The London School of Economics and Political Science

Essays in Information Economics

A thesis submitted to the Department of Economics
for the degree of Doctor of Philosophy

Alkiviadis Georgiadis-Harris

May 15^{th}, 2022

Acknowledgements

First and foremost I would like to deeply thank my advisor Balázs Szentes for his inexhaustible guidance, support, and patience. There is no doubt in my mind that without him my experience at LSE would have been painfully lacking.

I am also hugely indebted to Francesco Nava for supporting me throughout my time here, offering an abundance of advice and a wealth of good times. Also, to Matt Levy for always being willing to ponder about my questions and ideas, and for his invaluable help in preparing for the job-market.

I am also grateful to Deniz Aydin, Carlo Cabrera, Umut Cetin, Daniele Condorelli Angelos Diamantopoulos, Jason Roderick Donaldson, Andrew Ellis, Daniel Gottlieb, Maxi Guennewig, Gilat Levy, William Matcham, John Moore, Ronny Razin, and Christopher Sandmann for their many comments and discussions. Special thanks to Thomas Brzustowski for always offering me his strikingly thoughtful comments and advice.

Lastly, I would like to thank my partner Christiana, and my parents for making any of this possible.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 29,345 words.

Statement of conjoint work

I confirm that Chapter 2 was jointly co-authored with Thomas Brzustowski (LSE), and Balázs Szentes (LSE), and I contributed 33% of this work.

I confirm that Chapter 3 was jointly co-authored with Maximilian Guennewig (University of Bonn), and I contributed 50% of this work.

Abstract

This thesis comprises of three chapters.

The first chapter analyzes a dynamic model of information acquisition by a Bayesian decision-maker. In the model the decision-maker has full flexibility in choosing the fine qualitative features of the information they acquire, but faces a constraint on the 'per-period' quantity of information they can generate. Moreover, the decision-maker lacks control over the timing of their actions. At the optimum, the decision maker concentrates resources in generating a single piece of breakthrough news, contradicting their plan of action. In the absence of such news, the decision maker becomes more confident in their intentions. This leads them to sacrifice the frequency with which breakthroughs arrive in order to increase their impact on choice behaviour.

The second chapter (with T. Brzustowski, and B. Szentes), reconsiders the problem of a durable-good monopolist who cannot make intertemporal commitments. The buyer's valuation is binary and his private information. The seller has access to dynamic contracts and, in each period, decides whether deploy the previous period's contract or to replace it with a new one. Our main result is that the Coase Conjecture fails: the monopolist's payoff is bounded away from the low valuation irrespective of the discount factor.

Finally, the third chapter (with M. Guennewig) examines the efficacy of bail-ins in resolving the regulator's problem of committing not to conduct bail-outs. By endogenizing financing choices in a market which internalises the presence of the regulator, we find that debt maturity shortens. Creditors then respond to news on bank fundamentals leading to runs on loss-absorbing debt, which render bail-ins ineffective. The model provides an explanation why regulators impose minimum maturity requirements for bail-in debt and a motivation to treat short-term debt preferentially during intervention, which achieve constrained efficiency in our model.

Contents

1 Preparing to Act: Information Acquisition and the Timing of Actions **7**
1.1 Introduction . . . 7
1.2 Model . . . 13
1.3 Optimization Problem . . . 16
1.4 Appendix . . . 32
1.4.1 Proof of Proposition 1 . . . 32
1.4.2 Proof of Theorem 1 . . . 33
1.5 Online Appendix . . . 43
1.5.1 Proof of Lemma 1 . . . 43
1.5.2 Proof of Convexity . . . 44

2 Smart Contracts and the Coase Conjecture **50**
2.1 Introduction . . . 50
2.2 The Model . . . 54
2.3 Main Result . . . 57
2.4 Discussion . . . 65
2.5 Appendix . . . 72
2.5.1 Deployed Contracts . . . 72
2.5.2 Auxiliary Games . . . 80
2.5.3 Equilibrium Assessment . . . 81

3 On the Equilibrium Maturity of Bail-in Debt **83**
3.1 Introduction . . . 83
3.2 Environment . . . 86
3.2.1 Bankers . . . 86
3.2.2 Investors . . . 88
3.2.3 Depositors . . . 88
3.2.4 Government . . . 88
3.3 A simple notion of bailouts . . . 88
3.4 Bail-ins . . . 89
3.4.1 Long-term debt . . . 90
3.4.2 Short-term debt . . . 92
3.4.3 Equilibrium . . . 93
3.5 Entry into financing markets and efficiency . . . 95
3.6 Mixed maturity debt and long-term debt requirements . . . 96

3.6.1 Pari passu . . . 96
3.6.2 Subordination of long-term debt to short-term debt . . . 97
3.6.3 Discussion of current UK policy . . . 98
3.7 Disciplining effect of short-term debt . . . 99
3.8 Extensions . . . 100
3.8.1 Equity . . . 100
3.8.2 Market completeness . . . 100
3.9 Conclusion . . . 101

List of Figures

1.1 Typical optimal experiment. 19
1.2 Comparison with endogenous timing model. 20
1.3 The problem for arbitrary λ. 22
1.4 The problem at an optimal multiplier λ^*. 22
1.5 Left Panel: Potential jump locations. Right Panel: Optimal jump direction. 23
1.6 Checking for severe mistakes. 25
1.7 Multiple possibilities. 26
1.8 Increasing impact. 26
1.9 Change in the objective as beliefs vary. 27

2.1 Comparison of the seller's profits when the discount factor varies. 66

Chapter 1

Preparing to Act: Information Acquisition and the Timing of Actions

1.1 Introduction

The study of optimal learning in dynamic environments is a classical problem with a long history dating back to the works of Wald (1947) and Arrow, Blackwell and Girshick (1949). It is fair to say that the majority of existing work has considered this problem in environments where decision makers have full control over the timing of their actions. In such 'endogenous-timing' models, decision makers learn until they are confident about what to do, at which point they can stop and take an action.

However, in many environments opportunities to act are scarce and uncertain, and decision makers are deprived of control over *when* they take actions. The aim of this paper is to shed light on learning behaviour in such cases. In particular, this paper considers a Bayesian decision maker whose preferences over actions depend on an unknown state. To account for potential richness in information sources, we allow access to any dynamic experiment which they may employ to learn about the state. Although such a degree of flexibility is arguably extreme, it allows us to uncover the key forces at play without obscuring them with potentially superfluous assumptions about the available information. Importantly, the decision maker awaits a single opportunity to irreversibly act, which arrives at some *exogenous* random time—independently of the state. This assumption serves as a polar opposite to fully controlling the timing of actions, making this model one of what will be termed 'exogenous-timing'.

Many economic problems have this flavour. For instance, consider a pharmaceutical firm researching a new vaccine technology over a known existing technology. The problem they are facing is which technology to use in producing their next product. Before a pandemic occurs there may be little demand for their products and because of the limited term of patents, there may be little urgency to commit to any one technology.[1] Importantly, the onset of pandemics cannot be foreseen and firms will thus try to learn about the efficacy of the new technology in anticipation of the next outbreak. Understanding firms' optimal research agenda in such a context is essential, especially since such activities typically involve positive externalities one may wish to address. Similarly, we

[1]Studies of the intellectual property landscape surrounding mRNA vaccine technologies, suggest a clustering of patent filings of mRNA related technologies around the onset of the SARS-CoV-2 pandemic (Martin and Lowery, 2020).

may consider a government agency exploring different ways to tackle some unpredictable natural disaster, such as a volcano eruption or an asteroid strike. A good understanding of optimal learning can aid the design of research strategies for such problems, which can prove pivotal in their success.

Furthermore, control over the timing of actions is an important theoretical dimension along which dynamic information acquisition problems should be distinguished. Choosing 'information' amounts to choosing lotteries over beliefs. Consequently, when the timing of actions is controlled, a decision maker's preferences over time-lotteries take a central role in shaping their behaviour. As has been pointed out by DeJarnette et al. (2020), standard exponential discounting generates risk-taking behaviour on the time dimension. Zhong (2019) establishes the qualitative features of optimal information in an endogenous-timing model. He shows that the result is driven by the temporal-risk attitudes of the decision maker in the following way: the optimal experiment generates the maximal mean-preserving spread over decision times.[2] In contrast, the exogeneity of the decision time in this paper eliminates the role of preferences over time-lotteries, and we find that most qualitative aspects of the typical optimal experiment *reverse*.

Ultimately, in any decision problem there will be pieces of information whose timely arrival the decision maker values more. For instance, in the endogenous-timing setting the decision maker would like to know whether they should *stop* and take an action as quickly as possible. When the timing of actions is exogenous, what is of value is whether, and how, they should *change* their intended action; and when the opportunity to act is likely to be soon, they want to learn this promptly. An optimal dynamic experiment will try to deliver those relevant pieces of information as fast as feasibly possible.

Model Outline.— An opportunity to act arrives at some *exogenous* random time, and the decision maker's preferences over actions depend one of two unknown states. The future is discounted at some constant rate. Time is continuous and the decision maker can flexibly design a dynamic experiment to learn about the state prior to acting.

A dynamic experiment consists of a standard Blackwell experiment, and an information flow, which refines knowledge over time. Such experiments can be identified by the posterior process they induce, given a prior. Any dynamic experiment can be represented by a pair of characteristics. The first captures the speed at which 'incremental information' arrives. Such pieces of information generate erratic and local changes in the beliefs. The second characteristic governs the behaviour of 'breakthroughs'—that is, pieces of information which can generate abrupt and significant changes in beliefs. The decision maker can set up dynamic experiments delivering incremental information as well as breakthroughs of varying impact, arriving at different frequencies. Consequently, the set of experiments is rich enough to accommodate a large variety of dynamic learning patterns one may wish to explore.

Without any costs or constraints the decision maker would simply obtain a fully informative signal instantaneously, and then wait for the opportunity to act. We will therefore be imposing some information constraint. To describe it, we first define the total quantity of information in an experiment as a moment of the final posterior distribution it implements.[3] We can then specify a notion of information increment, or *flow-information*, generated by a dynamic experiment. As in Zhong (2019) and Hébert and Woodford (2021b), we specify a capacity constraint on the flow-

[2] In contrast, when there are linear delay costs, the decision maker is risk-neutral over time-lotteries, and the type of experiment becomes irrelevant.

[3] In the tradition of the literature on rational inattention, we consider a uniformly posterior separable measure of information.

information the experiments are allowed to generate.[4] Critically, such a constraint induces the decision maker to 'smooth' information acquisition over time.

This constraint can be intuitively understood as follows. As the decision maker learns their posterior changes; and, the more information they receive the more their beliefs 'move' around. Essentially, the constraint enforces that beliefs cannot 'move' too much in expectation over short periods of time.

Moreover, such a constraint introduces some important trade-offs and considerations the decision maker must resolve. The first is a trade-off between the frequency with which breakthroughs arrive, and their significance. Since the decision maker is constrained in the expected movements of their beliefs, if they want information *fast*, then its arrival cannot 'move' beliefs too much. For the same reason, if they want a very informative signal, generating potentially big movements in beliefs, then this cannot be arriving very frequently. Additionally, the decision maker needs to consider the direction and location, where beliefs move when impactful information arrives. For instance, at any point in time the decision maker will deem some action optimal given the information they have accumulated up to that point. They need to decide whether to seek evidence which confirms or contradicts this intended action; and if they seek to contradict it, which alternative actions should this new evidence favour.

Preview of Results.—To describe the main results, we define the *information gain* at each posterior belief as the difference between the decision maker's optimal payoff—their value, and their payoff from acting on the basis of their current information alone. The information gain quantifies for each belief the potential gains from learning. The optimal experiment is as follows. The decision maker concentrates resources in generating a single piece of breakthrough evidence. That is, the optimal experiment is a *single-jump* process. It is characterized at each point in time by its *frequency*—the rate at which breakthrough evidence arrive; its *impact*—the significance of such evidence, captured by the magnitude of the change in beliefs they induce; and finally, its *direction*—whether these evidence confirm or contradict the decision maker's prior conviction.

Close to a location of maximal information gain, the optimal experiment has high frequency and low impact. In the limit as the belief approaches a location of maximal information gain, frequency explodes and impact implodes, producing purely incremental information. As the information gain decreases, the frequency of breakthroughs decreases in favour of more impact. Furthermore, the jumps in beliefs induced by breakthrough evidence are in the *direction of increasing information gain*. Finally, along a sequence of breakthroughs the information gain is *constant*; that is, any jump in beliefs is at a location where the information gain has the same value as the starting belief.

In settings with finitely many actions, these properties combine to paint the following picture. At the optimum, this single piece of breakthrough evidence convinces the decision maker to change their action. Particularly, at each instant, either this piece of evidence arrives and beliefs jump to a posterior where they find some alternative action optimal; or, no evidence arrives leading to a small drift in beliefs, which *reaffirms* their current intentions. Moreover, frequency is increasing and impact is decreasing as the decision maker becomes less confident about which action is best.

Of course, an experiment which can *never* lead to a change in action is of no value. The content of the results above is that at the optimum the decision maker learns in a way which has the potential to falsify their intended action at *any* time-scale—no matter how small. For instance, it is not optimal to attempt to learn about the merits of alternative actions over time, by first jumping

[4] Analyzing a problem with some smooth convex cost function applied to the information increment is similar to the constrained problem where the capacity is endogenized, and varies with beliefs.

to beliefs where the current action is still dominant; regardless of whether it is more or less desirable at that new belief.

Whenever they fail to contradict their intentions, the decision maker becomes more confident, which makes them less eager to learn fast, and induces them to obtain more impactful information—which is necessary to falsify their preferred action, in which they are now more confident. If they do succeed in finding evidence against their current intention, it is substantial enough so that even more convincing evidence against the new preferred action is necessary to sway them in the future. Ultimately, the impact of breakthroughs is increasing over time at the optimum, along any path.

It is interesting to compare these qualitative features with those in the endogenous-timing model. In that model, the decision maker typically learns in a way which *confirms* their intentions, to the point they are confident enough to stop learning and take an action. Moreover, they acquire frequent signals when they are *confident* about what to do, because this can lead to early decision—saving on the delay costs.

When the timing is exogenous, there is no particular pressure to act fast, and hence there is little gain from a rare signal which tells the decision maker that their preferred action is actually good. Consequently, they opt to learn in a way which contradicts their intentions. Moreover, they acquire frequent signals when they are *confused* about what to do, because this breaks indifferences fast—raising the quality of decisions.

In contrast, when the decision maker is confused about what to do in the endogenous-timing model, they know that they will not benefit from acting soon, so they are willing to delay the decision and wait for some impactful information. Overall, the resolution of the key trade-offs goes in the opposite direction between the two models.

Related Research

The question of how to optimally gather information to decide on competing hypotheses about the world is probably as old as Statistics itself. In a series of seminal papers, Wald (1947) and Arrow, Blackwell and Girshick (1949) laid out the archetype sequential sampling model, in which a decision maker faces a dynamic problem of choosing between an irreversible decision whose consequences depend on an unknown state; or continuing to obtain information, in the form of independent draws from a distribution depending on the state. Costs are usually proportional to the number of observations.[5]

Although these early models recognize the importance of dynamics in statistical decision-making, they are limited in an important respect: the kinds of information sources available are beyond the decision maker's control.

Parallel to these models, the literature on 'experimentation' sought to merge information sources with payoff-generating activities: a decision maker chooses between competing options (bandits) of unknown quality whose payoffs are stochastic. The seminal papers in this literature are Robbins (1952), Weitzman (1979), and Gittins and Jones (1979).[6] Importantly, allocation of 'time' among these bandits produces dynamic information, and information costs arise as *opportunity costs* on the

[5] An interesting generalization of this type of model is pursued in Morris and Strack (2017) who allow for belief-dependent sampling costs. Their goal is establishing equivalences between such sequential models, and static information acquisition models with appropriate cost functions.

[6] For a survey see Bergemann and Vallimaki (2006). Karatzas (1984), and Bolton and Harris (1999) treat continuous-time versions of the problem in a diffusion setting, while Keller, Rady, and Cripps (2005) treat the Poisson framework.

foregone payoffs of unchosen options. There too, the exogeneity of the structure of bandits limits the decision maker in the kind of information they can obtain.

The success of these dynamic models in capturing relevant economic phenomena spurred interest in moving away from their restrictive assumptions. Moscarini and Smith (2001) study a model of a decision maker who learns via observations of a Gaussian process whose 'precision' they control at a convex cost. By endogenizing this aspect they uncover a monotonicity between the precision of learning and the cost of delay. Che and Mierendorff (2019) consider a decision maker who learns via observing Poisson processes, but who faces a constraint in allocating their time between them. They are able to study the choice of 'direction' in the decision maker's information by considering choices between perfectly revealing Poisson signals, and uncover rich patterns of behaviour.[7] Ke and Villas-Boas (2019) study the problem of a decision maker who learns before choosing between two alternatives by sampling Gaussian processes, at a heterogeneous cost which is linear in the time spent with each source. Liang, Mu, and Syrgkanis (2021) consider a model where the decision maker allocates a fixed budget of time among Gaussian signals. In all of these models the decision maker *controls* the timing of their action.[8]

The contribution of this paper vis à vis this line of research is two-fold. Firstly, it departs from the prevalent stopping-problem setting to consider other important environments. Namely: (i) settings where the timing of actions is exogenous; and (ii) settings with repeated actions and unobservable payoffs (see Footnote 11, after the presentation of the model, for a discussion of how this model can be viewed this way.) The former differ from the 'Wald-type' models above by shutting down any timing concerns. This introduces novel considerations orthogonal to the cost of delay in taking an action—which drives in large part the results in the aforementioned literature. The latter are distinct from experimentation models in that, although payoffs are generated by repeated actions, they are unobservable, and thus the decision maker must rely on additional information obtained elsewhere. This allows the study of information acquisition when decisions have incremental impact on payoffs, while retaining the flexibility in the design of experiments.

Schneider and Wolf (2020), also recognize the importance of departing from traditional timing assumptions. They develop a model where a decision maker experiments with exogenous exponential bandits until a fixed deadline. They show that 'time pressure' significantly alters the conclusions of the existing literature. In this paper, the 'deadline' is stochastic so the concept of time pressure is less pronounced. Rather, there is a possibility of having to act in each period.

Secondly, the present model allows for full freedom in what information the decision maker can generate, and thus is capable of capturing all subsets of the features of information studied in the literature. Methodologically, this paper is closely related to Zhong (2019) which studies flexible dynamic information acquisition in a 'Wald-type' model of optimal stopping. The recent literature on Bayesian Persuasion and Rational Inattention, and the techniques developed therein,[9] allow for flexibility in modelling information, which Zhong (2019) introduces to a dynamic model. In particular, Zhong (2019) associates to each dynamic signal a flow-information quantity, as the speed of reduction in an uncertainty measure. This is then used to specify costs of information which

[7] See also Mayskaya (2020) for a related model.

[8] A working paper version of Liang, Mu and Syrgkanis (2021), (Mu, Liang and Syrgkanis, 2017) also considers a framework which accommodates an exogenous deadline. Their focus is in establishing when myopically optimal behaviour is equivalent to the dynamically optimal one. Importantly, their exogenous Gaussian framework cannot be reconciled in the present model, where optimal information is almost never achievable by observing Gaussian sources.

[9] See Kamenica (2019) for a survey of Bayesian Persuasion, and Bergemann and Morris (2019) for a survey of Information Design.

are, importantly, super-additive in this flow-information. This forces the decision maker to smooth information acquisition over time generating non-trivial dynamics. Zhong (2019) characterizes the optimal signal as an (almost everywhere) single-jump Poisson process, with immediate stopping after arrival, and confirmatory direction. This paper adopts a similar specification in terms of costs of information, and the relationship between the results here and the stopping-problem in Zhong (2019) have been discussed in the Introduction.[10]

Furthermore, this paper shares features with the literature on dynamic rational inattention. Indeed, the information acquisition problem can be thought of as the 'information-processing' problem of a rationally inattentive agent. Following Sims' (2003) influential work, a vast literature emerged on how agents process information in the presence of processing costs. See Máckowiak, Matějka and Wiederholt (2021) for an extensive survey. Steiner, Stewart and Matějka (2017) develop a dynamic rational inattention model where in each period the decision maker obtains arbitrary information about a non-persistent state, prior to choosing an action. The costs of information are linear in the reduction of Shannon entropy. However, their flow utilities depend on the entire history of actions and states. Miao and Xing (2020), also study a dynamic rational inattention model but with more general uniformly posterior-separable information measures. In the language of the present paper, in both these models costs are linear in flow-information and hence the DM has no smoothing motive.

Moreover, Hébert and Woodford (2021b) pursue an optimal stopping-problem with flexible information acquisition as in Zhong (2019), but generalize the functional forms delivering the informational constraints by considering general divergences. They explore which properties of these divergences generate diffusion experiments and which pure-jump experiments.

Finally, any model seeking to allow flexibility in information choices must specify a notion of information quantity or information costs at the required level of generality. Defining a 'sensible' quantity of information is a very subtle and deep question, with a long history. Here we will be content with representing the total quantity of information as a moment of the final posterior distribution. As in, for example, Pomatto, Strack and Tamuz (2020), Mensch (2018), and Denti, Marinacci, and Rustichini (2021) we view the association between experiments and a numerical quantity of information as a completion of the Blackwell order, in that it generates a total order on final experiments, while respecting the Blackwell order. Although a representation as a prior-*dependent* moment is a consequence of a form of linearity of the information measure over Blackwell experiments,[11] the prior-independent representation pursued here, and in Zhong (2019), Hebért and Woodford (2021b), Steiner, Stewart and Matějka (2017), and Miao and Xing (2020), is harder to motivate—particularly as a 'physical' cost of information acquisition. This is the subject of recent debate in the literature on information costs and particularly the distinction between posterior-separable and *uniformly* posterior-separable costs. The specification in this paper falls in the latter class, which is nevertheless rich enough and includes widely used (or well-motivated) measures of information, such as Shannon entropy reduction of Sims (2003), Bayesian LLR costs of Pomatto, Strack and Tamuz (2020) and Bloedel and Zhong (2020), Neighbourhood-based Costs of Hébert and Woodford (2021a), and many others. Exploring a richer set of information measures is left for future work.

[10] See also the discussion following the binary example below.

[11] Versions of this result appear in characterizations of posterior-separable information costs in Caplin, Dean and Leahy (forthcoming); of cardinal measures of information in Mensch (2018), and Azrieli and Lehrer (2008); and most generally in the characterization of monotone affine functionals of experiments in Torgersen (1991).

1.2 Model

The Decision Problem.— There is a set of unknown states $\Theta = \{\theta_0, \theta_1\}$, and a Bayesian decision maker (DM) with prior belief $\pi = \Pr(\theta = \theta_1) \in (0,1)$. A perishable opportunity to irreversibly act arrives at some exogenous random time $\tau \sim \exp(\rho)$.[12] The DM has state-dependent preferences over their actions, described by some payoff function $u : \Theta \times A \to \mathbb{R}$, where $u(\theta, a)$ is the payoff from action $a \in A$ in state $\theta \in \Theta$. We assume A is compact and $u(\cdot, a)$ continuous in a, so that the maximum value function $U : [0,1] \to \mathbb{R}$, with:

$$U(\mu) = \max_{a \in A} \mathbb{E}_\mu[u(\theta, a)] = \max_{a \in A} \mu \cdot u(\theta_1, a) + (1-\mu) \cdot u(\theta_0, a)$$

is well-defined, and continuous. U is convex as the maximum of linear functions. The DM discounts the future at rate $r > 0$, so if they act at time t their payoffs are given by $e^{-rt} \cdot U(\mu_t)$.

Information Strategies and Costs.— While waiting for the opportunity to act, the DM may acquire information about the unknown state θ in the form of *dynamic* experiments, which we now describe. We identify a dynamic experiment by the posterior process $(\mu_t)_{t \geq 0}$ it induces, given some prior, and assume that this process is right-continuous and has left limits (càdlàg). Denote by $\langle \Omega^\pi, (\mathscr{F}_t)_{t \geq 0} \rangle$ the space of càdlàg paths starting at some prior $\pi \in (0,1)$, with the filtration generated by the observations of the realized path.[13] We identify the posterior with the coordinate process, $\mu : \mathbb{R}^+ \times \Omega^\pi \to [0,1]$, with $\mu_t(\omega) = \omega_t$. Then, a dynamic experiment amounts to a choice of probability measure $P \in \Delta(\Omega^\pi)$, under which μ is a martingale—reflecting Bayes' consistency. We denote such martingale laws by $\Delta_m(\Omega^\pi) \subseteq \Delta(\Omega^\pi)$.

We now describe the information constraint which will depend on a notion of flow-information generated by a dynamic experiment. To specify this, we start from a uniformly posterior-separable measure of information,[14] so that over a period of length $h > 0$, starting from time $t > 0$, the quantity of information generated by the experiment $P \in \Delta_m(\Omega^\pi)$ amounts to: $\mathbb{E}_P\left[G(\mu_{t+h}) - G(\mu_{t-}) \,\middle|\, \mathscr{F}_{t-}\right]$ for some convex function $G : [0,1] \to \mathbb{R}^+$. In the spirit of Zhong (2019) and Hébert and Woodford (2021b), we impose the following 'capacity' constraint on the amount of information that can be generated:

$$\mathbb{E}_P\left[G(\mu_{t+h}) - G(\mu_{t-}) \,\middle|\, \mathscr{F}_{t-}\right] \leq \kappa \cdot h \quad \text{for all } t \geq 0, \text{ and } h > 0 \tag{$\mathcal{I}$}$$

where $\kappa > 0$ is the capacity.[15] This constraint expresses the fact that the DM cannot generate

[12] Given the exponential assumption on τ, this model has an alternative interpretation. It captures the problem of a decision maker enjoying flow payoff from repeatedly taking an action, but who cannot learn from these payoffs. Such problems are economically meaningful. For example, consider a patient who has to take one of two life-saving treatments every day. In such a situation it is unlilkely they are able to observe the gradual effects of their choice, and will inevitably have to rely on other sources to learn about which treatment is best. This contrasts such an example to the 'experimentation' literature, where decision makers learn purely by choosing among different options yielding stochastic payoffs. See e.g. Gittins and Jones (1979), Bolton and Harris (1999) and Keller, Rady, and Cripps (2005). Optimal learning behaviour for such a problem is covered by the solution characterized here.

