Chunky Information and Adverse Selection*

Costas Cavounidis†

September 2, 2022

Abstract

A group of appraisers face a prospect in a random order. Not all prospects are good so, upon encountering a prospect, an appraiser may acquire costly information before deciding to accept or reject it. As each prospect can only be accepted by one appraiser, selectivity causes an adverse selection externality. I show that if appraisers may acquire any signal of the prospect's quality, all appraisers, regardless of their number, have positive expected payoffs. By contrast, when information is chunky, equilibrium adverse selection limits the number of appraisers who can profit by participating.

*I am grateful to Dan Bernhardt, Kevin Cooke, Sambuddha Ghosh, Shulamit Kahn, Kevin Lang, Bart Lipman, Calvin Luscombe, Albert Ma, Filip Matějka, Raghav Malhotra, Dilip Mookherjee, Benjamin Ogden, Juan Ortner, Doron Ravid, Juuso Toikka, and audiences at EARIE, Boston University, University of Georgia, University of Glasgow, University of Warwick, University of Sussex and University of St. Andrews. I am grateful to Hyungmin Park for excellent research assistance. I declare that I have no conflicts of interest relating to this paper.

†University of Warwick, c.cavounidis@warwick.ac.uk

1 Introduction

Consider a venture capitalist inspecting a start-up soliciting an investment. Ideally, she invests when the return on investment clears some benchmark; however, this is not immediately apparent. The venture capitalist may aid her decision by acquiring costly information about the start-up. The prior distribution of the start-up's profitability, naturally, informs the venture capitalist's decision of what information — if any — to purchase.

However, the venture capitalist must also consider that she may not be the first potential investor solicited. Could the start-up have already been turned down by other venture capitalists? What would those other venture capitalists have learned about the start-up, and how would it have influenced rejections? The answers to these questions affect the venture capitalist's beliefs about the start-up, and therefore her own information acquisition decisions. If other venture capitalists are being circumspect, perhaps this is cause to be circumspect oneself.

One may reason similarly of health insurers learning about applicants, employers inspecting job candidates, faculty taking time to learn about potential advisees, and consumers poking their fingers into that last avocado in the grocery aisle. More generally, information acquisition produces an interdependent adverse selection externality in common values settings. When agents screen in sequence, this manifests as a *solicitation curse*[1] analogous to the *winner's curse* in simultaneous setups. In these kinds of settings, an agent's incentives to acquire information are affected by the strategic information acquisition of other agents. An intriguing possibility arises: could equilibrium beliefs be so pessimistic that some agents give up on learning altogether? In a stylized environment, I will show that this depends on the fine details of information acquisition.

I develop a model with a set of *appraisers* that encounter a *prospect* in a uniform random *visit order*. As the realization of the visit order is unknown to appraisers, and bad prospects will in equilibrium be rejected more often by other appraisers, beliefs conditioned on encountering the prospect exhibit a solicitation curse.

When faced with the prospect, an appraiser can choose what information to learn about the prospect's quality. This is modeled as choosing a Blackwell experiment - a

[1] I use the term in the sense of Kim and Pease (2017), but the analogy is imperfect. In that paper, the solicitation curse results from worse prospects exerting more search effort, rather than their being rejected more often.

signal whose distribution varies with the prospect's quality. Experiments vary in cost, so a more informative one is not always preferable. An appraiser may, in fact, opt to spend nothing and receive no information. Following information acquisition, the appraiser must accept or reject the prospect. If the appraiser rejects the prospect she reaps neither loss nor benefit, and the prospect is encountered by the next appraiser.[2] If the appraiser accepts the prospect, she gains or loses utility on the basis of the prospect's quality, and the game ends.

I argue that this setting is ideal for studying the relationship between information acquisition and adverse selection externality in isolation. Information acquisition occurs at an 'interim' stage - the active appraiser holds full decision rights over the prospect at that time. Thus, information acquisition decisions are not affected by a concurrent competition for the prospect, as in the auction setting of Persico (2000) (in which information acquisition occurs 'ex ante'), but instead reflect current beliefs.[3] Furthermore, there is no additional source of adverse selection such as mutual consent to transactions, as in Lauermann and Wolinsky (2016).[4] Compared to models such as these, the present model considerably broadens the forms information acquisition may take. This is at the cost of dispensing with any notion of prices or bids, like Ely and Siegel (2013).

The model takes full advantage of recent advances in the literature on information acquisition. *Divisible* information entails the freedom to collect arbitrarily small increments of information and use them to inform further collection of information. A result from Bloedel and Zhong (2020) (henceforth, BZ) then allows me to represent divisible information with a Uniform Posterior-Separable (UPS) cost function over posteriors. Such functions are rapidly becoming ubiquitous since their introduction by Caplin and Dean (2013). When information is *chunky*, on the other hand, learning can still be sequenced, but information is not available in arbitrarily small increments. This is not to say chunky information is necessarily rigid - there may well be a rich

[2]Thus, as an appraiser may always opt to acquire no information and reject the prospect, 0 is a lower bound on the equilibrium expected payoff of each appraiser.

[3]In models of ex-ante information acquisition, a player's beliefs at the information acquisition stage are not affected by others' strategies. Thus, the costs of acquiring particular experiments are invariant to others' strategies. Instead, others' strategies affect a player's payoff from accepting a prospect. An interesting related case is Gershkov and Szentes (2009) in which the optimal mechanism asks players to acquire information in the interim in a uniform random visit order, but the active player does not know how her report will affect the decision at hand.

[4]The latter includes a form of (Groucho) Marxist adverse selection - one should be wary of being party to transactions that others would be party to.

menu of chunks.

In both cases, information acquisition strategies should be thought of as reduced-form descriptions of potentially complex dynamic plans. That is, appraisers are not merely "setting the bar"in a test, but also deciding on how long the test goes on for, as a function of the answers given.

Should information costs vary with beliefs? Mensch (2018) and Denti, Marinacci, and Rustichini (2019) demonstrate how information costs that depend arbitrarily on beliefs are inapt for strategic settings. On the other hand, BZ show that belief-invariant costs do not allow us to represent sequential procedures. It thus seems reasonable to at the least allow for information costs that are linear in beliefs, to accomodate Morris and Strack (2019)'s generalized Wald model. Rather than take a stand, I allow for generality. This requires using Geanakoplos, Pearce, and Stacchetti (1989)'s notion of a psychological equilibrium as the solution concept. In what I view as the canonical case of information costs linear in beliefs, however, the solution concept reduces to Bayes-Nash Equilibrium.

I show two main results. First, when information is divisible, all appraisers acquire information and each receives a positive expected payoff. This is due to the marginality of beliefs about rejected prospects. Appraisers acquire just enough information that their posterior beliefs about rejected prospects lie on a threshold between a region in which only rejection is optimal and one in which only further experimentation is optimal.[5] An appraiser encountering a prospect might be the first to do so, or might be later in the order. As a consequence, an appraiser assumed for contradiction not to acquire any information in equilibrium would hold beliefs equal to a convex combination of the prior for fresh prospects and the posterior for rejected ones. As the prior is above the threshold and the posterior for rejected prospects equals the threshold, these equilibrium beliefs would be above the threshold. But at such beliefs, the appraiser should acquire some information, a contradiction. Therefore, no matter how many appraisers are in the game, all acquire information in equilibrium; earning positive expected payoffs. Divisible information attenuates the informational externality due to each appraiser as the number of appraisers increases.[6]

[5]This follows from the fact UPS cost functions satisfy Caplin, Dean, and Leahy (2022)'s *Locally Invariant Posteriors* property. In particular, the distribution of posteriors following optimal information acquisition has support on the same two points - if feasible.

[6]The fact the visit order is uniformly distributed plays an important role here. If, say, a particular appraiser knew she was always last in the order, she would never acquire information, and would

Second, when information is chunky, the informational externality limits the number of appraisers reaping a nonzero expected payoff. There cannot be equilibria with too many appraisers gainfully acquiring and acting on information. If there were, either they would be spending too much on acquiring information, or the prospect would be accepted too often. In the former case some appraiser could gain by deviating to rejecting without information; in the latter case some appraiser could gain by screening more heavily. As information acquisition cannot be finely tuned in a chunky information market, posteriors for the rejected are strictly below the rejection/further experimentation threshold. Unlike in the divisible case, the externality due to this excessive information acquisition is powerful enough to limit market participation.

It is important to recognize that the latter result is not due to 'congestion', as in a model in which appraisers acquire information before encountering the prospect. In my model, all costs are paid at the *interim stage*, in the sense that the appraiser only pays for information when she has unilateral power to decide the fate of the prospect. Unlike in Bergemann, Shi, and Välimäki (2009), an increase in the number of appraisers does not affect the value of information via a decrease in 'market share'. Instead, beliefs become excessively pessimistic as the adverse selection externality intensifies. Section 4.1 shows how this force is blunted as information becomes 'more divisible'. It is important, however, that appraisers can always acquire *more* information; the example in Section 4.3 shows that the result is not valid when appraisers cannot sample repeatedly. Moreover, Section 4.4 illustrates the externality caused by over-acquiring information by way of a market with an appraiser with commitment power; she can commit to an experiment that keeps entrants out.

