lim_{x \to 0^+} \tilde{\mu}'_{pos}(x) - \lambda \mu'(\pi_0)$. For any finite λ, this limit is ∞, since $\lim_{x \to 0^+} \tilde{\mu}'_{pos}(x) = \infty$. On the other hand, $D(1) = \mu(1 - \pi_0) + \mu(-1) - \mu(-\pi_0) = \tilde{\mu}_{pos}(1 - \pi_0) - \lambda(\tilde{\mu}_{pos}(1) - \tilde{\mu}_{pos}(\pi_0))$. Since $\tilde{\mu}_{pos}(1) - \tilde{\mu}_{pos}(\pi_0) > 0$, we may find a large enough $\bar{\lambda} \geq 1$ so that $\tilde{\mu}_{pos}(1 - \pi_0) - \bar{\lambda}(\tilde{\mu}_{pos}(1) - \tilde{\mu}_{pos}(\pi_0)) < 0$. Whenever $\lambda \geq \bar{\lambda}$, we therefore get $D(\pi_0) = 0$, $\lim_{p \to \pi_0^+} D'(p) = \infty$, and $D(1) < 0$. By the intermediate value theorem applied to the continuous D, there exists some $p \in (\pi_0, 1)$ so that $D(p) = 0$. □

B.1.18 Proof of Proposition 22

Proof. Let J intermediate beliefs satisfying the hypotheses be given. We construct a gradual good news equilibrium where $p_t = q^{(t)}$ for $1 \leq t \leq J$, and $p_t = q^{(J)}$ for $J + 1 \leq t \leq T - 1$.

Let $M = \{g, b\}$ and consider the following strategy profile. In period $t \leq J$ where the public history so far h^{t-1} does not contain any b, let $\sigma(h^{t-1}; G)(g) = 1$, $\sigma(h^{t-1}; B)(g) = x$ where $x \in (0, 1)$ satisfies $\frac{p_{t-1}}{p_{t-1} + (1 - p_{t-1})x} = p_t$. But if public history contains at least one b, then $\sigma(h^{t-1}; G)(b) = 1$ and $\sigma(h^{t-1}; B)(b) = 1$. Finally, if the period is $t > J$, then $\sigma(h^{t-1}; G)(b) = 1$ and $\sigma(h^{t-1}; B)(b) = 1$. In terms of beliefs, suppose h^t has $t \leq J$ and every message so far has been g. Such histories are on-path and get assigned the Bayesian posterior belief. If h^t has $t \leq J$ and contains at least one b, then it gets assigned belief 0. Finally, if h^t has $t > J$, then h^t gets assigned the same belief as the subhistory constructed from its first J elements. It is easy to verify that these beliefs are derived from Bayes' rule whenever possible.

We verify that the sender has no incentive to deviate. Consider period $t \leq J$ with history h^{t-1} that does not contain any b. The receiver's current belief is p_{t-1} by construction.

In state B, we first calculate the sender's equilibrium payoff after sending g. The receiver will get some I periods of good news before the bad state is revealed, either by the sender or by nature in period T. That is, the equilibrium news utility with I periods of good news is given by

$$\sum_{i=1}^{I} \mu(p_{t-1+i} - p_{t-2+i}) + \mu(-p_{t-1+I}).$$

Since $p_{t-1+I} \in P^*(p_{t-2+I})$, we have $N_B(p_{t-1+I}; p_{t-2+I}) = N_B(0; p_{t-2+I})$, that is to say $\mu(p_{t-1+I} - p_{t-2+I}) + \mu(-p_{t-1+I}) = \mu(-p_{t-2+I})$. We may therefore rewrite the receiver's total news utility as $\sum_{i=1}^{I-1} \mu(p_{t-1+i} - p_{t-2+i}) + \mu(-p_{t-2+I})$. But by repeating this argument, we conclude that the receiver's total news utility is just $\mu(-p_{t-1})$. Since this result holds regardless of I's realization, the sender's expected total utility from sending g today is $\mu(-p_{t-1})$, which is the same as the news utility from sending b today. Thus, sender is indifferent between g and b and has no profitable deviation.

In state G, the sender gets at least $\mu(1 - p_{t-1})$ from following the equilibrium strategy. This is because the receiver's total news utility in the good state along the equilibrium path is given by $\sum_{i=1}^{J-(t-1)} \mu(p_{t-1+i} - p_{t-2+i}) + \mu(1 - p_{t-1+I})$. By sub-additivity in gains, this sum is strictly larger than $\mu(1 - p_{t-1})$. If the sender deviates to sending b today, then the receiver updates belief to 0 today and belief remains there until the exogenous revelation, when belief updates to 1. So this deviation gives the total news utility $\mu(-p_{t-1}) + \mu(1)$. We have

$$\begin{aligned} \mu(1) &< \mu(1 - p_{t-1}) + \mu(p_{t-1}) \\ &\leq \mu(1 - p_{t-1}) - \mu(-p_{t-1}), \end{aligned}$$

where the first inequality comes from sub-additivity in gains, and the second from weak loss aversion. This shows $\mu(-p_{t-1}) + \mu(1) < \mu(1 - p_{t-1})$, so the deviation is strictly worse than sending the equilibrium message.

Finally, at a history containing at least one b or a history with length K or longer, the receiver's belief is the same at all continuation histories. So the sender has no deviation incentives since no deviations affect future beliefs.

For the other direction, suppose by way of contradiction there exists a gradual good news equilibrium with the J intermediate beliefs $q^{(1)} < ... < q^{(J)}$. For a given $1 \leq j \leq J$, find the smallest t such that $p_t = q^{(k-1)}$ and $p_{t+1} = q^{(k)}$. At every on-path history $h^t \in H^t$ with $p^*(h^t) = p_t$, we must have $\sigma^*(h^t; B)$ inducing both 0 and $q^{(j)}$ with strictly positive probability. Since we are in equilibrium, we must have $\mu(-q^{(j-1)})$ being equal to $\mu(q^{(j)} - q^{(j-1)})$ plus the continuation payoff. If $j = J$, then this continuation payoff is $\mu(-q^{(j)})$ as the only other period of belief movement is in period T when the receiver learns the state is bad. If $j < J$, then find the smallest $\bar{t}$ so that $p_{\bar{t}+1} = q^{(j+1)}$. At any on-path $h^{\bar{t}} \in H^{\bar{t}}$ which is a continuation of h^t, we have $p^*(h^{\bar{t}}) = q^{(j)}$ and the receiver has not experienced any news utility in periods $t + 2, ..., \bar{t}$. Also, $\sigma^*(h^{\bar{t}}; B)$ assigns positive probability to inducing posterior belief 0, so the continuation payoff in question must be $\mu(-q^{(j)})$. So we have shown that $\mu(-q^{(j-1)}) = \mu(q^{(j)} - q^{(j-1)}) + \mu(-q^{(j)})$, that is $N_B(q^{(j)}; q^{(j-1)}) = N_B(0; q^{(j-1)})$. □

B.1.19 Proof of Corollary 5

Proof. We apply Proposition 22 to the case of quadratic. Recall the relevant indifference equation in the good state.

$$(!) \quad \mu(-q_t) = \mu(q_{t+1} - q_t) + \mu(-q_{t+1}).$$

Plugging in the quadratic specification and algebraic transformations lead to

$$0 = (\alpha_p - \alpha_n)(q_{t+1} - q_t) - \beta_p(q_{t+1} - q_t) + \beta_n(q_{t+1} - q_t)(q_{t+1} + q_t)$$

Define $r = q_{t+1} - q_t$. Then this relation can be written as

$$(\beta_p - \beta_n)r^2 + (\alpha_n - \alpha_p - 2\beta_n q_t)r = 0,$$

i.e. r is a zero of a second order polynomial. For P^* to be non-empty we need this root r to be in $(0, 1 - q_t)$. In particular the peak/trough $\bar{r}$ of the parabola defined by the second order polynomial should satisfy $\bar{r} \in (0, \frac{1-q_t}{2})$. Given that

$\bar{r} = \frac{2\beta_n q_t - (\alpha_n - \alpha_p)}{2(\beta_p - \beta_n)}$ for the case that $\beta_p \neq \beta_n$, we get the equivalent condition on the primitives

$$0 < \frac{2\beta_n q_t - (\alpha_n - \alpha_p)}{2(\beta_p - \beta_n)} < \frac{1 - q_t}{2}.$$

The root r itself is given by $r = \frac{2\beta_n q_t - (\alpha_n - \alpha_p)}{\beta_p - \beta_n}$, which leads to the recursion

$$(R) \quad q_{t+1} = q_t \frac{\beta_p + \beta_n}{\beta_p - \beta_n} - \frac{\alpha_n - \alpha_p}{\beta_p - \beta_n}.$$

This leads to the formula for $P^*(\pi)$ in part 1).

Case 1: When $\beta_p < \beta_n$ the coefficient in front of q_t is negative so that the recursion (R) leads to

$$(!!) \quad q_{t+1} - q_t = q_t \frac{2\beta_n}{\beta_p - \beta_n} - \frac{\alpha_n - \alpha_p}{\beta_p - \beta_n} < 0.$$

One also sees here that for the case that $\beta_p < \beta_n$ to give a gradual good news equilibrium of time-length 1, one needs a low enough prior: namely $\pi_0 < \frac{\alpha_n - \alpha_p}{2\beta_n} =: q^*$. For all priors larger or equal than q^*, there is no one-shot bad news partial good news equilibrium.

Case 2: When $\beta_p > \beta_n$ the slope in (R) is above 1 so that for all priors π_0 large enough we get an increasing sequence q_t which satisfies (!). It is also easy to see from (R) that

$$(q_{t+2} - q_{t+1}) - (q_{t+1} - q_t) = \left(\frac{\beta_p + \beta_n}{\beta_p - \beta_n} - 1 \right) > 0,$$

proving the statement in the text after the corollary.

That an equilibrium can exist where partial good news are released for more than two periods, is shown by the example in the main text following the statement of the Corollary (see Figure 2.4.2). □

B.1.20 Proof of Proposition 23

Proof. Since $N_B(p;\pi) - N_B(0;\pi) = 0$ for $p = \pi$ and $\frac{\partial}{\partial p} N_B(p;\pi)|_{p=\pi} > 0$, $N_B(p;\pi) - N_B(0;\pi)$ starts off positive for p slightly above π. Given that $|P^*(\pi)| \leq 1$, if we find some $p^{'} > \pi$ with $N_B(p^{'};\pi) - N_B(0;\pi) > 0$, then any solution to $N_B(p;\pi) - N_B(0;\pi) = 0$ in $(\pi, 0)$ must lie to the right of $p^{'}$.

If $q^{(j)}, q^{(j+1)}$ are intermediate beliefs in a GGN equilibrium, then by Proposition 22, $q^{(j)} \in P^*(q^{(j-1)})$ and $q^{(j+1)} \in P^*(q^{(j)})$. Let $p^{'} = q^{(j)} + (q^{(j)} - q^{(j-1)})$. Then,

$$\begin{aligned} N_B(p^{'};q^{(j)}) - N_B(0;q^{(j)}) &= \mu(p^{'} - q^{(j)}) + \mu(-p^{'}) - \mu(-q^{(j)}) \\ &= \mu(q^{(j)} - q^{(j-1)}) + \mu(-q^{(j)} - (q^{(j)} - q^{(j-1)})) - \mu(-q^{(j)}) \\ &> \mu(q^{(j)} - q^{(j-1)}) + \mu(-q^{(j-1)} - (q^{(j)} - q^{(j-1)})) - \mu(-q^{(j-1)}), \end{aligned}$$

where the last inequality comes from diminishing sensitivity. But, the final expression is $N_B(q^{(j)};q^{(j-1)}) - N_B(0;q^{(j-1)})$, which is 0 since $q^{(j)} \in P^*(q^{(j-1)})$. This shows we must have $q^{(j+1)} - q^{(j)} > q^{(j)} - q^{(j-1)}$. □

B.1.21 Proof of Corollary 6

Proof. We verify the sufficient condition in Proposition 23. We get $\frac{\partial}{\partial p} N_B(p;\pi) = \frac{\alpha}{(p-\pi)^{1-\alpha}} - \frac{\lambda\alpha}{p^{1-\alpha}}$, so $\frac{\partial}{\partial p} N_B(p;\pi)|_{p=\pi} = \infty$.