[13] For details on the construction of this canonical space, see Chapter VI.1 in Jacod and Shirayev (2013).

[14] See Caplin, Dean and Leahy (forthcoming) for a unified treatment of information costs.

[15] This allows for an analysis of the key trade-offs at play, in a relatively straightforward framework. Indeed, the case of smooth convex flow costs, can be analyzed in a similar manner by endogenizing the 'capacity constraint.'

arbitrarily large amounts of information over short time intervals. Distributions $P \in \Delta_m(\Omega^\pi)$ satisfying this constraint will be called feasible and we will write $P \in \mathcal{I}$.

We make the following assumptions:

Assumption 1. *The function $G : [0,1] \to \mathbb{R}^+$ satisfies:*

i. $G \in C^3(0,1)$ *with* $G''(x) > 0$ *for all* $x \in [0,1]$.

ii. $\lim_{x\to 0} G'(x) = -\infty$, *and* $\lim_{x\to 1} G'(x) = +\infty$.

iii. The mapping $x \longmapsto x^2(1-x)^2 \cdot G''(x)$ *is concave.*

The first assumption imparts smoothness to the problem, guaranteeing that the quantities to appear below are well-defined. The second prevents the DM from acquiring information leading to a jump which reveals the state, thereby guaranteeing interior solutions in the ensuing optimization problem. The third assumption is more substantial. Notice that the notion of information quantity specified above is defined over unconditional distributions of beliefs. One could equivalently specify it directly on the Blackwell experiment giving rise to that distribution. Because we insist on the function G to be independent of the prior, the induced information measure over Blackwell experiments must necessarily depend on the prior. That is, the information quantity generated by the same Blackwell experiment may differ depending on the prior.

The third assumption is equivalent to the induced information measure over Blackwell experiments being *concave in the prior* (see Lemma O1 in the Online Appendix). In particular, this implies that the DM can never be hurt by free information, and is sufficient for the value function being convex (see Lemma 1).[16] We view this last implication as fundamental to a Bayesian DM and hence enforce it with this assumption, while maintaining the uniform posterior separability of the information measure.

Examples of functions satisfying these assumptions include the Shannon entropy, with $G(x) = x \log x + (1-x) \log (1-x)$; and the Bayesian LLR costs of Pomatto, Strack and Tamuz (2020) and Bloedel and Zhong (2020), with $G(x) = \log \frac{x}{1-x} \cdot (\gamma_1 \cdot x - \gamma_0 \cdot (1-x))$, for positive constants γ_0, γ_1.

We now continue with determining the notion of flow-information. To any feasible distribution over paths $P \in \mathcal{I}$ we can associate a pair of (differential) characteristics (α^P, F^P), with the following interpretation.[17] For each $t \geq 0$, the term $\alpha_t^P \geq 0$ is the diffusion coefficient associated with the continuous part of the belief process. It captures the behaviour of 'incremental information' in the dynamic experiment. Such pieces of information generate rapid but small changes in beliefs. The second characteristic, F_t^P, is a positive measure associated to the jumps of the belief process. It

[16]See also Bloedel and Zhong (2020), who show that prior-concavity is almost necessary for convexity in static problems, in that if it fails there exist decision problems for which the DM is hurt by free information.

[17]The fact that such differential characteristics exist reflects the necessity to smooth information acquisition over time given the constraint. In other words, there cannot be any sudden lumpy changes in the quantity of information. That is not to say that beliefs cannot jump; just that they cannot jump in a predictable manner (quasi-left continuity).

governs the behaviour of 'breakthroughs'—that is, pieces of information which can generate abrupt and significant changes in beliefs. More specifically, each $F_t^P \in \mathcal{M}(\mu_{t-})$,[18] where

$$\mathcal{M}(x) = \left\{ F \ \text{ positive measure on } [0,1] : \int (y-x)^2 \mathrm{d}F(y) < +\infty, \ F(\{x\}) = 0 \right\}$$

That is, the measures F_t^P are positive, put no mass at the 'current' location μ_{t-}, and satisfy a square-integrability on the jump-sizes $(y - \mu_{t-})$.

Some illustrative examples are in order:

Example 1 (Gaussian)**.** *The pair (α_t, F_t) with $F_t \equiv 0$ and $\alpha_t = \alpha(\mu_t)$, where $\alpha(x) = \frac{1}{2}\sigma^2 x^2 \cdot (1-x)^2$, corresponds to the posterior process resulting from observing a Gaussian signal with variance σ^2, whose drift is the state $\theta \in \{0,1\}$.*

Example 2 (Fully-revealing Poisson)**.** *The pair (α_t, F_t) with $\alpha_t \equiv 0$ and $F_t = F(\cdot|\mu_{t-})$, where $F(\cdot|x) = \varphi x \cdot \delta_1$, corresponds to the posterior process resulting from observing a fully-revealing Poisson signal, which arrives at frequency $\varphi > 0$, only when the state is 1.*[19]

Example 3 (Non-directional jumps)**.** *The pair (α_t, F_t) with $\alpha_t \equiv 0$ and $F_t = \varphi \cdot Q$, where Q some centered probability distribution over posteriors with $\mathbb{E}_Q[y - x] = 0$, corresponds to the posterior process resulting from observing a 'static' experiment inducing posteriors distributed according to P, when this signal arrives at some random time $\tau \sim \exp(\varphi)$,* independently *of the state. That is, the arrival time is uninformative about the state.*

In what follows we will be extensively using the notation:

$$\mathcal{L}\psi(y,x) = \psi(y) - \psi(x) - \psi'(x) \cdot (y-x) \tag{$\mathcal{L}$}$$

for the deviation of a function $\psi(y)$ from its tangent line at x.

We have the following:

Proposition 1. *Let $P \in \mathcal{I}$ be feasible. Then, there exists a pair of predictable processes (α^P, F^P), such that for any $\psi \in C^2(0,1)$, and stopping time $\sigma < +\infty$:*

$$\mathbb{E}_P\left[\psi(\mu_\sigma) - \psi(\mu_{t-}) \,\middle|\, \mathscr{F}_{t-}\right] = \mathbb{E}_P\left[\int_t^\sigma \left(\alpha_s^P \cdot \psi''(\mu_{s-}) + \int \mathcal{L}\psi(y,\mu_{s-})\mathrm{d}F_s^P(y)\right)\mathrm{d}s \,\middle|\, \mathscr{F}_{t-}\right]$$

Moreover, let:

$$\mathcal{I}(x) = \left\{ (\alpha, F) \in \mathbb{R}^+ \times \mathcal{M}(x) \,\middle|\, \alpha \cdot G''(x) + \int \mathcal{L}G(y,x)\mathrm{d}F(y) \le \kappa \right\}$$

[18] All relationships such as $F_t^P \in \mathcal{M}(\mu_{t-})$ should be understood as holding $P \times \mathrm{d}t$-almost everywhere.

[19] Note that we have already assumed that the process is a martingale. Hence, strictly speaking there is a drift which 'compensates' the process to make this so. Since this drift depends only on F, we do not explicitly mention it.

Then,

$$\mathcal{I} = \left\{P \in \Delta_m(\Omega^\pi) : (\alpha^P, F^P) \in \mathcal{I}(\mu_{t-}),\ P \times \mathrm{d}t - \textit{almost everywhere}\right\}$$

Proof. Appendix 1.4.1 □

Proposition 1 provides a representation result for any arbitrary feasible information strategy. In particular, it allows us to describe feasible information strategies in terms of differential characteristics, and the information constraint ($\mathcal{I}$) as a point-wise constraint on those, thereby operationalizing the use of dynamic programming techniques.

In light of the formula for conditional expectations above applied to the information measure G, the term:

$$\alpha_t^P \cdot G''(\mu_{t-}) + \int \mathcal{L}G(y, \mu_{t-})\mathrm{d}F_t^P(y)$$

corresponds to the information increment or 'flow-information' generated by $P \in \mathcal{I}$ at time $t \geq 0$. The constraint restricts the flow-information to be below the capacity $\kappa > 0$.

Notice that the structure of the constraint is history-independent, in that the feasible distributions at time t, are those feasible distributions on the space $\Omega^{\mu_{t-}}$ of paths starting at μ_{t-}. Moreover, payoffs depend only on the value of beliefs at the time of decision. Consequently, one does not need to keep track of the entire belief history, but just of the current value μ_{t-}. That is, we may restrict the induced belief processes to be Markovian without loss of optimality.

1.3 Optimization Problem

With all the ingredients at hand we are ready to state the optimization problem. Given a problem with data $(r, \tilde{\rho}, \tilde{U})$ we consider the value function $\tilde{V} : [0,1] \to \mathbb{R}^+$, defined by:[20]

$$\tilde{V}(x) = \sup_{P \in \mathcal{I}} \mathbb{E}_P^x\left[e^{-r\tilde{\tau}}\tilde{U}(\mu_{\tilde{\tau}})\right]$$

where $\tilde{\tau} \sim \exp(\tilde{\rho})$, is the exogenous time at which the DM gets to act. The DM chooses a feasible law over belief paths to maximize their expected discounted payoff at the time of decision. It is easy to see that this problem is equivalent to a problem with no discounting, where $\tau \sim \exp(\tilde{\rho} + r)$ and the decision payoff is given by $U = \frac{\tilde{\rho}}{\tilde{\rho}+r} \cdot \tilde{U}$.

Hence, in what follows we will dispense with discounting and solve the general problem with data $(0, \rho, U)$:

$$V(x) = \sup_{P \in \mathcal{I}} \mathbb{E}_P^x\Big[U(\mu_\tau)\Big] = \sup_{P \in \mathcal{I}} \mathbb{E}_P^x\left[\int_0^{+\infty} e^{-\rho t}\rho U(\mu_t)\,\mathrm{d}t\right] \tag{V}$$

where $\tau \sim \exp(\rho)$, is the exogenous time at which the DM gets to act.

[20] We will use the notation $\mathbb{E}_P^x[\cdot] = \mathbb{E}_P[\cdot\ |\mu_t = x]$.

The second equality above follows from the independence of τ and the belief process, as well as the exponential distribution assumption. It reveals an alternative interpretation of this model: this is the objective of a DM who repeatedly acts receiving flow payoff ρU, but who does not learn from these payoffs. The two models are behaviourally equivalent in terms of the optimal information strategies they generate.

Characterization

We now proceed to characterize V and the optimal experiment.

Theorem 1. *The value function V is a C^1 solution to:*

$$\rho(V-U)(x) = \sup_{y\neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)} \cdot \kappa \tag{HJB}$$

with $V(1) = U(1)$ and $V(0) = U(0)$.

Moreover, there exists almost everywhere an optimal jump $y^*(x)$, such that:

i. Jumps leave the information gain unchanged*: $(V-U)(y^*(x)) = (V-U)(x)$. Moreover, x is an optimal jump from $y^*(x)$.*

ii. Jumps are in the direction *of increasing information gain: $(V-U)'(x)\cdot(y^*(x)-x) > 0$.*

iii. The impact $|y^*(x)-x|$ is increasing away from local maxima of the information gain.

Proof. Appendix 1.4.2. □

Before examining the properties of the optimal experiment we give a heuristic derivation of the Hamilton-Jacobi-Bellman equation (HJB) above. A typical heuristic limiting argument using the dynamic programming principle suggests that V solves, for all $x \neq 0,1$:

$$\rho(V-U)(x) = \sup_{(\alpha,F)\in\mathcal{I}(x)} \alpha \cdot V''(x) + \int \mathcal{L}V(y,x)\mathrm{d}F(y)$$

We firstly observe that the diffusion term $\alpha \geq 0$ can be feasibly approximated by the jump-measures F,[21] and consequently we may simplify the problem to:

$$\rho(V-U)(x) = \sup_{F\in\mathcal{I}(x)} \int \mathcal{L}V(y,x)\mathrm{d}F(y)$$

[21] Indeed, by setting $F^\varepsilon = \frac{\alpha}{\varepsilon^2}\cdot\delta_{x+\varepsilon}$, we have for any $\psi \in C^2(0,1)$,

$$\int \mathcal{L}\psi(y,x)\mathrm{d}F^\varepsilon(y) = \frac{\alpha}{\varepsilon^2}\cdot\big(\psi(x+\varepsilon)-\psi(x)-\psi'(x)\cdot\varepsilon\big) = \alpha\psi''(x) + \frac{o(\varepsilon^2)}{\varepsilon^2} \to \alpha\psi''(x) \quad \text{as } \varepsilon \to 0$$

Next note that dividing by the total mass $\varphi = F([0,1] \setminus \{x\})$, we may view the choice of F at each belief $x \in (0,1)$ as a choice of a pair (φ, Q), where $\varphi \geq 0$, and $Q \in \Delta(S^x)$, with $S^x = [0,1] \setminus \{x\}$. φ is the total frequency of a jump arriving, and Q is the distribution of jump locations conditional on such arrival. By writing out the constraint explicitly we arrive at the following problem:

$$\rho(V-U)(x) = \sup_{(\varphi,Q)} \varphi \cdot \int \mathcal{L}V(y,x)\mathrm{d}Q(y) \quad \text{subject to} \quad \varphi \cdot \int \mathcal{L}G(y,x)\mathrm{d}Q(y) \leq \kappa$$

Clearly, the DM will never leave any frequency $\varphi \geq 0$ unused and thus the constraint must always bind. By substituting from the binding constraint we get:

$$\rho(V-U)(x) = \sup_{Q\in\Delta(S^x)} \frac{\int \mathcal{L}V(y,x)\mathrm{d}Q(y)}{\int \mathcal{L}G(y,x)\mathrm{d}Q(y)} \cdot \kappa$$

The objective function on the RHS is a ratio of linear functionals of the choice variable Q. As such, it is quasi-convex. Since we are maximizing a quasi-convex functional it is enough to consider the extreme points of the choice set. One can see that:

$$\mathrm{ext}\Delta(S^x) = \big\{\delta_y : y \neq x\big\}$$

where δ_y is a point-mass at $y \neq x$, and consequently:

$$\rho(V-U)(x) = \sup_{y\neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)} \cdot \kappa$$

which is the (HJB) equation we were after.[22] Processes constructed this way correspond to single-jump processes.

We can recognize the LHS as the difference between the optimal payoff, V, and the payoff from taking the optimal action in the absence of further information, U. Their difference $V-U$ captures how much the DM gains from having the ability to learn optimally. We term this quantity the *information gain*.[23] It accrues at rate $\rho > 0$ which is the frequency with which the DM gets to take an action. The RHS is proportional to the ratio between the expected change in the value and the flow-information, optimized over single-jump experiments.[24] That ratio coincides with the *shadow-cost* of information. The HJB equation necessitates that, at the optimum, the information gain is

[22]It should be noted that exactly the same argument can be applied in the more general framework of Hébert and Woodford (2021b), where abstract divergences $D(y\|x)$ are used. Indeed, the information increment in their model is $\int D(y\|x)\mathrm{d}F(y)$ which is still linear in F. More broadly, whenever the information increment is a concave functional of the jump measure F, single-jump processes are optimal.

[23]Interpreting the solution to the original problem with data $(r, \rho, \tilde{U})$ with discounting in terms of the solution to the transformed problem, amounts to simply recognizing that the payoff in the absence of further information is given by $U(\mu) = \mathbb{E}[e^{-r\tilde{\tau}}]\tilde{U}(\mu) = \frac{\rho}{\rho+r}\tilde{U}(\mu)$.

[24]Of course, when the supremum is strict no single-jump experiment is optimal in which case the signal collapses to a diffusion.

exactly proportional to the shadow-cost of information. This monotonicity between the information gain and the shadow-cost of information is a key driver of the results.

Furthermore, the local characteristics of the optimal experiment are identified as follows. Let $\bar{V}''(x) = \limsup_{y \to x} \frac{\mathcal{L}V(y,x)}{(y-x)^2}$. If $\sup_{y \neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)} > \frac{\bar{V}''(x)}{G''(x)}$ then a pure-jump is optimal, and we recover it from:

$$y^*(x) \in \arg\max_{y \neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)} \neq \emptyset$$

and determine its frequency by the binding information constraint:

$$\varphi^*(x) = \frac{\kappa}{\mathcal{L}G(y^*(x), x)}$$

If $\sup_{y \neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)} = \frac{\bar{V}''(x)}{G''(x)}$, then a diffusion is uniquely optimal, and is characterized by binding the information constraint: $\alpha^*(x) = \frac{\kappa}{G''(x)}$.

Finally, to see why a diffusion cannot be optimal on an interval of beliefs, note that if it where, then the value function satisfies the equation: $\rho(v - U)(x) = \frac{v''(x)}{G''(x)} \cdot \kappa$ on the interior of that interval. However, for the solution v to this equation we can show that *some* jump must be optimal (see Lemma 4).

Binary Example.— Before exploring the properties in more detail, we present a simple example with two actions. Consider a pharmaceutical firm exploring different vaccine technologies. There are two 'risky' technologies: R and L whose payoffs depend on the state $\theta \in \{0, 1\}$. In state-1 technology R yields a payoff of 1 if developed, and in state-0 it yields a payoff of -1. Reciprocally, technology L yields a payoff of 1 in state-0, and -1 in state-1. Figure 1.1 below depicts the value V (solid) and the payoff from acting U (dotted), as functions of the belief $x \in [0, 1]$, that the state is 1.

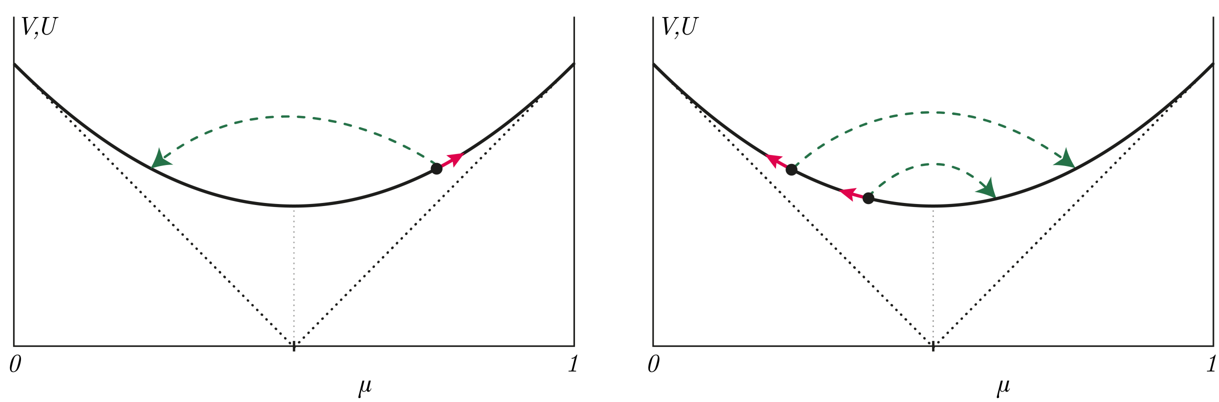

Figure 1.1: Typical optimal experiment.

From any starting belief it is clear that there is a unique location where the information gain is equal to its value at that point. By Property (i) in Theorem 1 this has to be the optimal jump, and the two points are coupled by the fact that the starting belief is also the uniquely optimal jump from that location. Moreover, optimal jumps are always contradictory in that they lead to a change in the intended action. Additionally, at starting locations close the the point of indifference—where the information gain is relatively higher—the impact of jumps decreases (see the right panel of Figure 1.1) which illustrates property (iii).

At this point it is interesting to briefly compare the qualitative properties of optimal learning described above, with those in the endogenous-timing model of Zhong (2019), and Hébert and Woodford (2021b).

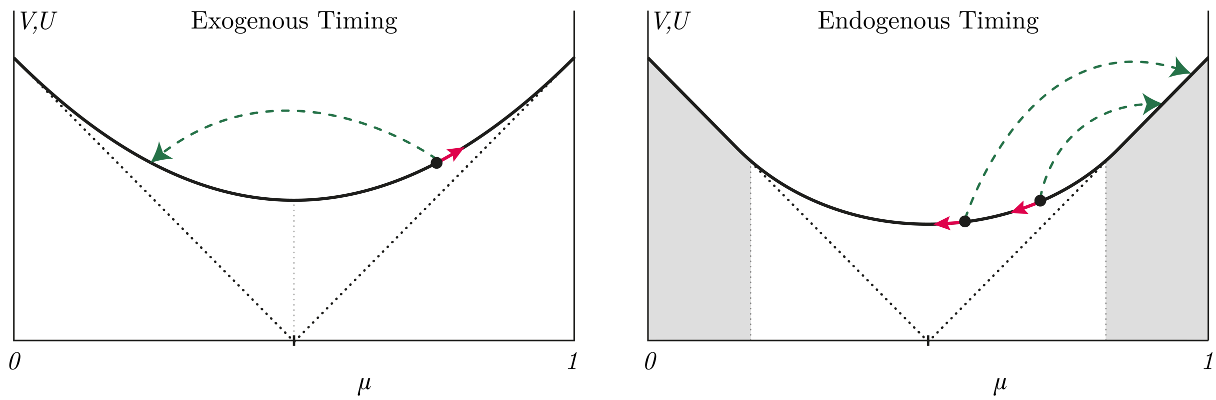

Figure 1.2: Comparison with endogenous timing model.

In the right panel of Figure 1.3 above, we see the value V (solid) and the payoff from acting U (dotted), as functions of the belief $x \in [0, 1]$, that the state is 1.[25] The shaded region corresponds to the stopping region where the firm stops learning and takes an immediate action. In the intermediate non-shaded region the firm learns about the state.

We see that the firm seeks to *confirm* their intensions: that is, they seek evidence whose arrival makes them more confident in the action they already find dominant. This evidence is furthermore impactful enough that it leads to immediate action. Moreover, frequency is higher when beliefs are closer to the stopping boundaries—that is, as the firm becomes relatively confident about the right action. These features are exactly reversed in the exogenous-timing model considered here.

[25]Note that the magnitude of the value V is not comparable across the two models in these graphs. Of course, the DM is always better-off in the endogenous-timing model.

Properties

We will now examine the key features of the optimal experiment, in terms of the *frequency* of jumps, their *direction*, and their *location*.

Location.— The first property in Theorem 1 states that optimal jumps must leave the information gain *unchanged*; namely, if $y^*(x)$ is an optimal jump starting from x, then $(V - U)(y^*(x)) = (V - U)(x)$. Decreasing the information gain is equivalent to capturing some of the outstanding gains from learning. The fact that jumps do not increase the information gain reflects the desire of the DM to use the significant pieces of information in the experiment to capture portions of the outstanding gains from learning. In a way, the DM wants to consolidate intermediate chucks of information into a single piece, until it delivers substantial returns. That jumps do not *strictly* decrease the information gain is quite surprising, and results from the fact that if $y^*(x)$ is optimal at x, then x is a feasible jump for the problem starting with belief $y^*(x)$; and moreover delivers a value exactly equal to the information gain at x. So the information gain at $y^*(x)$ must me weakly higher than at x, by optimality in the (HJB) equation.

To illustrate the nature of this property we will explore the optimization problem in the (HJB) equation in a bit more detail. In particular, consider the problem:

$$\lambda^*(x) = \sup_{y \neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)} \tag{λ^*}$$

which incidentally determines the value of the shadow-cost λ^*. A belief y^* is an optimal jump starting from x if and only if there exists some $\lambda^* > 0$ such that the following hold:

$$\mathcal{L}V(y,x) - \lambda^* \cdot \mathcal{L}G(y,x) \leq \mathcal{L}V(y^*,x) - \lambda^* \cdot \mathcal{L}G(y^*,x) = 0$$

By defining $J_\lambda(y) = V(y) - \lambda \cdot G(y)$, and noting the linearity of $\mathcal{L}$, the above can be re-written as:

$$\mathcal{L}J_{\lambda^*}(y,x) \overset{\text{(i)}}{\leq} \mathcal{L}J_{\lambda^*}(y^*,x) \overset{\text{(ii)}}{=} 0$$

Then, the inequality (i) above implies that y^* maximizes $\mathcal{L}J_{\lambda^*}(y,x)$, while the equality (ii) implies that this maximized value is 0. These can be envisioned as a game between the DM who takes λ as given and seeks to maximize $\mathcal{L}J_\lambda(y,x)$; and an adversarial Nature who chooses $\lambda > 0$ to induce $\mathcal{L}J_\lambda(y,x) = 0$, via the DM's optimal choice.

Figure 1.3 below illustrates the DM's problem for an arbitrary choice $\lambda > 0$. Starting from a current belief $x \in (0,1)$ the DM tries to maximize the deviation of the objective function $J_\lambda(y)$ from its tangent line at x. The tangent at any optimal point y^* must be parallel to the tangent at x. In the figure, the gap between J_λ and its tangent line at x is strictly positive at the maximizer y^*, which violates condition (ii), and hence Nature must adjust λ which alters the shape of the objective function until the DM sets $\mathcal{L}J_\lambda(y,x)=0$, optimally.