I argue that the contribution of this paper is threefold. First, it develops a model in which interdependent adverse selection can be studied in isolation. The model is applicable to several settings and yields tractable implications despite allowing agents wide discretion in information acquisition. Second, like Ravid, Roesler, and Szentes (2022), it shows that the minute details of the information acquisition technology affect macroscopic outcomes such as market structure. 'Cheap enough' information is not free, and its character not separable from market entry; moreover, the quanta of information matter. Finally, the paper illustrates the usefulness and limits of the modern information acquisition literature, at once applying its results to a new domain and showing that its stylization is not without loss.

receive a payoff of 0.

2 Model

There is a single *prospect*, characterized by quality θ drawn from $\Theta = \{\underline{\theta}, \bar{\theta}\}$ with the prior putting probability $p_0 \in (0, 1)$ on $\bar{\theta}$. There exist N identical *appraisers*, who will act as players in this game. The quality θ is the benefit to an appraiser of accepting the prospect. I assume $\underline{\theta} < 0 < \bar{\theta}$; appraisers want to accept some, but not all, prospects.

2.1 Timing

Nature chooses θ according to the prior, and also a *visit order* π uniformly from the set of permutations of the N appraisers. Then, the game proceeds in up to N stages, starting with stage 1.

If and when stage n is reached, appraiser $\pi(n)$ encounters the prospect. This appraiser may then choose to acquire information about θ using the available technology. Once information is acquired, the appraiser may choose to accept or reject the prospect. Following the appraiser's decision, if the prospect is accepted or $n = N$, the game ends. On the other hand, if the prospect is rejected and $n < N$, stage $n+1$ follows.

Each appraiser becomes active at a single information set. That is to say, appraisers are not aware of the prospect's history of encountering other appraisers. As a nonempty history has only rejections, which in equilibrium would convey bad news about the prospect's quality, a prospect (a startup, say) would not disclose previous encounters even if the model were augmented with a cheap-talk stage.

To summarize, the prospect visits appraisers in a (uniform) random order, until either an acceptance occurs, or the prospect has been rejected by every appraiser. When an appraiser encounters the prospect, the appraiser observes only that a prospect has arrived, not the history, the visit order, or the date.

2.2 Information Acquisition

When an appraiser encounters the prospect, that appraiser can then acquire information about the prospect's quality. Information acquisition is modeled as choice of a *Blackwell experiment*, a collection of quality-conditional distributions $\{\sigma(\cdot|\theta)\}_{\theta\in\Theta}$ for a signal on Polish space S_σ. The set of all Blackwell experiments is denoted $\mathcal{B}$. A null experiment σ_{null} carries no information, so that $\sigma_{\text{null}}(\cdot|\theta)$ does not depend on θ.[7]

Appraisers have access to a collection of experiments $\mathcal{G} \subseteq \mathcal{B}$, the *information menu*. This menu represents the procedures available to an appraiser in order to learn about prospects. The information menu is common to all appraisers. Experiments will vary in cost, so a more informative experiment will not always be preferable.

Unlike much of the literature, I do not always assume that $\mathcal{G}$ is equal to $\mathcal{B}$. This is crucial; I am primarily interested in varying the nature of $\mathcal{G}$ rather than the cost function c. I assume $\mathcal{G}$ contains a null experiment, so not learning is an option.

2.3 Strategies

Each appraiser n must choose two things: what (if any) information to acquire and how to use it. Thus, each appraiser n must choose a $\sigma_n \in \mathcal{G}$. The use of information reduces to a function $a_n : S_{\sigma_n} \to \{0, 1\}$, which denotes (as a function of the experiment's outcome) whether the appraiser accepts (1) or rejects (0) the prospect. Thus, the strategy space of each appraiser n is $\Delta\{(\sigma_n, a_n)|\sigma_n \in \mathcal{G},\ a_n : S_{\sigma_n} \to \{0, 1\}\}$. A typical (mixed) strategy for appraiser i is denoted φ_n.

2.4 Beliefs

To address payoffs, we first have to find the probability an appraiser encounters a prospect of each quality. In other words, we have to account for selection due to other appraisers' strategies.

When a prospect of quality θ visits appraiser m, the chance she *rejects* given her strategy φ_m is $\left(1 - E_{\varphi_m}\left[E_{\sigma_m(\cdot|\theta)}[a_m]\right]\right)$. Now, notice that for a given visit order π and profile of strategies for n's opponents φ_{-n}, the chance n encounters the prospect is the chance that *all* previous (with respect to π) appraisers reject the prospect; or[8]

$$\prod_{m:\pi(m)<\pi(n)} \left(1 - E_{\varphi_m}\left[E_{\sigma_m(\cdot|\theta)}[a_m]\right]\right). \tag{1}$$

As the visit order π is random, given a prospect of quality θ, to find the probability of a visit to appraiser n we take an expectation over the set of visit orders Π_N:

[7] Often, with some abuse of notation, σ_{null} will refer to a *generic* null experiment.

[8] Here, I ignore the distinction between n's beliefs about others' strategies and their actual strategies. As each appraiser acts at a single information set which is always on-path, this will not matter.

$$\sum_{\pi \in \Pi_N} \frac{1}{N!} \prod_{m:\pi(m)<\pi(n)} \left(1 - E_{\varphi_m}\left[E_{\sigma_m(\cdot|\theta)}[a_m]\right]\right). \tag{2}$$

Now, we wish to find appraiser n's *interim belief* - her belief that the prospect is quality $\bar{\theta}$ given that she encounters it. To find the interim belief, we weight by the prior and update on a visit to n occuring at all:[9]

$$p_n = \frac{p_0 \sum_{\pi \in \Pi_N} \prod_{m:\pi(m)<\pi(n)} \left(1 - E_{\varphi_m}\left[E_{\sigma_m(\cdot|\bar{\theta})}[a_m]\right]\right)}{E_{p_0}\left[\sum_{\pi \in \Pi_N} \prod_{m:\pi(m)<\pi(n)} \left(1 - E_{\varphi_m}\left[E_{\sigma_m(\cdot|\theta)}[a_m]\right]\right)\right]}. \tag{3}$$

The interim belief reflects the fact that other appraisers' strategies modulate the probability a prospect of each quality reaches n. We will show that in equilibrium, as appraisers accept good prospects more often than bad, all interim beliefs are (weakly) more pessimistic than the prior.

2.5 Payoffs

When appraiser n encounters a prospect of quality θ, that appraiser's probability of accepting given a pure strategy (σ_n, a_n) is

$$E_{\sigma_n(\cdot|\theta)}\left[a_n\right]. \tag{4}$$

On accepting a prospect of quality θ, an appraiser gets utility θ (rejections give zero). Thus, for any profile of strategies for other appraisers φ_{-n} and pure strategy of her own (σ_n, a_n), given a visit, appraiser n's payoff from accepted prospects is

$$E_{p_n}[\theta E_{\sigma_n(\cdot|\theta)}\left[a_n\right]] \tag{5}$$

where p_n is computed according to (3) using other appraisers' strategies.

Each experiment is assigned a cost, and this cost is allowed to vary with the experimenter's beliefs.[10] Formally, let $c : \mathcal{G} \times \text{int}(\Delta\Theta) \to \mathbb{R}_+$ denote the experiment

[9]Here, E_{p_0} is an abuse of notation; by this, I mean an expectation over qualities using the prior distribution, as parametrized by p_0.

[10]Experiments' costs varying with beliefs is reasonable, and in fact necessary, when experiments represent sequential procedures. For instance, when each flip of a coin has a constant cost, the

cost function, taking an experiment and a belief to a cost. It is assumed that, for any beliefs, experiments producing the same distribution over posterior likelihoods are equally costly, and that $c(\sigma, p) = 0$ iff σ is a null experiment.[11]

Thus, if appraiser n encounters a prospect and holds beliefs p_n, that appraiser computes her payoff from (σ_n, a_n) as

$$E_{p_n}[\theta E_{\sigma_n(\cdot|\theta)}[a_n]] - c(\sigma_n, p_n). \tag{6}$$

By computing p_n according to other appraisers' strategies φ_{-n} and (3), we can thus compute n's payoff on being visited given any strategy profile. As an appraiser gets 0 when she is not visited, and as being visited always has probability at least $1/N$, for the purpose of equilibrium it suffices to work with this interim utility function.