To show that $|P^*(\pi)| \leq 1$, it suffices to show that $\frac{\partial}{\partial p} N_B(p;\pi) = 0$ for at most one $p > \pi$. For the derivative to be zero, we need $(\frac{p}{p-\pi})^{1-\alpha} = \lambda$. As the LHS is decreasing for $p > \pi$, it can have at most one solution. □

B.1.22 Proof of Proposition 24

Proof. Consider the following operator φ on the space of continuous functions on $[0,1]$. For $V : [0,1] \to \mathbb{R}$, define $\varphi(V)(p) := \tilde{V}(p \mid p)$, where

$$\tilde{V}(\cdot \mid p) := \text{cav}_q[\mu(q-p) + \delta V(q) + (1-\delta)(q \cdot \mu(1-q) + (1-q) \cdot \mu(-q))].$$

We show that φ satisfies the Blackwell conditions and so is a contraction mapping.

Suppose that $V_2 \geq V_1$ pointwise. Then for any $p, q \in [0,1]$,

$$\mu(q-p) + \delta V_2(q) + (1-\delta)(q\mu(1-q) + (1-q)\mu(-q))$$
$$\geq (q-p) + \delta V_1(q) + ((1-\delta)(q\mu(1-q) + (1-q)\mu(-q))$$

therefore $\tilde{V}_2(\cdot \mid p) \geq \tilde{V}_1(\cdot \mid p)$ pointwise as well. In particular, $\tilde{V}_2(p \mid p) \geq \tilde{V}_1(p \mid p)$, that is $\varphi(V_2)(p) \geq \varphi(V_1)(p)$.

Also, let $k > 0$ be given and let $V_2 = V_1 + k$ pointwise. It is easy to see that $\tilde{V}_2(\cdot \mid p) = \tilde{V}_1(\cdot \mid p) + \delta k$ for every p, because the argument to the concavification operator will be pointwise higher by δk. So in particular, $\varphi(V_2)(p) = \varphi(V_1)(p) + \delta k$. By the Blackwell conditions, the operator φ is a contraction mapping on the metric space of continuous functions on $[0,1]$ with the supremum norm. Thus, the value function exists and is also unique.

To show pointwise monotonicity in δ, suppose $0 \leq \delta < \delta' < 1$. First, $V_\delta(0) = V_\delta(1) = 0$ for any $\delta \in [0,1)$. Now consider an environment where full revelation happens at the end of each period with probability $1-\delta$, and fix a prior $p \in (0,1)$. There exists some binary information structure with message space $M = \{0,1\}$, public histories $H^t = (M)^t$ for $t = 0, 1, \ldots$, and sender strategies $(\sigma_t)_{t=0}^{\infty}$ with $\sigma_t : H^t \times \Theta \to \Delta(M)$, such that (M, σ) induces expected news utility of $V_\delta(p)$ when starting at prior p.

We now construct a new information structure, $(\bar{M}, \bar{\sigma})$ to achieve expected news utility $V_\delta(p)$ when starting at prior p in an environment where full revelation happens at the end of each period with probability $1-\delta'$, with $\delta' > \delta$. Let $\bar{M} = \{0, 1, \varnothing\}$. The idea is that when full revelation has not happened, there is a $1 - \frac{\delta}{\delta'}$ probability each period that the sender enters into a babbling regime forever. When the sender enters the babbling regime at the start of period $t+1$, the receiver's expected utility going forward is the same as if full revelation happened at the start of $t+1$.

To implement this idea, after any history $h^t \in H^t$ not containing $\varnothing$, let

$$\bar{\sigma}_{t+1}(h^t;\theta) = \begin{cases} \varnothing & \text{w/p } 1 - \frac{\delta}{\delta'} \\ 1 & \text{w/p } \frac{\delta}{\delta'} \cdot \sigma_{t+1}(h^t;\theta)(1) \\ 0 & \text{w/p } \frac{\delta}{\delta'} \cdot \sigma_{t+1}(h^t;\theta)0) \end{cases}.$$

That is, conditional on not entering the babbling regime, $\bar{\sigma}$ behaves in the same way as σ. But, after any history $h^t \in H^t$ containing at least one $\varnothing$, $\bar{\sigma}_{t+1}(h^t;\theta) = \varnothing$ with probability 1. Once the sender enters the babbling regime, she babbles forever (until full revelation exogenously arrives at some random date). We need to verify that payoff from this strategy is indeed $V_\delta(p)$. Fix a history h^t not containing $\varnothing$ and a state θ, and suppose $p^*(h^t) = q$. Under $\bar{\sigma}_{t+1}$, with probability of $(1-\delta') + \delta'(1-\frac{\delta}{\delta'}) = 1-\delta$ the receiver gets the expected babbling payoff $q\mu(1-q) + (1-q)\mu(-q)$ in the period of state revelation. Analogously, under σ_{t+1}, there is probability $1-\delta$ that state revelation happens in period $t+1$ and the receiver gets $q\mu(1-q) + (1-q)\mu(-q)$ in expectation. With probability $\delta' \frac{\delta}{\delta'} = \delta$, the receiver facing $\bar{\sigma}$ gets the payoff induced by $\sigma_{t+1}(h^t;\theta)$ in period $t+1$ and the same distribution of continuation histories as under σ. The same argument applies to all these continuation histories, so $\bar{\sigma}$ must induce the same expected payoff as σ when starting at $(h^t;\theta)$. □

B.1.23 Proof of Proposition 25

Proof. We show first sufficiency. Consider the following strategy profile. In period t where the public history so far h^{t-1} does not contain any b, let $\sigma(h^{t-1};G)(g) = 1$, $\sigma(h^{t-1};B)(g) = x$ where $x \in (0,1)$ satisfies $\frac{p_{t-1}}{p_{t-1}+(1-p_{t-1})x} = p_t$. But if public history contains at least one b, then $\sigma(h^{t-1};G)(b) = 1$ and $\sigma(h^{t-1};B)(b) = 1$. In terms of beliefs, suppose h^t is so that every message so far has been g. Such histories are on-path and get assigned the Bayesian posterior belief. If h^t contains at least one b, then belief is 0. It is easy to verify that these beliefs are derived from Bayes' rule whenever possible.

We verify that the sender has no incentive to deviate. Consider period t with history h^{t-1} that does not contain any b. The receiver's current belief is p_{t-1} by construction.

In state B, we first calculate the sender's equilibrium payoff after sending g. For any realization of the exogenous revelation date, the receiver's total news utility in the good state along the equilibrium path is given by $\sum_{j=1}^{J} \mu(p_{t-1+j} - p_{t-2+j}) + \mu(-p_{t-1+J})$ for some integer $J \geq 1$. Since $p_{t-1+J} \in P^*(p_{t-2+J})$, we have $N_B(p_{t-1+J}; p_{t-2+J}) = N_B(0; p_{t-2+J})$, that is to say $\mu(p_{t-1+J} - p_{t-2+J}) + \mu(-p_{t-1+J}) = \mu(-p_{t-2+J})$. We may therefore rewrite the receiver's total news utility as $\sum_{j=1}^{J-1} \mu(p_{t-1+j} - p_{t-2+j}) + \mu(-p_{t-2+J})$. But by repeating this argument, we conclude that the receiver's total news utility is just $\mu(-p_{t-1})$. Since this result holds regardless of J, the sender's expected total utility from sending g today is $\mu(-p_{t-1})$, which is the same as the news utility from sending b today. Thus, sender is indifferent between g and b and has no profitable deviation.

In state G, the sender gets at least $\mu(1 - p_{t-1})$ from following the equilibrium strategy. This is because for any realization of the exogenous revelation date, the receiver's total news utility in the good state along the equilibrium path is given by $\sum_{j=1}^{J} \mu(p_{t-1+j} - p_{t-2+j}) + \mu(1 - p_{t-1+J})$ for some integer $J \geq 1$. By sub-additivity in gains, this sum is strictly larger than $\mu(1 - p_{t-1})$. If the sender deviates to sending b today, then the receiver updates belief to 0 today and belief remains there until the exogenous revelation, when belief updates to 1. So this deviation has the total news utility $\mu(-p_{t-1}) + \mu(1)$. We have

$$\begin{aligned}\mu(1) &< \mu(1 - p_{t-1}) + \mu(p_{t-1}) \\ &\leq \mu(1 - p_{t-1}) - \mu(-p_{t-1}),\end{aligned}$$

where the first inequality comes from sub-additivity in gains, and the second from weak loss aversion. This shows $\mu(-p_{t-1}) + \mu(1) < \mu(1 - p_{t-1})$, so the deviation is strictly worse than sending the equilibrium message.

Finally, at a history containing at least one b, the receiver's belief is the same at all continuation histories. So the sender has no deviation incentives since no deviations affect future beliefs.

We now show necessity. Suppose that we have a (possibly infinite) gradual good news equilibrium given by the sequence $p_0 < p_1 < \cdots < p_t < \ldots$. By Bayesian plausibility and because we are focusing on two-message equilibria the sender must be sending the messages $\{0, p_t\}$ in period t if the state is bad. The sender must thus be indifferent between these two posteriors in the bad state. Formally, $N_B(0; p_t) = N_B(p_{t+1}; p_t)$ for all $t \geq 0$, as long as there is no babbling. Written equivalently in the language of P^*: $p_{t+1} \in P^*(p_t)$ for all $t \geq 0$, as long as there's no babbling, where here $p_0 = \pi_0$. □

B.2 Residual Consumption Uncertainty

B.2.1 A Model of Residual Consumption Uncertainty

In the main text, we studied a model where the sender has perfect information about the receiver's final-period consumption level.

Now suppose the sender's information is imperfect. In state θ, the receiver will consume a random amount c in period $T+1$, drawn as $c \sim F_\theta$, deriving from it consumption utility $v(c)$. As before, v is a strictly increasing consumption-utility function. We interpret the state θ as the sender's private information about the receiver's future consumption, while the distribution F_θ captures the receiver's residual consumption uncertainty conditional on what the sender knows. The case where F_θ is degenerate for every $\theta \in \Theta$ nests the baseline model.

Assume that $\mathbb{E}_{c \in F_{\theta'}}[v(c)] \neq \mathbb{E}_{c \in F_{\theta''}}[v(c)]$ when $\theta' \neq \theta''$. We may without loss normalize $\min_{\theta \in \Theta} \mathbb{E}_{c \in F_\theta}[v(c)] = 0$, $\max_{\theta \in \Theta} \mathbb{E}_{c \in F_\theta}[v(c)] = 1$.

The mean-based news-utility function $N(\pi_t \mid \pi_{t-1})$ in this environment is the same as in the environment where the receiver always gets consumption utility $\mathbb{E}_{c \sim F_\theta}[v(c)]$ in state θ. This is because given a pair of beliefs $F_{\text{old}}, F_{\text{new}} \in \Delta(\Theta)$ about the state, the receiver derives news utility $N(F_{\text{new}} \mid F_{\text{old}})$ based on the difference in *expected* consumption utilities, $\mu(\mathbb{E}_{c \sim F_{\text{new}}}[v(c)] - \mathbb{E}_{c \sim F_{\text{old}}}[v(c)])$. So, all of the results in the paper concerning mean-based news utility immediately extend. The two results in the paper that are not specific to mean-based news utility, Propositions 11 and 12, apply to *any* functions $N(\pi_t \mid \pi_{t-1})$ satisfying the continuous differentiability condition stated in Section 2.2, without requiring any relationship between N and consumptions in different states.

We now define N using Kőszegi and Rabin [2009]'s percentile-based news-utility model with a power-function gain-loss utility, in an environment with residual consumption uncertainty. We apply Proposition 12 to the resulting N and show that one-shot resolution is strictly sub-optimal. This result applies for any $K \geq 2$.

Corollary 11. *Consider the percentile-based model with* $\mu(x) = \begin{cases} x^\alpha & x \geq 0 \\ -\lambda(-x)^\alpha & x < 0 \end{cases}$ *for* $0 < \alpha < 1$, $\lambda \geq 1$. *Suppose there are two states* $\theta_G, \theta_B \in \Theta$ *with distributions of consumption utilities* $v(F_{\theta_B}) = Unif[0, L]$, $v(F_{\theta_G}) = J + v(F_{\theta_B})$ *for some* $L, J > 0$. *One-shot resolution is strictly suboptimal for any finite* T.

Proof. We show that $\lim_{\varepsilon \to 0} \frac{N(1_G \mid (1-\varepsilon)1_G \oplus \varepsilon 1_B)}{\varepsilon} = \infty$ under this set of conditions. The argument behind Proposition 12 then implies some information structure involving perfect revelation of states other than θ_G, θ_B, one-shot bad news, partial good news for the two states θ_G, θ_B is strictly better than one-shot resolution.

For $r \in [0, 1]$, write F_r for the distribution of consumption utilities under the belief $r1_G \oplus (1 - r)1_B$.

Note we must have $\int_0^1 c_{F_1}(q) - c_{F_{1-\varepsilon}}(q)dq = J\varepsilon$, and that $c_{F_1}(q) - c_{F_{1-\varepsilon}}(q) \geq 0$ for all q.