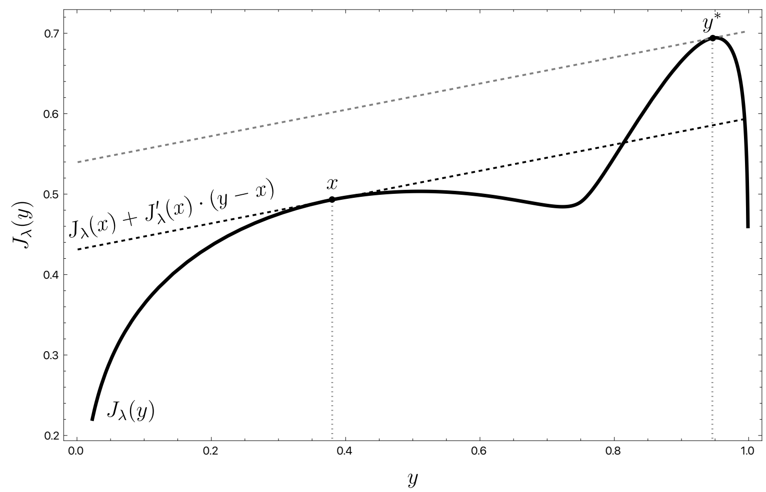

Figure 1.3: The problem for arbitrary λ.

In Figure 1.4 below we see a depiction of the problem at an optimal multiplier λ^*. There, y^* lies exactly on the tangent line that goes through x, which by optimality must also be tangent to J_{λ^*} at y^*. Figure 1.4 thus describes a 'concavification procedure' whereby the starting belief lies at one boundary of the concave envelope of $J_{\lambda*}$, and any optimal jump lies at some other boundary (see also Zhong (2019)).

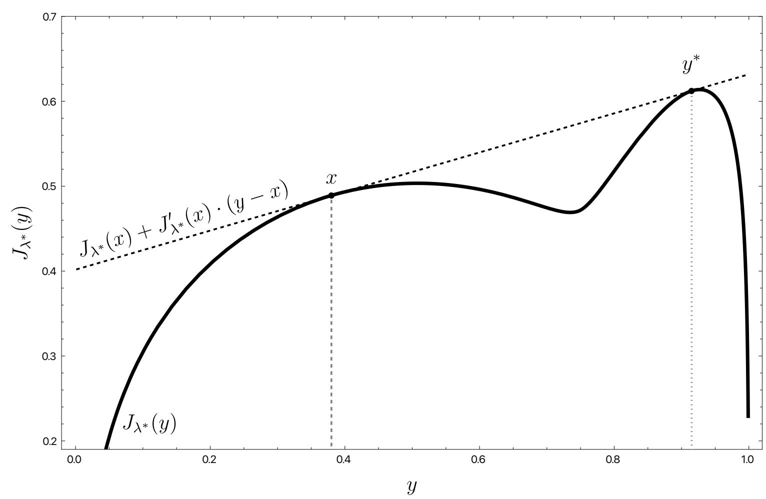

Figure 1.4: The problem at an optimal multiplier λ^*.

Now to understand property (i) in Theorem 1, we consider the problem of finding an optimal jump starting from y^*—the optimal jump starting from x in Figure 1.4 below. Note that the shape of the objective function in the figure depends only on the value of λ^*. Hence, by taking that same value of the multiplier for the problem starting from y^*, we see that the picture remains intact. But then clearly x maximizes the deviation $J_{\lambda^*}(y)$ from its tangent at y^*, and the value of this deviation at x is 0.

Consequently, the belief x is an optimal jump starting from y^*; that is, $x \in y^*(y^*(x))$. Moreover, the shadow-cost of information is equal at x and at $y^*(x)$; that is, $\lambda^*(x) = \lambda^*(y^*(x))$. Therefore, since by the HJB equation the shadow-cost is proportional to the information gain we must have $(V-U)(y^*(x)) = (V-U)(x)$. That is, the information gain is constant along optimal jumps, which is property (i).

As we saw in the binary example above, the fact that the information gain is constant on jumps, essentially determines the solution of every binary action problem. To illustrate the implications of this property further, and the ones to come, we will develop a simple running example with three actions by adding a third technology, S, to the example above. S is a safe technology which yields a sure payoff of 1/2.

Figure 1.5 below depicts the value V (solid) and the payoff from acting U (dotted), as functions of the belief $x \in [0,1]$, that the state is 1.

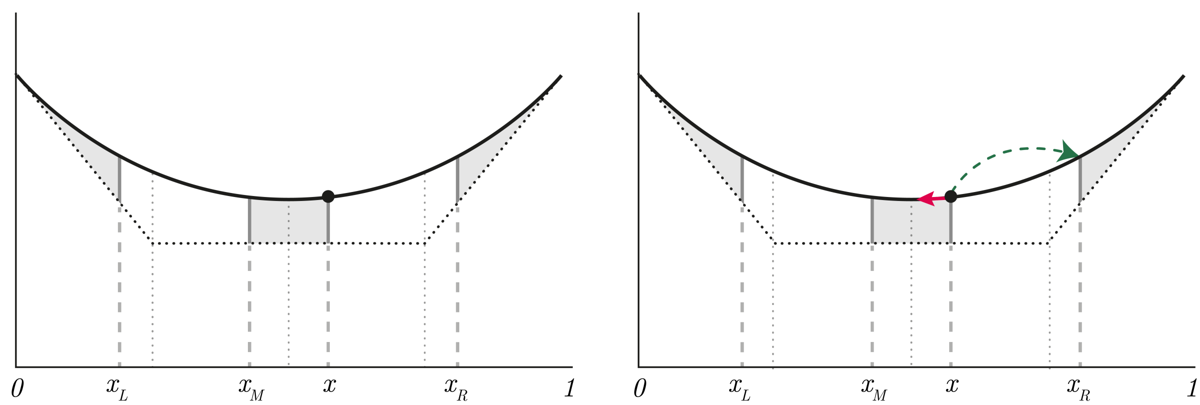

Figure 1.5: Left Panel: Potential jump locations. Right Panel: Optimal jump direction.

Starting from belief x in the the left panel of Figure 1.5, there are only three regions for which the information gain is lower than its value at x. These are the shaded regions in the left panel of Figure 1.5. By the first property in Theorem 1, the boundaries of these regions, x_L, x_M, and x_R are the only possible candidate locations for an optimal jump.

Direction.— The second property in Theorem 1 concerns the direction of optimal jumps. It states that optimal jumps are in the *direction of increasing information gain*. To see this consider the (HJB) equation at a belief $x \in (0,1)$, where an optimal jump exists. By the envelope theorem we have:

$$\rho(V-U)'(x) = \varphi^*(x) \cdot \Big(\lambda^*(x) \cdot G''(x) - V''(x)\Big) \cdot (y^*(x) - x)$$

where $\varphi^*(x) > 0$ is the optimal frequency. Recall that if an optimal jump exists, then $\lambda^*(x) > \frac{V''(x)}{G''(x)}$. Therefore, the first two terms in the RHS above are strictly positive. Consequently, $(V-U)'(x)$ has the same sign as $(y^*(x) - x)$. In other words, optimal jumps are in the direction where the information gain is increasing.

The right panel in Figure 1.5 shows the optimal jump for the running example. At belief x, the information gain is increasing locally, and therefore the optimal jump must be to the right. The only location to the right which satisfies the requirement of the first property, that information gain is constant, is belief x_R. It is therefore the optimal jump from x.

To gain some intuition for why this pattern of learning is optimal, notice that jumping in the direction of increasing information gain means that, in the absence of arrival of a jump, beliefs drift in the opposite direction. Consequently, at the optimum beliefs always drift in the direction of *decreasing* information gain. Coupled with the fact that jumps weakly decrease information gain by the first property, we see that this pattern of learning has good 'hedging' properties. Indeed, putting resources in generating a substantial piece of evidence runs the 'risk' of failing to generate it. By decreasing the information gain in this contingency the DM effectively hedges against that risk. The downside of this strategy is that in the future, more impactful information is needed to lower the information gain, at the expense of frequency. Nevertheless, this is accordance with the DM's intertemporal incentives, since at a lower information gain the desire for frequency is diminished—as described in property (iii).

Before we proceed, we will explore the richness of learning behaviour that is present in this model—even in a three-action example. In the context of the running example, Figure 1.5 suggests that the DM seeks to falsify the safe technology S in favour of the most promising alternative risky technology. One may be tempted to infer that it is always optimal to falsify your current intended action in favour of the 'closest' alternative action. However, this is not correct as Figure 1.6, below illustrates.

The orange vertical line corresponds to the minimal information gain achievable with a jump in the region where S is optimal. Since the information gain converges to 0 as the DM becomes certain about the state, while the minimal information gain at the interior is non-zero, there will always be extreme beliefs where the information gain is lower than the minimal information gain in the middle region. Such a point is x in Figure 1.6. Since a jump location of equal information gain cannot be found in the middle region, property (i) in Theorem 1 then prevents optimal jumps from landing there. Any optimal jump must travel all the way across at a location where the alternative risky technology is optimal. In such regions the firm finds it optimal to check whether they are making some 'severe mistake' by choosing the action they intend to.

Notice that in both cases illustrated by figures 1.5 and 1.6 optimal jumps increase the value.

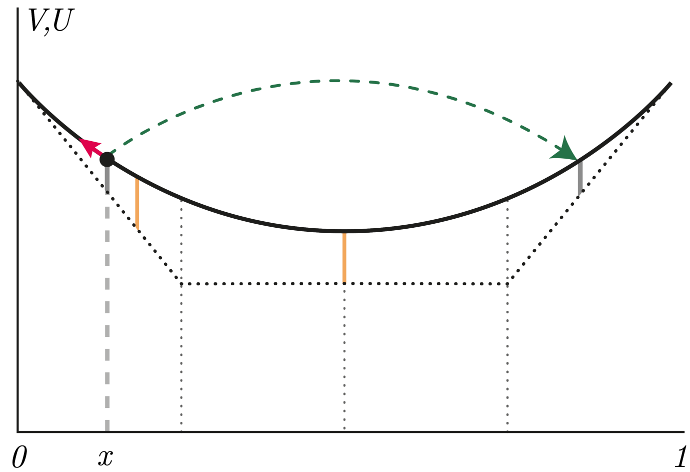

Figure 1.6: Checking for severe mistakes.

This, however, is not a general feature of the optimal experiment.[26] To complete the description of learning in the running example we need to describe optimal jumps from beliefs in the region to the right of the leftmost orange line until the kink where the optimal action changes. Consider belief x in Figure 1.7, below. The information gain at such a belief is larger than the minimal information gain at the middle region, and therefore multiple candidate locations of equal information gain exist. Moreover, all of the candidate locations are to the right of x which is the direction of increasing information gain.

Starting from the left, the second option can be eliminated as jumping back to x cannot be optimal, since x is in the direction of decreasing information gain from that location.[27] We are left with two options both of which satisfy the properties in Theorem 1, and thus we cannot determine which exact location is optimal solely on the basis of that result.

Nevertheless, we can intuitively expect that the closest candidate location should be optimal. The reason is that beliefs such as x in Figure 1.7 are close to points of indifference where the information gain is high. At such locations the DM values frequency more than impact and a jump to the closest location allows the DM achieve higher frequency. This intuition is confirmed by numerical results which suggest that optimal jumps are always at the closest candidate location.

Frequency and Impact.— The last property in Theorem 1 determines how the DM resolves the trade-off between the frequency of arrival of jumps, and their impact—that is, the magnitude of the change in beliefs that they induce. In particular, it states that the DM trades off frequency for

[26] By contrast, in the endogenous-timing models of Zhong (2019), and Hébert and Woodford (2021b) optimal jumps always increase the value.

[27] Recall that $x \in y^*(y^*(x))$.

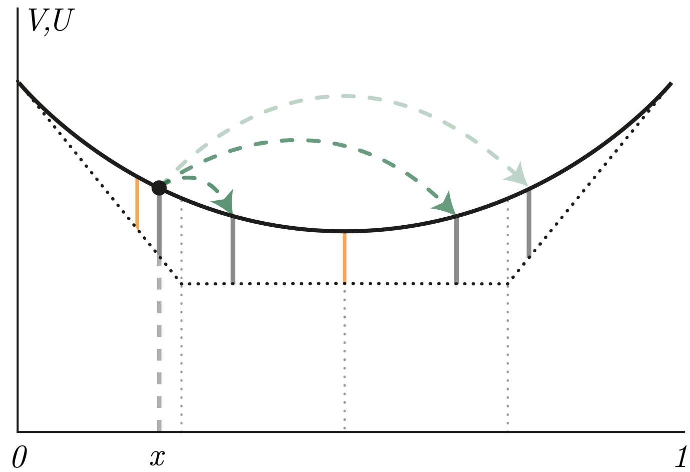

Figure 1.7: Multiple possibilities.

impact as beliefs move away from local maxima of the information gain. The implications of this for the running example are shown in Figure 1.8.

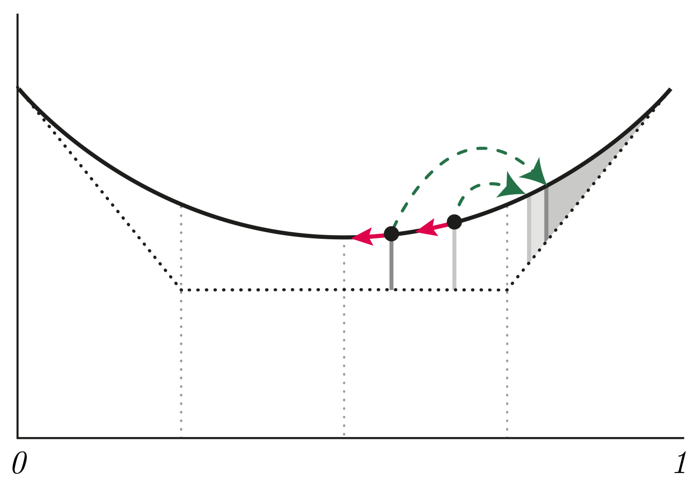

Figure 1.8: Increasing impact.

Intuitively, locations of high information gain correspond to beliefs where the DM values 'marginal' information a lot. In the context of the running example these are precisely the points where the firm is nearly indifferent between two technologies. In such a situation the firm seeks incremental information that arrives *fast*, which has the possibility of guiding them in some direction, even if the signal very uninformative. As they become more confident in their intended action, they become less eager to learn something fast, and they focus on obtaining more informative signals.

This property can be illustrated by looking at how the FOC implicit in the concavification problem described above, varies with beliefs. Figure 1.9 illustrates how the DM's objective changes as their belief moves from x_1 to x_2. Same coloured regions correspond to beliefs where the same action is optimal, and thus the same action is optimal at both x_1 and x_2. The boundary of these regions corresponds to the point of indifference between the two actions. At x_2, the DM is closer to being indifferent between the two actions, and thus less certain about which of the two is best. The red curves correspond to the objective and tangent at belief x_1; while the blue curves correspond to the same at belief x_2.

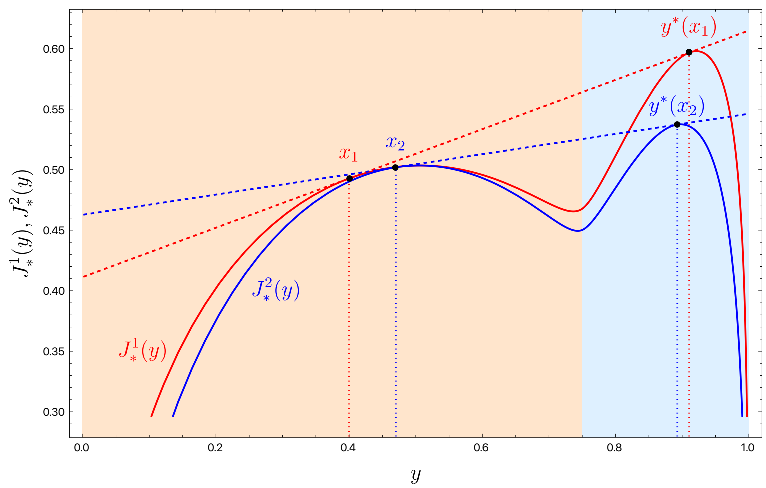

Figure 1.9: Change in the objective as beliefs vary.

It can be seen that at x_2, which is closer to the point of maximal information gain, the region below the concave envelope of J_*^2 contracts. That is, the boundaries come closer together, and correspondingly the difference between the belief x_2 and the optimal jump $y^*(x_2)$ becomes smaller than that between x_1 and $y^*(x_1)$. That is, impact decreases, and correspondingly frequency increases.

Finally, observe that under the optimal experiment, the information gain is decreasing over time with probability 1, since at any instant either a jump arrives leaving it unchanged; or beliefs drift in the direction in which it is decreasing. Because frequency is decreasing in the information gain, along any path, frequency must be decreasing. Equivalently, impact is increasing over time.

Finite Action Problems

We now combine the general properties described above to give a description of optimal learning in finite action problems. In such problems, the indirect utility U is piece-wise linear. Hence, any local maximum of the information gain must occur at a point of indifference between at least two actions.

This is intuitive as those are exactly the instances where the DM is most confused about what to do. In such situations the DM values imprecise signals, because by pumping up the frequency they can potentially 'outrun' the decision time, and raise the quality of the decision.

At the optimum, the DM concentrates resources in generating a single piece of breakthrough evidence which convinces them to change their action. At each instant, either this piece of evidence arrives, and beliefs jump to a posterior where they find some alternative action optimal; or, no evidence arrives, leading to a small drift in beliefs, which *reaffirms* their current intentions.

To see this recall that property (i) in Theorem 1 requires the optimal jumps to weakly decrease information gain, while property (ii) requires that these jumps are in the direction where it is increasing. The only way by which these two can be reconciled is if optimal jumps cross a local maximum of the information gain. Since these are located at points of indifference between actions, optimal jumps must land in regions where some alternative action is superior. From property (ii), in the absence of arrival beliefs move away from points of indifference between actions, and thus the DM becomes more confident in their current plan.

Of course, an experiment which can *never* lead to a change in action is of no value and can never be optimal. The content of the results above is that at the optimum the DM learns in a way which has the potential to change their action at *any* time-scale—no matter how small. For instance, it is not optimal to attempt to learn about the merits of alternative actions by first jumping to beliefs where the current action is still dominant; regardless of whether it is more or less desirable at the new belief.

The eagerness to be prepared to act when the time comes, induces the DM to *always* seek evidence contradicting their plans, giving them the opportunity to make drastic changes in time, if necessary. Moreover, they learn in a way in which, failing to falsify their current intentions, makes them more confident in them—which, in the future, motivates them to explore whether they are making some 'big mistake.' This requires a significantly informative experiment—at the expense of frequency, which the DM is happy to obtain, being more confident in the validity of their intended action.

Conclusion

This paper characterizes the optimal dynamic experiment in an environment where the decision maker lacks control over the timing of actions. The decision maker can flexibly choose all aspects of the experiment but faces a capacity constraint on the flow-information it can generate. In this environment resources are concentrated into generating single pieces of breakthrough evidence, which counter the decision maker's prior intentions, in a way in which absence of breakthroughs make the decision maker more confident in their action plan. This leads them to sacrifice frequency in the arrival of breakthroughs, to increase their significance. Over time, breakthroughs become more rare and more impactful.

Bibliography

[1] K. J. Arrow, D. Blackwell, and M. A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17(3/4):213–244, 1949. ISSN 00129682, 14680262. URL `http://www.jstor.org/stable/1905525`.

[2] Y. Azrieli and E. Lehrer. The value of a stochastic information structure. *Games and Economic Behavior*, 63(2):679–693, 2008.

[3] D. Bergemann and S. Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

[4] D. Bergemann and J. Välimäki. Bandit problems. 2006.

[5] A. W. Bloedel and W. Zhong. The cost of optimally-acquired information. *Unpublished Manuscript, November*, 2020.

[6] P. Bolton and C. Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.

[7] Y.-K. Che and K. Mierendorff. Optimal dynamic allocation of attention. *American Economic Review*, 109(8):2993–3029, 2019.

[8] P. DeJarnette, D. Dillenberger, D. Gottlieb, and P. Ortoleva. Time lotteries and stochastic impatience. 88(2):619–656, 2020. doi: 10.3982/ecta16427. URL `https://doi.org/10.3982%2Fecta16427`.

[9] T. Denti, M. Marinacci, A. Rustichini, et al. Experimental cost of information. Unpublished Manuscript, 2021.

[10] W. H. Fleming and H. M. Soner. *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media, 2006.

[11] J. C. Gittins and D. M. Jones. A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66(3):561–565, 12 1979. ISSN 0006-3444. doi: 10.1093/biomet/66.3.561. URL `https://doi.org/10.1093/biomet/66.3.561`.

[12] B. M. Hébert and M. Woodford. Neighborhood-based information costs. Technical report, National Bureau of Economic Research, 2021a.

[13] B. M. Hébert and M. Woodford. Rational inattention when decisions take time. Unpublished Manuscript, 2021b.

[14] J. Jacod and A. Shiryaev. *Limit theorems for stochastic processes*, volume 288. Springer Science & Business Media, 2013.

[15] E. Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11: 249–272, 2019.

[16] I. Karatzas. Gittins indices in the dynamic allocation problem for diffusion processes. *The Annals of Probability*, 12(1):173–192, 1984.

[17] T. T. Ke and J. M. Villas-Boas. Optimal learning before choice. *Journal of Economic Theory*, 180:383–437, 2019.

[18] G. Keller, S. Rady, and M. Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005.

[19] A. Liang, X. Mu, V. Syrgkanis, et al. Dynamic information acquisition from multiple sources. *arXiv preprint arXiv:1703.06367*, 58, 2017.

[20] A. Liang, X. Mu, and V. Syrgkanis. Dynamically aggregating diverse information. 2021.

[21] D. G. Luenberger. *Optimization by vector space methods.* John Wiley & Sons, 1997.

[22] B. Máckowiak, F. Matějka, and M. Wiederholt. Rational inattention: A review. Unpublished Manuscript, 2021.

[23] C. Martin and D. Lowery. mrna vaccines: intellectual property landscape. *Nature Reviews*, 2020.

[24] T. Mayskaya. Dynamic choice of information sources. *California Institute of Technology Social Science Working Paper, ICEF Working Paper WP9/2019/05*, 2020.

[25] J. Mensch. Cardinal representations of information. *Available at SSRN 3148954*, 2018.

[26] J. Miao and H. Xing. Dynamic discrete choice under rational inattention. Technical report, Working Paper, 2020.

[27] S. Morris and P. Strack. The wald problem and the equivalence of sequential sampling and static information costs. *Unpublished manuscript, June*, 2017.

[28] G. Moscarini and L. Smith. The optimal level of experimentation. *Econometrica*, 69(6):1629–1644, 2001.

[29] L. Pomatto, P. Strack, and O. Tamuz. The cost of information. *arXiv preprint arXiv:1812.04211*, 2020.

[30] H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.

[31] J. Schneider and C. Wolf. On risk and time pressure: When to think and when to do. *Available at SSRN 3490935*, 2020.

[32] C. A. Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690, 2003.

[33] J. Steiner, C. Stewart, and F. Matějka. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553, 2017.

[34] J. Stenger, F. Gamboa, and M. Keller. Optimization of quasi-convex function over product measure sets. *SIAM Journal on Optimization*, 31(1):425–447, 2021.

[35] E. Torgersen. *Comparison of statistical experiments*, volume 36. Cambridge University Press, 1991.

[36] A. Wald. Foundations of a general theory of sequential decision functions. *Econometrica*, 15 (4):279–313, 1947. ISSN 00129682, 14680262. URL `http://www.jstor.org/stable/1905331`.

[37] M. L. Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.

[38] W. Zhong. Optimal dynamic information acquisition. unpublished manuscript, 2019.

1.4 Appendix

1.4.1 Proof of Proposition 1

Fix a feasible $P \in \mathcal{I}$. From Proposition II.2.9 in Jacod and Shirayev (2013), there exists a unique predictable triple $(c^P, K^P; A^P)$, where: $c_t^P(\omega) \geq 0$, $K_{t,\omega}^P(dy)$ a transition kernel, with: $K_{t,\omega}^P(\{\omega_{t-}\}) = 0$ and $\int (y - \omega_{t-})^2 K_{t,\omega}^P(\mathrm{d}y) \leq 1$; and $A_t^P(\omega)$ an increasing process, P-a.s., such that:

$$\mathbb{E}_P\Big[G(\mu_{t+h}) - G(\mu_{t-})\big|\mathscr{F}_{t-}\Big] = \mathbb{E}_P\left[\int_t^{t+h}\left(c_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})K_s^P(\mathrm{d}y)\right)\mathrm{d}A_s^P \;|\mathscr{F}_{t-}\right]$$

We will show that the constraint necessitates that A_t^P is P-a.s. absolutely continuous with respect to the Lebesgue measure. Fix $t \geq 0$, and let $\tilde{\Omega} \subseteq \Omega$ such that A^P is not absolutely continuous, for $\omega \in \tilde{\Omega}$ at $t \geq 0$. That is, there is some $\varepsilon > 0$, such that $A_{t+h}^P(\omega) - A_t^P(\omega) \geq \varepsilon > 0$ for all $h > 0$. Let

$$\mathcal{G}_*^h(\omega) = \inf_{s\in[t,t+h]} c_s^P(\omega) \cdot G''(\omega_{s-}) + \int \mathcal{L}G(y,\omega_{s-})K_{s,\omega}^P(\mathrm{d}y) \geq 0$$

Then, we have:

$$\mathcal{G}_*^h(\omega) \cdot \left[A_{t+h}^P(\omega) - A_t^P(\omega)\right] \leq \int_t^{t+h}\left(c_s^P(\omega) \cdot G''(\omega_{s-}) + \int \mathcal{L}G(y,\omega_{s-})K_{s,\omega}^P(\mathrm{d}y)\right)\mathrm{d}A_s^P(\omega)$$

so in particular:

$$\varepsilon \cdot \mathcal{G}_*^h \cdot \mathbb{1}_{\tilde{\Omega}} \leq \mathcal{G}_*^h \cdot \left[A_{t+h}^P - A_t^P\right] \cdot \mathbb{1}_{\tilde{\Omega}} \leq \left(\int_t^{t+h} c_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})K_s^P(\mathrm{d}y)\right)\mathrm{d}A_s^P \cdot \mathbb{1}_{\tilde{\Omega}}$$

and hence:

$$\liminf_{h\to0+} \varepsilon \cdot \mathcal{G}_*^h \cdot \mathbb{1}_{\tilde{\Omega}} \leq \liminf_{h\to0+}\left(\int_t^{t+h} c_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})K_s^P(\mathrm{d}y)\right)\mathrm{d}A_s^P \cdot \mathbb{1}_{\tilde{\Omega}}$$

From Fatou's lemma and the fact that all variables are positive since G is convex, we have:

$$\begin{aligned}
0 \leq \mathbb{E}_P\left[\liminf_{h\to0+} \varepsilon \cdot \mathcal{G}_*^h \cdot \mathbb{1}_{\tilde{\Omega}} \;|\mathscr{F}_{t-}\right] &\leq \mathbb{E}_P\left[\liminf_{h\to0+}\left(\int_t^{t+h}\left(c_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})K_s^P(\mathrm{d}y)\right)\mathrm{d}A_s\right) \cdot \mathbb{1}_{\tilde{\Omega}} \;|\mathscr{F}_{t-}\right] \\
&\leq \mathbb{E}_P\left[\liminf_{h\to0+}\int_t^{t+h}\left(c_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})K_s^P(\mathrm{d}y)\right)\mathrm{d}A_s\big|\mathscr{F}_{t-}\right] \\
&\leq \liminf_{h\to0+} \mathbb{E}_P\left[\int_t^{t+h}\left(c_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})K_s^P(\mathrm{d}y)\right)\mathrm{d}A_s\big|\mathscr{F}_{t-}\right] \\
&\leq \lim_{h\to0+} \kappa \cdot h = 0
\end{aligned}$$

where the last inequality follows form the fact that P is feasible. When some strictly positive amount of information is acquired over $[t, t+h]$ for all $h > 0$, $\liminf_{h\to0+} \mathcal{G}_*^h > 0$, and the inequalities above necessitate $P(\tilde{\Omega}) = 0$. That is, A^P is P-a.s. absolutely continuous to the Lebesgue measure on $\mathbb{R}^+$.