Following BZ, given an experiment $\sigma \in \mathcal{B}$ and a function that maps its outcomes into experiments, $\sigma' : S_\sigma \to \mathcal{B}$, I denote by $\sigma * \sigma'$ a *compound experiment* that performs σ first, and then for each outcome s performs the experiment $\sigma'(s)$. We say that c exhibits *indifference to sequential learning* if whenever σ and $\sigma'(s)$ for each s are in $\mathcal{G}$, and an experiment Blackwell-equivalent to $\sigma * \sigma'$ (wlog $\sigma * \sigma'$ itself) is also in $\mathcal{G}$, we have

$$c(\sigma * \sigma', p) = c(\sigma, p) + E_\sigma [c(\sigma'(s), p(s))] \tag{7}$$

where $p(s)$ is the posterior on Θ induced by experiment σ's outcome s.

Indifference to sequential learning requires that compound experiments cost as much as the sum of the expected costs of their constituent parts, where those costs are computed at the induced posteriors. For instance, if our venture capitalist chooses to conduct interviews first and then proceed with market analysis only if she likes what she hears, the cost of this procedure is equal to the cost of only the interview, plus the probability of 'liking what she hears' times the cost of only the market analysis at the updated beliefs.

Though BZ study indifference to sequential learning as a characteristic of UPS cost functions on $\mathcal{B}$, I will also find it handy when we talk about chunky information

procedure 'flip this coin until it lands on heads' has a total cost that is linear in beliefs about the coin's bias. Linear costs correspond to the 'Wald cost function' of Morris and Strack (2019). To accommodate yet more information costs, such as those based on Shannon Entropy, I go beyond linearity.

[11]It is well-known that experiments produce the same distribution over posterior likelihood functions iff they are Blackwell equivalent – see, for instance, Torgersen (1991).

acquisition. Briefly, it allows me to identify what appraisers do when they reach certain beliefs by experimenting, based on what is done by appraisers who start with those same beliefs.

2.6 Equilibrium

I regard as a solution to the game a psychological equilibrium strategy profile φ^*. That is, each φ^*_n maximizes appraiser n's expected payoff (6), given p_n is computed from φ^*_{-n}. Note that when c is linear in beliefs, the solution concept collapses to BNE.

2.7 Market Size

The main question I'll be asking is, how does the information available for acquisition affects the number of appraisers who can gainfully participate in the market? Accordingly, I should define the benchmarks for market size results.

I'll say that *information always limits entry* if there exists a number of appraisers $\bar{N}$ such that, if the number of appraisers is any $N > \bar{N}$, at most $\bar{N}$ of them can earn a nonzero expected payoff. Suppose a rough model of appraisers contemplating entry. When information always limits entry, even if entry is free, there are always equilibria with as few as $\bar{N}$ appraisers. Moreover, the number of entrants (in pure equilibria of the entry game) will be bounded by $\bar{N}$ if there is an entry cost, no matter how small.

I'll say that *information never limits entry* if, for any number of appraisers N, in every equilibrium of the game, every appraiser gets a positive expected payoff. When information never limits entry, every appraiser that can enter the market freely will. This requires that in all equilibria, each appraiser occasionally accepts the prospect.

These two cases, while not exhaustive, will suffice for my results.

3 Equilibria of divisible- and chunky-information markets.

First, let's agree to say the market is *trivial* if a null experiment maximizes (6) when p_n is set to p_0. A 'monopsonist' would not acquire information in a trivial market. She might choose to accept or reject, depending on p_0, but will do so without acquiring information. In such a market, there is always an equilibrium in which all appraisers use null experiments. Additionally, any equilibrium with experimentation would reflect pure coordination failure: appraisers acquiring information only because other appraisers are. Neither chasing equilibrium refinements nor producing policy that alleviates coordination failure is germane to the topic at hand; so we won't dwell on trivial markets.

Before we proceed to examine divisible and chunky information separately, we can immediately state a rather trivial general finding.

Lemma 1. *In any equilibrium of the game φ, for each n, $p_n \leq p_0$.*

Proof. From (3), the contrary would imply

$$\sum_{\pi \in \Pi_N} \prod_{m:\pi(m)<\pi(n)} \left(1 - E_{\varphi_m}\left[E_{\sigma_m(\cdot|\bar{\theta})}[a_m]\right]\right) > \sum_{\pi \in \Pi_N} \prod_{m:\pi(m)<\pi(n)} \left(1 - E_{\varphi_m}\left[E_{\sigma_m(\cdot|\underline{\theta})}[a_m]\right]\right). \tag{8}$$

This would in turn require that for some appraiser m, $E_{\sigma_m(\cdot|\underline{\theta})}[a_m] > E_{\sigma_m(\cdot|\bar{\theta})}[a_m]$, or, in other words, that bad visiting prospects are accepted by m more frequently than good ones. This is incompatible with optimization. □

With this in hand, we commence our search for equilibria, beginning with the case of divisible information.

3.1 Markets with divisible information

Divisible information, under which appraisers have the freedom to acquire any experiment in $\mathcal{B}$, most closely matches the models of information acquisition used in the literature. In particular, I will exploit what is known about Uniform Posterior-Separable (UPS) cost functions.

A cost function $c : \mathcal{B} \times (0,1)$ is UPS if there exists a strictly convex potential function V such that

$$c(\sigma_n, p_n) = E_{\sigma_n,p_n} V(q) - V(p_n) \tag{9}$$

where $q(s)$ is the posterior on Θ induced by the experimental outcome s.

The first step is to invoke BZ to establish that in the divisible information setting, under a mild continuity condition, c is UPS.[12]

[12] The exact type of continuity required is that when a sequence of experiments' distributions of induced posteriors weak*-converge, their cost converges to the cost of an experiment that produces the limiting posterior distribution.

Theorem (Bloedel and Zhong (2020)). *Suppose $\mathcal{G} = \mathcal{B}$, and that c is continuous and exhibits indifference to sequential learning. Then, c is Uniform Posterior-Separable.*

From here, we can identify the solution to the information acquisition problem for a single appraiser, for a given p_n. Lemmata 2, 3, and 4 retread some ground from the rational inattention literature (see, for instance, Caplin, Dean, and Leahy (2022)), adapted to the present setting. They are presented here both for completeness, and for utility to an application-minded audience. First, Lemma 2 shows that with a UPS cost function, optimally acquired information results in at most as many posteriors as there are actions - in our case, just two.

Lemma 2. *Suppose $\mathcal{G} = \mathcal{B}$ and that c is UPS. For a given interim belief p_n, an optimal experiment leads to at most two posteriors.*

A proof of Lemma 2 appears in Appendix A.1. Briefly, if two experimental outcomes inducing different posteriors led to the same action (acceptance or rejection), the appraiser could merge them. This would coarsen the experiment, and thus by the convexity of V cheapen it. This would not change the distribution of actions conditional on the quality θ; so the same distribution of outcomes is achieved, while spending less on information. Now, Lemma 3 can use this fact to show that the optimal experiment is essentially unique.

Lemma 3. *Suppose $\mathcal{G} = \mathcal{B}$ and that c is UPS. For a given interim belief p_n, the optimal experiment is unique (up to Blackwell-equivalence).*

The proof, in Appendix A.2, leverages the fact that the cost of a posterior distribution is mixture-linear. Thus, if there were two distinct optima, any combination of them would also be optimal. But such a combination must necessarily have at least three points in its support, contradicting Lemma 2.

Now, we show behavior is uniform - in terms of the posteriors of optimal experiments - even as p_n is allowed to vary. Formally, Lemma 4 shows that there is an interval of interim beliefs for which appraisers will acquire information, and that the purpose of optimally acquired information is to push posteriors to the endpoints of that same interval. This lemma retrieves the *Locally Invariant Posteriors* property shown to be common to UPS functions in Caplin, Dean, and Leahy (2022).

Lemma 4. *Suppose $\mathcal{G} = \mathcal{B}$, that c is UPS, and that the market is non-trivial. There exist beliefs q_a^*, q_r^* such that n's optimal experiment is non-null iff $p_n \in (q_a^*, q_r^*)$; if so, it produces only the posteriors q_a^*, q_r^*. An optimal acceptance policy is pure.*

The proof, to be found in Appendix A.3, simply shows that whenever the optimal experiment splits the interim belief into two posteriors, those two posteriors do not vary with the interim belief. Non-triviality then ensures there is an open interval of interim beliefs where this occurs. When the agent's interim beliefs do not lie in the interval, the agent will acquire no information.

Figure 1 illustrates how acceptance rates vary with interim beliefs under optimally-acquired information, as well as beliefs about the rejected. From Lemma 4, any information acquired by an appraiser produces a posterior q_r^* on a rejection. Thus, all rejections are marginal: an appraiser that rejects a prospect is just indifferent between acquiring more information about them or not. Furthermore, an appraiser that had as their interim beliefs exactly the rejection belief, $p_n = q_r^*$, would be *just on the fence* between acquiring some information or not. The fact that these thresholds coincide is crucial. It is a consequence of the fact UPS cost functions all satisfy indifference to sequential learning: an appraiser who just encountered a prospect and holds some interim belief q is in the same position as an appraiser who got to that belief by way of an experiment's realization.