Let $q^* = \min(\varepsilon \cdot J/L, \varepsilon)$. It is the quantile at which $c_{F_{1-\varepsilon}}(q^*) = J$.

For all $q \geq q^*$, $c_{F_1}(q) - c_{F_{1-\varepsilon}}(q) \leq \varepsilon L$.

Case 1: $J \geq L$, so $q^* = \varepsilon$.

$$\int_0^{q^*} c_{F_1}(q) - c_{F_{1-\varepsilon}}(q)dq = \int_0^{\varepsilon} J - q \cdot \frac{1}{\varepsilon} \cdot ((1-\varepsilon)L)dq$$
$$= J\varepsilon - \frac{1}{2}\varepsilon(1-\varepsilon)L.$$

This implies $\int_{q^*}^1 c_{F_1}(q) - c_{F_{1-\varepsilon}}(q)dq = \frac{1}{2}\varepsilon(1-\varepsilon)L$.

The worst case is when the difference is εL on some q-interval, and 0 elsewhere. For small $\varepsilon < 0$ so that $\varepsilon L < 1$,

$$\int_{q^*}^1 (c_{F_1}(q) - c_{F_{1-\varepsilon}}(q))^{\alpha} dq \geq (\varepsilon L)^{\alpha} \cdot \frac{(1/2) \cdot \varepsilon(1-\varepsilon)L}{\varepsilon L}$$
$$= \frac{1}{2}(\varepsilon L)^{\alpha}(1-\varepsilon).$$

Therefore, for small $\varepsilon > 0$, $\frac{N(1_G|(1-\varepsilon)1_G \oplus \varepsilon 1_B)}{\varepsilon} = \frac{1}{2}\frac{1}{\varepsilon^{1-\alpha}}L^{\alpha}(1-\varepsilon)$, which diverges to ∞ as $\varepsilon \to 0$.

Case 2: $J < L$, so $q^* = \varepsilon J/L$.

$$\int_0^{\varepsilon J/L} c_{F_1}(q) - c_{F_{1-\varepsilon}}(q)dq = \int_0^{\varepsilon J/L} J - q \cdot \frac{1}{\varepsilon J/L}(J - \frac{J}{L}\varepsilon \cdot L)dq$$
$$= \frac{1}{2}\frac{J^2}{L}\varepsilon + \frac{1}{2}\frac{J^2}{L^2}\varepsilon^2 L$$
$$< \frac{1}{2}J\varepsilon + \frac{1}{2}L\varepsilon^2$$

using $J < L$. This then implies $\int_{q^*}^1 c_{F_1}(q) - c_{F_{1-\varepsilon}}(q)dq > \frac{1}{2}J\varepsilon - \frac{1}{2}L\varepsilon^2$.

So, again using the worst-case of the difference being εL on some q-interval, and 0 elsewhere,

$$\frac{N(1_G \mid (1-\varepsilon)1_G \oplus \varepsilon 1_B)}{\varepsilon} > \frac{1}{\varepsilon}(\varepsilon L)^{\alpha} \cdot \frac{\frac{1}{2}J\varepsilon - \frac{1}{2}L\varepsilon^2}{\varepsilon L}$$
$$= \frac{1}{\varepsilon^{1-\alpha}}L^{\alpha} \cdot \left(\frac{1}{2}J/L - \frac{1}{2}\varepsilon\right).$$

As $\varepsilon \to 0$, RHS converges to ∞. □

B.2.2 A Calibration Comparing Percentile-Based News Utility and Mean-Based News Utility

Since Proposition 11's procedure for computing the optimal information structure applies to general N, including both the percentile-based and the mean-based news-utility functions in an environment with residual consumption uncertainty, we can compare the solutions to the sender's problem for these two models.

Consider two states of the world, $\Theta = \{G, B\}$. For some $\sigma > 0$, suppose consumption is distributed normally conditional on θ with $F_G = \mathcal{N}(1, \sigma^2)$, $F_B = \mathcal{N}(0, \sigma^2)$, consumption utility is $v(x) = x$, and gain-loss utility (over consumption) is $\mu(x) = \sqrt{x}$ for $x \geq 0$, $\mu(x) = -1.5\sqrt{-x}$ for $x < 0$. We calculated the optimal information structure for the mean-based model in an analogous environment, as reported in Figure 2.3.1.

With the percentile-based model, an agent who believes $\mathbb{P}[\theta = G] = \pi$ has a belief over final consumption given by a mixture normal distribution, $\pi F_G \oplus (1 - \pi) F_B$, illustrated in Figure B.2.1.

We plot in Figure B.2.2 the optimal information structures for $T = 5$, $\sigma = 1$. The optimal information structures for $\sigma = 0.1, 1, 10$ all involve gradual good news, one-shot bad news. Table B.2.1 lists the optimal disclosure of good news over time. Not only are the shapes of the concavification problems qualitatively similar to those of the mean-based model, but the resulting optimal information structures also bear striking quantitative similarities.

From Table B.2.1, it appears that percentile-based and mean-based models deliver more similar results for larger σ^2. We provide an analytic result consistent with the idea that these two models generate similar amounts of news utility when the state-dependent consumption utility distributions have large variances.

Proposition 43. *Suppose $\Theta = \{B, G\}$ and the distributions of consumption utilities in states B and G are Unif$[0, L]$ and Unif$[d, L + d]$ respectively, for $L, d > 0$. Let*

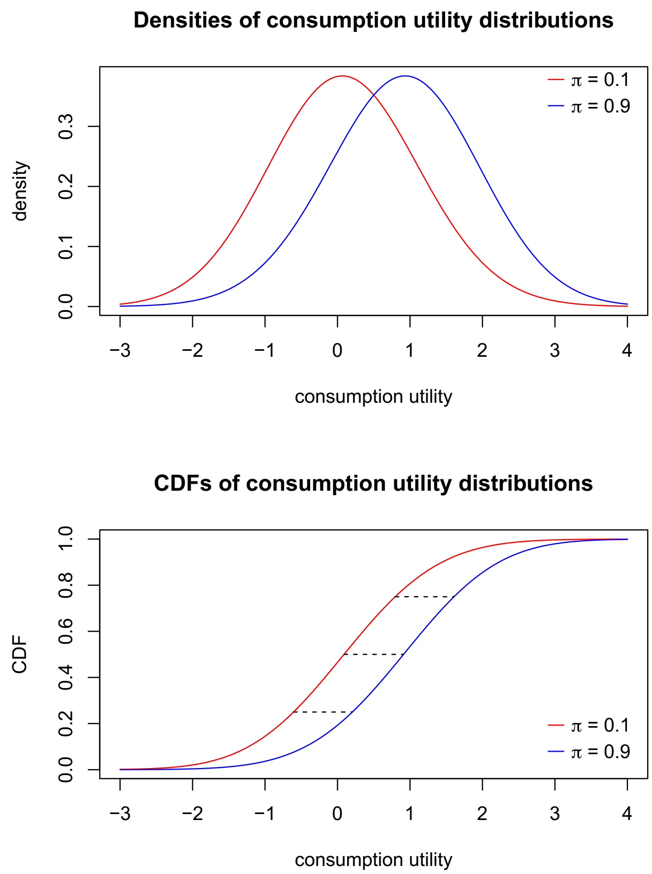

Figure B.2.1: The densities and CDFs of final consumption utility distributions under two different beliefs.

$\mathbb{P}[\theta = G]$, $\pi = 0.1$ and $\pi = 0.9$. The dashed black lines in the CDFs plot show the differences in consumption utilities at the 25th percentile, 50th percentile, and 75th percentile levels between these two beliefs. The news utility associated with updating belief from $\pi = 0.1$ to $\pi = 0.9$ in the percentile-based model is calculated by applying a gain-loss function μ to all these differences in consumption utilities at various quantiles, then integrating over all quantiles levels in $[0, 1]$.

$N^{perc}(p_2 \mid p_1)$ *be the news utility associated with changing belief in* $\theta = G$ *from* p_1 *to* p_2

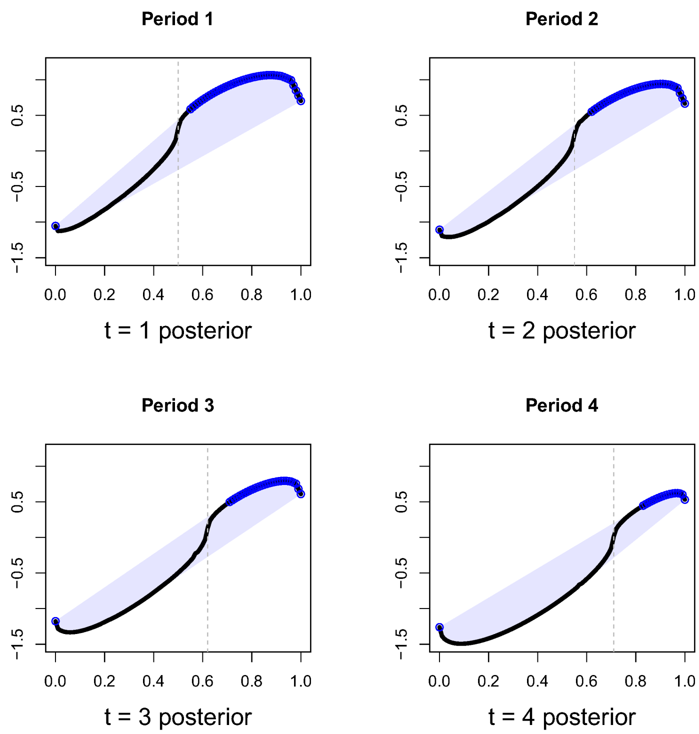

Figure B.2.2: The concavifications giving the optimal information structure with Kőszegi and Rabin [2009]'s percentile-based

More precisely, we consider a Gaussian environment. horizon $T = 5$, gain-loss function $\mu(x) = \begin{cases} \sqrt{x} & \text{for } x \geq 0 \\ -1.5\sqrt{-x} & \text{for } x < 0 \end{cases}$, prior $\pi_0 = 0.5$, and $\sigma = 1$. The y-axis in each graph shows the sum of news utility this period and the value function of entering next period with a certain belief.

in a percentile-based news-utility model with a continuous gain-loss utility μ. Then,

$$\lim_{L\to\infty}\left(\sup_{0\leq p_1,p_2\leq 1}|N^{perc}(p_2 \mid p_1) - \mu[(p_2 - p_1)d]|\right) = 0.$$

In a uniform environment, if there is enough unresolved consumption risk even conditional on the state θ, then the difference between percentile-based news utility and mean-based news utility goes to zero uniformly across all

	$t=0$	$t=1$	$t=2$	$t=3$	$t=4$	$t=5$
percentile-based, $\sigma=0.1$	0.50	0.55	0.61	0.69	0.80	1.00
percentile-based, $\sigma=1$	0.50	0.55	0.62	0.71	0.83	1.00
percentile-based, $\sigma=10$	0.50	0.56	0.63	0.72	0.84	1.00
mean-based, any σ	0.500	0.556	0.626	0.715	0.834	1.000

Table B.2.1: Optimal disclosure of good news. The optimal information structure under a square-root gain-loss function with $\lambda = 1.5$ takes the form of gradual good news, one-shot bad news both in the mean-based model and the percentile-based model for $T = 5$, $\sigma = 0.1, 1, 10$. The table shows belief movements conditional on the good state in different periods.

possible belief changes.[1]

Proof. Let $F_p(x)$ be the distribution function of the mixed distribution $p \cdot \text{Unif}[d, L+d] \oplus (1-p) \cdot \text{Unif}[0, L]$, and $F_p^{-1}(q)$ its quantile function for $q \in [0,1]$. By a simple calculation, $F_p^{-1}(d/L) = d + pd$ and $F_p^{-1}(1-d/L) = L + pd - d$. At the same time, for $d/L \leq q \leq 1 - d/L$ where $q = d/L + y$, we have $F_p^{-1}(q) = d + pd + yL$.

This shows that over the intermediate quantile values between d/L and $1-d/L$,

$$\int_{d/L}^{1-d/L} \mu\left[F_{p_2}^{-1}(q) - F_{p_1}^{-1}(q)\right] dq = \int_{d/L}^{1-d/L} \mu\left[(p_2 - p_1)d\right] dq = (1-2d/L)\cdot\mu[(p_2-p_1)d].$$

For the lower part of the quantile integral $[0, d/L]$, using the fact that $F_p^{-1}(d/L) = d + pd$, we have the uniform bound $0 \leq F_p^{-1}(q) \leq 2d$ for all $p \in [0,1]$ and $q \leq d/L$. So,

$$\left|\int_0^{d/L} \mu\left[F_{p_2}^{-1}(q) - F_{p_1}^{-1}(q)\right] dq\right| \leq \frac{d}{L} \cdot \max_{x\in[-2d,2d]} |\mu(x)|.$$

[1]Lemma 3 in the Online Appendix of Kőszegi and Rabin [2009] states a similar result, but for a different order of limits.