Consequently, we may consider the pair of differential characteristics (α^P, F^P), where $(\alpha_t^P, F_t^P) \in$

$\mathbb{R}^+ \times \mathcal{M}(\mu_{t-})$, $P \times \mathrm{d}t$-a.s. From Theorem II.2.42 in Jacod and Shirayev (2013), we can write the constraint at time $t \geq 0$ as:

$$\mathbb{E}_P\left[\int_t^{t+h}\left(\alpha_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})\,\mathrm{d}F_s^P(y)\right)\mathrm{d}s\,\Big|\mathscr{F}_{t-}\right] \leq \kappa \cdot h, \quad \text{for all } h > 0$$

Having $\alpha_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})\,\mathrm{d}F_s^P(y) \leq \kappa$, for $P \times \mathrm{d}t$ a.a. $s \geq 0$ is clearly sufficient for the constraint to hold. We show it is also necessary.

First, by the mean value theorem:

$$\frac{1}{h}\int_t^{t+h}\left(\alpha_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})\,\mathrm{d}F_s^P(y)\right)\mathrm{d}s \to \alpha_t^P \cdot G''(\mu_{t-}) + \int \mathcal{L}G(y,\mu_{t-})\,\mathrm{d}F_t^P(y)$$

By applying Taylor's theorem with Langrange remainder to $\mathcal{L}G(y,x)$, we have, $\mathcal{L}G(y,x) \leq \frac{1}{2}\sup_{z\in[\eta,1-\eta]} G''(z)(y-x)^2 = B < \infty$ by continuity of G'', and the compactness of $[\eta, 1-\eta]$. Hence, we have:

$$\begin{aligned}\frac{1}{h}\int_t^{t+h}\left|\alpha_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})\,\mathrm{d}F_s^P(y)\right|\mathrm{d}s &\leq \frac{1}{h}\int_t^{t+h}\left(\alpha_s^P + \int (y-\mu_{s-})^2\,\mathrm{d}F_s^P(y)\right)\cdot B\,\mathrm{d}s \\ &\leq \frac{1}{h}\int_t^{t+h} C \cdot B\,\mathrm{d}s = C \cdot B\end{aligned}$$

where C is a bound on $\int (y-\mu_{s-})^2\,\mathrm{d}F_s^P(y)$ and α_s^P, where one may need to stop first before the process exits $[\eta, 1-\eta]$, for B to apply. Thus the sequence is also bounded. By the dominated convergence theorem we have:

$$\alpha_t^P \cdot G''(\mu_{t-}) + \int \mathcal{L}G(y,\mu_{t-})\,\mathrm{d}F_t^P(y) = \lim_{h\to 0+}\frac{1}{h}\mathbb{E}_P\left[\int_t^{t+h}\left(\alpha_s^P \cdot G''(\mu_{s-}) + \int \mathcal{L}G(y,\mu_{s-})\,\mathrm{d}F_s^P(y)\right)\mathrm{d}s\,\Big|\mathscr{F}_{t-}\right] \leq \kappa$$

which proves necessity.

1.4.2 Proof of Theorem 1

We split the proof in various steps, starting with a determination of the (HJB) equation, and proceeding to prove the properties of the optimal jumps.

Before we give the characterization we collect some properties of the value function.

Lemma 1. *The value function V is convex, and continuously differentiable on* $[0,1]$.

Proof. Online Appendix. □

Next, we give the definition of a viscosity solution to:

$$\rho(V-U)(x) - \sup_{(\alpha,F)\in\mathcal{I}(x)} \alpha \cdot V''(x) + \int \mathcal{L}V(y,x)\,\mathrm{d}F(y) = 0 \quad \text{and} \quad V = U \quad \text{on } \{0,1\}$$

A continuous function $V : [0,1] \to \mathbb{R}^+$ is a viscosity solution to the above if:[28]

[28]See Fleming and Soner (2006), Chapter II.4.

- For all $v \in C^2(0,1)$, such that $V \geq v$ and $V - v$ has a global minimum of 0, at $\hat{x} \in (0,1)$:

$$\rho(v-U)(\hat{x}) - \sup_{(\alpha,F)\in\mathcal{I}(x)} \alpha \cdot v''(\hat{x}) + \int \mathcal{L}v(y,\hat{x})\,\mathrm{d}F(y) \geq 0 \tag{super}$$

- For all $v \in C^2(0,1)$, such that $V \leq v$ and $V - v$ has a global maximum of 0, at $\hat{x} \in (0,1)$:

$$\rho(v-U)(\hat{x}) - \sup_{(\alpha,F)\in\mathcal{I}(x)} \alpha \cdot v''(\hat{x}) + \int \mathcal{L}v(y,\hat{x})\,\mathrm{d}F(y) \leq 0 \tag{sub}$$

We note the following lemma, which implies in particular the compactness of the constraint-set:

Lemma 2. *For any $x \in (0,1)$, we have:*

$$\sup_{(\alpha,F)\in\mathcal{I}(x)} \left\{\alpha + \int \min\left\{|y-x|,(y-x)^2\right\}\mathrm{d}F(y)\right\} \leq K < +\infty \tag{K}$$

Proof. Since $G \in C^3(0,1)$, we can apply Taylor's Theorem with the Langrange form for the remainder to $\mathcal{L}G(y,x)$ to get:

$$\mathcal{L}G(y,x) = \frac{1}{2}G''(\zeta_y)\cdot(y-x)^2 \quad \text{for some } \zeta_y \text{ between } x \text{ and } y$$

We let $B = \inf_{z\in(0,1)} G''(z) > 0$. For $(\alpha,F) \in \mathcal{I}(x)$ we have:

$$\begin{aligned}\kappa \geq \alpha\cdot G''(x) + \int \mathcal{L}G(y,x)\mathrm{d}F(y) &= \alpha\cdot G''(x) + \frac{1}{2}\int G''(\zeta_y)(y-x)^2\mathrm{d}F(y)\\ &\geq \alpha\cdot B + \frac{1}{2}B\int (y-x)^2\mathrm{d}F(y)\\ &\geq \frac{1}{2}B\cdot\left(\alpha + \int (y-x)^2\mathrm{d}F(y)\right)\end{aligned}$$

Consequently, for $K = \frac{2\cdot\kappa}{B}$ we have:

$$\alpha + \int (y-x)^2\mathrm{d}F(y) \leq K < +\infty \quad \text{for all } (\alpha,F)\in\mathcal{I}(x)$$

The claim follows since $|y-x| < 1$.

□

We now show the viscosity solution property.

Proposition 2. *The value function V, is a viscosity solution to:*

$$\rho(V-U)(x) = \sup_{(\alpha,F)\in\mathcal{I}(x)} \alpha\cdot V''(x) + \int \mathcal{L}V(y,x)\,\mathrm{d}F(y)$$

with boundary conditions $V(1) = U(1)$ and $V(0) = U(0)$.

Proof. We proceed to show the two properties:

(i) Super-solution: Fix $x_0 \in (0,1)$ and take $v \in C^2(0,1)$ such that $V \geq v$ and $V(x_0) = v(x_0)$. From the definition of the value V we have for any stopping-time $\sigma > 0$,

$$\begin{aligned}
0 &= \sup_{P\in\mathcal{I}} \mathbb{E}_P\left[e^{-\rho\sigma}V(\mu_h) - V(x) + \int_0^\sigma e^{-\rho t}\rho U(\mu_{t-})\mathrm{d}t\right] \\
&\geq \sup_{P\in\mathcal{I}} \mathbb{E}_P\left[e^{-\rho\sigma}v(\mu_h) - v(x) + \int_0^\sigma e^{-\rho t}\rho U(\mu_{t-})\mathrm{d}t\right] \\
&= \sup_{P\in\mathcal{I}} \mathbb{E}_P\left[\int_0^\sigma e^{-\rho r t}v'(\mu_{t-})\mathrm{d}\mu_t\right] \\
&+ \mathbb{E}_P\left[\int_0^\sigma e^{-\rho t}\rho\big(U(\mu_{t-}) - v(\mu_{t-})\big)\mathrm{d}t\right] \\
&+ \mathbb{E}_P\left[\int_0^\sigma e^{-\rho t}v''(\mu_{t-})\cdot\alpha_t^P\mathrm{d}t\right] \\
&+ \mathbb{E}_P\left[\int_0^\sigma e^{-\rho t}\int \mathcal{L}v(y,\mu_{t-})\mathrm{d}F_t^P(y)\ \mathrm{d}t\right]
\end{aligned}$$

We set $\sigma = \sigma_\varepsilon \wedge h$, where $\sigma_\varepsilon = \inf\{t|\mu_t \notin [\varepsilon, 1-\varepsilon]\}$ for $\varepsilon > 0$, and $h > 0$. The first term is 0 since the integrand is a true martingale, and the fact that v' is bounded on $[\varepsilon, 1-\varepsilon]$. We have that for all $P \in \mathcal{I}$:

$$\begin{aligned}
0 &\geq \frac{1}{h}\mathbb{E}_P\left[\int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}\rho\big(U(\mu_t) - v(\mu_{t-})\big)\mathrm{d}t\right] \\
&+ \frac{1}{h}\mathbb{E}_P\left[\int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}v''(\mu_{t-})\cdot\alpha_t^P + \int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}\int \mathcal{L}v(y,\mu_{t-})\mathrm{d}F_t^P(y)\ \mathrm{d}t\right]
\end{aligned} \tag{I}$$

For sufficiently small $h > 0$, $\sigma_\varepsilon \wedge h = h$ a.s. Hence, the random variables:

$$\begin{aligned}
&\frac{1}{h}\int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}\rho\big(U(\mu_{t-}) - v(\mu_{t-})\big) + v''(\mu_{t-})\cdot\alpha_t^P + \int \mathcal{L}v(y,\mu_{t-})\mathrm{d}F_t^P(y)\ \mathrm{d}t \\
&\to \rho\big(U(\mu_{0-}) - v(\mu_{0-})\big) + v''(\mu_{0-})\cdot\alpha_0^P + \int \mathcal{L}v(y,\mu_{0-})\mathrm{d}F_0^P(y)\ \text{ a.s.}
\end{aligned}$$

Next, we show that they are bounded. Indeed, we have:

$$
\begin{aligned}
&\left|\frac{1}{h}\int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}\rho\big(U(\mu_{t-})-v(\mu_{t-})\big)+v''(\mu_{t-})\cdot\alpha_t^P+\int \mathcal{L}v(y,\mu_{t-})\mathrm{d}F_t^P(y)\;\mathrm{d}t\right|\\
&\le \frac{1}{h}\int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}\rho\cdot\sup{}_{z\in[\varepsilon,1-\varepsilon]}\big|U(z)-v(z)\big|+\sup_{z\in[\varepsilon,1-\varepsilon]}|v''(z)|\cdot\alpha_t^P+\int|\mathcal{L}v(y,\mu_{t-})|\mathrm{d}F_t^P(y)\;\mathrm{d}t\\
&\le \frac{1}{h}\int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}\rho\cdot B_1+B_2\cdot\alpha_t^P+\int\frac{1}{2}B_2\cdot(y-\mu_{t-})^2\mathrm{d}F_t^P(y)\;\mathrm{d}t\\
&\le \frac{1}{h}\int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}\rho\cdot B_1+B_2\cdot\sup_{(\alpha,F)\in\mathcal{I}(\mu_{t-})}\left(\alpha+\int\cdot(y-\mu_{t-})^2\mathrm{d}F(y)\right)\;\mathrm{d}t\\
&\le \frac{1}{h}\int_0^{\sigma_\varepsilon\wedge h} e^{-\rho t}\rho\cdot B_1+B_2\cdot K\;\mathrm{d}t\\
&\le \frac{1}{h}\int_0^{h}\rho\cdot B_1+B_2\cdot K\;\mathrm{d}t=\rho\cdot B_1+B_2\cdot K<+\infty
\end{aligned}
$$

Thus, by taking limits in (I) and using the dominated convergence theorem we have:

$$
0\ge\rho\big(U(x_0)-v(x_0)\big)+v''(x_0)\cdot\alpha_0^P+\int\mathcal{L}v(y,x_0)\mathrm{d}F_0^P(y)
$$

Taking supremum over the (α_0^P,F_0^P) and re-arranging we get:

$$
\rho\big(v(x_0)-U(x_0)\big)-\sup_{(\alpha,F)\in\mathcal{I}(x_0)}v''(x_0)\cdot\alpha+\int\mathcal{L}v(y,x_0)\mathrm{d}F(y)\ge 0
$$

(ii) Sub-solution: Fix $x_0\in(0,1)$ and take $v\in C^2(0,1)$ such that $V\le v$ and $V(x_0)=v(x_0)$. From the definition of the value, and arguing as above we have for any stopping-time $\sigma>0$:

$$
\begin{aligned}
0&=\sup_{P\in\mathcal{I}}\mathbb{E}_P\left[e^{-\rho\sigma}V(\mu_h)-V(x)+\int_0^\sigma e^{-\rho t}\rho U(\mu_{t-})\mathrm{d}t\right]\\
&\le\sup_{P\in\mathcal{I}}\mathbb{E}_P\left[e^{-\rho\sigma}v(\mu_h)-v(x)+\int_0^\sigma e^{-\rho t}\rho U(\mu_{t-})\mathrm{d}t\right]\\
&=\sup_{P\in\mathcal{I}}\mathbb{E}_P\left[\int_0^\sigma e^{-\rho t}\rho\big(U(\mu_{t-})-v(\mu_{t-})\big)\mathrm{d}t\right]\\
&+\mathbb{E}_P\left[\int_0^\sigma e^{-\rho t}v''(\mu_{t-})\cdot\alpha_t^P\mathrm{d}t\right]\\
&+\mathbb{E}_P\left[\int_0^\sigma e^{-\rho t}\int\mathcal{L}v(y,\mu_{t-})\mathrm{d}F_t^P(y)\;\mathrm{d}t\right]
\end{aligned}
$$

We suppose that the sub-solution property fails so that:

$$
\rho\big(v(x_0)-U(x_0)\big)-\sup_{(\alpha,F)\in\mathcal{I}(x_0)}v''(x_0)\cdot\alpha+\int\mathcal{L}v(y,x_0)\mathrm{d}F(y)>0
$$

From Berge's maximum theorem,[29] the LHS is continuous at x_0 and we can find a radius $\eta > 0$ and an $\varepsilon > 0$ such that:

$$\rho(v-U)(x) - \sup_{(\alpha,F)\in\mathcal{I}(x)} v''(x)\cdot\alpha + \int \mathcal{L}v(y,x)\mathrm{d}F(y) \geq \varepsilon \quad \forall x : |x_0 - x| < \eta$$

But then for $\sigma = \inf\{t : |\mu_t - x_0| \geq \eta\}$, denoting $\mathcal{A}^{(\alpha,F)}v(x) = v''(x)\cdot\alpha + \int \mathcal{L}v(y,x)\mathrm{d}F(y)$, we have:

$$\begin{aligned}
0 &\leq \sup_{P\in\mathcal{I}} \mathbb{E}_P\left[\int_0^\sigma e^{-\rho t}\Big[-\rho(v-U)(\mu_{t-}) + \mathcal{A}^{(\alpha_t^P,F_t^P)}v(\mu_{t-})\Big]\mathrm{d}t\right] \\
&\leq \sup_{P\in\mathcal{I}} \mathbb{E}_P^{\hat{x}}\left[\int_0^\sigma e^{-\rho t}\Big[-\varepsilon + \mathcal{A}^{(\alpha_t^P,F_t^P)}v(\mu_t) - \sup_{(\alpha,F)\in\mathcal{I}(\mu_{t-})}\mathcal{A}^{(\alpha,F)}v(\mu_t)\Big]\mathrm{d}t\right] \\
&\leq \sup_{P\in\mathcal{I}} -\varepsilon\mathbb{E}_P^{\hat{x}}\left[\int_0^\sigma e^{-\rho t}\mathrm{d}t\right] < 0
\end{aligned}$$

where the last inequality follows from the fact that no instantaneous jump causing $P(\sigma = 0) = 1$ is feasible. This is a contradiction yielding the claim.

□

Next we prove the following lemma. To economize on notation we view the choice of parameters $(\alpha, F) \in \mathcal{I}(x)$ as a choice of an operator $\mathcal{A} \in \mathcal{I}(x)$. Moreover, we define $\mathcal{E}(x) = \mathbb{R}^+ \times \mathcal{M}(x)$.

Lemma 3. *For any test-function $v \in C^2(0,1)$,*

$$\sup_{\mathcal{A}\in\mathcal{I}(x)} \mathcal{A}v(x) := \sup_{(\alpha,F)\in\mathcal{I}(x)}\left\{\alpha\cdot v''(\hat{x}) + \int \mathcal{L}v(y,\hat{x})\, dF(y)\right\} = \sup_{y\neq x}\frac{\mathcal{L}v(y,x)}{\mathcal{L}G(y,x)}\cdot\kappa$$

Proof. To prove this we first appeal to the saddle-point characterization of the optimal solution to this problem. First, $\mathcal{A}\to\mathcal{A}v(x)$ and $\mathcal{A}\to\mathcal{A}G(x)$ are linear and the set $\mathcal{E}(x)$ is convex. Moreover, there is clearly an element $\mathcal{A}\in\mathcal{E}(x)$ for which $\mathcal{A}G(x) < \kappa$, and the cone in $\mathbb{R}$ is closed with non-empty interior. By the Langrange Multiplier Theorems in Luenberger (1997, p.219 and 221),

$$\sup_{\mathcal{A}\in\mathcal{I}(x)} \mathcal{A}v(x) = L(\mathcal{A}^*, \lambda^*, x)$$

where $(\mathcal{A}^*, \lambda^*)$ forms a saddle-point of the functional over the domain $\mathcal{E}(x)\times\mathbb{R}^+$:

$$L(\mathcal{A},\lambda,x) = \mathcal{A}v(x) - \lambda\mathcal{A}G(x) + \lambda\cdot\kappa$$

That is,

$$L(\mathcal{A},\lambda^*,x) \leq L(\mathcal{A}^*,\lambda^*,x) \leq L(\mathcal{A}^*,\lambda,x) \quad \text{ for all } \mathcal{A}\in\mathcal{E}(x), \lambda\geq 0$$

Next, we provide such a saddle-point. Firstly, define $\lambda^*(x) = \sup_{\mathcal{A}}\frac{\mathcal{A}v(x)}{\mathcal{A}G(x)}$. We will show that:

[29] The correspondence $\mathcal{I} : (0,1) \rightrightarrows \mathbb{R}^+ \times \mathcal{M}$ is continuous and compact-valued. Compactness follows from Lemma 2 and the results in Liu and Neufeld (2018). See that paper for the definition of the topology on $\mathcal{M}(x)$.

$$\lambda^*(x) = \sup_{y \neq x} \frac{\mathcal{L}v(y,x)}{\mathcal{L}G(y,x)}$$

that is, the supremum can be taken over single-jump processes. Firstly, because we can always approximate a diffusion by an appropriate choice of jump measures, we have:

$$\lambda^*(x) = \sup_{(0,F)\in\mathcal{E}(x)} \frac{\int \mathcal{L}v(y,x)\mathrm{d}F(y)}{\int \mathcal{L}G(y,x)\mathrm{d}F(y)} = \sup_{(0,P)\in\mathcal{E}(x)} \frac{\int \mathcal{L}v(y,x)\mathrm{d}P(y)}{\int \mathcal{L}G(y,x)\mathrm{d}P(y)}$$

where P are probability measures. The second equality holds because the value of the ratio is independent of the total mass of the measures F.

Next, let $S^x = [0,x)\cup(x,1]$, and $\Delta(S^x)$ the set of Borel probability measures on S^x. The space S^x is Polish as a G_δ subset of a Polish space. The extreme points are given by $\mathrm{ext}\Delta(S^x) = \{\delta_y : y \neq x\}$.[30] Since the ratio above is a quasiconvex, lower semicontinuous functional of $P \in \Delta(S^x)$, we have by Theorem 3.2 in Stenger, Gamboa, and Keller (2021):

$$\sup_{P\in\Delta(S^x)} \frac{\int \mathcal{L}v(y,x)\mathrm{d}P(y)}{\int \mathcal{L}G(y,x)\mathrm{d}P(y)} = \sup_{P\in\mathrm{ext}\Delta(S^x)} \frac{\int \mathcal{L}v(y,x)\mathrm{d}P(y)}{\int \mathcal{L}G(y,x)\mathrm{d}P(y)} = \sup_{y\neq x} \frac{\mathcal{L}v(y,x)}{\mathcal{L}G(y,x)}$$

which proves the claim.

Next, define:

$$\mathcal{A}^* = \begin{cases} \varphi^* \cdot \delta_{y^*} \text{ with } y^* \in \arg\max_{y\neq x} \frac{\mathcal{L}v(y,x)}{\mathcal{L}G(y,x)} \text{ and } \varphi^* = \frac{\kappa}{\mathcal{L}G(y^*,x)}, & \text{if } \lambda^*(x) > \frac{v''(x)}{G''(x)} \\ \\ \alpha^* = \frac{\kappa}{G''(x)}, & \text{if } \lambda^*(x) = \frac{v''(x)}{G''(x)} \end{cases}$$

Notice that $\mathcal{A}^*G(x) = \kappa$ and $\mathcal{A}^*v(x) - \lambda^*(x)\mathcal{A}^*G(x) = 0$.[31] Therefore, $L(\mathcal{A}^*,\lambda^*,x) = \lambda^*(x)\cdot\kappa$. We argue that $(\mathcal{A}^*,\lambda^*)$ is a saddle-point.

For arbitrary $\mathcal{A}$, from the definition of λ^* (we omit the dependence on x for clarity):

$$L(\mathcal{A},\lambda^*) = \mathcal{A}v - \lambda^*\cdot\mathcal{A}G + \lambda^*\kappa \leq \lambda^*\cdot\kappa = L(\mathcal{A}^*,\lambda^*,x)$$

Moreover, for arbitrary λ:

$$\begin{aligned} L(\mathcal{A}^*,\lambda) = \mathcal{A}^*v - \lambda\mathcal{A}G + \lambda\cdot\kappa &= \mathcal{A}^*v - \lambda^*\mathcal{A}G + \lambda\cdot(\kappa - \mathcal{A}^*G) + \lambda^*\mathcal{A}^*G \\ &= \lambda\cdot(\kappa - \mathcal{A}^*G) + \lambda^*\mathcal{A}^*G \\ &= \lambda^*\kappa \end{aligned}$$

Consequently, the pair $(\mathcal{A}^*,\lambda^*)$ defined above is a saddle-point for L.

□

To conclude with the characterization of the value, we put Lemma 1, Proposition 2, and Lemma 3 together to conclude that the value is a C^1 solution to:

[30] See Stenger, Gamboa, and Keller (2021), Theorem 2.1. Here we are taking the trivial moment class.