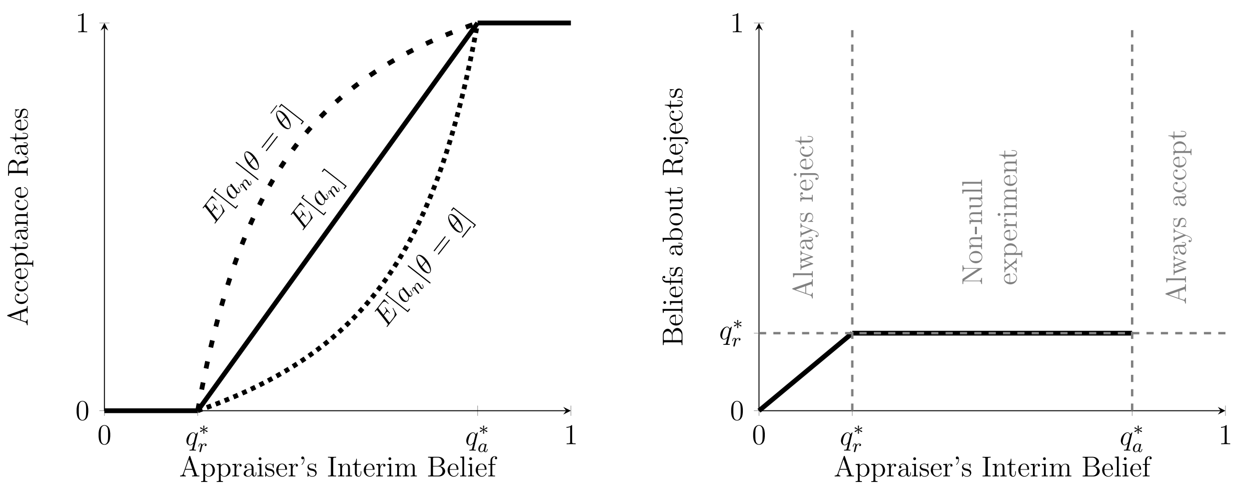

Figure 1: Optimal 'divisible information' acquisition. Left panel: acceptance rates for good and bad prospects, as well as the average rate, as a function of interim beliefs. Right panel: beliefs about rejected prospects, as a function of interim beliefs.

Thus, if the interim belief of an appraiser were *any higher* than the beliefs held by appraisers who acquired a 'bad' experiment realization, that appraiser would acquire information. The work done by the next proposition is to show that the possibility of a fresh prospect - one no other appraiser has screened yet - is always sufficient, in equilibrium, to produce interim beliefs for all appraisers above the q_r^* threshold. The reader is reminded that information never limits entry if, for any number of appraisers N, in every equilibrium of the game, every appraiser gets a positive expected payoff.

Proposition 1. *Suppose that $\mathcal{G} = \mathcal{B}$, that c is continuous, and that the market is not trivial. Then, information never limits entry. Moreover, in every equilibrium, every appraiser uses a non-null experiment.*

Proof. First, we call upon Bloedel and Zhong's Theorem to establish that c is UPS so we can use our lemmata. To begin, notice that using a null experiment and rejecting following any outcome of the experiment yields a payoff of 0. From Lemma 3 the optimal experiment is essentially unique. Therefore if in equilibrium an appraiser receives an expected payoff of 0, that appraiser must be acquiring a null experiment and always rejecting. So, the task at hand consists of finding that there are no such 'inactive' appraisers in equilibrium. We will proceed by contradiction.

First, suppose that there is an equilibrium φ of the game with some N appraisers featuring an inactive appraiser (wlog) N. An inactive appraiser never accepts, or in other words sets $E_\varphi[a_N|\theta] = 0$. Notice that the inactive appraiser drops out of (3). Thus, $(\varphi_n)_{n \leq N-1}$ is an equilibrium of the game with $N-1$ appraisers as for each appraiser $n < N$ the interim belief p_n computed from (3) is the same as when appraiser N is present, and thus the best responses of n coincide in the two scenaria. So, if the game with N appraisers features an equilirbium in which $k \geq 1$ are inactive, then the game with $N-(k-1)$ appraisers features an equilibrium in which just one is inactive. Therefore, for our purposes it will suffice to show that there is no equilibrium of the game for any N in which *one* appraiser is inactive.

Now, suppose there is an equilibrium φ of the game with some N appraisers such that appraiser (wlog) N is the sole inactive one. From non-triviality, $N > 1$. From the above, φ_{-N} is an equilibrium of the game with $N-1$ appraisers. Fix an appraiser $n \neq N$ and consider an auxiliary game with N appraisers in which we alter the visit order by forcing N to always come exactly after n. The conditioned distribution is, of course, a uniform distribution over $\{1, ..., n-1, (n, N), n+1, ..., N-1\}$.

I claim φ is an equilibrium of the auxiliary game. Optimality for $m \neq N$ comes from the fact that N drops out of the calculation of p_n and therefore incentives are as in the $N-1$ appraiser game. Optimality for N being inactive flows from the fact that N's interim belief is computable as p_n updated for a rejection by n, which coincide with n's beliefs when rejecting. From Lemma 1, $p_n < p_0$; from non-triviality $p_0 < q_a^*$; thus, n acquires a non-null experiment. Using Lemma 4, the beliefs of n on rejection are q_r^*. At these beliefs, indeed N finds it optimal to reject.

Now, notice the relation between N's beliefs in the auxiliary games and the original N-appraiser game while holding fixed the strategy profile φ. We can rewrite N's interim beliefs in the original game as a probability-weighted combination of her beliefs when she knows she comes right after n, for each n, and her beliefs when she knows she comes first. In the former cases her beliefs are q_r^*; if she knew she were first in the order, her beliefs would be p_0. As non-triviality implies that $q_r^* < p_0$, and coming first has probability at least $1/N$, we have that $p_N > q_r^*$; but this contradicts Lemma 2. □

Proposition 1 relies crucially on the fact that the visit order is uniform.[13] Thus, an inactive appraiser can be ignored in calculations by other appraisers. The fact that an inactive appraiser's interim beliefs updated on the event of acting after some other appraiser are exactly equal to that other appraiser's beliefs on rejection is contingent on uniformity.

This underlies the key argument in the proof: an inactive appraiser who knew she was acting after another appraiser would have interim beliefs equal to the rejection beliefs q_r^*. An appraiser who knew she's first in the order would have beliefs equal to the prior. An inactive appraiser who is uncertain about which is the case, should have beliefs equal to a convex combination of the two, and hence better than the rejection beliefs. Figure 2 illustrates this - for any number of other appraisers who are best-responding, the inactive appraiser has interim beliefs that make it profitable for her to acquire a non-null experiment. Thus, our inactive appraiser should be acting after all!

[13]There is a rather immediate counterexample to Proposition 1 when the visit order is not uniform. For any visit order distribution with a deterministic first appraiser, that appraiser acts as the monopsonist would, and all other appraisers have beliefs equal to q_r^* in equilibrium, and hence are inactive.

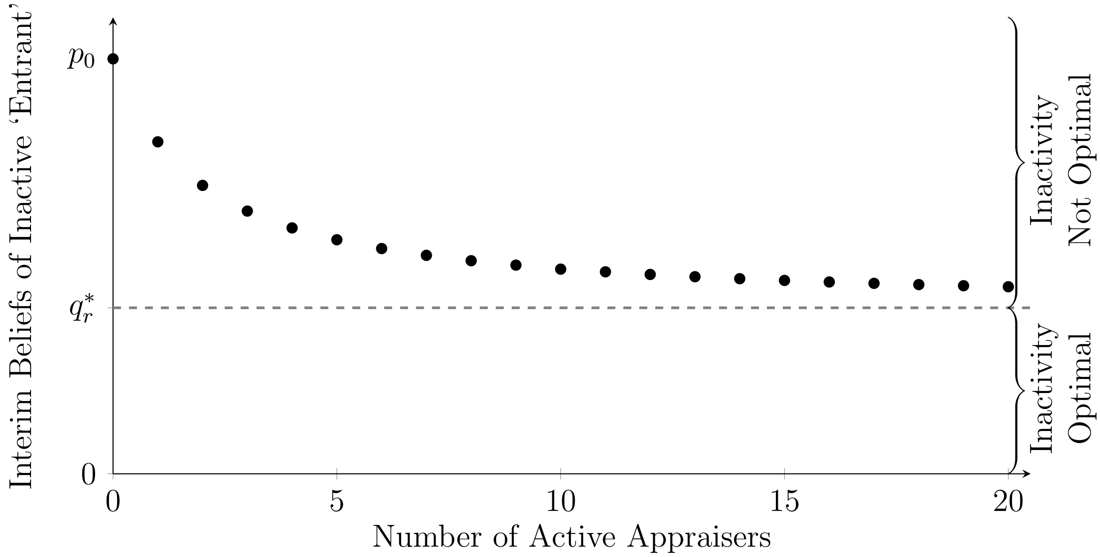

Figure 2: An inactive $(N+1)$st appraiser computing interim beliefs when others play a symmetric equilibrium of the N-appraiser game. For any N, the inactive appraiser's beliefs indicate inactivity is not, in fact, a best-response.