By an analogous argument,

$$\left| \int_{1-d/L}^{1} \mu \left[F_{p_2}^{-1}(q) - F_{p_1}^{-1}(q) \right] dq \right| \leq \frac{d}{L} \cdot \max_{x \in [-2d, 2d]} |\mu(x)|.$$

So for any $0 \leq p_1, p_2 \leq 1$,

$$|N^{\text{perc}}(p_2 \mid p_1) - \mu[(p_2 - p_1)d]| \leq \frac{2d}{L} \max_{x \in [d,d]} |\mu(x)| + \frac{2d}{L} \max_{x \in [-2d,2d]} |\mu(x)|,$$

an expression not depending on p_1, p_2. The max terms are seen to be finite by applying extreme value theorem to the continuous μ, so the RHS tends to 0 as $L \to \infty$. □

C

Appendix to Chapter 3

C.1 Properties of the value function of the agent

We note here properties of the value functions $A \to \tilde{W}_0(A, \pi_0)$. Most of the proofs are routine checking, except for the second part of 5) which is established by an example.

Lemma 24. *(Axioms from De Oliveira et al. [2017]) Fix the sequential experiment* $\mathcal{E}$.

For both cases SeSa-GD and SeSa-LC the map
$\tilde{W}_0 : \mathcal{A} \times \Delta(S) \to \mathbb{R}, \quad (A, \pi) \to \tilde{W}_0(A, \pi)$ *mapping a menu A and a prior* π *to the ex-ante value of the decision problem has the following properties.*

1) it is continuous,[1]

[1]Here we equip $\mathcal{A} \times \Delta(\Delta(S))$ with the product topology of the Hausdorff topology on $\mathcal{A}$ and the weak convergence topology on $\Delta(\Delta(S))$.

2) it is convex in the first argument,

3) it is monotonic with respect to set inclusion in the first argument,

4) it satisfies Dominance with respect to the taste u:

(Dominance) for every information structure π, menu A and act g such that there exists $f \in A$ with $u(f(s)) \geq u(g(s))$ for all $s \in S$ it holds
$\tilde{W}_0(A, \pi) = \tilde{W}_0(A \cup \{g\}, \pi)$,

5) in case B. it satisfies Independence of Degenerate Decisions:

(IDD) for every fixed information structure π, pair of acts $h, h' \in \mathbb{F}$ and menus $A, B \in \mathcal{A}$ it holds

$$\tilde{W}_0(\lambda A + (1 - \lambda)\{h\}, \pi) \geq \tilde{W}_0(\lambda B + (1 - \lambda)\{h\}, \pi)$$
$$\Longleftrightarrow$$
$$\tilde{W}_0(\lambda A + (1 - \lambda)\{h'\}, \pi) \geq \tilde{W}_0(\lambda B + (1 - \lambda)\{h'\}, \pi).$$

In general IDD is not satisfied for the case SeSa-GD.

Proposition 24 implies that the SeSa-LC model satisfies the axioms in De Oliveira et al. [2017] and so is a rational inattention model in their terminology. In a general rational inattention model the costs of information enter the overall utility of the agent in an additive separable way. The following example shows that IDD may be violated under SeSa-GD.

Example 1 (Example 1-A: IDD may be violated for SeSa-GD)**.** *Take $Z = \mathbb{R}_+$ and an agent with taste $u(z) = z$ and state space $S = \{s_1, s_2\}$ as well as uniform prior $\pi_0 = (\frac{1}{2}, \frac{1}{2})$. If the agent decides to acquire information she learns the true state of the world. Finally, let $\delta \in (\frac{1}{2}, 1)$. Consider the menus $A = \{(1, 0), (0, 1)\}$ and $B = \{(1 - \varepsilon, 1 - \varepsilon)\}$ with $\varepsilon \in (0, \frac{1}{2})$. Thus menu B offers full insurance to the agent but at a cost, whereas menu A exposes her to full uncertainty. Since the agent is relatively patient she will decide to learn for menu A. She will never learn for menu B. Now consider for $a \in (0, 1)$ the menus $aA = \{(a, 0), (0, a)\}$ and $aB = \{(a(1 - \varepsilon), a(1 - \varepsilon))\}$. These are just rescalings of the menus A and B.*

One can easily show that

i. $\tilde{W}_0(A, \pi) > \tilde{W}_0(B, \pi)$,

ii. $\tilde{W}_0(\alpha A, \pi) > \tilde{W}_0(\alpha B, \pi)$ *(mixing with* $\{\mathbf{0}\}$*)*,

iii. $\tilde{W}_0(\alpha A + (1-\alpha)\mathbf{1}, \pi) < \tilde{W}_0(\alpha B + (1-\alpha)\mathbf{1}, \pi)$ *(mixing with* $\{\mathbf{1}\}$*)*, *whenever* α *satisfies* $\frac{\delta - \frac{1}{2}}{1-\delta} > \frac{1-\alpha}{\alpha} > \frac{\delta-(1-\varepsilon)}{1-\delta}$.

Thus IDD is violated in the case of impatience costs.

Remark 13. *Part 2) of Lemma 24 holds true under the model assumptions for every value function* $A \mapsto \tilde{W}_t(A, e^t)$ *for all* e^t *and* $t \in \{1, \ldots, T\}$*. This is proven as in Lemma 24 through induction over* $T - t$ *(induction start is Lemma 24 with* $T = 0$*).*

C.2 Continuity axiom and menus without ties

We start with a Lemma establishing that the collection of menus where the agent has to break ties is 'small' relative to the collection of all menus $\mathcal{A}$.

Lemma 25. *Suppose the model is true, i.e. either of the cases SeSa-GD or SeSa-LC are true. Then the set of menus* $\hat{\mathcal{A}} \subset \mathcal{A}$ *such that the agent has no ties w.r.t. either stopping decision or choice from menu upon stopping is dense in the set* $\mathcal{A}$.

Proof. We use utility acts for the proof. If the SeSa models are true, which is the premise of the Lemma, this is w.l.o.g. Recall that we are assuming a worst prize $w \in Z$ whose utility value we take to be 0.

There are two sorts of ties in the model. First, there may be ties due to choice out of menu upon stopping. Since the tree of posteriors of the agent is finite there are only finitely many possible SEUs at which the agent can stop. The set of menus out of $\mathcal{A}$ for which each SEU exhibits ties is nowhere dense in $\mathcal{A}$ and also closed and convex. Taking the union of all these sets we get a nowhere dense set, because the union of finitely many closed, nowhere dense sets is nowhere dense and closed. Therefore, the collection of menus for which the agent has ties on choice out of menus upon stopping is a subset of a nowhere dense set and thus itself a nowhere dense set. Denote this collection of menus by $\mathcal{A}^{n,s}$. Note that this argument does not depend on whether we are in case SeSa-GD or SeSa-LC.

The second kind of ties are ties coming from the stopping decision, i.e. menus where the agent at some node corresponding to a history of experimental outcomes (and a related history of posteriors) is indifferent between continuing and stopping.

Case A - discounting.

We show that for any menu A and any open neighborhood $\mathcal{N}_A$ of A in the Hausdorff topology of $\mathcal{A}$ such that A has ties about continuing or stopping in a period $t \in \{0, \dots, T-1\}$ there exists a menu A' in the neighborhood such that the agent strictly wants to stop in every node where she was indifferent before. Moreover, A' will be such that otherwise, for the other history nodes the agent will take the same decision continue/stop decision as for A. Fix in the following a deterministic prize $z_0 > 0$ and assume that e^t is a first (in terms of time t) node where the agent is indifferent between stopping and continuing.

This means that for all shorter histories e^s, $s < t$ in previous periods as well as for nodes $e'^t \in E_1 \times \cdots \times E_t$, where the agent has strict incentives for the decision to stop/continue we have

$$V(A, \pi(e^s)) \neq \delta W_{s+1}(A, e^s). \tag{C.1}$$

Now due to continuity of the value functions it holds that the relation (C.1) holds true in an open neighborhood $\mathcal{N}_A$ of A (in the Hausdorff topology on $\mathcal{A}$) for all nodes e^s of depth $s \leq t$ where the agent has strict incentives for the continue/stop decision with the same inequality direction as for A. This uses continuity of the value functions $V(A, \pi(e^s))$, $W_{s+1}(A, e^s)$ in the menu A.

Now note that it holds for $B = \{z_0\}$ that $V(B, \pi(e^t)) = W(B, e^t)$. Moreover, due to convexity of $W(\cdot, e^t)$ and the fact that $\delta < 1$ it also holds for every $\lambda \in (0, 1)$ that

$$V(\lambda A + (1-\lambda)B, \pi(e^t)) = \lambda V(A, \pi(e^t)) + (1-\lambda)V(B, \pi(e^t)) \tag{C.2}$$
$$> \delta \left[\lambda W_{t+1}(A, e^t) + (1-\lambda)W_{t+1}(B, e^t)\right] \geq \delta W_{t+1}(\lambda A + (1-\lambda)B, e^t).$$

The first (strict) inequality is due to $\delta < 1$ and the fact that $V(A, \pi(e^t)) = W_{t+1}(A, e^t)$ by assumption, whereas the second inequality follows from convexity of $W_{t+1}(\cdot, e^t)$.

Now pick an $\lambda(e^t) \in (0, 1)$ s.t. for all $\lambda \in (\lambda(e^t), 1)$ it holds $\lambda A + (1-\lambda)B \in \mathcal{N}_A$. If there are other $e'^t \in E_1 \times \cdots \times E_t$ with ties about the continuation decision then pick $\lambda_t := \max \lambda(e'^t) \in (0, 1)$ where the max runs over these history nodes of depth t that have ties in the continue/stop decision.

In addition, pick $\lambda_{<t} \in (0, 1)$ such that for all $\lambda \in (\lambda_{<t}, 1)$ it holds $\lambda A + (1-\lambda)B \in \mathcal{N}_A$ and also that $\lambda A + (1-\lambda)B$ has no ties in the continuation/stopping decision at *all* history nodes of depth $s < t$. Recall, that by minimality choice of t at the start this is always possible. Finally, set $\hat{\lambda} = \max\{\lambda_t, \lambda_{<t}\}$.

We have then that for all $\lambda \in (\hat{\lambda}, 1)$ the continuation/stopping decision at times $s \leq t$ is always strict for menus of the type $\lambda A + (1-\lambda)B$ and that all these menus are in $\mathcal{N}_A$. Now pick such a mixed menu and denote it by $A(\lambda)$. If for $A(\lambda)$ the agent becomes indifferent at some future period $t' > t$ then let $t' > t$ be the earliest such period and repeat the procedure above for $A(\lambda)$ with the same B to arrive at some new $A(\lambda')$ with some $\lambda' \in (\lambda, 1)$. One goes on like this inductively till the last period is reached. The process ends because of the finiteness of the tree of posteriors.

In this way one constructs a sequence of menus $A_n \rightarrow A$ such that all of the A_n have no ties regarding stopping or continuation decisions at any node.[2] We also note here that the sequence A_n we construct in the proof has at most $|A|$ elements.[3]

[2] Recall that $\mathcal{A}$ is a metric space with the Hausdorff topology.

[3] This is important for the discussion right after the proof.

Denote $\mathcal{A}^{n,c}$ the set of menus where the agent has ties w.r.t. the continuing/stopping decision. We just showed that the set $\mathcal{A}^{n,c}$ is nowhere dense in $\mathcal{A}$. It follows that the set $\mathcal{A}^{n,c} \cup \mathcal{A}^{n,s}$ is also nowhere dense. Take $\hat{\mathcal{A}} := \mathcal{A} \setminus (\mathcal{A}^{n,s} \cup \mathcal{A}^{n,c})$ to conclude.

Case B - additive costs.

This is analogous to case A. with the only change being that instead of (C.2) we use the chain of inequalities

$$V(\lambda A + (1-\lambda)B, \pi(e^t)) = \lambda V(A, \pi(e^t)) + (1-\lambda)V(B, \pi(e^t))$$
$$> [\lambda(W_{t+1}(A, e^t) - c) + (1-\lambda)(W_{t+1}(B, e^t) - c)] \geq W_{t+1}(\lambda A + (1-\lambda)B, e^t) - c.$$

Note that here we have used $c > 0$.

□

Recovering $\hat{\mathcal{A}}$ from a RCDT.