[31] Note that the constructed $\mathcal{A}^*$ achieves the supremum defining λ^*.

$$\rho(V-U)(x) = \sup_{y\neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)} \cdot \kappa, \quad \text{with } V = U \text{ on } \{0,1\}$$

Properties

We proceed to prove the properties of the optimal jumps. Let $\bar{V}''(x) = \limsup_{y\to x} \frac{\mathcal{L}V(y,x)}{(y-x)^2}$, and note that a diffusion is optimal if and only if $\frac{\bar{V}''(x)}{G''(x)} = \sup_{y\neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)}$. We first have:

Lemma 4. *Diffusion can only be optimal on a set of Lebesgue measure zero, i.e.: the set:*

$$\mathcal{D} = \left\{ x : \rho(V-U)(x) = \frac{\bar{V}''(x)}{G''(x)} \cdot \kappa \right\} \quad \textit{contains no intervals}$$

Proof. Suppose that diffusion is optimal on some interval. This interval will contain a further subinterval on which either $(V-U)' > 0$ or $(V-U)' < 0$. We consider the case $(V-U)' > 0$ since the other is similar, and denote such a subinterval by (a,b). Then the value V coincides with the solution to the equation:

$$\rho(v-U)(x) = \frac{v''(x)}{G''(x)} \cdot \kappa \qquad v(a) = V(a), v(b) = V(b)$$

This equation has a C^3-solution in the interior (a,b), and by differentiating we get:

$$\rho(v-U)'(x) = \frac{\kappa}{G''(x)^2} \cdot \Big(v'''(x) \cdot G''(x) - G'''(x) \cdot v''(x) \Big) > 0 \tag{1.1}$$

From the hypothesis that no optimal jump exists on (a,b) we get that for all $\varepsilon > 0$:

$$\begin{aligned}
\frac{v(x+\varepsilon) - v(x) - v'(x)\cdot\varepsilon}{G(x+\varepsilon) - G(x) - G'(x)\cdot\varepsilon} &< \frac{v''(x)}{G''(x)} \Rightarrow \\
\frac{v(x+\varepsilon) - v(x) - v'(x)\cdot\varepsilon}{v''(x)} &< \frac{G(x+\varepsilon) - G(x) - G'(x)\cdot\varepsilon}{G''(x)} \Rightarrow \\
\frac{\frac{1}{2}v''(x)\cdot\varepsilon^2 + \frac{1}{6}v'''(x)\cdot\varepsilon^3 + o(\varepsilon^3)}{v''(x)} &< \frac{\frac{1}{2}G''(x)\cdot\varepsilon^2 + \frac{1}{6}G'''(x)\cdot\varepsilon^3 + o(\varepsilon^3)}{G''(x)} \Rightarrow \\
\frac{1}{2}\varepsilon^2 + \frac{1}{6}\frac{v'''(x)}{v''(x)}\cdot\varepsilon^3 + o(\varepsilon^3) &< \frac{1}{2}\varepsilon^2 + \frac{1}{6}\frac{G'''(x)}{G''(x)}\cdot\varepsilon^3 + o(\varepsilon^3) \Rightarrow \\
\frac{v'''(x)}{v''(x)} + \frac{o(\varepsilon^3)}{\varepsilon^3} &< \frac{G'''(x)}{G''(x)} + \frac{o(\varepsilon^3)}{\varepsilon^3} \Rightarrow \\
\frac{v'''(x)}{v''(x)} &\leq \frac{G'''(x)}{G''(x)}
\end{aligned}$$

This contradicts (1.1). In the case where the $(V-U)' < 0$ on that interval, we repeat the above for $x - \varepsilon$, to get a contradiction.

$\square$

Next we will prove the first property (i) in Theorem 1.

Proposition 3 (location)**.** *Optimal jumps leave the information gain unchanged:*

$$\rho(V-U)(y^*(x)) = \rho(V-U)(x)$$

Proof. We will show that $\lambda^*(y^*(x)) = \lambda^*(x)$, which by the (HJB) equation is equivalent to the claim. Firstly, we show that $\lambda^*(y^*(x)) \geq \lambda^*(x)$. Since $y^*(x) \neq x$, jumping to x is feasible from $y^*(x)$ and we must have:

$$\lambda^*(y^*(x)) = \sup_{y \neq y^*(x)} \frac{\mathcal{L}V(y, y^*(x))}{\mathcal{L}G(y, y^*(x))} \geq \frac{\mathcal{L}V(x, y^*(x))}{\mathcal{L}G(x, y^*(x))}$$

Moreover, for any differentiable function ψ, we have:

$$\mathcal{L}\psi(x,y) + \mathcal{L}\psi(y,x) = [\psi'(y) - \psi'(x)] \cdot (y-x)$$

Consequently,

$$\frac{\mathcal{L}V(x, y^*(x))}{\mathcal{L}G(x, y^*(x))} = \frac{-\mathcal{L}V(y^*(x), x) + [V'(y^*(x)) - V'(x)] \cdot (y^*(x) - x)}{-\mathcal{L}G(y^*(x), x) + [G'(y^*(x)) - G'(x)] \cdot (y^*(x) - x)} \tag{*}$$

Next, we consider optimizing:

$$\max_{y \neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)}$$

Since an optimal jump exists at x, the FOC necessitates that any maximizer $y^*(x)$ satisfies:

$$\big(V'(y^*(x)) - V'(x)\big) - \big(G'(y^*(x)) - G'(x)\big) \cdot \frac{\mathcal{L}V(y^*(x), x)}{\mathcal{L}G(y^*(x), x)} = 0 \Rightarrow$$

$$\big(V'(y^*(x)) - V'(x)\big) = \lambda^*(x) \cdot \big(G'(y^*(x)) - G'(x)\big) \tag{FOC}$$

From the (FOC), and the optimality of $y^*(x)$ yielding, $\mathcal{L}V(y^*(x), x) = \lambda^*(x) \cdot \mathcal{L}G(y^*(x), x)$ we may re-write (*) as:

$$\frac{\mathcal{L}V(x, y^*(x))}{\mathcal{L}G(x, y^*(x))} = \frac{-\lambda^*(x) \cdot \mathcal{L}G(y^*(x), x) + \lambda^*(x)[G'(y^*(x)) - G'(x)] \cdot (y^*(x) - x)}{-\mathcal{L}G(y^*(x), x) + [G'(y^*(x)) - G'(x)] \cdot (y^*(x) - x)} = \lambda^*(x)$$

We therefore have $\lambda^*(y^*(x)) \geq \lambda^*(x)$.

Now suppose that $\lambda^*(y^*(x)) > \lambda^*(x)$. In other words,

$$\sup_{y \neq y^*(x)} \frac{\mathcal{L}V(y, y^*(x))}{\mathcal{L}G(y, y^*(x))} = \lambda^*\big(y^*(x)\big) > \lambda^*(x) = \sup_{y \neq x} \frac{\mathcal{L}V(y,x)}{\mathcal{L}G(y,x)}$$

We will show that the original jump to $y^*(x)$ was not optimal. First, there exists an $\varepsilon > 0$ and some $\hat{y} \neq y^*(x)$ such that:

$$\frac{\mathcal{L}V(\hat{y}, y^*(x))}{\mathcal{L}G(\hat{y}, y^*(x))} > \lambda^*\big(y^*(x)\big) - \varepsilon > \lambda^*(x) \tag{**}$$

We will show that $\hat{y}$ constitutes a profitable deviation to $y^*(x)$ starting from x.

Indeed, for any differentiable ψ:

$$\mathcal{L}\psi(\hat{y},x) = \mathcal{L}\psi(\hat{y},y^*(x)) + \mathcal{L}\psi(y^*(x),x) + \big(\psi'(y^*(x)) - \psi'(x)\big)\cdot(\hat{y}-y^*(x))$$

and therefore:

$$\frac{\mathcal{L}V(\hat{y},x)}{\mathcal{L}G(\hat{y},x)} = \frac{\mathcal{L}V(\hat{y},y^*(x)) + \mathcal{L}V(y^*(x),x) + \big(V'(y^*(x)) - V'(x)\big)\cdot(\hat{y}-y^*(x))}{\mathcal{L}G(\hat{y},y^*(x)) + \mathcal{L}G(y^*(x),x) + \big(G'(y^*(x)) - G'(x)\big)\cdot(\hat{y}-y^*(x))} \qquad (***)$$

Consequently, by using the (FOC) we can re-write (***) as:

$$\begin{aligned}\frac{\mathcal{L}V(\hat{y},x)}{\mathcal{L}G(\hat{y},x)} &= \frac{\frac{\mathcal{L}V(\hat{y},y^*(x))}{\mathcal{L}G(\hat{y},y^*(x))}\cdot\mathcal{L}G(\hat{y},y^*(x)) + \lambda^*(x)\cdot\mathcal{L}G(y^*(x),x) + \lambda^*(x)\big(G'(y^*(x)) - G'(x)\big)\cdot(\hat{y}-y^*(x))}{\mathcal{L}G(\hat{y},y^*(x)) + \mathcal{L}G(y^*(x),x) + \big(G'(y^*(x)) - G'(x)\big)\cdot(\hat{y}-y^*(x))}\\ &> \frac{\lambda^*(x)\cdot\mathcal{L}G(\hat{y},y^*(x)) + \lambda^*(x)\cdot\mathcal{L}G(y^*(x),x) + \lambda^*(x)\big(G'(y^*(x)) - G'(x)\big)\cdot(\hat{y}-y^*(x))}{\mathcal{L}G(\hat{y},y^*(x)) + \mathcal{L}G(y^*(x),x) + \big(G'(y^*(x)) - G'(x)\big)\cdot(\hat{y}-y^*(x))}\\ &= \lambda^*(x)\end{aligned}$$

where the inequality follows from (**). This contradicts the optimality of $y^*(x)$, and we must therefore have $\lambda^*(y^*(x)) = \lambda^*(x)$.

□

We proceed to the next property:

Proposition 4 (direction)**.** *If an optimal jump exists at x, then there is an optimal jump $y^*(x)$ in the direction of increasing information gain:*

$$(V-U)'(x) > 0 \Rightarrow y^*(x) - x > 0$$

and

$$(V-U)'(x) < 0 \Rightarrow y^*(x) - x < 0$$

Moreover, if $(V-U)'(x) = 0$, and an optimal jump exists, then there must exist optimal jumps on either side of x.

Proof. Suppose an optimal jump exists at $x \in (0,1)$, and let $\mathcal{Y}^*(x)$ be the set of maximizers. From Clarke's envelope theorem (Clarke, 1983, Corollary 1, p. 242) and the discussion at the bottom of p. 245,[32]

$$\rho(V-U)'(x) \in \text{co}\left\{\bigcup_{y^*\in\mathcal{Y}^*(x)} \frac{\kappa}{\mathcal{L}G(y^*(x),x)}\cdot\Big(\lambda^*(x)\cdot G''(x) - Q\Big)\cdot(y^*-x) : Q \in \partial V'(x)\right\} \qquad \text{(ET)}$$

where $\partial V'(x)$ is the generalized gradient of V' at x:

$$\partial V'(x) = \big\{\lim_n V''(x_n) \mid x_n \to x, \text{ and } x_n \notin \mathcal{K}_{V'}, \forall n \in \mathbb{N}\big\}$$

[32]See also Morand, Reffett, and Tarafdar (2015).

where $\mathcal{K}_{V'}$ is the set of points where V' is not differentiable.

Furthermore, we show that for each $Q \in \partial V'(x)$, $\lambda^*(x) \cdot G''(x) - Q \geq 0$. Indeed, let $(x_n)_n$ be the sequence such that: $\lim_n V''(x_n) = Q$. For each n we have:

$$\lambda^*(x_n) \geq \frac{\mathcal{L}V(x_n+\varepsilon, x_n)}{\mathcal{L}G(x_n+\varepsilon, x_n)} = \frac{V''(x_n) + o(\varepsilon^2)/\varepsilon}{G''(x_n) + o(\varepsilon^2)/\varepsilon} \Rightarrow \lambda^*(x_n) \cdot G''(x_n) \geq V''(x_n)$$

By taking limits and as $n \to +\infty$ we have: $\lambda^*(x) \cdot G''(x) \geq Q$.

Now suppose that $(V-U)'(x) > 0$. (ET) implies that there must exist $y^* \in \mathcal{Y}^*(x)$ and $Q \in \partial V'(x)$, such that $y^* - x > 0$ and $\lambda^*(x) \cdot G''(x) - Q > 0$. We conclude similarly for the case $(V-U)'(x) < 0$. Finally, if $(V-U)'(x) = 0$, not all optimal jumps can be in the same direction, otherwise 0 cannot be in the convex hull.

□

Finally, we show the last property.

Proposition 5 (frequency vs impact)**.** *Let $\varphi^*(x)$ be the optimal frequency. Then φ^* is increasing in the information gain:*

$$(V-U)(x_1) \leq (V-U)(x_2) \Rightarrow \varphi^*(x_1) \leq \varphi^*(x_2)$$

Proof. In each neighbourhood where an optimal jump exists and is continuous, the FOC holds and we get:

$$\big(V'(y^*(x)) - V'(x)\big) = \lambda^*(x) \cdot \big(G'(y^*(x)) - G'(x)\big)$$

We use the Implicit Function Theorem to obtain:

$$V''(y^*(x)) \cdot \frac{dy^*(x)}{dx} - V''(x) = \lambda^*(x) \cdot \left(G''(y^*(x))\frac{dy^*(x)}{dx} - G''(x)\right) + (\lambda^*)'(x) \cdot \big(G'(y^*(x)) - G'(x)\big)$$

which we re-arrange to get:

$$\frac{dy^*(x)}{dx} \cdot \Big(V''(y^*(x)) - \lambda^*(x) \cdot G''(y^*(x))\Big) = V''(x) - \lambda^*(x)G''(x) + (\lambda^*)'(x) \cdot \big(G'(y^*(x)) - G'(x)\big) \quad (*)$$

Moreover, from the Envelope Theorem:

$$(\lambda^*)'(x) = \frac{1}{\mathcal{L}G(y^*(x), x)} \cdot \Big(\lambda^*(x) \cdot G''(x) - V''(x)\Big) \cdot (y^*(x) - x) \Rightarrow$$

$$V''(x) - \lambda^*(x) \cdot G''(x) = -(\lambda^*)'(x) \cdot \frac{\mathcal{L}G(y^*(x), x)}{y^*(x) - x}$$

Substituting into (*) yields:

$$\frac{dy^*(x)}{dx} \cdot \Big(V''(y^*(x)) - \lambda^*(x) \cdot G''(y^*(x))\Big) = (\lambda^*)'(x) \cdot \left(G'(y^*(x)) - G'(x) - \frac{\mathcal{L}G(y^*(x), x)}{y^*(x) - x}\right) \Rightarrow$$

$$\frac{dy^*(x)}{dx} \cdot \Big(V''(y^*(x)) - \lambda^*(x) \cdot G''(y^*(x))\Big) \cdot (y^*(x) - x) = (\lambda^*)'(x) \cdot \Big(G'(x) - G'(y^*(x)) - G'(y^*(x)) \cdot \big(x - y^*(x)\big)\Big)$$

Re-arranging, and using the expression of λ^* in terms of the information gain we get:

$$-\frac{dy^*(x)}{dx} \cdot \big(y^*(x) - x\big) = \frac{\rho}{\kappa}(V-U)'(x) \cdot \mathcal{L}G\big(x, y^*(x)\big) \cdot \Big(\lambda^*(x) \cdot G''(y^*(x)) - V''(y^*(x))\Big)^{-1}$$

The term $\mathcal{L}G\big(x, y^*(x)\big) > 0$ by convexity of G. Additionally, the SOC necessitates:

$$V''(y^*(x)) - \lambda^*(x) \cdot G''(y^*(x)) \leq 0$$

Hence, we express the derivative as:

$$-\frac{dy^*(x)}{dx} = \frac{(V-U)'(x)}{y^*(x) - x} \cdot K(x, y^*(x)) \tag{**}$$

with $K(x, y^*(x)) > 0$.

From Proposition 4, $(V-U)'(x)$ has the same sign as $y^*(x) - x$. Therefore, the RHS in (**) is always positive, which implies that:

$$\frac{dy^*(x)}{dx} < 0$$

We use this to show that impact $|y^*(x) - x|$ is increasing in the distance from a local maximum of the information gain. Let $\hat{x}$ be a local maximum of $(V-U)$. Consider, $x_1 < x_2 < \hat{x}$. We have $y^*(x_1) > y^*(x_2)$, and:

$$|y^*(x_1) - x_1| = y^*(x_1) - x_1 > y^*(x_2) - x_1 > y^*(x_2) - x_2 = |y^*(x_2) - x_2|$$

where the first equality follows from Proposition 4. Similarly, for $\hat{x} < x_2 < x_1$, $y^*(x_2) > y^*(x_1)$, and

$$|y^*(x_1) - x_1| = x_1 - y^*(x_1) > x_1 - y^*(x_2) > x_2 - y^*(x_2) = |y^*(x_2) - x_2|$$

This proves the assertion. Since frequency and impact are inversely related at the optimum we have that:

$$(V-U)(x_1) < (V-U)(x_2) \Rightarrow |\hat{x} - x_1| > |\hat{x} - x_2| \Rightarrow \varphi^*(x_1) < \varphi^*(x_2)$$

which completes the proof.

□

1.5 Online Appendix

1.5.1 Proof of Lemma 1

Convexity follows from the fact that in a hypothetical scenario where the DM observes some free information generating a spread $\lambda x_1 + (1-\lambda)x_2 = x$ over their beliefs, the DM can still wait to observe which belief materializes and then implement the same experiment they found optimal at belief x. The prior-concavity of the information measure is sufficient for this strategy to be feasible. The formal proof proceeds by approximation from discrete-time problems and is postponed to Appendix 1.5.2.

We now prove differentiability. Fix some $x \in (0, 1)$. Consider the following experiment: a signal arrives at some exponential time σ independently of the state, and implements a 50:50 spread:

$F = \frac{1}{2} \cdot \delta_{x+\varepsilon} + \frac{1}{2} \cdot \delta_{x-\varepsilon}$. The frequency of $\varphi(x, \varepsilon)$ of the arrival time σ is selected to bind the constraint, that is:

$$\varphi(x, \varepsilon) = \frac{\kappa}{\frac{1}{2}\Big[G(x+\varepsilon) - G(x)\Big] + \frac{1}{2}\Big[G(x-\varepsilon) - G(x)\Big]}$$

Call the resulting distribution of beliefs Q. For any time interval $h > 0$, we have:

$$V(x) \geq \mathbb{E}_Q\left[\int_0^h e^{-\rho t}\rho U(\mu_t)\,\mathrm{d}t + e^{-\rho h}V(\mu_h)\right]$$

since implementing this experiment until h time passes and then following the optimal experiment must be unprofitable. We re-arrange and expand to get for h small enough:

$$0 \geq \int_0^h e^{-\rho t}\rho U(x) - \rho V(x) + \varphi(x, \varepsilon) \cdot \left(\frac{1}{2} \cdot V(x+\varepsilon) + \frac{1}{2} \cdot V(x+\varepsilon) - V(x)\right)\mathrm{d}t$$

Now dividing by h and sending $h \to 0$, and substituting for $\varphi(x, \varepsilon)$ we have:

$$\frac{\rho}{\kappa}(V-U)(x) \cdot \Big[G(x+\varepsilon) - G(x) + G(x-\varepsilon) - G(x)\Big] \geq \Big[V(x+\varepsilon) - V(x) + V(x-\varepsilon) - V(x)\Big]$$

Dividing by $\varepsilon > 0$ we get:

$$\frac{\rho}{\kappa}(V-U)(x) \cdot \left[\frac{G(x+\varepsilon) - G(x)}{\varepsilon} + \frac{G(x-\varepsilon) - G(x)}{\varepsilon}\right] \geq \left[\frac{V(x+\varepsilon) - V(x)}{\varepsilon} + \frac{V(x-\varepsilon) - V(x)}{\varepsilon}\right]$$

We take limits $\varepsilon \to 0+$ on both sides to get:

$$\frac{\rho}{\kappa}(V-U)(x) \cdot \Big[G'_+(x) - G'_-(x)\Big] \geq \Big[V'_+(x) - V'_-(x)\Big]$$

where we notice that the limits exist by convexity of V and G.

Since G is differentiable the LHS is zero and we have:

$$V'_-(x) \geq V'_+(x)$$

Consequently, V is differentiable at x. As a convex function we have that it is also continuously differentiable, which completes the proof.

1.5.2 Proof of Convexity

First, we show the following:

Suppose G satisfies

$$x \longmapsto x^2(1-x)^2 \cdot G''(x) \quad \text{concave} \tag{C}$$

Then, the induced information measure $I(\Sigma|x)$ over Blackwell experiments $\Sigma = (P^1, P^0)$, is concave in the prior $x \in (0, 1)$ for all Σ. Conversely, if $I(\Sigma \mid x)$ is concave in the prior for all Σ, then (C) holds.

Proof. Note that any Blackwell experiment Σ is equivalent to a choice of distribution F of the relative density $Z = \frac{\mathrm{d}P^0}{\mathrm{d}P^1}$, with $\mathbb{E}_F[Z] = 1$.

By definition $I(\Sigma \,|\, x) = \mathbb{E}_P[G(\mu) - G(x)]$ where P is the unconditional distribution of posteriors generated by Σ. We have:

$$\begin{aligned} I(\Sigma \mid x) = \mathbb{E}_P[G(\mu) - G(x)] &= x \cdot \mathbb{E}^1[G(\mu) - G(x)] + (1-x) \cdot \mathbb{E}^0[G(\mu) - G(x)] \\ &= x \cdot \mathbb{E}^1[G(\mu) - G(x)] + (1-x) \cdot \mathbb{E}^1[Z \cdot \big(G(\mu) - G(x)\big)] \\ &= \mathbb{E}^1\left[\Big(G(\mu) - G(x)\Big) \cdot \Big(x + (1-x) \cdot Z\Big)\right] \\ &= \mathbb{E}_F\left[\left(G\left(\frac{x}{x + (1-x) \cdot Z}\right) - G(x)\right) \cdot \Big(x + (1-x) \cdot Z\Big)\right] \end{aligned}$$

where we have used the fact that $\mu = \mu(x, Z) := \frac{x}{x+(1-x)\cdot Z}$. We now differentiate twice with respect to the prior $x \in (0, 1)$:

$$\begin{aligned} I_{xx}(\Sigma \mid x) &= \mathbb{E}_F\left[G''\left(\frac{x}{x + (1-x) \cdot Z}\right) \cdot \frac{Z^2}{\big(x + (1-x) \cdot Z\big)^3} - \big(x + (1-x) \cdot Z\big) \cdot G''(x) + 2 \cdot \big(Z - 1\big) \cdot G'(x)\right] \\ &= \mathbb{E}_F\left[G''\left(\frac{x}{x + (1-x) \cdot Z}\right) \cdot \frac{Z^2}{\big(x + (1-x) \cdot Z\big)^3} - G''(x)\right] \end{aligned}$$

where the equality follows from $\mathbb{E}_F[Z] = 1$. Insisting on $I_{xx}(\Sigma \,|\, x) \leq 0$ for all Blackwell experiments Σ, we must have for all F, with $\mathbb{E}_F[Z] = 1$:

$$G''(x) \geq \mathbb{E}_F\left[G''\left(\frac{x}{x + (1-x) \cdot Z}\right) \cdot \frac{Z^2}{\big(x + (1-x) \cdot Z\big)^3}\right]$$

We now express this back in terms of the unconditional distribution over beliefs. We multiply through by $x^2 \cdot (1-x)^2$ to get:

$$
\begin{aligned}
x^2 \cdot (1-x)^2 \cdot G''(x) &\geq \mathbb{E}_F \left[G''\left(\frac{x}{x+(1-x)\cdot Z}\right) \cdot \frac{Z^2 x^2 \cdot (1-x)^2}{\big(x+(1-x)\cdot Z\big)^3} \right] \\
&= \mathbb{E}_F \left[G''\left(\frac{x}{x+(1-x)\cdot Z}\right) \cdot \frac{Z^2 x^2}{\big(x+(1-x)\cdot Z\big)^2} \cdot \frac{(1-x)^2}{\big(x+(1-x)\cdot Z\big)} \right] \\
&= \mathbb{E}_F \left[G''\Big(\mu(x,Z)\Big) \cdot \mu(x,Z)^2 \cdot \frac{(1-x)^2}{\big(x+(1-x)\cdot Z\big)} \right] \\
&= \mathbb{E}_F \left[G''\Big(\mu(x,Z)\Big) \cdot \mu(x,Z)^2 \cdot \frac{(1-x)^2 \cdot \big(x+(1-x)\cdot Z\big)}{\big(x+(1-x)\cdot Z\big)^2} \right] \\
&= x \cdot \mathbb{E}_F \left[G''\Big(\mu(x,Z)\Big) \cdot \mu(x,Z)^2 \cdot \frac{(1-x)^2}{\big(x+(1-x)\cdot Z\big)^2} \right] + \\
&+ (1-x) \cdot \mathbb{E}_F \left[G''\Big(\mu(x,Z)\Big) \cdot \mu(x,Z)^2 \cdot \frac{(1-x)^2}{\big(x+(1-x)\cdot Z\big)^2} \cdot Z \right] \\
&= x \cdot \mathbb{E}_F \left[G''\Big(\mu(x,Z)\Big) \cdot \mu(x,Z)^2 \cdot \big(1-\mu(x,Z)\big)^2 \right] + \\
&+ (1-x) \cdot \mathbb{E}_F \left[G''\Big(\mu(x,Z)\Big) \cdot \mu(x,Z)^2 \cdot \big(1-\mu(x,Z)\big)^2 \cdot Z \right] \\
&= \mathbb{E}_P \left[G''(\mu) \cdot \mu^2 \cdot (1-\mu)^2 \right]
\end{aligned}
$$

where we have repeatedly expressed the posteriors in terms of the relative density Z and the prior x, as $\mu(x, Z)$; and transformed back to the unconditional distribution over posteriors P. We conclude that prior-concavity for each Blackwell experiment Σ is equivalent to:

$$
x^2 \cdot (1-x)^2 \cdot G''(x) \geq \sup_P \mathbb{E}_P \left[\mu^2 \cdot (1-\mu)^2 \cdot G''(\mu) \right] = \operatorname{cav}\big[\mu^2 \cdot (1-\mu)^2 \cdot G''(\mu)\big](x)
$$

Consequently, since $x^2 \cdot (1-x)^2 \cdot G''(x)$ must be always above its concave envelope, it must itself be concave.