3.2 Markets with chunky information

To talk sensibly about information being "chunky" but retaining the experimental nature of the exercise, I want to ensure that appraisers are free to run experiments repeatedly or in whatever sequence they choose. To that end, we'll say that $\mathcal{G}$ is *procedurally closed* if whenever non-null σ^1 is in $\mathcal{G}$, and for each $s \in S_{\sigma^1}$, $\sigma^2(s)$ is in $\mathcal{G}$, we have that $\sigma^1 * \sigma^2 \in \mathcal{G}$.

Intuitively, this captures the notion that appraisers may always perform more tests, and the set of available experiments does not depend on prior experiments.[14]

Information is said to be *chunky* when $\inf\{c(\sigma, p_n) | \text{non-null } \sigma \in \mathcal{G}, p_n \in (0,1)\} > 0$, that is, when the cost of non-null experiments is bounded away from zero. This infimal non-null experimentation cost is denoted $\underline{c}$.

As an example, we could start with finitely many 'primitive' experiments and then repeatedly apply the definition of procedurally closed to obtain a procedurally closed, chunky $\mathcal{G}$. Any non-null experiment in $\mathcal{G}$, in this case, can be interpreted as the

[14]In principle, the results of this section extend even if we allow the menu to vary with beliefs, modifying our definition of procedural closure. Procedural closure would (only) require composition with experiments that are available at the induced posterior beliefs. The definition of chunkiness would also be edited to only check the cost of available experiments.

reduced form of a procedure that at each stage, depending on the history of experimental output, either orders another primitive experiment or terminates. This can model acquiring information via a lab which prices each medical test independently.

For a different example, one can start with a cost function c over $\mathcal{B} \times (0, 1)$ and a constant $\underline{c}$ and produce a restricted menu $\mathcal{B}_{\underline{c}}$ by choosing only those non-null experiments that c always assigns a cost of at least $\underline{c}$. In this case as well, arbitrarily cheap non-null experiments are not available; but, above the cutoff, experiments may well be designed very flexibly. An example is a consultant who can be hired to do all sorts of market research and charges by the minute, but won't accept tasks worth less than $\underline{c}$.

Unfortunately, procedural closure may only allow us to link the behavior of appraisers at very few collections beliefs. For instance, the set of interim beliefs at which an appraiser rejects outright could be disconnected. Intuitively, we'd like it to be the case that an optimistic appraiser accepts outright, a pessimistic appraiser rejects outright, and one in between pays for information. To this end, we are forced to make an assumption directly on behavior. We say that the market satisfies *connected behavior* if the set of interim beliefs at which outright rejection maximizes (6) and the set of beliefs at which outright acceptance maximizes (6) are closed intervals. In Appendix B, I show that several standard cases, including costs linear in beliefs, and UPS costs restricted to $\mathcal{B}_{\underline{c}}$, satisfy connected behavior.[15]

Lemma 5 establishes that any profiting appraiser is acquiring a non-null experiment, and therefore imposes an externality on other appraisers.

Lemma 5. *In every equilibrium of a non-trivial market with chunky information and connected behavior, any appraiser making a profit acquires a non-null experiment.*

Proof. From Lemma 1 we have that that for each n, $p_n \leq p_0$ in every equilibrium. From non-triviality and connected behavior, we have that $p_0 < \bar{p}$, where $\bar{p}$ is the least belief at which outright acceptance is optimal. Together, they give us $p_n < \bar{p}$ in every equilibrium. Thus, in equilibrium either $p_n \leq \underline{p}$ and n gets a payoff of zero (as immediate rejection is optimal), or $\underline{p} < p_n \leq p_0 < \bar{p}$ and acquiring a null experiment and rejecting is *not* optimal. Thus, the appraiser's payoff in this second case must be greater than 0. So, whenever the appraiser's expected payoff is positive in equilibrium, a null experiment is not optimal. □

[15]In particular, linear - or Wald - costs include what I regard as the canonical case: costs that vary only with the actual quality of the prospect.

Together, the assumption of connected behavior and Lemma 5 will allow us to show that, as the number of appraisers grows, along any sequence of equilibria in which the acceptance probability of the market as a whole goes to 1, eventually some appraiser must be accepting when her beliefs on acceptance are below $\bar{p}$. This would imply that she has a profitable deviation due to indifference to sequential learning and the procedural closure of the information menu. On the flip side, if the acceptance probability of the market as a whole limits to something other than 1 along a sequence of equilibria in which all appraisers profit, Lemma 5 implies they must be acquiring information. Then, the chunkiness of information implies the expected market-wide total expenditure on experiments is going to infinity, so that some appraiser is getting a negative payoff and should deviate to being inactive. Finally, a sequential compactness argument shows one kind of subsequence or another exists if there are equilibria in which arbitrarily large numbers of appraisers profit. Thus, such a sequence of equilibria cannot exist, and there must be a maximal number of profitable appraisers; in other words, *information always limits entry.* Proposition 2 formalizes this argument.

Proposition 2. *Suppose that $\mathcal{G}$ is procedurally closed, that information is chunky, and that the market is not trivial and exhibits connected behavior. Then, information always limits entry.*

Proof. To prove the proposition, I proceed by contradiction. Suppose that there is a strictly increasing sequence of natural numbers $(N_i)_{i\in\mathbb{N}}$ such that for each i, there is an equilibrium φ^i of the game with some $N_i' \geq N_i$ appraisers where N_i appraisers profit. For each such equilibrium, we consider the *total rejection rate* for each quality θ. The probability that every appraiser rejects a prospect of the given quality is

$$R_i(\theta) = \prod_{n=1}^{N_i'} \left(1 - E_{\varphi_n^i}\left[E_{\sigma_n^i(\cdot|\theta)}[a_n^i]\right]\right). \tag{10}$$

For all i, the profile of total rejection rates $R_i = (R_i(\bar{\theta}), R_i(\underline{\theta}))$ is in $[0,1]^2$, which is a sequentially compact space. Hence, there is some accumulation point R^∞ which is the limit point of some subsequence $(R_{i_k})_{k\in\mathbb{N}}$. I will show that whatever this limit is, we can find some φ^{i_k} that is not an equilibrium.

First, suppose $R^\infty \neq (0,0)$. Then, there is some $\varepsilon > 0$, some $\theta \in \Theta$, and some $k^* \in \mathbb{N}$ such that for all $k \geq k^*$, $R_{i_k}(\theta) > \varepsilon$. Thus, with enough appraisers, there is some quality that is rejected by all appraisers with probability at least ε. Quality θ is assigned probability at least $\min\{p_0, 1-p_0\}$ by the prior. Choose some $k > k^*$ that satisfies

$$N_{i_k} > \frac{\bar{\theta}}{\varepsilon \min\{p_0, 1-p_0\} \cdot \underline{c}}. \tag{11}$$

Thus, there is at least a $\varepsilon \min\{p_0, 1-p_0\}$ chance that the prospect is quality θ and visits every appraiser in the equilibrium φ^{i_k}. From Lemma 5 every appraiser with a positive utility must conduct a non-null experiment, so the expected cost of all experiments in the market is at least $N_{i_k} \cdot \varepsilon \min\{p_0, 1-p_0\} \cdot \underline{c}$. On the other hand, the highest expected gains from acceptances can be across the market is $\bar{\theta}$. Using (11), $N_{i_k} \cdot \varepsilon \min\{p_0, 1-p_0\} \cdot \underline{c} > \bar{\theta}$, and therefore the total expected payoffs in the equilibrium are negative. Thus, at least one appraiser has negative expected utility under φ^{i_k}; deviating to null experimentation and rejection, she could get 0. Therefore φ^{i_k} could not be an equilibrium, and hence we reject the hypothesis that $R^\infty \neq (0,0)$.