Definition 11. Say that a menu $A \in \mathcal{A}$ *is without ties* if there exists a neighborhood $\mathcal{N}_A$ of A in $\mathcal{A}$ such that for every sequence $A_n \in \mathcal{N}_A$ with $A_n \to A$ and $|A_n| = |A|$ it holds uniformly for $t \in \{0, \dots, T\}$ that $P_{A_n}(f_n, t) \to P_A(f, t)$ if $A_n \ni f_n \to f \in A$.

Say that a menu *has ties* if it is not a menu without ties.

Say that a menu has no ties with respect to the stopping decision if and only if there exists a neighborhood $\mathcal{N}_A$ of A in $\mathcal{A}$ such that for every sequence $A_n \in \mathcal{N}_A$ with $A_n \to A$ it holds uniformly for $t \in \{0, \dots, T\}$ that $P_{A_n}(\tau = t) \to P_A(\tau = t)$.

Imposing the following Axiom on RCDT data ensures that we can focus on menus without ties throughout the analysis of the model.

Axiom: Continuity The set of menus without ties as defined in Definition 11 is dense in $\mathcal{A}$.

If the model is correct then the set of menus without ties corresponds to $\hat{\mathcal{A}}$. To see this, note that if $A \in \hat{\mathcal{A}}$ then it is trivial to show that A has no ties according to

Definition 11, given the finiteness of the model.

On the other hand, if $A \in \mathcal{A}^{n,c}$ then there exists a first time t where a tie in the stop/continue decision appears. Take then a perturbation A_n of A as in the proof of Lemma 25. It follows that $P_{A_n}(\tau = t) \not\rightarrow P_A(\tau = t)$ even though $A_n \rightarrow A$ and $|A_n| = |A|$ for all large enough n. A cannot have no ties in the sense of Definition 11 then, because if it had no ties there would be convergence of the marginal of decision time for the menus A_n as well.

Let on the other hand $A \in \mathcal{A}^{n,s} \setminus \mathcal{A}^{n,c}$. Let t be a first time where a tie in choice appears. In particular the agent stops with positive probability in period t. Looking at $P_A(\cdot|\tau = t)$, this is a SCF as in Duraj [2018a] or also full observability SCF in the terminology of Duraj and Lin [2019]).[4] In particular, there is a tie in choice at time t if and only if there exists a $f \in A$ and a sequence $f_n \rightarrow^m f$ as well as $B_n \rightarrow^m A \setminus \{f\}$[5] with $\lim_{n\rightarrow\infty} P_{B_n \cup \{f_n\}}(f_n|\tau = t) \neq P_A(f|\tau = t)$. But because A has no ties w.r.t. decision time it follows that A will have ties according to Definition 11 since it will violate the definition of no-ties there at the tuple (A, f, t) for the sequence $B_n \cup \{f_n\}$ which again has at most $|A|$ elements (and ultimately also exactly $|A|$ elements).

C.3 Proofs for sections 3.3, 3.4 and 3.6

Proof of Theorem 4. We work with utility acts throughout. We look first at the discounting case. Note that a linear combination of absolutely continuous functions is absolutely continuous. In particular, even more is true: any linear combination of Lipschitz functions is again Lipschitz. Recall also that a Lipschitz function is absolutely continuous. Let $G \in \Delta\left(\Delta(S)\right) \times \mathcal{T})$, write $G_t \in \Delta(\Delta(S))$ for the conditional distribution $G(\cdot \subset \Delta(S)|\tau = t)$ on stopping at time t[6] and

[4] The concept of SCF for menus of Anscombe-Aumann acts first appears in Lu [2016].

[5] See Definition 11 in Duraj [2018a] for the concept of mixture convergence $\rightarrow^m$.

[6] We denote with τ the projection $\tau : \Delta(S) \times \mathcal{T} \rightarrow \mathcal{T}$ with $\tau(\pi, t) = t$.

define

$$f(G,\lambda) = \sum_{t=0}^{T} G(\tau = t)\delta^t \mathbb{E}_{\pi\sim G_t}\left[V(A\cup\{b-\lambda\},\pi)\right], \quad \lambda\in[0,b].$$

For fixed belief π the function $\lambda \to V(A\cup\{b-\lambda\},\pi)$ is Lipschitz in λ with Lipschitz-constant 1 and so is then also $\lambda\to\mathbb{E}_{\pi\sim G_t}\left[V(A\cup\{b-\lambda\},\pi)\right]$ upon integration. In particular, the latter is also absolutely continuous.

One can adapt Lemma 3 in Lin [2018] easily (this is Lemma A2 in the jmp version as of January 22) and get for the derivative

$$\frac{d}{d\lambda}\mathbb{E}_{\pi\sim G_t}\left[V(\pi, A\cup\{b-\lambda\})\right] = -1 + \rho^{G_t}_{A\cup\{b-\lambda\}}(A), \text{ a.e. } \lambda\in(0,b),$$

(alternatively this follows from results in Lu [2016]). Here ρ^{G_t} is a stochastic choice function (Lu [2016], Duraj [2018a]) corresponding to the subjective learning model (Dillenberger et al. [2018]) with posterior distribution G_t. Overall it follows that $f(G,\cdot)$ is absolutely continuous and we have

$$\frac{d}{d\lambda}f(G,\lambda) = \sum_{t=0}^{T}\delta^t G(\tau=t)\left(-1+\rho^{G_t}_{A\cup\{b-\lambda\}}(A)\right) \tag{C.3}$$

and

$$\left|\frac{d}{d\lambda}f(G,\lambda)\right| \le 2\sum_{t=0}^{T}\delta^t < \infty.$$

Note that the agent cannot choose arbitrary $G\in\Delta(\Delta(S)\times\mathcal{T})$. G has to correspond to the distribution over $\Delta(S)\times\mathcal{T}$ of a random variable pair (π_τ,τ) where τ is a (randomized) stopping time as introduced in subsection 3.2.2 and π is the process of martingale beliefs induced through the sequential experiment $\mathcal{E}$ and the prior π_0. Recall that we denoted by $\mu(\pi_0)$ the measure of this belief process over the space of experimental outcomes. For further use denote the collection of such distributions G over $\Delta(S)\times\mathcal{T}$ resulting from (randomized) stopping the martingale process with measure $\mu(\pi_0)$ by Π'. One can show easily that this set is compact in the topology of weak convergence of probability measures; it is namely closed in that topology and also tight, since the measures have bounded support.

Conditions for using Theorems 1 and 2 from Milgrom and Segal [2002] are satisfied so that applying them and using (C.3) as well as the fact that $V(\{b\}) = b$ gives

$$V_\delta(A) = b - \int_0^b \sum_{t=0}^{T} \delta^t P_{A\cup\{\lambda\}}(\tau = t) d\lambda + \int_0^b \sum_{t=0}^{T} \delta^t P_{A\cup\{\lambda\}}(A, t) d\lambda = b - \int_0^b \sum_{t=0}^{T} \delta^t P_{A\cup\{\lambda\}}(\lambda, t) d\lambda.$$

Overall we get for any $b \geq b(A)$ (recall that $b(A)$ denotes the best prize feasible under acts from menu A)

$$V_\delta(A) = \int_0^b \left(1 - \sum_{t=0}^{T} \delta^t P_{A\cup\{\lambda\}}(\lambda, t)\right) d\lambda.$$

Now take any $b > b(A)$. Then it follows that $P_{A\cup\{\lambda\}}(\tau = t) = \delta_{\{0\}}(t)$ for the optimal choice of G for $A \cup \{\lambda\}$ by the agent so that the integrand is zero for such b and we can let $b \to \infty$ and get the result.[7]

We now turn to the case of linear costs. Now the function f takes the form

$$f(G, \lambda) = \sum_{t=0}^{T} G(\tau = t) \left(\mathbb{E}_{\pi\sim G_t}\left[V(A \cup \{b - \lambda\}, \pi)\right] - ct\right), \quad \lambda \in [0, b].$$

It follows just as above that $f(G, \cdot)$ is absolutely continuous with

$$\frac{d}{d\lambda} f(G, \lambda) = \sum_{t=0}^{T} G(\tau = t) \left(-1 + \rho^{G_t}_{A\cup\{b-\lambda\}}(A)\right)$$

and

$$\left|\frac{d}{d\lambda} f(G, \lambda)\right| \leq 2T < \infty.$$

Using again the results from Milgrom and Segal [2002] and the fact that

[7] We use the usual notation $\delta_{\{0\}}(t) = 0$ unless $t = 0$ in which case it becomes 1.

$V(\{b\}) = b$ leads to

$$V_c(A) = b - \int_0^b \sum_{t=0}^{T} P_{A\cup\{\lambda\}}(\tau = t)d\lambda + \int_0^b \sum_{t=0}^{T} P_{A\cup\{\lambda\}}(A, t)d\lambda = \int_0^b \sum_{t=0}^{T} P_{A\cup\{\lambda\}}(A, t)d\lambda.$$

Now take again any $b > b(A)$. Then it follows that $P_{A\cup\{\lambda\}}(\lambda, \tau = 0) = 1$ so that we can let $b \to \infty$ and get the result. □

Proof of Proposition 26. We focus again on utility acts and we look at variations of a menu A of the form $A + k$, $k \geq 0$, where k is a constant act giving the constant prize of $k \geq 0$ in every state of the world. Namely, we want to justify writing

$$V_\delta(A + k) = V_\delta(A) + \int_0^k \frac{d}{d\lambda} V_\delta(A + \lambda)d\lambda. \tag{C.4}$$

In line with the notation from Milgrom and Segal [2002] we now define for a given $G \in \Pi'$

$$f(G, \lambda) = \sum_{t=0}^{T} G(\tau = t)\delta^t \mathbb{E}_{\pi\sim G_t}\left[V(A + \lambda, \pi)\right], \quad \lambda \in [0, k],$$

where $G_t := G(\pi \in \cdot | \tau = t)$. It holds now that

$$V(A + \lambda, \pi) = \lambda + V(A, \pi).$$

Therefore it follows that $\lambda \to \mathbb{E}_{\pi\sim G_t}\left[V(A + \lambda, \pi)\right]$ is trivially Lipschitz with constant 1 and so is $\lambda \to f(G, \lambda)$ for every $G \in \Pi'$ as well. In particular, it is absolutely continuous with derivative $\frac{d}{d\lambda}\mathbb{E}_{\pi\sim G_t}\left[V(A + \lambda, \pi)\right] = 1$.

It follows easily that

$$\frac{d}{d\lambda} f(G, \lambda) = \sum_{t=0}^{T} \delta^t G(\tau = t) = \mathbb{E}_{\tau\sim marg_{\{0,\dots,T\}} G}[\delta^\tau]$$

and

$$\left|\frac{d}{d\lambda}f(G,\lambda)\right| \leq \sum_{t=0}^{T}\delta^t < \infty.$$

All conditions from Theorem 2 of Milgrom and Segal [2002] are thus satisfied whenever we assume that k comes from a bounded set. We arrive at the following representation.

$$V_\delta(A+k) = V_\delta(A) + \sum_{t=0}^{T}\delta^t \int_0^k P_{A+\lambda}(\tau = t)d\lambda. \tag{C.5}$$

If instead k is negative, we can write the following after the obvious changes to the argument above.

$$V_\delta(A) = V_\delta(A-k) + \sum_{t=0}^{T}\delta^t \int_0^{-k} P_{A-k+\lambda}(\tau = t)d\lambda. \tag{C.6}$$

□

Proof of Proposition 27. Looking at (3.4) we define

$$f(G,k) = \sum_{t=0}^{T} G(\tau = t)\delta^t \mathbb{E}_{\pi\sim G_t}\left[V(A,\pi)\right] + k\cdot\sum_{t=0}^{T} G(\tau = t)\frac{1-\delta^{t+1}}{1-\delta},$$

for some $G \in \Pi'$. For fixed G this is clearly a Lipschitz function of k. Moreover, we have the derivative

$$\frac{d}{dk}f(G,k) = \sum_{t=0}^{T} G(\tau = t)\frac{1-\delta^{t+1}}{1-\delta},$$

which is clearly bounded uniformly for all $G \in \Pi'$. In particular, all conditions of the Theorems 1 and 2 in Milgrom and Segal [2002] are satisfied so that the application of the abstract envelope theorems is warranted. From this the statements of the Proposition follow immediately.