□

Consider the constrained problem where the DM can only change the information strategy after a time-interval of length $h > 0$, and moreover, can only act at the end of that interval even if the opportunity arises in the interim. The Bellman equation for this problem is given by:

$$
V_h(\mu_t) = \sup_P \ \mathbb{E}[p_h \cdot V_h(\mu_{t+h}) + (1-p_h) \cdot U(\mu_{t+h})] \quad \text{subject to } \ \mathbb{E}_P[G(\mu_{t+h}) - G(\mu_t)] \leq \kappa \cdot h \quad (\text{HJB}_h)
$$

where $p_h = e^{-\rho \cdot h}$ is the probability of no decision in period $t+h$, and $V_h(\cdot)$, is the value conditional on no decision. The problem in (HJB_h) is seemingly *less* constrained on the information strategies as the capacity can be violated somewhere on the interval $[t, t+h]$. However, for each distribution P which is feasible in (HJB_h), there exists a continuous-time information strategy which is feasible, and implements the distribution P at $t+h$, which follows from Zhong (2019). Indeed, we have:

Let $P \in \Delta([0,1])$, with $\mathbb{E}_P[\mu] = x$, such that:

$$\mathbb{E}_P[G(\mu) - G(x)] \leq \kappa \cdot h$$

Then, there exists a posterior process $(\mu_t)_{0 \leq t \leq h}$, such that $\mu_0 = x$ and $\mu_h \sim P$ and:

$$\alpha_t \cdot G(\mu_{t-}) + \int \mathcal{L}G(y, \mu_{t-}) \mathrm{d}F_t(y) \leq \kappa \quad \text{for all } t \in [0,h]$$

Proof. Straightforward application of Lemma S.3 in Zhong (2019). □

In what follows, any choice of P in (HJB_h) will be understood being implemented feasibly as per the lemma above. Such a restriction does not impact optimal payoffs in (HJB_h).

Moreover, the choice variable in problem (HJB_h) is a Bayes' consistent distribution of posteriors $P \in \Delta[0,1]$, with $\mathbb{E}_P[\mu_{t+h}] = \mu_t$. Equivalently, one can view this problem as choosing a Blackwell experiment $\Sigma = (P^1, P^0)$, that is, a pair of conditional distributions over signals. We can re-write the objective as:

$$\begin{aligned} &\mu_t \cdot \mathbb{E}^1[p_h \cdot V_h(\mu_{t+h}) + (1-p_h) \cdot U(\mu_{t+h})] + (1-\mu_t) \cdot \mathbb{E}^0[p_h \cdot V_h(\mu_{t+h}) + (1-p_h) \cdot U(\mu_{t+h})] = \\ &= \mathbb{E}^1\Big[\big(p_h \cdot V_h(\mu_{t+h}) + (1-p_h) \cdot U(\mu_{t+h})\big) \cdot \big(\mu_t + (1-\mu_t) \cdot Z\big)\Big] \end{aligned}$$

where $Z = \frac{\mathrm{d}P^0}{\mathrm{d}P^1}$ is the relative density corresponding to the experiment. Notice, that choosing a Blackwell experiment is equivalent to choosing a distribution F for the variable Z, under the constraint $\mathbb{E}_F[Z] = 1$. Observing that $\mu_{t+h} = \frac{\mu_t}{\mu_t + (1-\mu_t)\cdot Z}$, we re-write the Bellman equation as:

$$V_h(\mu_t) = \sup_F \ \mathbb{E}_F\Big[W_h\big(\mu_t, Z \mid V_h\big)\Big] \quad \text{subject to } I_h(F|\mu_t) \leq \kappa \cdot h \tag{HJB_h^*}$$

where $I_h(F|\mu_t)$ is the induced information measure over Blackwell experiments implied by G, and for any function $f : [0,1] \to \mathbb{R}$,

$$W_h(\mu, z \mid f) = \left[p_h \cdot f\left(\frac{\mu}{\mu + (1-\mu)\cdot z)}\right) + (1-p_h) \cdot U\left(\frac{\mu}{\mu + (1-\mu)\cdot z)}\right)\right] \cdot \big(\mu + (1-\mu) \cdot z\big)$$

We construct the Bellman operator $T_h : B \to B$ for (HJB_h^*).

$$T_h f(\mu) = \sup_F \ \mathbb{E}_F\Big[W_h\big(\mu, Z \mid f\big)\Big] \quad \text{subject to } I_h(F|\mu) \leq \kappa \cdot h$$

We have the following:

The operator T_h is a contraction for all $h > 0$.

Proof. We check Blackwell's sufficient conditions: First, monotonicity $f \geq g \Rightarrow T_h f \geq T_h g$ is evident since $W_h\big(\mu, z \mid f\big) \geq W_h\big(\mu, z \mid g\big)$. Second, discounting, $T_h(f+a) \leq T_h f + \beta \cdot a$, for $\beta < 1$, works for any $\beta \leq p_h < 1$.

□

As a consequence we get the existence of a unique fixed-point V_h, which is the value of this h-constrained problem. We wish to prove that prior-concavity of $I_h(F|\mu)$ in μ, is sufficient for convexity of the value.

T_h maps convex to convex functions.

Proof. Pick a convex function f. Firstly, we have that $W_h(\mu, z \mid f)$ is convex in μ. Indeed, for a convex function g, the function:

$$\mu \longmapsto g\left(\frac{\mu}{\mu + (1-\mu)\cdot z}\right)\cdot\big(\mu + (1-\mu)\cdot z\big)$$

is convex in μ,[33] and $W_h\big(\mu, z \mid f\big)$ is a convex combination of such terms when f is convex.

Now, suppose that starting from μ the DM first observes a *free* signal σ, generating a spread $\mu = \xi\cdot\mu_1 + (1-\xi)\cdot\mu_2$, and let F^* be the optimal experiment at belief μ. We consider the strategy $\sigma\circ F^*$, where the DM ignores that free information and chooses F^* regardless of the realization of σ.

The quantity of information generated by this strategy (composite experiment) is given by:

$$I_h(\sigma|\mu) + \xi\cdot I_h(F^*|\mu_1) + (1-\xi)\cdot I_h(F^*|\mu_2)$$

and since σ is free the constraint applies only to the second part:

$$\xi\cdot I_h(F^*|\mu_1) + (1-\xi)\cdot I_h(F^*|\mu_2) \le \kappa\cdot h$$

Since $I_h(F|\cdot)$ is concave we have:

$$\xi\cdot I_h(F^*|\mu_1) + (1-\xi)\cdot I_h(F^*|\mu_2) \le I_h(F^*|\xi\cdot\mu_1 + (1-\xi)\cdot\mu_2) = I_h(F^*|\mu) \le \kappa\cdot h$$

where the last inequality follows from the feasibility of F^* at μ. Hence, $\sigma\circ F^*$ is feasible whenever σ is free.

Evaluating the payoff from $\sigma\circ F^*$ we have:

$$\begin{aligned}
&\xi\cdot\mathbb{E}_{F^*}\Big[W_h\big(\mu_1, Z\mid f\big)\Big] + (1-\xi)\cdot\mathbb{E}_{F^*}\Big[W\big(\mu_2, Z\mid f\big)\Big]\\
&\ge \mathbb{E}_{F^*}\Big[W_h\big(\xi\cdot\mu_1 + (1-\xi)\cdot\mu_2, Z\mid f\big)\Big]\\
&= \mathbb{E}_{F^*}\Big[W_h\big(\mu, Z\mid f\big)\Big] = T_h f(\mu)
\end{aligned}$$

where the inequality follows from convexity of $W(\cdot, z\mid f)$. Since the DM can generally optimize after observing σ, we have:

$$\xi\cdot T_h f(\mu_1) + (1-\xi)\cdot T_h f(\mu_2) \ge T_h f(\mu)$$

Since, ξ, μ, μ_1, μ_2 were arbitrary, we conclude that $T_h f$ is convex.

[33] For instance, for a twice-differentiable g we have that the second derivative of the mapping is:

$$\mu \longmapsto \frac{z^2\cdot g''\left(\frac{\mu}{\mu+(1-\mu)\cdot z}\right)}{(\mu + (1-\mu)\cdot z)^3} \ge 0$$

□

From Lemma 1.5.2 we have that T_h maps convex to convex functions, and since the set of convex functions is closed, we must have that the fixed point of T_h is convex. That is, the h-constrained value, V_h, is convex.

Finally, we show uniform convergence of $V_h \to V$ as $h \to 0+$. Notice, that since the DM cannot act exactly when they want, and decreasing h relaxes their problem, $V_h \geq V_{h'}$ whenever $h \leq h'$. That is, the family of value functions: $\{V_h\}_{h>0}$ is monotone-decreasing in h. We have the following:

We have $V_h \to V$ uniformly as $h \to 0+$. Moreover, V is convex.

Proof. Firstly, we note that $V(x) = \sup_{h>0} V_h(x)$, by imagining the DM first optimizes the information strategy for any arbitrary delay $h > 0$, and then picks the delay optimally. Next, for each $x \in [0, 1]$, the family of real numbers $\{V_h(x)\}_{h>0}$ is bounded and increasing as $h \to 0+$. Therefore, $V_h(x) \to \sup_{h>0} V_h(x) = V(x)$; that is, V_h converges pointwise to V. Since V is continuous, from Dini's Theorem the convergence is also uniform. Finally, we have that V is convex as the limit of convex functions.

□

Chapter 2

Smart Contracts and the Coase Conjecture

2.1 Introduction

The Coase Conjecture is a manifestation of the striking consequences of the lack of intertemporal commitment power: if the monopolist can post prices frequently, she clears the market quickly, at prices close to the lowest possible willingness-to-pay even when most consumers have high valuation for the good.[1] The goal of this paper is to examine the extent to which this conclusion is robust to considering more complex selling mechanisms than just price-posting. Motivated by *smart contracts* used in digital markets, we allow the seller to offer general dynamic contracts. Our main result is that if the monopolist has access to such contracts, the Coase Conjecture no longer holds.

In our model, there is a seller of a single good and a buyer. The buyer's valuation for the good is binary, high or low, and it is his private information[2]. We consider the case where the probability of high valuation is large enough for the static monopoly price to be the high valuation. Time is discrete and both parties discount the future at the same rate. In the initial period, the seller offers a contract from a space described below. If the buyer accepts the contract, it determines the probabilities of trade and the transfers in subsequent periods until it is replaced. At the beginning of each period, the seller decides whether to proceed with the current contract or to void it and offer a new one.

We use the analogy of a mediator to describe a typical contract from our contract space. In each period, both the seller and the buyer may send messages to the mediator. In turn, the mediator sends private (and possibly public) signals to the contracting parties and implements an allocation. Perhaps the most notable feature of such a contract is that the mediator can possess information which the seller does not. When the seller abandons a contract, she loses that information. This feature is shown to be the driving force of our main result.

Some aspects of our contract space are reminiscent of the technologies developed in relation to the aforementioned smart contracts. First, smart contracts are automated in the sense that they

[1] This phenomenon was first described by Coase (1972), and later formalized by Stokey (1982), Fudenberg, Levine and Tirole (1985), and Gul, Sonnenschein and Wilson (1986).

[2] We focus on the "gap case" and assume that the seller's production cost, normalized to be zero, is smaller than the low valuation. In the "no-gap case", Ausubel and Deneckere (1989) show that the Coase Conjecture fails even with posted prices.

execute trades without further consents from the contracting parties. Similarly, in our model, the allocation proposed by the mediator is implemented and cannot be renegotiated. Second, smart contracts in practice can be, and often are switched off just like the seller can abandon her current contract in our model. One of the reasons that contracts in digital markets are designed so that they can be switched off is to avoid the execution of unlawful transactions, for example, due to bankruptcy procedure against a contracting party.[3] Finally, we note that cryptographic encoding of a party's input can prevent the other contracting party to recover that input even if she has access to the contract's code. Such encoding also plays an important role in digital markets: smart contracts deployed on blockchain networks use cryptographically signed transactions.[4]

Since the seller's commitment power is limited, she may benefit from a small contract space. The reason is that removing contracts from the seller's action space makes the set of possible deviations shrink which, in turn, may enable the seller to stick with contracts which are advantageous from the ex ante perspectives. This can be seen most vividly by considering the scenario when each contract available to the seller specifies trading at the high valuation. In this case, the seller could achieve the full-commitment profit because, even though she maybe tempted to lower the price if there is no trade, she cannot do so. So, in order to model the consequences of limited commitment in a meaningful way, the contract space should be rich enough. To this end, we assume that the seller has access to all *simple and direct* contracts defined as follows. A contract is called simple and direct if the contract elicits the buyer's valuation in the initial period of its deployment and does not communicate with the contracting parties ever after. Our main result holds as long as the seller's contract space is rich enough to include all such contracts.

We note that our model assumes a certain amount of commitment power of the monopolist. Namely, if the contract in place determines an allocation, that allocation will be implemented and the seller cannot take further actions. The same assumption is maintained in the standard Coasian model: if the buyer is willing to buy at the posted price, trade will take place and the seller cannot reneg on the price.[5]

Our main result is that the monopolist's payoff is bounded away from the low valuation irrespective of the discount factor. We prove this result by constructing a simple and direct but suboptimal mechanism which never reveals any information to the seller. In the initial period, the buyer reports his valuation and the high type trades at a price less than his valuation with probability less than one. The low-type buyer does not trade in the first period. Even though the seller receives no signal from the mechanism, she updates her prior about the buyer's type whenever trade does not occur. In every subsequent period, the probability of trade is constant and does not depend on the buyer's type, so the seller's posterior remains the same unless there is sale. Furthermore, the price is the buyer's valuation. This means that the low-value buyer's payoff is zero and the high-value buyer earns rent only in the initial period. Finally, from the second period onwards, the discounted present value of the seller's payoff is larger than the low valuation.

If the seller abandons this contract, she loses its information content. The trading probabilities

[3]Another context in which a contracting party retains her right to void smart contracts is where the issuer deploys these contracts in her private blockchains. Examples for such issuers include Walmart, Comcast, Spotify, DHL, JPMorgan and MetLife.

[4]While communications in many digital trading platforms are public, there are examples for protocols which allow for private communications. We discuss these examples and the implementation of such private communications in the concluding section.

[5]McAdams and Schwarz (2007) and Akbarpour and Li (2020) consider static mechanism design problems where the principal has even less committment power and she cannot credibly promise to follow the rules of her own mechanism.

in this mechanism are specified so that the optimal full-commitment mechanism in all but the initial period is clearing the market at the low valuation. Since the seller's expected payoff is larger than the low valuation, the constraint guaranteeing that she does not abandon the mechanism is satisfied in each future period. The mechanism we construct may not be optimal: In the initial period, the seller might prefer to choose a different mechanism. But that would only imply that her equilibrium payoff is even larger than the one generated by the mechanism described above. Therefore, since the seller's expected profit is bounded away from the low valuation in our mechanism, her equilibrium payoff is also larger than the low valuation. That is, the Coase Conjecture fails.

From the methodoligical viewpoint, a contribution of our paper is the introduction of the aforementioned contracts into a dynamic principal-agent problem without intertemporal commitment. Following the tradition of Mechanism Design, we do not impose restrictions on the contract space and consider general dynamic contracts similar to those in the full-commitment benchmark. Of course, when the principal can commit to long-term contracts, the information revealed to her by the mechanism about the agent's prior communications is irrelevant. In contrast, when the principal lacks commitment power and re-optimizes in each period, such information is detrimental in shaping the future relationship with the agent. Consequently, the information revealed by the contract should be part of the optimal contract design problem and hence, we allow for contracts that store more information than the principal has access to. We speculate that this approach may turn out to be useful in analyzing dynamic mechanism design problems with limited commitment in various environments. The application of this idea to the problem of a durable-good monopolist merely clarifies that the Coase Conjecture is not only due to the seller's lack of commitment power but also to her restricted contract space.

Literature Review

The literature on dynamic contracting in the absence of commitment probably started with the papers by Laffont and Tirole (1988 and 1990). The authors offer two related yet distinct approaches to model such environments. The first one is to consider one-period contracts. In each period, the principal offers a contract which, if accepted by the agent, determines the allocation in that period as a function of contractible variables.[6] The second approach is to allow dynamic contracts which can be voided and replaced if both parties wish to do so. That is, equilibrium contracts must be *renegotiation-proof*.[7]

The methodological contribution of our paper is to put forward another approach of modeling limited commitment which appears to be new. In order to explore the consequences of the absence of commitment in the context of a principal-agent relationship, it is desirable to consider a setting which differs from the full-commitment benchmark only in the assumption regarding the principal's commitment power. In the full-commitment benchmark, the principal has access to dynamic contracts and has full bargaining power. Therefore, our model combines the two approaches of Laffont and Tirole (1988 and 1990) in the following way. On the one hand, the set of mechanisms is not restricted to be one-period ones and the principal has access to infinite-horizon dynamic contracts.

[6] Examples for recent papers analysing dynamic screening problems with short-term contracts include Gerardi and Maestri (2020), Beccuti and Möller (2018), Acharya and Ortner (2017) and Tirole (2016).

[7] Among others, Battaglini (2007) and Maestri (2017) generalize the results of Laffont and Tirole (1990) in various ways. Strulovici (2017) provides a foundation for renegotiation-proof contracts in a bargaining environment. Hart and Tirole (1988) and Breig (2019) compare the two modelling approaches in a dynamic buyer-seller relationship.

On the other hand, the principal can offer new contracts in each period and the agent's consent is not required to abandon the previous contract.

Doval and Skreta (2020a) also consider mechanism design problems with limited commitment. They generalize the approach in Laffont and Tirole (1988) and consider one-period contracts. Their mechanisms do not only determine allocations but can also reveal public information. The authors develop a Revelation Principle and show that the information revealed by a mechanism can be assumed to be the principal's posterior about the agent's type.[8] In their companion paper, Doval and Skreta (2020b) show that the Coase Conjecture still holds with such a contract space. Indeed, the authors demonstrate that in a Coasian environment, the seller optimally posts prices in each period.[9]

Our paper also contributes to the literature documenting failures of the Coase Conjecture in the 'gap case'. With multiple atomic buyers, Bagnoli, Salant and Swierzbinski (1989), von der Fehr and Kuhn (1995) and Montez (2013) show that the seller can maintain high posted prices until a trade occurs. Feinberg and Skrzypacz (2005) show that higher-order uncertainty can generate delay. Other papers demonstrate that the Coase Conjecture is not robust to the assumption that the seller's marginal cost of production is constant, see, for example, Kahn (1986), McAfee and Wiseman (2008), Karp (1993), and Ortner (2017).[10] Bulow (1982) argues that the monopolist benefits from renting the good rather than selling it.[11]

Another approach to break the Coase Conjecture is to allow the seller to intratemporally screen, e.g. by producing a variety, see Wang (1998), Takeyama (2002), Hahn (2006), Inderst (2008), or Board and Pycia (2014). A notable contribution by Nava and Schiraldi (2019) demonstrates that all these results are consistent with the Coasian logic in the following sense. The seller's limit payoff is the maximal static monopoly profit subject to the market-clearing condition.

The Coase conjecture has been also proved to fail when market deterioration is prevented by the arrival of new buyers or stochastically changing values. Important contributions include Sobel (1991), Biehl (2001), and Fuchs and Skrzypacz (2010).

Our work is also related to the literature on smart contracts. The term 'smart contract' was first coined by Nick Szabo in the mid-90's, whose prototypical example of a vending machine highlights the ideas of automatic execution and immutability. Since then, with the advent of bitcoin and the popularization of blockchain technologies such as Ethereum, interest in smart contracts has heightened. For some recent papers on the blockchain, see Huberman, Leshno and Moallemi (2021) who provide an insightful analysis of the Bitcoin Payment System and Abadi and Brunnermeier (2018) who study the impossibility of any distributed ledger to satisfy certain desiderata.

Recent research on smart contracts has explored how these contracts can enlarge the space of implementable economic outcomes. Cong and He (2019) study the effects of smart contracts on industrial organization, while Tinn (2018) studies how financial contracting may be affected.

[8] Bester and Strausz (2001) also develop a Revelation Principle in finite-horizon environments and finite type spaces.

[9] Lomys and Yamashita (2021) introduce a mediator into the model of Doval and Skreta (2020a) who controls the communication between the contracting parties. The mediator cannot be replaced by the principal and can possess information which the principal does not have. The authors demonstrate that such a mediator expands the set of implementable allocations of Doval and Skreta (2020a). See also Fanning (2021a, 2021b) who explores how a mediator who can withhold information can improve equilibrium outcomes in a reputational bargaining problem.

[10] Also related is the literature on obsolescence or imperfect durability of the good, see Bulow (1986), Waldman (1993), and Fudenberg and Tirole (1998).

[11] Hart and Tirole (1988) point out that the arguments of Bulow (1982) rely on buyer-anonymity and show and that renting may make the seller worse-off.

Bakos and Halaburda (2020) delineate the effects of enhanced information generation of technologies dubbed the Internet-of-Things, and the automatic execution offered by smart contracts in a simple contracting game. Finally, Holden and Malani (2018) examine the use of smart contracts in the context of the hold-up problem. Two key properties of smart contracts underpin all of the above papers: (i) enhanced commitment power—for example, through lowering enforcement costs via automatic execution, or preventing renegotiation of terms altogether; and (ii) better information—for example, by reducing state-verification costs.

While we recognize that restoring some commitment power is possibly the main reason for the popularity of smart contracts, our paper intends to provide a different perspective. We take the view of Laffont and Tirole (1988) that the lack of intertemporal commitment is a form of contractual incompleteness. In other words, contracting parties may refrain from signing long-term, binding contracts due to potential unforeseen or non-contractible contingencies even if such contracts were feasible.[12] Although we do not model these contingencies, our assumption that the seller cannot commit not to switch off a deployed contract embodies the idea that she prefers a contract allowing for discretion in the future.[13] Our main result suggests that smart contracts may turn out to be useful even in such environments because they can store information securely.

2.2 The Model

There is a seller of a durable, indivisible good and a buyer whose willingness-to-pay for the good is his private information. The buyer's valuation is either high, v_h, or low, v_l so that $v_h > v_l > 0$. The probability of high valuation, μ, is common knowledge. We assume that $\mu v_h > v_l$, so the static monopoly price is v_h. Time is discrete and indexed by 0, 1,... In the initial period, the seller offers a contract from the set $\mathcal{C}$ described below. This contract then determines the allocation, i.e., the probability of trade and the transfer, in every period unless it is replaced. In each subsequent period, the seller decides whether to proceed with the previous period's contract or to deploy a new one. If the seller deploys a new contract, then it will determine the allocation in that period as well in every future period until it is replaced. The game ends when the good is sold. We assume that both parties discount the future according to the common factor $\delta\,(\in(0,1))$. If the buyer's valuation is $v\,(\in\{v_l,v_h\})$, trade occurs in period T and the transfer is p_t at time t, then the payoffs of the buyer and seller are

$$\delta^T v - \sum_{t=0}^{\infty} \delta^t p_t \text{ and } \sum_{t=0}^{\infty} \delta^t p_t,$$

respectively. Moreover, both parties maximize their expected payoffs.

The Contract Space $\mathcal{C}$.— We describe a typical contract, c, from the seller's contract space $\mathcal{C}$. The contract specifies both the communication and the implemented allocation in each period when the contract is deployed. Formally, $c = \left(M_T^b, M_T^s, S_T^b, S_T^s, \mathbf{x}_T, \mathbf{p}_T, \rho_T\right)_{T=0}^{\infty}$, where M_T^b and M_T^s are the messages available to the buyer and the seller in a given period if the contract was already

[12] Unexpected software security vulnerabilities, bugs, novel types of attack threats, the need for upgrades, and regulatory risk, are among the reasons one may willingly retain discretion over aspects of a smart contract, preventing its absolute immutability in practice.

[13] Such control can be exercised via 'admin keys' retained by the issuer. It is worth noting that 12 out of the 15 most popular Decentralized Finance protocols, governed by smart contracts, have such 'admin keys' (https://cointelegraph.com/news/how-many-defi-projects-still-have-god-mode-admin-keys-more-than-you-think).

deployed T consecutive periods immediately preceding that period.[14] The sets S_T^b and S_T^s are the set of signals the buyer and the seller may receive privately. The functions $\mathbf{x}_T : \left(M_\gamma^b, M_\gamma^s\right)_{\gamma=0}^T \times \left(S_\gamma^b, S_\gamma^s\right)_{\gamma=0}^{T-1} \to [0,1]$ and $\mathbf{p}_T : \left(M_\gamma^b, M_\gamma^s\right)_{\gamma=0}^T \times \left(S_\gamma^b, S_\gamma^s\right)_{\gamma=0}^{T-1} \to \mathbb{R}$ specify the probability of trade and the transfer conditional on sale[15] as a function of histories of messages and signals. Finally, the function $\rho_T = \left(\rho_T^b, \rho_T^s\right) : \left(M_\gamma^b, M_\gamma^s\right)_{\gamma=0}^T \times \left(S_\gamma^b, S_\gamma^s\right)_{\gamma=0}^{T-1} \to \Delta\left(S_T^b, S_T^s\right)$ specifies the distributions of the signals revealed to the buyer and the seller as a function of the history of reports. To model the buyer's participation decision, we assume that, for each T, the buyer's message space, M_T^b, includes a special message, r, which triggers no trade. Sending this message is interpreted as *rejecting the contract*. If the buyer rejects the contract, $m_T^b = r$, then $\mathbf{x}_T = \mathbf{p}_T = 0$.[16] We say that the contract c is *actively deployed* in a given period, if the seller deploys c and the buyer does not reject it in that period. The seller's contract space $\mathcal{C}$ is a set of contracts described above.

Note first that the signals revealed to the contracting parties are assumed to be private. However, when the signals are perfectly correlated, they are effectively public. In fact, contracts are defined to be general enough to also allow the mixture of private and public communication; signals may have both private and public components. Second, despite the seller having no private information to start with, it is important to allow a contract to condition on the seller's messages. The reason is that the seller learns over time and when she decides to deploy a contract, she may benefit from inputting her posterior and making the implemented allocations dependent on it.