Hence, suppose now that $R^\infty = (0,0)$ - along this sequence of equilibria, eventually nearly every prospect is accepted by some appraiser. In the market with N'_{i_k} appraisers, in equilibrium φ^{i_k}, the proportion of prospects accepted by the market as a whole that are of quality $\bar{\theta}$ is

$$p_a^{i_k} = \frac{p_0(1 - R_{i_k}(\bar{\theta}))}{p_0(1 - R_{i_k}(\bar{\theta})) + (1-p_0)(1 - R_{i_k}(\underline{\theta}))}. \tag{12}$$

Consider the beliefs at which acceptance occurs given φ^{i_k}; $p_a^{i_k}$ is their average. As $k \to \infty$, $(R_{i_k}(\underline{\theta}), R_{i_k}(\bar{\theta})) \to (0,0)$; thus $p_a^{i_k} \to p_0$. From non-triviality and the assumption of connected behavior, we have $p_0 < \bar{p}$, where $\bar{p}$ is the least belief at which outright acceptance is optimal. Thus from $p_a^{i_k} \to p_0$ there is a $k \in \mathbb{N}$ such that in the game with N'_{i_k} appraisers, there is an equilibrium in which some appraiser n accepts with positive probability and, on average, at a belief below $\bar{p}$; call this belief p^*. From Lemma 5, an appraiser with beliefs p^* would not find a null experiment optimal; instead, call (σ^*, a^*) a strategy that outperforms null experiments. Now, consider a deviation for n that compounds experiment σ^* on any outcome of n's original experiment that leads to acceptance, and instead uses acceptance rule a^* at the conclusion of this compounded component. From the fact $\mathcal{G}$ is procedurally closed, this compound experiment is available. From the indifference to sequential learning assumption, this deviation is profitable. Thus we are again led to a contradiction. □

The proposition implies that regardless of the available information, so long as the information menu $\mathcal{G}$ does not have arbitrarily cheap experiments on offer, the solicitation bias eventually prevents the entry of additional appraisers. We may, in some sense, regard $\bar{N}$ as the natural size of the market - the largest number of profitable appraisers.

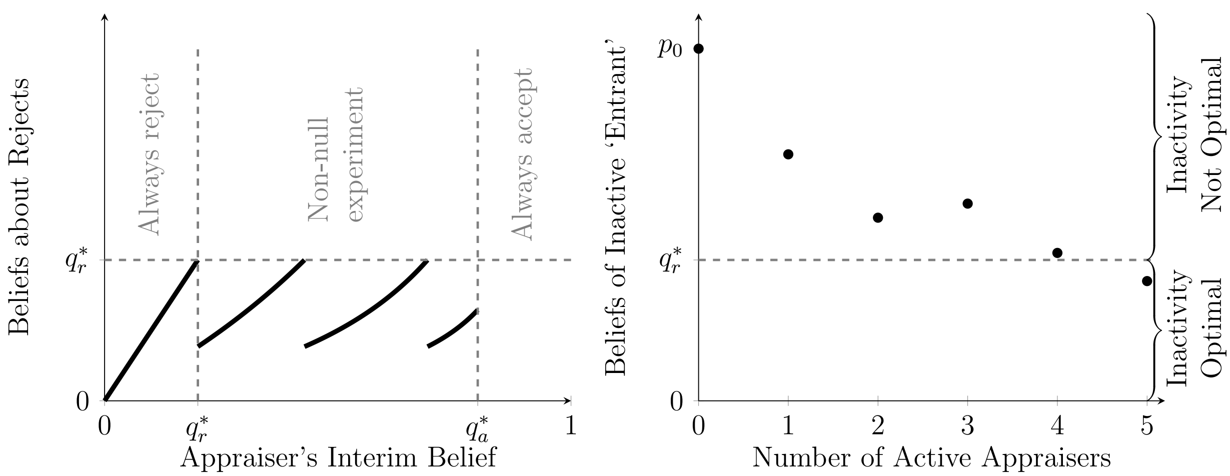

Figure 3: Chunky information example. Left panel: beliefs about rejected prospects, as a function of interim beliefs. Right panel: An inactive $(N + 1)$st appraiser computing interim beliefs when others play a symmetric equilibrium of the N-appraiser game. Five appraisers are enough to keep the sixth inactive.

Figure 3 illustrates the 'excessively negative' posteriors in a chunky information market. When appraisers reject, their beliefs at that stage are below q_r^*. As the number of appraisers increases, this eventually leads to equilibria with inactive appraisers.

It is important to note that this is not a 'congestion' result in any way. An appraiser that is not visited never pays for information; all costs are effectively marginal, as an appraiser has exclusive transaction rights at the time of information acquisition. Therefore, the number of competitors only matters via adverse selection and not congestion. Moreover, this result hinges crucially on the procedural closure of $\mathcal{G}$; see Section 4.3 for a counterexample to the result when procedural closure fails.

4 Discussions and extensions

4.1 Reconciling the two Propositions

How are Proposition 1 and Proposition 2 to be reconciled? The former states that divisible information never limits entry: all appraisers, no matter their number, profit. The latter states that if information is the least bit chunky, information always limits entry: all but a small number of appraisers are excluded by the informational externality. We will see how these results are related by starting with a chunky menu, and progressively making it more divisible. For a given number of appraisers N, eventually equilibria in which all appraisers profit appear, and equilibria in which appraisers are excluded disappear.

Start with a market $(c, \mathcal{B})$ satisfying the assumptions of Proposition 1 and also that for all non-null $\sigma \in \mathcal{B}$, $\inf_p(\sigma, p) > 0$. Now, take a positive monotonically decreasing sequence $(k_i)_{i\in\mathbb{N}}$ with $lim_{i\to\infty}k_i = 0$. Let

$$\mathcal{G}_i = \{\sigma \in \mathcal{B} | \sigma \text{ is null or } \inf_p c(\sigma, p) \geq k_i\}. \tag{13}$$

Each $\mathcal{G}_i$ is by construction procedurally closed and therefore satisfies the assumptions of Proposition 2. Also, whenever $j > i$, $\mathcal{G}_i \subset \mathcal{G}_j$ from monotonicity. Therefore, for any $\sigma \in \mathcal{B}$, there is an i such that $\sigma \in \mathcal{G}_j$ for all $j > i$. We'll be fixing a number of appraisers N and inspecting equilibria in the game with $(c, \mathcal{G}_i)$ as i increases.

Returning to $(c, \mathcal{B})$, from Proposition 1 any equilibrium of the game with $N \in \mathbb{N}$ appraisers φ^* has each appraiser using a strategy in which she undertakes non-null, and hence costly, experimentation. Therefore, there is an i after which φ^* is an equilibrium strategy profile in $(c, \mathcal{G}_i)$. Therefore, for any *fixed* number of appraisers N, if the information is *divisible enough* (even if not perfectly divisible) there is an equilibrium in which all N appraisers profit.

Moreover, fix an $i \in \mathbb{N}$ and suppose the game with $(c, \mathcal{G}_i)$ and N appraisers has an equilibrium φ' in which some appraiser receives a payoff of 0. From Proposition 1, φ' is not an equilibrium in the game given by $(c, \mathcal{B})$. As a consequence, there is some appraiser $m \leq N$, a $\sigma_m \in \mathcal{B} \setminus \mathcal{G}_i$ and some acceptance rule a_m, such that (σ_m, a_m) would be a profitable deviation for m if σ_m were available to m. But recall that $\inf_p c(\sigma_m, p) > 0$. Thus, there is a j so that $l \geq j \implies \sigma_m \in \mathcal{G}_l$. Hence, such an equilibrium disappears when information becomes divisible enough.

4.2 Generalizing Propositions 1 and 2

One may wonder in what ways we can generalize the statements of Propositions 1 and 2. The key to the answer is that each uses indifference to sequential learning in a different way, whereas neither proposition alone needs its full strength.

Proposition 1 relies on the fact that with smooth information, if there is an active appraiser whose rejection belief lies below another inactive interim belief, that second appraiser will optimally acquire a non-null signal. This comes from the fact that we can write the first appraiser's experiment as a compound one, whose second component demands a non-null experiment at the second appraiser's beliefs. Indifference to sequential learning then shows this experiment is as cheap for the second appraiser, and therefore optimal. The implication goes through if we make this second experiment cheaper to the second appraiser still, by replacing indifference to sequential learning with BZ's preference for sequential learning:

$$c(\sigma * \sigma', p) \geq c(\sigma, p) + E_\sigma \left[c(\sigma'(s), p(s)) \right]. \tag{14}$$

In this setting, this plays a role akin to 'decreasing returns to scale' for information generation.[16] Under this prism, it makes sense that the market can always support more entrants, *so long as information is available in fine enough increments.*

On the other hand, Proposition 2 relies on the fact that an appraiser will never acquire an experiment that induces any posterior q^* such that a null experiment is not optimal for an appraiser with an interim belief of q^*. This is because indifference to sequential learning implies that the original appraiser could then tack on an additional experiment that is more profitable than a null one. Plainly, an appraiser wants to continue experimenting at any belief at which it would be optimal to begin experimenting. In other words, the proposition uses BZ's preference for one-shot learning

$$c(\sigma * \sigma', p) \leq c(\sigma, p) + E_\sigma \left[c(\sigma'(s), p(s)) \right]. \tag{15}$$

rather than the full force of indifference to sequential learning. With this in place, we might relax the definition of chunky information. Rather than non-null 'cheap enough' experiments not being available, it suffices that non-null *cheap enough experiments are*

[16]This is not the same notion of returns to scale as in **??**.

never optimal at any interim belief. It could be, for instance, that the informativeness of experiments goes to zero faster than the cost.[17]

4.3 Chunkiness without procedural closure

This section illustrates how procedural closure is key to Proposition 2. Procedural closure requires that if an experiment is available to an appraiser, so is running that experiment, then conducting another experiment, or not, depending on the outcome. To show that it's essential, we'll study a market in which there is only one non-null experiment available.