□

Proof of Proposition 28. Looking at (3.5) we define

$$f(G,c) = \sum_{t=0}^{T} G(\tau = t)\left(\mathbb{E}_{\pi\sim G_t}\left[V(A,\pi)\right] - ct\right),$$

for some $G \in \Pi'$. For fixed G this is clearly a Lipschitz function of c. Moreover, we have the derivative

$$\frac{d}{dc}f(G,c) = -\sum_{t=0}^{T} G(\tau = t)t,$$

which is clearly bounded uniformly for all $G \in \Pi'$. In particular, all conditions of the Theorems 1 and 2 in Milgrom and Segal [2002] are satisfied so that the application of the abstract envelope theorems is warranted. From this the statements of the Proposition follow immediately. □

Proof of Claim 2. This is similar to the proof of Claim 1 in text. We have

$$V_\delta^k(A+\lambda) = \mathbb{E}[\delta^{\tau_{A,k}} V(A,\mu_{\tau_{A,k}})] + k\mathbb{E}\left[\frac{1-\delta^{\tau_{A,k}+1}}{1-\delta}\right].$$

Here $\tau_{A,k}$ is an optimal stopping strategy for the menu A when the duration of the experiment is subsidized by k and expectations are w.r.t. the random realizations of $\tau_{A,k}$ and $\mu_{\tau_{A,k}}$. It holds by revealed preference for $k \neq k'$, both positive numbers, that

$$\mathbb{E}[\delta^{\tau_{A,k}} V(A,\mu_{\tau_{A,k}})] + kE\left[\frac{1-\delta^{\tau_{A,k}+1}}{1-\delta}\right] \geq \mathbb{E}[\delta^{\tau_{A,k'}+1} V(A,\mu_{\tau_{A,k'}})] + k\mathbb{E}\left[\frac{1-\delta^{\tau_{A,k'}}}{1-\delta}\right]. \tag{C.7}$$

By combining the inequality (C.7) for the optimality of $\tau_{A,k}$ in the case of a subsidy of k with its analogue for the optimality of $\tau_{A,k'}$ in the case of a subsidy of k' we get

$$
\begin{aligned}
&\frac{1-\delta}{\delta k}\left(\mathbb{E}[\delta^{\tau_{A,k}} V(A, \mu_{\tau_{A,k}})] - \mathbb{E}[\delta^{\tau_{A,k'}} V(A, \mu_{\tau_{A,k'}})]\right) \\
&\qquad \geq \mathbb{E}[\delta^{\tau_{A,k}}] - \mathbb{E}[\delta^{\tau_{A,k'}}] \\
&\geq \frac{1-\delta}{\delta k'}\left(\mathbb{E}[\delta^{\tau_{A,k}} V(A, \mu_{\tau_{A,k}})] - \mathbb{E}[\delta^{\tau_{A,k'}} V(A, \mu_{\tau_{A,k'}})]\right).
\end{aligned}
$$

From here Claim 2 easily follows. □

Proof of Claim 3 in section 3.3 and of (3.16). One makes use of the fact that the set $\mathbb{T}$ of randomized stopping times is a lattice and shows easily that the map $\mathbb{R} \times \mathbb{T} \ni \mathbb{R}, (k, \tau) \mapsto k\mathbb{E}[\tau]$ is supermodular in (k, τ). Moreover, the function $\mathbb{R} \times \mathbb{T} \to \mathbb{R}, (k, \tau) \mapsto \mathbb{E}[V(A, \pi_\tau)]$ is also supermodular in (k, τ). This is because π is a Martingale and so $\mathbb{E}[V(A, \pi_\tau)]$ is increasing in τ. In particular, the function $\mathbb{R}_+ \times \mathbb{T} \to \mathbb{R}, (c, \tau) \mapsto \mathbb{E}[V(A, \pi_\tau)] - c\mathbb{E}[\tau]$ is supermodular in $(-c, \tau)$. It follows from Theorem 2.8.2 in Topkis [1998] that τ_A^c is pointwise decreasing in c. The last statement is a trivial implication of the FOSD-monotonicity of the expectation operator.

Proof of (3.16): Since $s \to \sum_{t=0}^{T} tP_A^s(\tau = t)$ is weakly decreasing in the costs s it follows that

$$
\begin{aligned}
& c\sum_{t=0}^{T} tP_A^c(\tau = t) - c'\sum_{t=0}^{T} tP_A^{c'}(\tau = t) \\
& = (c - c')\sum_{t=0}^{T} tP_A^c(\tau = t) + c'\left(\sum_{t=0}^{T} tP_A^c(\tau = t) - \sum_{t=0}^{T} tP_A^{c'}(\tau = t)\right) \\
& \leq (c - c')\sum_{t=0}^{T} tP_A^c(\tau = t) \leq \int_{c'}^{c} \sum_{t=0}^{T} tP_A^s(\tau = t) ds \\
& = V_{c'}(A) - V_c(A).
\end{aligned}
$$

We have used monotonicity of the expected stopping time in costs in the two inequalities. Note that we get a strict inequality in

$$c\sum_{t=0}^{T} tP_A^c(\tau = t) - c'\sum_{t=0}^{T} tP_A^{c'}(\tau = t) \leq V_{c'}(A) - V_c(A),$$

whenever $s \rightarrow \sum_{t=0}^{T} tP_A^s(\tau = t)$ is strictly decreasing in some range of (c', c). □

Proof of Theorem 39. This is done in the text of chapter 3. □

Proof of Theorem 6. The axioms are clearly necessary. Sufficiency is also clear, once one notes that axioms 1-4 in both case A. and B. ensure that the identification procedure from section 3.4 works to deliver a unique cost of learning in the respective models. Namely, Axiom 1 and 2 yield the prior and the Bernoulli utility of the agent in a classical way. Moreover, Axiom 3 (Taste stationarity) ensures that the risk preferences of the agent remain ordinally the same across time. This allows the construction of utility acts and their associated menus. Axiom 4-δ ensures that the identification procedure in SeSa-GD for the discount factor works for at least one menu $A \in \mathcal{A}_t$ and yields a discount factor strictly in $(0, 1)$. Axiom 4-c implies that costs of information enter utility of the agent additively, that they are proportional to the expected decision time and that the marginal costs of an additional experiment are menu-independent. It also ensures that the identification of the additive costs does not depend on the menu $A \in \mathcal{A}_t$ without ties chosen. Axiom 5 for both versions of the model then complete rationalizability by showing that the agent $\mathbb{A} = (u, \delta, \mathcal{E}, \pi_0)$ in the case of SeSa-GD or $\mathbb{A} = (u, c, \mathcal{E}, \pi_0)$ in the case of SeSa-LC, where taste, prior and costs of information are identified from Axioms 1-4, match the data of the RCDT for menus without ties where the agent has strict incentives to learn. □

C.4 Calculations for section 3.5

Calculations for SeSa-GD example.

Identification. This can be solved by backwards induction. Suppose that her discount factor is $\delta = \frac{4}{5}$. Consider choosing from menu $\{f_s, r\}$ for $r \in [0, 1]$. Without loss of generality, suppose $s = s_1$. The optimal learning strategy can be derived by backward induction. Below, we do not deal with the situation where indifference in acts or stopping strategy occurs.

At $t = 1$, choose f_s only if s_1 occurs.

At $t = 1$, suppose that $\{s_1, s_2\}$ occurs. If $r \in (\frac{1}{4}, \frac{2}{3})$, then delay choice and the value of A is $\frac{2}{5}(1 + r)$. If $r < \frac{1}{4}$, then choose f_s to enjoy expected utility $\frac{1}{2}$. If $r > \frac{2}{3}$, then choose r. Suppose instead $\{s_1, s_2\}$ does not occur. Then always choose r.

At $t = 0$, if $r < \frac{1}{8}$ or $r > \frac{4}{11}$, then choose f_s or r respectively. If $r \in (\frac{1}{8}, \frac{4}{11})$, then waiting is optimal. This leads to Table 1.

Welfare. SeSa-GD model with a lump-sum tax.

A. Suppose that $S = \{s_1, \cdots, s_4\}$ and $T = 2$. At t=1, it is revealed whether event $\{s_1, s_2\}$ occurs or not. At $t = 2$, the true state of nature is disclosed.

B. The prior belief over S is uniform ($\pi_0(s) = \frac{1}{4}$ for all $s \in S$). The discount factor is $\delta = \frac{4}{5}$.

C. Let $A = \{e_1, \frac{1}{3}\}$. Let's solve the learning behavior for menu $A + k$ given various value of k.

D. Note that if the agent continues at $t = 0$ and learns $\{s_3, s_4\}$ at $t = 1$, then it is optimal to stop at $t = 1$ for any $k > 0$ because $f_1 + k$ is never optimal. Therefore, we can focus only on the following three stopping strategies:

τ: Continue at $t = 0$. Continue at $t = 1$ if and only if $\{s_1, s_2\}$ occurs.

τ': Continue at $t = 0$. Stop at $t = 1$ no matter which event happens.

τ'': Stop at $t = 0$.

E. First, we derive the value of k for which τ is optimal. When agent learns

$\{s_1, s_2\}$ at $t = 1$, continuing is optimal if and only if

$$\frac{1}{2} + k \leq \frac{4}{5}\left[\frac{1}{2}(1+k) + \frac{1}{2}\left(\frac{1}{3}+k\right)\right] \Leftrightarrow k \leq \frac{1}{6}.$$

Given $k \leq \frac{1}{6}$, continuing at $t = 0$ is optimal if and only if

$$\frac{1}{3}+k \leq \frac{4}{5}\left\{\frac{1}{2} \times \frac{4}{5}\left[\frac{1}{2}(1+k) + \frac{1}{2}\left(\frac{1}{3}+k\right)\right] + \frac{1}{2}\left(\frac{1}{3}+k\right)\right\} \Leftrightarrow k \leq \frac{1}{21}.$$

Thus, τ is optimal if and only if $k \leq \frac{1}{21}$.

F. Second, we derive the value of k for which τ' is optimal. When agent learns $\{s_1, s_2\}$ at $t = 1$, stopping is optimal if and only if

$$\frac{1}{2} + k \geq \frac{4}{5}\left[\frac{1}{2}(1+k) + \frac{1}{2}\left(\frac{1}{3}+k\right)\right] \Leftrightarrow k \geq \frac{1}{6}.$$

Given $k \geq \frac{1}{6}$, continuing at $t = 0$ is optimal if and only if

$$\frac{1}{3} + k \leq \frac{4}{5}\left[\frac{1}{2}\left(\frac{1}{2}+k\right) + \frac{1}{2}\left(\frac{1}{3}+k\right)\right] \Leftrightarrow k \leq 0.$$

Thus, τ' is not optimal for any k.

G. According to the previous discussion, τ'' is optimal if and only if $k \geq \frac{1}{21}$.

H. Suppose that $k = 1$. Now, we want to compare the indirect utilities from A and $A + k$. These utilities can be computed from behavior.

I. Consider menu $(A + 1) \cup \{r\} = \{f_1 + 1, \frac{4}{3}, r\}$ where $r > 0$. When $r \leq \frac{4}{3}$, this menu is equivalent to $A + 1$ and so the agent chooses $\frac{4}{3}$ at $t = 0$. When $r > \frac{4}{3}$, this menu is equivalent to $\{f_1, r - 1\} + 1$. It can be shown that stopping at $t = 0$ is still optimal, and r will be chosen. The following table summarizes.

$P_{(A+1)\cup\{r\}}$	$(A+1,0)$	$(r,0)$	$(A+1,1)$	$(r,1)$	$(A+1,2)$	$(r,2)$
$r \in [0, \frac{4}{3})$	1	0	0	0	0	0
$r \in (\frac{4}{3}, \infty)$	0	1	0	0	0	0

Table C.4.1: RCDT for the SeSa-GD example - part 1.

J. Next consider menu $A \cup \{r\}$. When $r < \frac{1}{3}$, this menu is equivalent to A. When $r > \frac{1}{3}$, this menu is equivalent to $\{f_1, r\}$. Thus, we have the following table (recall that we have solved for this menu before).

$P_{A\cup\{r\}}$	$(A,0)$	$(r,0)$	$(A,1)$	$(r,1)$	$(A,2)$	$(r,2)$
$r \in [0, \frac{1}{3})$	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0
$r \in (\frac{1}{3}, \frac{4}{11})$	0	0	0	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$
$r \in (\frac{4}{11}, 1]$	0	1	0	0	0	0

Table C.4.2: RCDT for the SeSa-GD example - part 2.

K. Suppose that an analyst has already identified $\delta = \frac{4}{5}$. Then, using Theorem 1 and Table 7, an analyst can compute the indirect utility of A:

$$\int_0^\infty \left(1 - \sum_{t\in\mathbb{T}} \delta^t P_{A\cup r}(r,t)\right) dr = \frac{1}{3} + \left(\frac{4}{11} - \frac{1}{3}\right)\left(1 - \frac{4}{5} \times \frac{1}{2} - \frac{16}{25} \times \frac{1}{4}\right) = \frac{26}{75}.$$

If we compute the value of A directly from the model, we get

$$\frac{4}{5}\left[\frac{1}{2} \times \frac{4}{5}\left(\frac{1}{2} \times 1 + \frac{1}{2} \times \frac{1}{3}\right) + \frac{1}{2} \times \frac{1}{3}\right] = \frac{26}{75}.$$

This verifies Theorem 4 once again.