Simple and Direct Contracts.— We are not making any additional assumption on the contract space except that it contains all those mechanisms which ask the buyer to report her valuation in the initial period of deployment but involve no additional meaningful communication. We call such contracts *simple and direct* and define them formally below. Again, let us describe a typical simple and direct contract, d. First, if d is deployed repeatedly then sending the message r only triggers a one period of delay. So, the easiest way to describe d is to index the message and signal spaces as well as the allocations defining these contracts by the number of those consecutive periods of deployment in which the contract d was not rejected. More precisely, at each history, let τ denote the number of previous periods in which d was actively deployed since a different contract was deployed.[17] Then, with a slight abuse of notation, the contract d is defined by the collection $(\mathbf{x}_\tau, \mathbf{p}_\tau)_{\tau=0}^\infty$, where $\mathbf{x}_\tau : \{v_l, v_h\} \to [0,1]$ and $\mathbf{p}_\tau : \{v_l, v_h\} \to \mathbb{R}$. In the initial period of deployment, and in every other period in which $\tau = 0$, the buyer is asked to report his valuation, so $M_0^b = \{v_l, v_h, r\}$. If the buyer reports $v \in \{v_l, v_h\}$ then trade occurs with probability $\mathbf{x}_0(v)$ at price $\mathbf{p}_0(v)$. If the buyer sends the message r and the seller deploys d in the next period, the buyer's message space is again $\{v_l, v_h, r\}$ and the allocation is determined by $(\mathbf{x}_0, \mathbf{p}_0)$. After the buyer does not reject d and reports a valuation, he can only accept or reject the contract, that is, $M_\tau^b = \{a, r\}$ for all $\tau > 0$. The seller is only informed whether or not the buyer rejected the contract, that is, $S_\tau^s = \{a, r\}$ for all τ and $\rho_\tau^s\left(m_\tau^b\right) = r$ if, and only if, $m_\tau^b = r$. The seller does not communicate to the contract and the buyer does not receive any information, so the seller's message spaces and the buyer's signal spaces are singletons. The set of such simple and direct contracts is denoted by $\mathcal{D}$ and we assume that $\mathcal{D} \subset \mathcal{C}$.

[14]For an example, suppose that c is deployed at $t = 0, 2, 3$ but not at $t = 1$. Then, $T = 0$ at $t = 0, 2$ and $T = 1$ at $t = 3$.

[15]For notational simplicity, we assume that transfers are deterministic and paid only if there is trade. Allowing random transfers has no impact on our results.

[16]One may also find it natural to assume that the seller is informed about the buyer's rejection of the contract. Our main result holds irrespective of such an assumption.

[17]For example, if d was deployed at $t = 0, 1, 2, 3$ and was rejected only at $t = 1$, then $\tau = 2$ at $t = 3$.

We point out that the set $\mathcal{D}$ is different from the set of contracts one may wish to call *direct* in our environment. In general, a contract should be defined to be direct if its message spaces in each period of its deployment are rich enough to allow the seller and the buyer to report their private information. Since such a contract may send signals to both parties and information may also evolve in those periods when the contract is not deployed, a direct contract must allow the reporting of hierarchies of beliefs. For example, the seller's type includes his posterior about the buyer's valuation, her belief about the buyer's belief about her posterior, etc.

Equilibrium Concept and Existence.— We focus on Weak Perfect Bayesian Equilibria. That is, an equilibrium is defined as an assessment: a pair of system of beliefs and a (possibly mixed) strategy profile. The belief system specifies for each information set of the game a probability distribution over the set of nodes in that set, which is then interpreted as the belief of the contracting party who moves at that information set. An assessment is a Weak Perfect Bayesian Equilibrium if (i) the strategy profile is sequentially rational at each information set and (ii) beliefs are derived by Bayes' rule at those information sets which are reached with positive probability.

The concept of Weak Perfect Bayesian Equilibrium places little restrictions on the players' out-of-equilibrium beliefs. Since we provide a lower bound on the seller's equilibrium payoffs, one may suspect that this result is supported by constructing beliefs off the equilibrium path which may appear unreasonable. For example, if the seller believes that the buyer's willingness-to-pay is surely v_h whenever he rejects a contract, she would rationally offer a contract which specifies trade only at price v_h in subsequent periods. In fact, the seller may maintain this belief even after the buyer rejects contracts arbitrarily many times. This, in turn, may deter the buyer to reject an otherwise unattractive contract in the first place if it generates non-negative payoffs. We emphasize that our analysis does not rely on such arguments and we impose two further restrictions on the seller's off-equilibrium beliefs. First, we require the assessment to satisfy the "*no-signaling-what-you-don't know*" condition. In particular, the seller's posterior regarding the buyer's type cannot change after her own deviation. Second, in the spirit of the concept of Sequential Equilibrium,[18] special care is taken to construct the seller's beliefs so that they are limit points of beliefs derived by Bayes' rule along a sequence of totally mixed strategy profiles converging to the equilibrium strategy profile.

It is not hard to show that equilibria exist in a discretized version of our model, i.e., the set of contracts, the message and signal spaces are all finite[19]. We also prove existence for the case when the seller only has access to simple and direct contracts, that is, $\mathcal{C} = \mathcal{D}$, see the Online Appendix[20]. In the rest of this paper, we assume that an equilibrium exists irrespective of the discount factor.

Revelation Principle.— The concept of the Revelation Principle in our environment is slightly more subtle than in the case of full-commitment. In the latter case, standard revelation principles state that the contract space can be restricted to be a canonical class without the loss of generality. That is, irrespective of the contract space, the seller's payoff can never exceed that generated by the optimal canonical contract. As mentioned in the Introduction, when commitment power is limited, the seller may benefit from reducing her contract space. Consequently, whether or not one may restrict attention to a certain canonical class of contracts depends on the original space. Nevertheless, it is not hard to show that if the contract space is rich enough, a Revelation Principle similar to the standard one still holds. More precisely, if the contract space includes all direct

[18]Myerson and Reny (2020) discuss the difficulty to extend the definition of Sequential Equilibrium to games with infinite sets of signals and actions, and propose the new concept of Perfect Conditional ε-Equilibrium.

[19]See Fudenberg and Levine (1983).

[20]Available at: tinyurl.com/36tzxh3k.

contracts, each equilibrium outcome can be reproduced by an incentive compatible direct contract which the seller deploys in each period and the buyer never rejects it.

By no means does the aforementioned Revelation Principle imply that the incentive compatible direct contract can also be assumed to be simple. That is, the equilibrium direct contract may involve communication with the contracting parties in addition to eliciting the buyer's willingness-to-pay in the initial period. One may naively suspect that such additional communication can be dispensed with because the seller offers the same direct contract in each period irrespective of the signals she received. So, it appears that unifying all the information nodes in each period at which the seller offers the same contract would make no difference. Similarly, revealing information to the buyer would only make it harder to satisfy his incentive compatibility constraints. Unfortunately, this reasoning is incomplete and, in general, the equilibrium contract cannot be assumed to be simple and direct. The reason is that the seller may benefit from receiving signals because the change in her beliefs due to these signals may deter her from deviating to other contracts. This is because knowing that the seller received information about the buyer's willingness-to-pay, the buyer would reject certain off-equilibrium contracts in the future which, in turn, can enable the seller to stick with the equilibrium contract. Since the direct contract guaranteed by the Revelation Principle may involve complex communication and reporting hierarchies of beliefs, this concept does not seem to be operational in characterizing equilibrium outcomes. Our forthcoming analysis will not rely on this concept and hence, we do not present formal proofs for the statements above.

2.3 Main Result

In order to state our main theorem, let $\pi(\mathcal{C}, \delta)$ denote the supremum of the seller's payoff across all equilibria if the contract space is $\mathcal{C}$ and the discount factor is δ.

There exists a $\underline{\pi} > v_l$, such that for all $\delta \in (0, 1)$,

$$\pi(\mathcal{C}, \delta) \geq \underline{\pi}.$$

We remark that this theorem implies the failure of the Coase Conjecture: no matter how close the discount factor is to one, the largest equilibrium payoff of the seller is bounded from below by a constant, $\underline{\pi}$, which is larger than the low valuation, v_l.

The key to the arguments leading to the statement of Theorem 2.3 is to analyze a particular set of simple and direct contracts, coined as *abiding contracts*. The identifying feature of these contracts is that if they are actively deployed forever then (i) the buyer's continuation payoff is weakly positive in each period irrespective of his type and (ii) the seller's expected continuation payoff is larger than her full-commitment profit in all but the initial periods. The proof of the theorem consists of two steps. We first show that the seller's largest equilibrium payoff cannot be smaller than her payoff generated by any of the abiding contract. The second step is to construct an abiding mechanism which generates a payoff to the seller which is larger than v_l and does not depend on the discount factor.

Next, we define incentive compatible and abiding contracts formally.

Incentive Compatible Simple and Direct Contracts.— Whether the buyer has incentive to report his willingness-to-pay truthfully after accepting a simple and direct contract depends on what contracts he expects to be deployed in the future. Moreover, it also depends on the discount factor, δ.

In what follows, we define incentive compatibility conditional on the same contract being deployed forever. Before presenting the formal definition, recall that the allocation determined by a simple and direct contract, $(\mathbf{x}_\tau, \mathbf{p}_\tau)_0^\infty$, depends only on the initial report of the buyer. Observe that if a simple and direct contract is actively deployed in each period, the buyer's report, v, determines the unconditional probability of trade, $X_\tau(v)$, and the expected transfer, $P_\tau(v)$, in each period by

$$X_\tau(v) = \mathbf{x}_\tau(v)\,\Pi_{t=0}^{\tau-1}\left(1 - \mathbf{x}(v)_t\right) \text{ and } P_\tau(v) = \mathbf{p}_\tau(v)\,\mathbf{x}_\tau(v)\,\Pi_{t=0}^{\tau-1}\left(1 - \mathbf{x}(v)_t\right).$$

Vice versa, each simple and direct contract $d \in \mathcal{D}$ can be described by $(X_\tau, P_\tau)_{\tau=0}^\infty$, where $X_\tau : \{v_l, v_h\} \to [0,1]$ denotes the probability that trade occurs in period τ conditional d being actively deployed in each period and $P_\tau : \{v_l, v_h\} \to \mathbb{R}$ is the expected transfer in that period.

Note that if a simple and direct contract is deployed forever, the buyer may maximize his payoff by misreporting his type in the initial period of deployment and optimizing with respect to his rejection-acceptance strategy in the future. Let $U(v, \widehat{v}, d, \delta)$ denote the buyer's value if the contract d is deployed forever, the discount factor is δ, the buyer's valuation is v and he reported $\widehat{v}$ in the initial period. Recall that if the buyer rejects a simple and direct contract, he only induces a one-period delay. Therefore, he only rejects the contract if his continuation payoff is negative, in which case, he would reject it forever. Consequently,

$$U(v, \widehat{v}, d, \delta) = \sup_{T \geq 0} \sum_{t=0}^{T} \delta^t \left[X_t(\widehat{v})\, v - P_t(\widehat{v})\right],$$

where T denotes the time period after which the buyer rejects the contract forever. We are now ready to define incentive compatibility.

The contract $d = (X_\tau, P_\tau)_{\tau=0}^\infty \in \mathcal{D}$ is δ-incentive compatible if for $v \in \{v_l, v_h\}$

$$v \in \arg\max_{\widehat{v} \in \{v_l, v_h\}} U(v, \widehat{v}, d, \delta).$$

Abiding Contracts.— As mentioned before, we intend to call a contract abiding if, conditional on the contract being actively deployed forever, the buyer's continuation payoff is non-negative and the seller's continuation payoff exceeds her full-commitment profit in each period. More precisely, we require that an abiding contract specifies trading probabilities with each type of the buyer so that, after the initial period, the static monopoly price becomes the low valuation. That is, conditional on not trading, the seller becomes so pessimistic regarding the buyer's willingness-to-pay that she would optimally clear the market at price v_l. Before providing the formal definition, let us introduce an additional piece of notation. If an incentive compatible, simple and direct contract is actively deployed in each period then $\mu_t(d)$ denotes the posterior probability that the buyer's willingness-to-pay is v_h in period t.

The contract $d = (X_\tau, P_\tau)_{\tau=0}^\infty \in \mathcal{D}$ is δ-abiding if it is δ-incentive compatible and, in addition,
(i) $\sum_{t=T}^\infty \delta^{t-T}\left[X_t(v)\, v - P_t(v)\right] \geq 0$ for all $v \in \{v_l, v_h\}$, $T \geq 0$,
(ii) $\mu_t(d) \leq v_l / v_h$ for all $t \geq 1$, and
(iii) $\mu_T(d) \sum_{t=T}^\infty \delta^{t-T} P_t(v_h) + (1 - \mu_T(d)) \sum_{t=T}^\infty \delta^{t-T} P_t(v_l) \geq v_l$ for all $T \geq 1$.

Condition (i) implies that if a δ-abiding contract is deployed forever, accepting the contract in each period is an optimal strategy of the buyer if his discount factor is δ. Conditions (ii) and (iii) require that the static monopoly price is v_l and the seller's continuation value is larger than v_l in all but the initial period if d is actively deployed forever.

Let $v(d, \delta)$ denote the seller's payoff if the incentive compatible, simple and direct contract $d = (X_\tau, P_\tau)_{\tau=0}^{\infty} \in \mathcal{D}$ is actively deployed forever, that is,

$$v(d, \delta) = \mu \sum_{t=0}^{\infty} \delta^t P_t(v_h) + (1-\mu) \sum_{t=0}^{\infty} \delta^t P_t(v_l).$$

We are ready to state that the seller's value generated by any abiding contract is a lower bound on her largest equilibrium payoff.

Lemma 2.3.1. *Suppose that $d \in \mathcal{D}$ is a δ-abiding contract. Then $\pi(\mathcal{C}, \delta) \geq v(d, \delta)$.*

Let us explain the main arguments leading to this result. If the statement was false, the seller's payoff in each equilibrium would be strictly less than $v(d, \delta)$. Therefore, to prove the lemma, it is enough to argue that each such equilibrium can be modified so that, in the new equilibrium, the contract d is actively deployed forever. On the modified equilibrium path, the seller always deploys d and the buyer always accepts it. Off the equilibrium path the new equilibrium assessment is constructed based on the original equilibrium. In particular, the seller's payoff from offering a contract different from d in the initial period is the same as from offering that contract in the original equilibrium. Since the seller's payoff from offering any contract in the initial period is smaller than $v(d, \delta)$ in the original equilibrium, such deviations are not profitable. In subsequent periods, when d was already deployed a number of times, the seller's continuation payoff from offering it exceeds the full-commitment profit because d is abiding. Of course, if the seller abandons d and loses its information content, her continuation payoff cannot exceed the full-commitment profit. So even in later periods, the seller has no incentive to deviate from offering d. If the buyer ever rejects the contract, the seller's posterior belief remains the same.[21] Given this belief and that the buyer is expected to accept d, the seller rationally offers this contract even after many periods of rejection. In turn, knowing this, the buyer best-responds by accepting the seller's offer because d is abiding so it provides him with a non-negative continuation payoff irrespective of his willingness-to-pay.

Proof. We prove this lemma by contradiction. Suppose that the seller's payoff in each equilibrium is strictly smaller than $v(d, \delta)$. In what follows, we fix such an equilibrium and, by modifying it, we construct a new equilibrium so that the contract d is actively deployed forever and, consequently, the seller's payoff is $v(d, \delta)$, yielding a contradiction.

Let us first define the new equilibrium assessment at those information sets which are reached by paths along which no contract was offered but d. The seller always offers d and the buyer never rejects it. Moreover, in the initial period, the buyer reports his type truthfully. So, the equilibrium path, and hence payoffs, are determined by the repeated active deployment of d. If the seller moves at such an information set, her belief is defined to be $\mu_\tau(d)$ if d was actively deployed τ times before reaching that information set, irrespective of the number of times the buyer rejected the contract. In other words, when the buyer rejects d along a path where no other contract was offered, the seller does not update her belief.

Next, we define the assessment at each information set which is reached by a path along which a contract $c \neq d$ is offered. Observe that if d is actively deployed τ times before the seller deviates for

[21] This belief is the limit of beliefs derived by Bayes' rule along a sequence of mixed strategy of the buyer over rejecting and accepting the contract, along which the probability of rejection goes to zero and does not depend on the buyer's valuation.

the first time, her posterior is $\mu_\tau(d)$. Next, we show that even in the original equilibrium assessment there are information sets at which the seller's posterior is exactly $\mu_\tau(d)$. We accomplish this by demonstrating the existence of a simple and direct contract $c(d,\tau) = (X_\tau, P_\tau)_{\tau=0}^{\infty}$, with the following properties. In each equilibrium,

(i) the buyer accepts $c(d,\tau)$ in the initial period,

(ii) the buyer truthfully reports his type in the initial period if $c(d,\tau)$ is deployed and

(iii) the seller's posterior belief is $\mu_\tau(d)$ after the initial period if there is no trade.

To this end, let $P_0(v_l) = -v_h$, $X_0(v_l) = 1/2$, $P_0(v_h) = -v_h + q(v_l + \varepsilon)$ and $X_0(v_h) = 1/2 + q$, so that

$$\mu_\tau(d) = \frac{\mu\left(\frac{1}{2} - q\right)}{\mu\left(\frac{1}{2} - q\right) + (1-\mu)\frac{1}{2}}$$

and $\varepsilon\,(>0)$ is small enough so that $v_h - (v_l + \varepsilon) > \delta(v_h - v_l)$. Moreover, let $P_\tau(v) = X_\tau(v) = 0$ for all $\tau > 0$ and $v \in \{v_l, v_h\}$. One interpretation of this contract is that each type trades with probability half at price $-2v_h$ in the initial period. If the buyer reports v_h, he trades with an additional probability of q at a price just above v_l. After the initial period, the contract prescribes autarky. Note that accepting this contract generates an instantaneous payoff of at least v_h to the buyer. The sequential rationality of the seller implies that the expected continuation payoff of the buyer cannot exceed v_h, so the buyer accepts this contract in every equilibrium, yielding (i). To see part (ii), first recall that reporting v_h triggers trade with an additional probability of q at $v_l + \varepsilon$. Observe that the sequential rationality of the seller implies that the object is never sold at a price lower than v_l so the high-value buyer is better off trading at $v_l + \varepsilon$ with an additional probability of q whereas the low-value buyer is not. Hence, the buyer reports his value truthfully. To obtain part (iii), observe that q is defined so that the seller's posterior is computed by Bayes' rule and is $\mu_\tau(d)$.

We are now ready to specify the new assessment at those information sets which are reached by paths along which a contract $c \neq d$ is offered. To this end, consider an information set at which c is offered and along the paths reaching this set the contract d was actively deployed τ times and no other contract was ever offered. Therefore, the seller's posterior belief when offering c is $\mu_\tau(d)$. Of course, the continuation game starting at this information set is isomorphic to the continuation game in which the seller offers $c(d,\tau)$ in the initial period, the buyer accepts it and reports his type truthfully, trade does not occur and the seller offers c in the next period. Indeed, the seller's posterior is also $\mu_\tau(d)$ by the definition of $c(d,\tau)$. Therefore, we define the equilibrium assessment in the continuation game starting from offering c to be the same as the original equilibrium assessment in the continuation game in which the seller offers $c(d,\tau)$ and c in the first and second periods, respectively.

It remains to prove that the new assessment defined above is indeed an equilibrium assessment. We first argue that players are sequentially rational at each information set. In the initial period, the seller's payoff from offering $c\,(\neq d)$ is at most as large as her payoff in the original equilibrium. Since, $v(d,\delta)$ is larger than that, the seller rationally offers d. At those information sets which are reached by paths along which only d was offered, the seller's continuation payoff is larger than her full-commitment payoff given her posterior. So, even at those information sets, the seller rationally offers d. At any other information set, the seller's strategy is sequentially rational because it is defined by the sequentially rational original equilibrium assessment in the corresponding isomorphic continuation game. The buyer's strategy is also sequentially rational at those information sets which are reached by those paths along which no contract other than d was offered. The reason is that d provides the buyer with a non-negative payoff and rejecting d would only delay those payoffs given that the seller offers it again after any number of rejections. Since d is incentive compatible, the buyer rationally reports his type truthfully in the first period. At any other information set, the buyer's strategy is sequentially rational because it is defined by the original equilibrium assessment in the corresponding isomorphic continuation game. Also note that the seller's belief is defined by Bayes' rule at each information set which is reached with positive probability. Indeed, the seller's belief after the contract d was actively deployed τ times is $\mu_\tau(d)$. □

Having established Lemma 2.3.1, in order to prove Theorem 2.3, we need to demonstrate the existence of a δ-abiding contract for each δ which generates a payoff to the seller which is larger than a bound which is bigger than v_l. The next lemma states that such contracts exist.

Lemma 2.3.2. *For all $\delta \in (0,1)$, there exists a δ-abiding contract $d_\delta \in \mathcal{D}$ so that $v(d_\delta, \delta) \geq \underline{\pi} > v_l$.*

In what follows, we construct a contract for each δ satisfying this lemma's statement. This contract will have the following properties. In its initial period of deployment, if the buyer reports v_h, the contract specifies a positive probability of trade, α, at price $p \in [v_l, v_h]$ and the buyer does not trade if he reports v_l. In any subsequent period, trade occurs with probability β with both types at a price equal to the buyer's report. In other words, this simple and direct contract depends on three parameters $(\alpha, \beta, p) \in [0,1]^2 \times [v_l, v_h]$ and can be formally defined as follows. In the initial period, $x_0(v_h) = \alpha$, $x_0(v_l) = 0$ and $p_0(v_h) = p_0(v_l) = p$.[22] For each $\tau > 0$, $x_\tau(v) = \beta$ and $p_\tau(v) = v$ for $v \in \{v_l, v_h\}$. In what follows, we express the constraints guaranteeing that the contract is not only incentive compatible but also abiding in terms of these parameters.

Let us first discuss incentive compatibility. First, note that the buyer with type v_l weakly prefers to report his willingness-to-pay irrespective of the parameter values. The reason is that if he does so, he always trades at price v_l and hence, his expected payoff is zero. On the other hand, if he reports v_h, the price is always weakly larger than v_l and hence, his expected payoff is non-positive. Consider now the buyer whose valuation is v_h. If he reports his true valuation, his payoff is $\alpha(v_h - p)$ because he trades with probability α at price p in the initial period and, any time in the future, the price is v_h. If he reports v_l, he does not trade in the initial period and, conditional on not trading before, he trades with probability β at price v_l in every period in the future. Therefore, if the high-type buyer misreports his type, the expected discounted present value of his payoff is

$$\delta \sum_{t=0}^{\infty} \delta^t (1-\beta)^t \beta (v_h - v_l) = \frac{\beta\delta}{1-\delta+\beta\delta}(v_h - v_l).$$

So, the incentive constraint of the buyer with type v_h is satisfied if

$$\alpha(v_h - p) \geq \frac{\beta\delta}{1-\delta+\beta\delta}(v_h - v_l). \tag{2.1}$$

Next, we investigate the set of those parameters for which the contract is abiding. First, note that whenever the buyer trades, the price is weakly smaller than his willingness-to pay. Therefore, if such a contract is actively deployed forever, the continuation value of each type is weakly larger than zero, so the contract satisfies part (i) of Definition 2.3 for any parameters. Let us now describe the constraint corresponding to part (ii) of Definition 2.3. That is, we describe conditions under which the seller's posterior in each future period is such that the full-commitment monopoly price

[22] Since $x_0(v_l) = 0$, $p_0(v_l)$ can be defined arbitrarily.

is v_l and her continuation payoff exceeds v_l if the contract is deployed forever. To this end, observe that, conditional on no trade, the seller's posterior remains the same after the initial period because probability of trade in future periods, β, does not depend on the buyer's report. This posterior depends only on the probability of trade in the initial period, α, and we denote it by $\widetilde{\mu}(\alpha)$. It can be computed by Bayes' Rule,

$$\widetilde{\mu}(\alpha) = \frac{(1-\alpha)\mu}{1-\mu+(1-\alpha)\mu}. \tag{2.2}$$

So, condition (ii) of Definition 2.3 holds and the static monopoly price is v_l if, and only if,

$$v_l \geq \widetilde{\mu}(\alpha)\, v_h. \tag{2.3}$$

We now compute the seller's continuation payoff in each period after the first deployment of the contract. Since neither the posterior distribution of types nor the probability of trade depend on time, this continuation payoff is also independent of time and can be expressed as

$$\sum_{t=T}^{\infty} \delta^{t-T}(1-\beta)^{t-T}\beta[\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l] = \frac{\beta}{1-\delta+\beta\delta}[\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l].$$

So, part (iii) of Definition 2.3 holds if this payoff is larger than v_l. In fact, we will construct parameter values so that that the seller's continuation value also exceeds his payoff from deploying this contract for one more period and selling the good at v_l immediately if the contract does not recommend trade, that is,

$$\frac{\beta}{1-\delta+\beta\delta}[\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l] \geq \beta[\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l] + (1-\beta)\, v_l. \tag{2.4}$$

In order to prove Lemma 2.3.2, for each large enough δ, it is enough to show the existence of a triple, $(\alpha^*, \beta^*, p^*) \in [0,1]^2 \times [v_l, v_h]$, such that the constraints (2.1), (2.3) and (2.4) are satisfied. Furthermore, we need to demonstrate that the seller's payoff is bounded away from v_l uniformly.