Take a market in which $\mathcal{G} = \{\sigma_{\text{null}}, \sigma\}$ where σ has two outcomes $\{0, 1\}$, with $\sigma(s = 1|\theta = 1) = \sigma(s = 0|\theta = 0) = k > .5$. Assume $c(\sigma, \cdot)$ is constant; the experiment's cost does not vary with beliefs. Then the interim beliefs at which it is optimal to acquire σ and condition acceptance or rejection on it satisfy

$$\max\{p_n\bar{\theta} + (1 - p_n)\underline{\theta}, 0\} \leq p_n k\bar{\theta} + (1 - p_n)(1 - k)\underline{\theta} - c. \tag{16}$$

so that an appraiser with interim beliefs $p_n \in [\frac{-(1-k)\underline{\theta}+c}{k\bar{\theta}-(1-k)\underline{\theta}}, \frac{-k\underline{\theta}-c}{(1-k)\bar{\theta}-k\underline{\theta}}]$ will acquire information, with strict incentives in the interior. The interval has an interior iff the market satisfies non-triviality.

Now, suppose there are N appraisers and each acquires σ and conditions acceptance on the outcome. By (3), the interim beliefs of each appraiser n are given by

$$p_n^{(N)} = \frac{p_0 \sum_{m=0}^{N-1}(1-k)^m}{p_0 \sum_{m=0}^{N-1}(1-k)^m + (1-p_0)\sum_{m=0}^{N-1} k^m}. \tag{17}$$

Information never limits entry if and only if $p_n^{(N)} \geq \frac{-(1-k)\underline{\theta}+c}{k\bar{\theta}-(1-k)\underline{\theta}}$ for all N. As

$$p_n^{(N)} - p_{n+1}^{(N)} = \frac{\frac{p_0(1-p_0)}{k(1-k)}\left[(1-k)k(k^{n-1} - (1-k)^{n-1}) + k^n(1-k)^n(2k-1)\right]}{\left(p_0\frac{1-(1-k)^n}{k} + (1-p_0)\frac{1-k^n}{1-k}\right)\left(p_0\frac{1-(1-k)^{n+1}}{k} + (1-p_0)\frac{1-k^{n+1}}{1-k}\right)}$$

is always positive (as $k > .5$), $p_n^{(\cdot)}$ is monotonically decreasing. Thus, to check if information never limits entry, it suffices to inspect $\lim_{N\to\infty} p_n^{(N)}$. Using (17) we have

[17] For a particular formalization that makes the cost function differentiable, we might state this as the derivative of the cost function exploding around the null experiment.

that $\lim_{N\to\infty} p_n^{(N)} = \frac{\frac{p_0}{k}}{\frac{p_0}{k}+\frac{(1-p_0)}{1-k}}$. Thus, the relevant comparison is

$$\frac{\frac{p_0}{k}}{\frac{p_0}{k}+\frac{(1-p_0)}{1-k}} \lessgtr \frac{-(1-k)\underline{\theta}+c}{k\bar{\theta}-(1-k)\underline{\theta}}. \tag{18}$$

Depending on the parameters, this inequality can go either way. Thus, absent procedural closure, the general features of the information acquisition technology alone do not determine the market structure.

4.4 Chunky commitments

Consider an appraiser with the ability to commit to acquiring particular information. This may be the result of a permanent information acquisition unit, or a contract with a third party. Such an appraiser may well benefit from committing to acquiring excessive information. The reason is simple: in a symmetric equilibrium with N appraisers, the most each appraiser can get is $1/N$ times the monopsonist profit. But committing to over-acquire information can keep other appraisers out. This example illustrates the point.

Suppose $\Theta = \{-1, 1\}$ with a prior $p_0 > .5$. Take a binary experiment with precision k: $\sigma(s = 1|\theta = 1) = \sigma(s = -1|\theta = -1) = k$, and set its cost to be belief-independent c. Take $\mathcal{G}$ to contain a null experiment, σ, and its procedural closure.

Breaking with the 'non-triviality' condition, here let's assume that immediate acceptance is preferable to acquiring σ before making a decision: $\hat{p}\bar{\theta} + (1-\hat{p})\underline{\theta} > \hat{p}k\bar{\theta} + (1-\hat{p})(1-k)\underline{\theta} - c$ for each $\hat{p} \geq .5$. As both sides are linear, this reduces to

$$.5 > k - c. \tag{19}$$

This makes it so that acquiring σ is, in fact, never a best-response in the simultaneous game, as it is preferable to accept with a null signal above a belief of .5 and preferable to reject with a null signal below .5. Assuming that $c = 1/8$ will make it so that it is never optimal to acquire a compound experiment, either. In equilibrium without an appraiser who can commit, it is immediate that every appraiser acquires a null experiment and accepts. Equilibrium beliefs are equal to the prior for each appraiser. Each appraiser thus reaps a payoff of $\frac{1}{N} \cdot (p_0 \cdot \bar{\theta} + (1-p_0)\underline{\theta}) = \frac{2p_0-1}{N}$ in equilibrium.

By committing to acquiring σ, and then accepting following a good outcome and rejecting otherwise, could appraiser 1 benefit? For this to be preferable to no commitment, it has to be that

$$p_0 k - (1-p_0)(1-k) - c > \frac{2p_0 - 1}{N}. \tag{20}$$

To see if there is an equilibrium in which appraisers $\{2, .., N\}$ reject all prospects, we first compute the induced interim belief p_n of appraiser $n \neq 1$,

$$p_n = \frac{p_0(1+(1-k))}{p_0(1+(1-k)) + (1-p_0)(1+k)}. \tag{21}$$

For n's immediate rejection to be a best-response, we need that $p_n\bar{\theta} + (1-p_n)\underline{\theta} \leq 0$, or

$$p_0(2-k) - (1-p_0)(1+k) \leq 0 \tag{22}$$

which reduces to simply

$$p_0 \leq \frac{1+k}{3}. \tag{23}$$

Thus, to summarize, by committing to acquiring experiment σ, an appraiser can keep other appraisers inactive if equation (23) holds. When (20) holds, it is profitable for an appraiser with commitment power to do this.

Therefore, it is possible for an appraiser with commitment power to become a monopsonist by acquiring *too much* information, enough that the informational externality keeps other appraisers out. This is a novel form of entry deterrence, and should give policymakers pause. In this example the market without commitments was more efficient, but this is not necessarily the case in general, complicating the matter.

5 Conclusion

I hope I have illustrated that the fine details of information acquisition matter. Even though information bears only interim and not ex-ante costs, the viscosity of information determines the emergent market structure. I show that agents constrained to chunky information produce 'excessive' informational externalities, and that in my setting this leads to informational oligopolies. This should make us think more broadly about informational mechanisms that stymie market entry, and the sensitivity of our models to the details of related assumptions.

Additionally, I bring recent work on information acquisition to bear on an applied setting and achieve some qualitative results without large compromises on abstraction. Admittedly, the setting studied is quite stylized, lacking such ordinary things as prices. Although this is a shortcoming, it allows for a clearer focus on interdependent information acquisition by shutting down strategic responses in other domains.

Bibliography

Dirk Bergemann, Xianwen Shi, and Juuso Välimäki. Information acquisition in interdependent value auctions. *Journal of the European Economic Association*, 7(1): 61–89, 2009.

Alexander W Bloedel and Weijie Zhong. The cost of optimally-acquired information. *Unpublished Manuscript, November*, 2020.

Andrew Caplin and Mark Dean. Behavioral implications of rational inattention with shannon entropy. Technical report, National Bureau of Economic Research, 2013.

Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130(6): 1676–1715, 2022.

Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. Experimental cost of information. Working Papers 657, IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University, 2019.

Jeffrey C. Ely and Ron Siegel. Adverse selection and unraveling in common-value labor markets. *Theoretical Economics*, 8(3):801–827, 2013.

John Geanakoplos, David Pearce, and Ennio Stacchetti. Psychological games and sequential rationality. *Games and economic Behavior*, 1(1):60–79, 1989.

Alex Gershkov and Balázs Szentes. Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, 144(1):36–68, 2009. ISSN 0022-0531. doi: https://doi.org/10.1016/j.jet.2008.02.004. URL https://www.sciencedirect.com/science/article/pii/S002205310800046X.

Kyungmin Kim and Marilyn Pease. Costly search with adverse selection: solicitation curse versus acceleration blessing. *The RAND Journal of Economics*, 48(2):526–553, 2017.