L. Similarly, from the data in part I., an analyst can compute the indirect utility of $A+1$, which is $\frac{4}{3}$. Then this analyst concludes that

$$V_\delta(A+1) - V_\delta(A) = \frac{4}{3} - \frac{26}{75} = \frac{74}{75} < 1.$$

Thus, if a social planner subsidize menu A with 1 util, there is a deadweight loss $\frac{1}{75}$.

Calculations for SeSa-LC example.

Identification We study first the belief process given the prior $\pi_0 = (\frac{1}{2}, \frac{1}{2})$ and the i.i.d. experiments prescribed in the example.

Suppose that $j \neq i$. By Bayes' rule,

$$\pi_1(s_i|s_i) = a; \ \pi_1(s_i|s_j) = 1-a.$$

And

$$\pi_2(s_i|s_i, s_i) = \frac{a^2}{1-2a+2a^2} \qquad \pi_2(s_i|s_j, s_j) = \frac{(1-a)^2}{1-2a+2a^2},$$

$$\pi_2(s_i|s_i, s_j) = \pi_2(s_i|s_j, s_i) = \frac{1}{2}.$$

Let's consider SeSa-LC model. Consider menu $A = \{f_1, \frac{7}{12}\}$. Suppose that $a = \frac{2}{3}$. Let q_t denote the belief on state s_1. The random process $\{q_t\}$ is the following:

$$q_0 = \frac{1}{2}; q_1 = \begin{cases} \frac{2}{3} & \text{with prob. } \frac{1}{2}, \\ \frac{1}{3} & \text{with prob. } \frac{1}{2}; \end{cases}$$

$$q_2|_{e_1=s_1} = \begin{cases} \frac{4}{5} & \text{with prob. } \frac{5}{9}, \\ \frac{1}{2} & \text{with prob. } \frac{4}{9}; \end{cases} ; q_2|_{e_1=s_2} = \begin{cases} \frac{1}{2} & \text{with prob. } \frac{4}{9}, \\ \frac{1}{5} & \text{with prob. } \frac{5}{9}; \end{cases}$$

When facing A, f_1 is optimal only if the agent observes only one or two consecutive signals s_1. For various learning cost c, only three stopping strategies may be optimal. They are

τ_0: Stop at $t = 0$.

τ_1: Continue at $t = 0$; stop at $t = 1$.

τ_2: Continue at $t = 0$; continue at $t = 1$ if and only if $q_1 = \frac{2}{3}$.

Let $V(A, \tau, c)$ denote the expected utility from A under stopping strategy τ and learning cost c. Then

$$
\begin{aligned}
V(A, \tau_0, c) &= \frac{7}{12};\\
V(A, \tau_1, c) &= \frac{1}{2} \times \frac{2}{3} + \frac{1}{2} \times \frac{7}{12} - c = \frac{5}{8} - c;\\
V(A, \tau_2, c) &= \frac{5}{18} \times \frac{4}{5} + \frac{13}{18} \times \frac{7}{12} - \left(\frac{1}{2} \times 2c + \frac{1}{2}c\right) = \frac{139}{216} - \frac{3}{2}c.
\end{aligned}
$$

Thus, τ_0 is optimal if and only if

$$\frac{7}{12} \geq \max\left\{\frac{139}{216} - \frac{3}{2}c, \frac{5}{8} - c\right\} \Leftrightarrow c \geq \frac{1}{24}.$$

Strategy τ_1 is optimal if and only if

$$\frac{5}{8} - c \geq \max\left\{\frac{139}{216} - \frac{3}{2}c, \frac{7}{12}\right\} \Leftrightarrow \frac{1}{27} \leq c \leq \frac{1}{24}.$$

Strategy τ_2 is optimal if and only if

$$\frac{139}{216} - \frac{3}{2}c \geq \max\left\{\frac{7}{12}, \frac{5}{8} - c\right\} \Leftrightarrow c \leq \frac{1}{27}.$$

Intuitively, when c gets smaller, the agent learns more information.

Now we work backwards and assume that $c = \frac{1}{36}$. Thus τ_2 is optimal for A, and the value of A is $\frac{139}{216} - \frac{3}{2} \times \frac{1}{36} = \frac{65}{108}$.

Consider facing menu $A \cup r$. When $r \leq \frac{7}{12}$, the agent behaves as if facing A because r is dominated. Thus τ_2 is optimal. When $r > \frac{4}{5}$, the agent will not learn because f_1 is never optimal.

Consider $r \in (\frac{2}{3}, \frac{4}{5})$. Then f_1 will be chosen only when observing two signals s_1. In other words, under τ_1, r is always chosen and so τ_1 is worse than τ_0. The expected utility from $A \cup r$ under τ_2 is

$$V(A \cup r, \tau_2, c) = \frac{5}{18} \times \frac{4}{5} + \frac{13}{18} \times r - \frac{3}{2} \times \frac{1}{36} = \frac{13}{72} + \frac{13}{18} r.$$

The utility is r under τ_0. Thus τ_0 is optimal for all $r \in (\frac{2}{3}, \frac{4}{5})$.

Consider $r \in (\frac{7}{12}, \frac{2}{3})$.

Then

$$\begin{aligned} V(A \cup r, \tau_0, c) &= r; \\ V(A \cup r, \tau_1, c) &= \frac{1}{2} \times \frac{2}{3} + \frac{1}{2} \times r - \frac{1}{36} = \frac{11}{36} + \frac{1}{2} r; \\ V(A \cup r, \tau_2, c) &= \frac{5}{18} \times \frac{4}{5} + \frac{13}{18} \times r - \frac{3}{2} \times \frac{1}{36} = \frac{13}{72} + \frac{13}{18} r. \end{aligned}$$

Thus, τ_2 is optimal when $r \in (\frac{7}{12}, \frac{13}{20})$; τ_0 is optimal when $r \in (\frac{13}{20}, \frac{2}{3})$. In sum, an analyst, who does not know c, observes the table below.

$P_{A\cup r}$	$(A, 0)$	$(r, 0)$	$(A, 1)$	$(r, 1)$	$(A, 2)$	$(r, 2)$
$r \in [0, \frac{7}{12})$	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0
$r \in (\frac{7}{12}, \frac{13}{20})$	0	0	0	$\frac{1}{2}$	$\frac{5}{18}$	$\frac{4}{18}$
$r \in (\frac{13}{20}, \infty)$	0	1	0	0	0	0

Table C.4.3: RCDT for the SeSa-LC example - part 1.

Then consider menu $\alpha A + (1 - \alpha)(A \cup r)$. When α is close to 1 enough, the agent's choice from this mixed menu reveals how she would choose from $A \cup r$ if she follows the same stopping strategy as facing A. An analyst observes the following table.

$\tau_{\{A,r\}}$	$(A,0)$	$(r,0)$	$(A,1)$	$(r,1)$	$(A,2)$	$(r,2)$
$r \in [0, \frac{7}{12})$	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0
$r \in (\frac{7}{12}, \frac{4}{5})$	0	0	0	$\frac{1}{2}$	$\frac{5}{18}$	$\frac{4}{18}$
$r \in (\frac{4}{5}, \infty)$	0	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$

Table C.4.4: RCDT for the SeSa-LC example - part 2.

An analyst given above two tables can compute the ex-ante value of A:

$$\int_0^1 \sum_{t\in\mathbb{T}} P_{A\cup r}(A,t)dr = \frac{7}{12} \times 1 + (\frac{13}{20} - \frac{7}{12}) \times \frac{5}{18} = \frac{65}{108},$$

compute the expected utility gain from menu A:

$$\int_0^1 \tau_{A\cup r}(A)dr = \frac{7}{12} \times 1 + (\frac{4}{5} - \frac{7}{12}) \times \frac{5}{18} = \frac{139}{216},$$

and compute the average decision time when facing A:

$$\sum_{t\in\mathbb{T}} t \times P_A(A,t) = \frac{1}{2} \times 1 + \frac{1}{2} \times 2 = \frac{3}{2}.$$

Consequently, this analyst recovers the flow cost of time:

$$c = \frac{\frac{139}{216} - \frac{65}{108}}{\frac{3}{2}} = \frac{9}{216} \times \frac{2}{3} = \frac{1}{36}.$$

Welfare. Consider an information market. An agent who has to choose from menu A would like to learn information about the state of the world. A scientist can run the experiment at a fixed fee. The price (in utils) of the experiment is c. That is, if the agent wants to run the experiment k times, she has to pay $k \times c$ to the scientist.

Suppose that $c = \frac{1}{36}$. As we have shown before, the agent adopts learning strategy τ_2.

Now suppose that the government intends to tax the experiment. Specifically, if the agent runs the experiment k times, she has to pay the government $k \times (\frac{1}{25} - \frac{1}{36})$. Thus, the effective cost of learning becomes $\frac{1}{25}$. The agent will adopt stopping strategy τ_1 instead, learning less information. Recall that expected decision time is equal to 1.

Before tax, the welfare of the government is 0. The welfare of the scientist is $\frac{1}{36}E_{\tau_2}[t] = \frac{1}{36} \times \frac{3}{2}$. The welfare of the agent is $V_c(A)|_{c=\frac{1}{36}} = \frac{65}{108}$. Total welfare in the economy is $\frac{139}{216}$.

After tax, the welfare of the government is $\frac{1}{25} - \frac{1}{36}$. The welfare of the scientist is $\frac{1}{36}$. The welfare of the agent is $V_c(A)|_{c=\frac{1}{25}} = \frac{5}{8} - \frac{1}{25} = \frac{117}{200}$. The total welfare is $\frac{5}{8}$.

The agent is worse off, and the total welfare also decreases. Note that

$$V_c(A)|_{c=\frac{1}{36}} - V_c(A)|_{c=\frac{1}{25}} = \frac{91}{5400} > \frac{1}{36} \times \frac{3}{2} - \frac{1}{25} = \frac{1}{600}.$$

Therefore we have a dead-weight loss of

$$\frac{91}{5400} - \frac{1}{600} = \frac{41}{2700}.$$

C.5 Identification of the additive costs in SeSa-LC without decision time data

Here we use again utility acts w.r.t. a Bernoulli utility u as well as the identification result from Lin [2018]. From there we know that the minimal canonical costs of a rational inattention model are given by

$$c(\pi) = \sup_{A \in \mathcal{A}} \left[\int_{\Delta(S)} \max_{f \in A}(p \cdot f)\pi(dp) - \int_0^\infty P_{A \cup \{a\}}(A) da \right]. \tag{C.8}$$

Now for the sequential sampling model with linear costs c per unit of time, if we allow for randomized stopping times and the agent's belief martingale is the $\Delta(S)$-valued process μ_t, $t \leq T$ we know that the set of posterior distributions achieved is given as follows.

$$\mathcal{F}(\mu) = \{\pi \in \Delta(\Delta(S)) : \mu_\tau =^d \pi, \tau \text{ randomized stopping time}\}.$$

This is a convex, compact subset of $\Delta(\Delta(S))$. If we look at the decision problem of the agent in (3.5) we see that it corresponds to a rational inattention model with costs

$$C(\pi) = c \min_{\tau \text{ randomized}, \mu_\tau =^d \pi} \mathbb{E}[\tau]. \tag{C.9}$$

Note that the costs here are also normalized: $C(\delta_{\pi_0}) = 0$ (using the prior is costfree).

Due to Theorem 2 in Lin [2018] the two costs (C.8) and (C.9) should be equal for every $\pi \in \mathcal{F}(\mu)$ that is realized for some menu $A \in \mathcal{A}$. But since for the other π-s we can change the canonical costs in any way we want it follows overall that we can use the following formula to identify the costs c:

$$c = \frac{\sup_{A \in \mathcal{A}} \left[\int_{\Delta(S)} \max_{f \in A}(p \cdot f)\pi(dp) - \int_0^\infty P_{A \cup \{a\}}(A) da\right]}{\min_{\tau \text{ randomized}, \mu_\tau =^d \pi} \mathbb{E}[\tau]}, \text{ for some } \pi \in \mathcal{F}(\mu).$$

References

D. Abreu and F. Gul. Bargaining and reputation. *Econometrica*, 68:85–118, 2000.

V.V. Acharya, P. DeMarzo, and I. Kremer. Endogenous information flows and the clustering of announcements. *American Economic Review*, 101(1):2955–2979, 2011.