Stephan Lauermann and Asher Wolinsky. Search with adverse selection. *Econometrica*, 84(1):243–315, 2016.

Jeffrey Mensch. Cardinal representations of information. Working papers, 2018.

Stephen Morris and Philipp Strack. The wald problem and the relation of sequential sampling and ex-ante information costs. *Available at SSRN 2991567*, 2019.

Nicola Persico. Information acquisition in auctions. *Econometrica*, 68(1):135–148, 2000.

Doron Ravid, Anne-Katrin Roesler, and Balázs Szentes. Learning before trading: On the inefficiency of ignoring free information. *Journal of Political Economy*, 130(2): 346–387, 2022. doi: 10.1086/717350. URL https://doi.org/10.1086/717350.

Erik Torgersen. *Comparison of statistical experiments*. Cambridge University Press, 1991.

A Proofs Omitted from the Text

A.1 Proof of Lemma 2

Proof. Denoting by $\alpha : [0,1] \to \{0,1\}$ the acceptance policy as a function of induced posteriors, and by $F \in \Delta[0,1]$ the appraiser's distribution over posteriors q induced by information acquisition, the appraiser's problem can be rewritten as

$$\max_{F\in\Delta[0,1],\ \alpha:[0,1]\to\{0,1\}} E_F[\alpha(q)(q\bar{\theta} + (1-q)\underline{\theta}) - (V(q) - V(p_n))] \tag{24}$$

$$\text{subject to } E_F[q] = p_n. \tag{25}$$

Optimally, we have $\alpha(q) = 0$ if $q \leq -\frac{\underline{\theta}}{\bar{\theta}-\underline{\theta}}$ and $\alpha(q) = 1$ if $q \geq -\frac{\underline{\theta}}{\bar{\theta}-\underline{\theta}}$. Thus we can rewrite the problem of choosing F as

$$\max_{F\in\Delta[0,1]} E_F[\max\{q\bar{\theta} + (1-q)\underline{\theta}, 0\} - V(q)] \tag{26}$$

$$\text{subject to } E_F[q] = p_n. \tag{27}$$

Suppose F's support contains more than a single $q \leq -\frac{\underline{\theta}}{\bar{\theta}-\underline{\theta}}$. In this region, the maximand is strictly concave in q, and thus reallocating all this mass to $E_F[q|q \leq -\frac{\underline{\theta}}{\bar{\theta}-\underline{\theta}}]$ is strictly preferable. The same applies for those $q \geq -\frac{\underline{\theta}}{\bar{\theta}-\underline{\theta}}$. Thus an optimal experiment either induces two posteriors, one leading to acceptance q_a and one to rejection q_r, satisfying $q_a > p_n > q_r$, or the experiment is null and produces only one posterior, p_n. □

A.2 Proof of Lemma 3

Proof. Suppose F^*, $F^{*\prime}$ are optimal and not identical. Let $\hat{F}$ be any mixture of F^* and $F^{*\prime}$. By the fact the maximand (28) and constraint (25) are linear in F, the distribution $\hat{F}$ is feasible and also achieves the maximum. But such a mixture will feature at least three points in its support, a contradiction to Lemma 2. Thus the maximizer F must be unique. Any two experiments generating the same distribution of posteriors are Blackwell-equivalent, so the optimal experiment is unique up to Blackwell-equivalence. □

A.3 Proof of Lemma 4

Proof. Suppose the optimal experiment is non-null. From Lemma 2 an optimal non-null experiment splits p_n into an acceptance posterior q_a^* and a rejection posterior q_r^*. Let's manipulate the constraint, $E_F[q] = p_n$. For an experiment producing only two posteriors q_a and q_r with probabilities z and $1-z$ respectively, the constraint reduces to $zq_a + (1-z)q_r = p_n$. Solving for z, we get that the probability of a signal corresponding to posterior q_a is $z = \frac{p_n - q_r}{q_a - q_r}$. Thus, we can rewrite the objective as

$$\max_{q_a,q_r} \left[\frac{p_n - q_r}{q_a - q_r}[q_a\bar{\theta} + (1-q_a)\underline{\theta}] - \frac{p_n - q_r}{q_a - q_r}V(q_a) - \frac{q_a - p_n}{q_a - q_r}V(q_r)\right]. \tag{28}$$

Differentiating, we arrive at the first-order conditions

$$-q_r^*\bar{\theta} - (1 - q_r^*)\underline{\theta} + V(q_a^*) - V(q_r^*) - (q_a^* - q_r^*)V'(q_a^*) = 0 \tag{29}$$

$$-q_a^*\bar{\theta} - (1 - q_a^*)\underline{\theta} + V(q_a^*) - V(q_r^*) - (q_a^* - q_r^*)V'(q_r^*) = 0 \tag{30}$$

which are independent of p_n. Furthermore, the SOSC confirms that if a solution exists and is feasible (requiring $p_n \in (q_r^*, q_a^*)$), it is a maximum. A corner solution (with $q_r^* = 0$, $q_a^* = 1$, or both) may instead obtain; but from Lemma 3 the solution is unique. If neither an interior or a corner optimum exists, a null experiment is optimal. Notice that if a solution (q_r^*, q_a^*) does not exist, then a null experiment is optimal for *any* $p_n \in [0, 1]$. Thus even a monopsonist with $p_n = p_0$ would decline to acquire information. This contradicts non-triviality, so there exists a solution q_a^*, q_r^*, proving the first part of the lemma.

For the second part of the lemma, it suffices to see that since from non-triviality $0 \leq q_r^* < p_0 < q_a^* \leq 1$, (i) if $p_n \leq q_r^*$ then a null experiment is chosen and the acceptance policy rejects; (ii) if $p_n \geq q_a^*$ a null experiment is chosen and the prospect is accepted; and (iii) if $q_r^* < p_n < q_a^*$ then a non-null experiment which only induces the posteriors q_r^*, q_a^* is chosen with rejection at the former posterior and acceptance at the latter. □

B Connected Behavior

In this appendix, I show that two prominent cases of chunky information satisfy connected behavior, as introduced in section 3.2.

B.1 Linear costs lead to connected behavior

Suppose that for any $\sigma \in \mathcal{G}$, the cost $c(\sigma, p_n)$ is linear in beliefs p_n. A strategy of acquiring a null experiment and accepting is optimal when the interim belief puts probability 1 on good quality. This yields a payoff of (near) $\bar{\theta}$. Conversely, at this belief any other pure strategy can give at most 0 if it acquires a null experiment and rejects, and $\bar{\theta} - \underline{c}$ if it involves a non-null experiment. From this and the fact that for any strategy φ_n, appraiser n's payoff (6) is linear in p_n, the set of beliefs where acquiring a null experiment and accepting is optimal is some interval $[\bar{p}, 1]$. Similarly, acquiring no information but rejecting is optimal on some interval $[0, \underline{p}]$. It might still be the case that $\underline{p} = \bar{p}$. Non-triviality asserts that p_0 is a belief in which acquiring a null experiment is not optimal; as a consequence, $\underline{p} < p_0 < \bar{p}$. Thus, it is strictly better to acquire a non-null experiment for beliefs in $(\underline{p}, \bar{p})$.[18]

B.2 UPS costs with a $\underline{c}$-restricted menu lead to connected behavior

As above, at low enough and high enough beliefs, a null experiment is optimal. Now, to show that the set of beliefs at which a non-null experiment is acquired is convex. First, notice that maximizing (6) under $c(\sigma, p_n) > \underline{c}$ still results in two posteriors for a UPS c. The reason for this is simple: if c is UPS, so is $(\lambda + 1)c$, where λ is the Lagrange multiplier on the constraint.

Let $p_i < p_n$ and suppose that non-null experiment σ^i is acquired at p_i. Let the induced posteriors be q_a^{*i}, q_r^{*i}.

Suppose, for contradiction, that it is optimal to acquire a null experiment and reject at p_n. Then, $q_a^{*i} > p_n$ or else the supposed strategy at p_i would give negative payoff. We can factor σ^i into an experiment σ^{i1} that splits p^i into q_r^{*i} and p^n, and an experiment σ^{i2} that splits p^n into q_a^{*i} and q_r^{*i}. From indifference to sequential learning, the composition experiment is assigned a cost by c strictly more than $\sigma^{i1}(p^n)$ times $c(\sigma^{i2}, p_n)$. Then, at p_n running σ^{i2} (or, if $c(\sigma^{i2} < \underline{c}$, an experiment that costs $\underline{c}$ that is Blackwell-more informative) and accepting at the posterior q_a^{*i} yields strictly more payoff than the indicated strategy at p_i, which in turn must be weakly positive. Thus, outright rejection is not optimal at p_n. A similar argument shows that if $p_n < p_k$ and it is optimal to accept outright at p_n, the same is true at p_k.

[18] However, we have not made assumptions guaranteeing the existence of optimal experiments in this range. This is part of why Proposition 2 concerns properties of equilibria, not whether they exist.