G. Akerlof. The market for lemons: Quality uncertainty and the market mechanism. *Quarterly Jorunal of Economics*, 89, 1970.

K.J. Arrow, D. Blackwell, and M. A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17(3):213–244, 1949.

Ned Augenblick and Matthew Rabin. Belief movement, uncertainty reduction, and rational updating. *Working Paper*, 2018.

R. J. Aumann and M. B. Maschler. *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press, 1995.

L.M. Ausubel, P. Cramton, and R.J. Deneckere. Bargaining with incomplete information. In *Handbook of Game Theory*, volume 3, chapter 50, pages 1897–1945. Elsevier Science B.V., 2002.

J. Banks and J. Sobel. Equilibrium selection in signalling games. *Econometrica*, 55 (3):647–661, 1987.

Nicholas Barberis and Wei Xiong. Realization utility. *Journal of Financial Economics*, 104(2):251–271, 2012.

D. Bergemann and A. Bonatti. Markets for information: An introduction. *Cowles discussion paper No. 2142*, 2018.

D. Bergemann and J. Välimaki. Dynamic mechanism design: An introduction. *Journal of Economic Literature*, 57:235–274, 2019.

S. Board and M. Pycia. Outside options and the failure of the coase conjecture. *American Economic Review*, 104(2):656–671, 2014.

David Bowman, Deborah Minehart, and Matthew Rabin. Loss aversion in a consumption–savings model. *Journal of Economic Behavior and Organization*, 38(2):155–178, 1999.

Markus K Brunnermeier and Jonathan A Parker. Optimal expectations. *American Economic Review*, 95(4):1092–1118, 2005.

A. Caplin and M. Dean. Revealed preference, rational inattention and information acquisition. *American Economic Review*, 105(7):2183–2203, 2011.

H. Chade and E. E. Schlee. Bargaining with arrival of new traders. *Journal of Economic Theory*, 107:421–452, 2002.

K. Chaterjee and L. Samuelson. Bargaining with two-sided incomplete information. *Review of Economic Studies*, 54(2):175–192, 1987.

Y.-K. Che and K. Mierendorff. Optimal sequential decision with limited attention. *working paper, December 2018 version*, 2018.

P.C. Cramton. Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty. *Review of Economic Studies*, 51(4): 579–593, 1984.

J. Crémer and F. Khalil. Gathering information before signing a contract. *American Economic Review*, 82(3):566–578, 1992.

J. Crémer, F. Khalil, and J.-C. Rochet. Contracts and productive information gathering. *Games and Economic Behavior*, 25(2):174–193, 1998.

B. Daley and B. Green. Bargaining and news. *American Economic Review, forthcoming*, 2019.

H. De Oliveira, T. Denti, M. Mihm, and K. Ozbek. Rationally inattentive preferences and hidden information costs. *Theoretical Economics*, 12:621–654, 2017.

P.M DeMarzo, I. Kremer, and A. Skrzypacz. Test design and minimum standarts. *American Economic Review, forthcoming*, 2019.

R.J. Deneckere and M.-Y. Liang. Bargaining with interdependent values. *Econometrica*, 74(5):1309–1364, 2006.

T. Denti. Posterior-separable cost of information. *working paper 2018*, 2018.

D. Dillenberger, J.S. Lleras, P. Sadowski, and N. Takeoka. A theory of subjective learning. *Journal of Economic Theory*, 173:56–94, 2018.

David Dillenberger and Collin Raymond. Additive-belief-based preferences. *Working Paper*, 2018.

D. Dranove and G. Z. Jin. Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4):935–963, 2010.

J. Duraj. Dynamic random subjective expected utility. *working paper, August 2018 version*, 2018a.

J. Duraj and Y.-H. Lin. Costly information and random choice. *working paper 2018*, 2019.

Jetlir Duraj. Mechanism design with news utility. *Working Paper*, 2018b.

R. Durrett. *Probability: Theory and Examples*. Cambridge University Press, 4 edition, 2010.

R. Dye. Disclosure of non-proprietary information. *Journal of Accounting Research*, 23(3), 1985.

F. Echenique and K. Saito. Response time and utility. *Journal of Economic Behavior and Organization*, 139:49–59, 2017.

Kfir Eliaz and Ran Spiegler. Can anticipatory feelings explain anomalous choices of information sources? *Games and Economic Behavior*, 56(1):87–104, 2006.

Jeffrey Ely, Alexander Frankel, and Emir Kamenica. Suspense and surprise. *Journal of Political Economy*, 123(1):215–260, 2015.

Jeffrey C Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.

L.G. Epstein. Stationary cardinal utility and optimal growth under uncertainty. *Journal of Economic Theory*, 31:133–152, 1983.

P. Esö and C. Wallace. Information and evidence in bargaining. *Economic Theory Bulletin*, 2(1):23–32, 2014.

P. Esö and C. Wallace. Bargaining with hard evidence. *The Economic Journal*, 129: 2039–2063, 2019.

L. C. Evans. *Partial Differential Equations, 2nd edition*. AMS: Graduate Studies in Mathematics, 2010.

Y. Feinberg and A. Skrzypacz. Uncertainty about uncertainty and delay in bargaining. *Econometrica*, 73(1):69–91, 2005.

P.C. Fishburn and A. Rubinstein. Time preference. *International Economic Review*, 23(3):677–694, 1982.

M. Frick, R. Iijima, and T. Strzalecki. Dynamic random expected utility. *working paper 2018*, 2018.

W. Fuchs and A. Skrzypacz. Bargaining with arrival of new traders. *American Economic Review*, 100:802–836, 2010.

D. Fudenberg and J. Tirole. Sequential bargaining with incomplete information. *Review of Economic Studies*, 50(2):221–247, 1983.

D. Fudenberg and J. Tirole. *Game Theory*. Cambridge, MA: MIT Press, 1991.

D. Fudenberg, D. K. Levine, and J. Tirole. Infinite-horizon models of bargaining with one-sided incomplete information. In *Game-theoretic Models of Bargaining*, chapter 5. Cambridge University Press, 1985.

D. Fudenberg, D. K. Levine, and J. Tirole. Incomplete information bargaining with outside opportunities. *Quarterly Journal of Economics*, 102(1):37–50, 1987.

D. Fudenberg, P. Strack, and T. Strzalecki. Speed, accuracy and the optimal timing of choices. *American Economic Review*, 108(12):3651–3684, 2018.

P.A. Gompers, W. Gornall, S.N. Kaplan, and I.A. Strebulaev. How do venture capitalists make decisions? *Journal of Financial Economics, forthcoming*, 2019.

S. Grossman. The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, 24:461–489, 1981.

F. Gul, H. Sonnenschein, and R. Wilson. Foundations of dynamic monopoly and the Coase conjecture. *Journal of Economic Theory*, 39:155–190, 1986.

I. Guttman, I. Kremer, and A. Skrzypacz. Not only what but also when: A theory of dynamic voluntary disclosure. *American Economic Review*, 104(8): 2955–2979, 2014.

B. Hebert and M. Woodford. Information costs and sequential information sampling. *working paper, November 2018 version*, 2018a.

B. Hebert and M. Woodford. Rational inattention in continuous time. *working paper, November 2018 version*, 2018b.

Vicky Henderson. Prospect theory, liquidation, and the disposition effect. *Management Science*, 58(2):445–460, 2012.

I. Hwang. A theory of bargaining deadlock. *Games and Economic Behavior*, 109: 501–522, 2018a.

I. Hwang. Dynamic trading with developing adverse selection. *Journal of Economic Theory*, 176:761–802, 2018b.

I. Hwang and F. Li. Transparency of outside options in bargaining. *Journal of Economic Theory*, 167:116–147, 2017.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.

E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

R. Kirpalani and E. Madsen. Information acquisition and strategic investment timing. *working paper August 2019 version*, 2019.

T. C. Koopmans. Stationary ordinal utility and impatience. *Econometrica*, 28 (287-309), 1960.

Botond Kőszegi. Emotional agency. *Quarterly Journal of Economics*, 121(1): 121–155, 2006.

Botond Kőszegi and Matthew Rabin. Reference-dependent consumption plans. *American Economic Review*, 99(3):909–36, 2009.

Albert S Kyle, Hui Ou-Yang, and Wei Xiong. Prospect theory and liquidation decisions. *Journal of Economic Theory*, 129(1):273–288, 2006.

S. Lang. *Undergraduate Analysis*. Springer Science+Business Media, Inc., 1997.

Fei Li and Peter Norman. Sequential persuasion. *Working Paper*, 2018.

Y. Li. Efficient mechanisms with information acquisition. *Journal of Economic Theory*, 182:279–328, 2019.

Y.-H. Lin. Stochastic choice and rational inattention. *working paper, November 2018 version*, 2018.

Elliot Lipnowski and Laurent Mathevet. Disclosure to a psychological audience. *American Economic Journal: Microeconomics*, 10(4):67–93, 2018.

N. Lomys. Learning while bargaining: Experimentation and coasean dynamics. *working paper 2017*, 2017.

J. Lu. Random choice and private information. *Econometrica*, 84(6):1983–2027, 2016.

Rosario Macera. Dynamic beliefs. *Games and Economic Behavior*, 87:1–18, 2014.

T. Magnac and D. Thesmar. Identifying dynamic discrete decision processes. *Econometrica*, 70(2):801–816, 2002.

A. Mas-Collel, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. New York: Oxford University Press, 1995.

P. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12:380–391, 1981.

P. Milgrom. What the seller won't tell you: Persuasion and disclosure in markets. *Journal of Economic Perspectives*, 22(2):115–131, 2008.

P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.

Ted O'Donoghue and Charles Sprenger. Reference-dependent preferences. *Handbook of Behavioral Economics-Foundations and Applications 1*, page 1, 2018.

J. Ortner. Durable goods monopoly with stochastic costs. *Theoretical Economics*, 12(2):817–861, 2017.

J. Ortner. A continuous-time model of bilateral bargaining. *Games and Economic Behavior*, 113:720–733, 2019.

Michaela Pagel. Expectations-based reference-dependent preferences and asset pricing. *Journal of the European Economic Association*, 14(2):468–514, 2016.

Michaela Pagel. Expectations-based reference-dependent life-cycle consumption. *Review of Economic Studies*, 84(2):885–934, 2017.

Michaela Pagel. A news-utility theory for inattention and delegation in portfolio choice. *Econometrica*, 86(2):491–522, 2018.

Matthew Rabin and Georg Weizsäcker. Narrow bracketing and dominated choices. *American Economic Review*, 99(4):1508–43, 2009.

R. Radner and J. Stiglitz. A nonconcavity in the value of information. In *Bayesian Models of Economic Theory*, pages 33–52. Elsevier, B.V., 1984.

D. Ravid. Bargaining with rational inattention. *working paper, March 2019 version*, 2019.

A. Rubinstein. A bargaining model with incomplete information about time preferences. *Econometrica*, 53(5):1151–1172, 1985.

Nikolaus Schweizer and Nora Szech. Optimal revelation of life-changing information. *Management Science*, 64(11):5250–5262, 2018.

Hersh Shefrin and Meir Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of finance*, 40(3): 777–790, 1985.

Xianwen Shi. Optimal auctions with information acquisition. *Games and Economic Behavior*, 74(2):666–686, 2012.

J. Sobel and I. Takahashi. A multistage model of bargaining. *Review of Economic Studies*, 50(3):411–426, 1983.

D.M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 1998.

Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323, 1992.

A. Wald. Foundations of a general theory of sequential decision functions. *Econometrica*, 15(4):279–313, 1947.

Wenhao Wu. Sequential bayesian persuasion. *Working Paper*, 2018.

M. Yildiz. Waiting to persuade. *Quarterly Journal of Economics*, 119(1):223–248, 2004.

C. Zeisberger, M. Prahl, and B. White. *Private Equity in Action: Case Studies from Developed and Emerging Markets*. Wiley & Sons, 2017a.

C. Zeisberger, M. Prahl, and B. White. *Mastering Private Equity: Transformation via Venture Capital, Minority Investments and Buyouts*. Wiley & Sons, 2017b.

W. Zhong. Optimal dynamic information acquisition. *working paper, January 2019 version*, 2019.

Colophon

THIS THESIS WAS TYPESET using LaTeX, originally developed by Leslie Lamport and based on Donald Knuth's TeX. The body text is set in 11 point Arno Pro, designed by Robert Slimbach in the style of book types from the Aldine Press in Venice, and issued by Adobe in 2007. A template, which can be used to format a PhD thesis with this look and feel, has been released under the permissive MIT (X11) license, and can be found online at github.com/suchow/ or from the author at suchow@post.harvard.edu.