Proof of Lemma 2.3.2. First, we construct a δ-abiding contracts for small discount factors. For each δ, consider the contract which specifies trade with the high-type buyer in the initial period at a price $v_h - \delta(v_h - v_l)$ and specifies trade with the low-type buyer in the next period at price v_l. This contract corresponds to the parameter triple where $\alpha = \beta = 1$ and $p = v_h - \delta(v_h - v_l)$. Note that this price makes the high-type buyer indifferent between buying the good immediately and trading at v_l a period later and hence, the constraint (2.1) is satisfied. Furthermore, since $\alpha = \beta = 1$ and $\widetilde{\mu}(1) = 0$, this contract is obviously δ-abiding and satisfies the constraint (2.4). Let π_δ denote the seller's value generated by this contract and note that

$$\pi_\delta = \mu[v_h - \delta(v_h - v_l)] + (1-\mu)\,\delta v_l = \mu v_h(1-\delta) + \delta v_l > v_l(1-\delta) + \delta v_l = v_l,$$

where the inequality follows from $\mu > v_l/v_h$. Moreover, observe that π_δ decreases in δ and converges to v_l as δ goes to one. Therefore, it is enough to prove the lemma's statement for large δ's. That is, we show that there exists a $\overline{\delta} \in (0,1)$, such that for all $\delta \geq \overline{\delta}$, there exists a δ-abiding contract $d_\delta \in \mathcal{D}$ so that $v(d_\delta, \delta) \geq \widehat{\pi} > v_l$. Then, setting $\underline{\pi}$ to be $\underline{\pi} = \min\{\pi_{\overline{\delta}}, \widehat{\pi}\}$, the lemma follows.

Let us explain how we construct the aforementioned triple of parameters for large δ. First, for each α we define $\widetilde{\beta}(\alpha)$ so that the abiding constraint (2.4) evaluated at $\beta = \widetilde{\beta}(\alpha)$ binds and ignore the constraint that $\widetilde{\beta}(\alpha)$ must be a probability. Second, we define $\widetilde{p}(\alpha)$ so that the incentive constraint (2.1) evaluated at $(\beta, p) = \left(\widetilde{\beta}(\alpha), \widetilde{p}(\alpha)\right)$ binds and ignore the constraint that $\widetilde{p}(\alpha) \in [v_l, v_h]$. Then, we consider the functional form of the seller's payoff at $\left(\alpha, \widetilde{\beta}(\alpha), \widetilde{p}(\alpha)\right)$, maximize it with respect to α subject to the constraint (2.3) and define α^* to be the maximizer. Finally, we show that, if δ is large enough, the parameters $\widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)$ are feasible, that is, $\left(\widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right) \in [0,1] \times [v_l, v_h]$. Moreover, the seller's payoff generated by the contract corresponding to $\left(\alpha^*, \widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right)$ is strictly larger than v_l and does not depend on δ.

For each $\alpha \in [0,1]$, let $\widetilde{\beta}(\alpha)$ be defined so that the constraint (2.4) binds, that is,

$$\widetilde{\beta}(\alpha) = \beta = \frac{1-\delta}{\delta} \cdot \frac{v_l}{\widetilde{\mu}(\alpha)(v_h - v_l)}. \tag{2.5}$$

In addition, let us define $\widetilde{p}(\alpha)$ for each $\alpha \in [0,1]$ so that the high-type buyer's incentive constraint, (2.1) binds, that is

$$\alpha(v_h - \widetilde{p}(\alpha)) = \frac{\widetilde{\beta}(\alpha)\,\delta}{1-\delta+\widetilde{\beta}(\alpha)\,\delta}(v_h - v_l). \tag{2.6}$$

We now turn our attention to the seller's payoff generated by the contract corresponding to the triple $\left(\alpha, \widetilde{\beta}(\alpha), \widetilde{p}(\alpha)\right)$. We first compute the seller's continuation payoff in each period $t > 0$. As mentioned above, this continuation payoff does not depend on t. Since the abiding constraint (2.4) binds at $\beta = \widetilde{\beta}(\alpha)$, this payoff can be computed by plugging $\widetilde{\beta}(\alpha)$ into the right-hand side of this constraint,

$$\begin{aligned} & \frac{1-\delta}{\delta} \cdot \frac{v_l}{\widetilde{\mu}(\alpha)(v_h - v_l)}[\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l] + \left(1 - \frac{1-\delta}{\delta} \cdot \frac{v_l}{\widetilde{\mu}(\alpha)(v_h - v_l)}\right) v_l \\ = \; & \frac{1-\delta}{\delta} \cdot \frac{v_l v_h}{(v_h - v_l)} - \frac{1-\delta}{\delta} \cdot \frac{v_l^2}{(v_h - v_l)} + v_l = \frac{v_l}{\delta}. \end{aligned}$$

We are now ready to compute the seller's payoff generated by the contract defined by $\left(\alpha, \widetilde{\beta}(\alpha), \widetilde{p}(\alpha)\right)$. Let $\nu(\alpha)$ denote this payoff. Observe that, in the initial period, the seller receives $\widetilde{p}(\alpha)$ with probability $\mu\alpha$ and, in the next period, her continuation payoff is v_l/δ. Therefore,

$$\nu(\alpha) = \mu\alpha\widetilde{p}(\alpha) + \delta(1-\mu\alpha)\frac{v_l}{\delta} = \mu\alpha\widetilde{p}(\alpha) + (1-\mu\alpha)\, v_l. \tag{2.7}$$

Substituting $\widetilde{p}(\alpha)$ from equation (2.6) and using equation (2.2) yield

$$\nu(\alpha) = v_l + (v_h - v_l)\left(1 - \frac{1-\mu}{1-\widetilde{\mu}(\alpha)} - \frac{\mu v_l}{\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l}\right). \tag{2.8}$$

Finally, we define α^* to maximize $\nu(\alpha)$, subject to the constraint (2.3). That is, α^* solves

$$\max\left\{\nu(\alpha) : \alpha \in [0,1],\, \widetilde{\mu}(\alpha) \le v_l/v_h\right\}. \tag{2.9}$$

We now show that α^* is uniquely determined. To this end, note that ν depends on α only through $\widetilde{\mu}(\alpha)$. Also note that, by (2.2), the function $\widetilde{\mu}$ is continuous, strictly decreasing in α and, in addition, $\widetilde{\mu}(0) = \mu$ and $\widetilde{\mu}(1) = 0$. Let $\Pi(\hat{\mu})$ denote $\nu\left(\widetilde{\mu}^{-1}(\hat{\mu})\right)$. In what follows, we characterize the unique solution, $\hat{\mu}^*$, of the following maximization problem

$$\max_{\hat{\mu} \in [0, v_l/v_h]} \Pi(\hat{\mu}).$$

Then, it follows that $\alpha^* = \widetilde{\mu}^{-1}(\hat{\mu}^*)$ is the unique solution of the problem (2.9). Note that

$$\Pi'(\hat{\mu}) = -(v_h - v_l)\left(\frac{1-\mu}{(1-\hat{\mu})^2} - \frac{\mu v_l (v_h - v_l)}{(\hat{\mu} v_h + (1-\hat{\mu}) v_l)^2}\right),$$

so $\Pi'(\hat{\mu}) \geq 0$ if, and only if, $\hat{\mu} \leq \left[\sqrt{\frac{\mu}{1-\mu}} - \sqrt{\frac{v_l}{v_h - v_l}}\right] / \left[\sqrt{\frac{\mu}{1-\mu}} + \sqrt{\frac{v_h - v_l}{v_l}}\right]$. Therefore,

$$\hat{\mu}^* = \min\left\{\frac{\sqrt{\frac{\mu}{1-\mu}} - \sqrt{\frac{v_l}{v_h - v_l}}}{\sqrt{\frac{\mu}{1-\mu}} + \sqrt{\frac{v_h - v_l}{v_l}}}, \frac{v_l}{v_h}\right\} \tag{2.10}$$

and note that $\hat{\mu}^* \in (0, v_l/v_h]$ because $\mu \in (v_l/v_h, 1)$.

Let us now return to examine whether $\left(\widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right) \in [0,1] \times [v_l, v_h]$ if $\alpha^* = \widetilde{\mu}^{-1}(\hat{\mu}^*)$. Observe that, by equation (2.5), $\widetilde{\beta}(\alpha^*) \in [0,1]$ if, and only if,

$$\hat{\mu}^* \left(= \widetilde{\mu}(\alpha^*)\right) \geq \frac{1-\delta}{\delta} \cdot \frac{v_l}{v_h - v_l}.$$

Since the right-hand side is decreasing in δ and converges to zero as δ goes to one, there exists $\overline{\delta}$ such that $\widetilde{\beta}(\alpha^*) \in [0,1]$ whenever $\delta \in \left(\overline{\delta}, 1\right)$. Let us turn our attention to the first period's transfer, $\widetilde{p}(\alpha^*)$. By the definition of the function $\widetilde{p}$, it follows that $\widetilde{p}(\alpha^*) \leq v_h$ for all $\alpha \in [0,1]$. Furthermore, equation (2.7) implies that the seller's payoff, $\nu(\alpha)$, can be expressed as a convex combination of $\widetilde{p}(\alpha^*)$ and v_l. Therefore, in order to establish that $\widetilde{p}(\alpha^*) \in [v_l, v_h]$ we only need to show that $\nu(\alpha^*) > v_l$, what we will do next.

Before proceeding, we note that the construction of the parameters depends on the prior distribution of types, $\mu \in (v_l/v_h, 1)$. We now make this dependency explicit and express the seller's payoff induced by the contact constructed above as a function of μ. To this end, let us write the seller's posterior defined by (2.10) as a function of μ, $\hat{\mu}^*(\mu)$. Now, observe that, by equation (2.8), the seller's payoff can be written as

$$V(\mu) = v_l + (v_h - v_l)\left(1 - \frac{1-\mu}{1-\hat{\mu}^*(\mu)} - \frac{\mu v_l}{\hat{\mu}^*(\mu) v_h + (1-\hat{\mu}^*(\mu)) v_l}\right). \tag{2.11}$$

In order to prove that $\nu(\alpha^*) > v_l$, it is enough to show that V is strictly increasing on $(v_l/v_h, 1)$ and $\lim_{\mu \to v_l/v_h} V(\mu) = v_l$. To this end, note that V is continuous on $(v_l/v_h, 1)$. From equation (2.10), it follows that there is a cutoff value of μ, $\bar{\mu} \in (v_l/v_h, 1)$[23], such that $\hat{\mu}^*(\mu) = v_l/v_h$ whenever $\mu \in (\bar{\mu}, 1)$. On this domain, $V(\mu) = \mu v_h^2/(2v_h - v_l)$, which is indeed strictly increasing. Since $\hat{\mu}^*$

[23]It can be shown that

$$\bar{\mu} = \frac{v_l(2v_h - v_l)^2}{v_l(2v_h - v_l)^2 + (v_h - v_l)^3}.$$

was chosen to maximize the seller's payoff, on the domain $(v_l/v_h, \bar{\mu})$, the Envelope Theorem implies that

$$\begin{aligned} V'(\mu) &= (v_h - v_l)\left[\frac{1}{1-\hat{\mu}^*(\mu)} - \frac{v_l}{\hat{\mu}^*(\mu)v_h + (1-\hat{\mu}^*(\mu))\, v_l}\right] \\ &= (v_h - v_l)\left[1 - 2\frac{v_l}{v_h} + \sqrt{\frac{v_l}{v_h}\left(1-\frac{v_l}{v_h}\right)}\left(\sqrt{\frac{\mu}{1-\mu}} - \sqrt{\frac{1-\mu}{\mu}}\right)\right]. \end{aligned}$$

It is clear from inspecting the expression in the second line that V' is strictly increasing on $(v_l/v_h, \bar{\mu})$. Furthermore, since $\lim_{\mu\to v_l/v_h} \hat{\mu}^*(\mu) = 0$ by (2.10), the first line of the previous equality chain implies that $\lim_{\mu\to v_l/v_h} V'(\mu) = 0$. Therefore, V' is strictly positive on $(v_l/v_h, \bar{\mu})$. Recall that V' is also strictly positive on $(\bar{\mu}, 1)$ and continuous on $(v_l/v_h, 1)$. Then, by noting that $\lim_{\mu\to v_l/v_h} V(\mu) = v_l$, we conclude that $V > v_l$ on $(v_l/v_h, 1)$.

To summarize, we have constructed a triple of parameters, $(\alpha^*, \beta^*, p^*) = \left(\alpha^*, \widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right)$. We have demonstrated the existence of $\overline{\delta}$ such that $\left(\widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right) \in [0,1] \times [v_l, v_h]$, so these parameters indeed define a contract. By equations (2.5) and (2.7), this contract is incentive compatible and abiding. Finally, we have proved that the seller's value from deploying this contract forever is strictly larger than v_l. To conclude the lemma's statement, all is left to do is to argue that, provided that $\delta > \overline{\delta}$, the seller's value does not depend on δ. This, however, is evident from equations (2.11) and (2.10). □

We are ready to argue that the statement of Theorem 2.3 follows from Lemmas 2.3.1 and 2.3.2.

Proof of Theorem 2.3. Recall that Lemma 2.3.2 guarantees the existence of a δ-abiding contract $d_\delta \in \mathcal{D}$ for each $\delta \in (0,1)$ such that seller's value generated by d_δ is bounded away from v_l, that is, $v(d_\delta, \delta) \geq \underline{\pi} > v_l$. Then Lemma 2.3.1 implies that, for all $\delta \in (0,1)$, the seller's largest equilibrium payoff exceeds $\underline{\pi}$, that is, $\pi(\mathcal{C}, \delta) \geq \underline{\pi}$. □

2.4 Discussion

Optimal Contracts.— Theorem 2.3 states that the seller's largest equilibrium payoff is bounded away from v_l but it provides no further information about this payoff. In fact, we do not know what the seller's optimal contract is generating her largest equilibrium profit. However, when we prove equilibrium existence for the case of $\mathcal{C} = \mathcal{D}$, we construct an equilibrium contract which induces a payoff to the seller which is significantly larger than the bound provided by Lemma 2.3.2 (see the Online Appendix[24]). For each δ, this equilibrium contract specifies trade before a certain date, $T(\delta)$, with probability one. In the initial period, only the buyer with valuation v_h trades with a positive probability at a price in (v_l, v_h). Ever after, the price is always the reported valuation, just like in the case of the contract described in the proof of Lemma 2.3.2. In early time periods, only the high-type buyer trades and the seller is becoming more and more pessimistic. The low-type buyer only trades in the last and in the penultimate periods. Of course, as δ goes to one, $T(\delta)$ converges to infinity. Figure 1 plots the seller's payoff generated by this contract as a function of δ for the example where $v_l = 1$, $v_h = 3$ and $\mu = .95$.

[24] Available at: tinyurl.com/36tzxh3k.

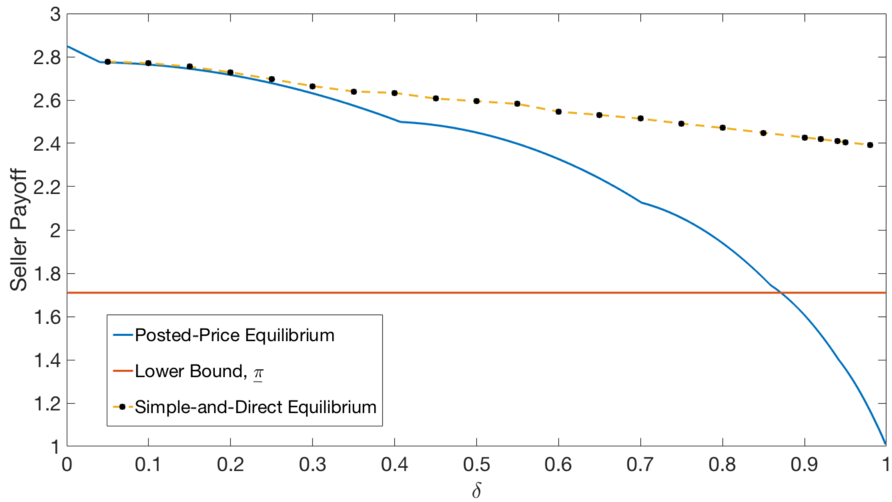

Figure 2.1: Comparison of the seller's profits when the discount factor varies.

What happens if the seller's contract space is larger than $\mathcal{D}$? As mentioned above, it is not hard to show that in different principal-agent models, the principal benefits from having access to contracts which reveal information to both contracting parties. It is also possible to construct examples where the principal is worse off if her contract space includes such contracts. Unfortunately, we were unable to establish whether the seller benefits or is hurt by enlarging the set $\mathcal{D}$.

One-period Contracts.— As mentioned in the Introduction, a common approach to model the lack of intertemporal commitment is to restrict the contract space to be the set of one-period contracts. Doval and Skreta (2020b) pursue this approach in the context of a durable-good monopolist. In their setup, a contract of a given period determines the probability of trade and transfer in that period and reveals a public signal which can be assumed to be the seller's posterior. The authors show that the largest equilibrium payoff the seller can get, can be generated by a sequence of posted prices. Moreover, the seller's equilibrium profit converges to v_l as the discount factor goes to one, that is, the Coase Conjecture holds. Our main theorem highlights that the Coase Conjecture in Doval and Skreta (2020b) is not only the consequence of the seller's limited commitment power, but also of the restricted contract space.

Figure 1 also plots the seller's largest equilibrium profit in the model of Doval and Skreta (2020b) as a function of the discount factor. Note that when the discount factor is small, this profit level is larger than the one generated by the stationary contract described in Lemma 2.3.2. However, the profit induced by the aforementioned equilibrium contract is larger than the maximum profit in Doval and Skreta (2020b) irrespective of δ.

Side-Contracts.— If the seller decides not to proceed with the previous period's contract, the information content of the contract is lost. This feature enables the seller to redeploy contracts which implement allocations which are dominated from the viewpoints of both contracting parties. For example, in the context of the simple and direct contract of Lemma 2.3.2, the buyer trades with probability $\beta\,(< 1)$ at the price of the reported valuation in all but the initial periods. Of course, both the seller and the buyer would be (weakly) better off if this probability was larger. However, if the seller wants to replace the contract with another one with larger trading probabilities, she would need to pay information rent to the buyer again which makes such a deviation non-profitable. A possible way to circumvent such ex-post inefficiencies associated to a contract would be to consider the possibility of writing side-contracts. That is, the seller continues to redeploy the previous period's contract but can offer a side-contract which conditions on the outcome of the redeployed contract. Despite the fact that the seller cannot offer such contracts in our model, we next argue that our main result is robust to the introduction of side-contracts.

To this end, consider again the contract of Lemma 2.3.2 and note that the only information this contract ever reveals is whether there is trade. Of course, if the contract specifies trade, the game ends. So, a side-contract must condition on the event of no trade. Recall that the trading probability in the initial period is specified so that the static monopoly price is v_l in any subsequent period. Therefore, the seller's optimal side-contract would implement trade at price v_l if the outcome of the original contract is no trade. Note that the seller's payoff from offering such a side-contract is the right-hand side of equation (2.4). In other words, the abiding constraint (2.4) guarantees that the seller's continuation value from deploying the original contract exceeds her payoff from such a side-contract.

Implementation.— As explained in the Introduction, the contracts considered in our model have features resembling those of smart contracts used in digital markets. For our main result to hold, it is essential that the buyer communicates his willingness-to-pay to the contract privately. While communications on blockchain-based software platforms, such as Ethereum, are typically public, there are numerous examples for smart contracts involving private communication[25]. Implementing such communication is not hard using the cryptographic technology already employed in those markets. One way to do this is to encrypt the buyer's messages and let the buyer retain the decryption key. Then, in each period, the buyer can input the decryption key and an allocation is determined. Inputting an incorrect key would simply be treated as rejecting the contract.[26]

[25] Examples for digital protocols where such private communication is implemented in practice include Tornado.Cash and Aztec.Network.

[26] For discussions of implementing private communications on public blockchains, see Kerber, Kiayias, and Kohlweiss (2021), Steffen et al. (2019) or Bünz et al. (2020).

Bibliography

[1] Abadi, Joseph, and Markus K. Brunnermeier. *Blockchain Economics*. Working Paper, 2019.

[2] Acharya, Avidit, and Ortner, Juan. *Progressive Learning*. Econometrica 85 (6), 1965-1990, 2017.

[3] Akbarpour, Mohammad, and Li, Shengwu. *Credible Auctions: A Trilemma*. Econometrica 88 (2), 425–467, 2020.

[4] Ausubel, Lawrence M., and Deneckere, Raymond J. *Reputation in Bargaining and Durable Goods Monopoly*. Econometrica 57(3), 511–531, 1989.

[5] Bagnoli, Mark, Salant, Stephen, and Swierzbinski, Joseph E. *Durable-Goods Monopoly with Discrete Demand*. Journal of Political Economy 97(6), 1459–1478, 1989.

[6] Bakos, Yannis and Halaburda, Hanna. *Smart Contracts, IoT Sensors and Efficiency: Automated Execution vs. Better Information*. Working Paper, Available at SSRN: https://ssrn.com/abstract=3394546, 2020.

[7] Battaglini, Marco. *Optimality and Renegotiation in Dynamic Contracting*. Games and Economic Behavior 60(2), 213–246, 2007.

[8] Battigalli, Pierpaolo, and Martin Dufwenberg *Dynamic Psychological Games*. Journal of Economic Theory 144 (1), 1-35, 2009.

[9] Beccuti, Juan, and Möller, Marc. *Dynamic adverse selection with a patient seller*. Journal of Economic Theory 173 (C), 95–117, 2018.

[10] Bester, Helmut, and Strausz, Roland. *Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case*. Econometrica: 69(4), 1077-1098, 2004.

[11] Biehl, Andrew R.. *Durable-Goods Monopoly with Stochastic Values*. The RAND Journal of Economics: 32 (3), 565–577, 2001.

[12] Board, Simon, and Pycia, Marek. *Outside Options and the Failure of the Coase Conjecture*. American Economic Review: 104(2), 656–671, 2014.

[13] Breig, Zachary. *Endogenous and Exogenous Commitment*. Economics Letters: 183 (C), 1–1, 2019.

[14] Bulow, Jeremy I. *Durable-Goods Monopolists*. Journal of Political Economy 90 (2): 314–332, 1982.

[15] Bulow, Jeremy I. *An Economic Theory of Planned Obsolescence.* The Quarterly Journal of Economics 101 (4): 729-749, 1986.

[16] Bünz, Benedikt, Agrawal, Shashank, Zamani, Mahdi, and Boneh, Dan. *Zether: Towards Privacy in a Smart Contract World.* In: Bonneau J., Heninger N. (eds) Financial Cryptography and Data Security. FC 2020. Lecture Notes in Computer Science, vol 12059. Springer, Cham.

[17] Coase, Ronald H. *Durability and monopoly.* Journal of Law and Economics 15: 143–149, 1972.

[18] Cong, Lin William, and He, Zhiguo. *Blockchain Disruption and Smart Contracts.* The Review of Financial Studies 32 (5), 1754–1797, 2019.

[19] Doval, Laura, and Skreta, Vasiliki. *Mechanism Design with Limited Commitment.* Working Paper, Available at SSRN: https://ssrn.com/abstract=3281132, 2020a.

[20] Doval, Laura and Skreta, Vasiliki. *Optimal Mechanism for the Sale of a Durable Good.* Working Paper, Available at Arxiv: https://arxiv.org/abs/1904.07456, 2020b.

[21] Fanning, Jack. *Mediation in Reputational Bargaining.* American Economic Review, 111(8): 2444–2472, 2021a.

[22] Fanning, Jack. *Optimal Dynamic Mediation.* Working Paper, 2021b.

[23] Feinberg,Yossi, and Skrzypacz, Andrzej. *Uncertainty about Uncertainty and Delay in Bargaining.* Econometrica 73 (1): 69–91, 2005.

[24] Fuchs, William, and Skrzypacz, Andrzej. *Bargaining with Arrival of New Traders.* American Economic Review: 100 (3), 802-836, 2010.

[25] Fudenberg, Drew and Levine, David. *Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games.* Journal of Economic Theory 31 (2): 251-268, 1993.

[26] Fudenberg, Drew and Tirole, Jean. *Upgrades, Tradeins, and Buybacks.* The RAND Journal of Economics 29 (2): 235–258, 1998.

[27] Fudenberg, Drew, Levine, David, and Tirole, Jean. *Infinite Horizon Models of Bargaining with One-sided Incomplete Information*, in Bargaining with Incomplete Information, (Alvin Roth, Ed.), London/New York, Cambridge Univ. Press: 73–98, 1985.

[28] Gerardi, Dino, and Maestri, Lucas. *Dynamic Contracting with Limited Commitment and the Ratchet Effect.* Theoretical Economics 15 (2): 583-623, 2020.

[29] Gul, Faruk, and Sonnenschein, Hugo, and Wilson, Robert. *Foundations of Dynamic Monopoly and the Coase Conjecture.* Journal of Economic Theory 39: 155–190, 1986.

[30] Hahn, Jong-Hee. *Damaged Durable Goods.* The RAND Journal of Economics 37 (1): 121–133, 2006.

[31] Hart, Oliver D., and Jean Tirole. *Contract Renegotiation and Coasian Dynamics.* The Review of Economic Studies 55 (4): 509-540, 1988.

[32] Holden, Richard and Malani, Anup. *Can Blockchain Solve the Holdup Problem in Contracts?*. Working Paper, Available at SSRN: https://ssrn.com/abstract=3093879, 2017.

[33] Huberman, Gur, Leshno, Jacob D., and Moallemi, Ciamac. *Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System*. The Review of Economic Studies, 0 1-30, 2021.

[34] Inderst, Roman. *Durable Goods with Quality Differentiation*. Economics Letters 100 (2): 173–177, 2008.

[35] ahn, Charles. *The Durable Goods Monopolist and Consistency with Increasing Costs*. Econometrica: 275–294, 1986.

[36] Karp, Larry S. *Monopoly Extraction of a Durable Non-Renewable Resource: Failure of the Coase Conjecture*. Economica: 1–11, 1993.

[37] Kerber, Thomas, Kiayias, Aggelos and Kohlweiss, Markulf. *Kachina – Foundations of Private Smart Contracts*. IEEE 34th Computer Security Foundations Symposium: 1–16, 2021.

[38] Laffont, Jean-Jacques, and Tirole, Jean. *The Dynamics of Incentive Contracts*. Econometrica 56 (5): 1153–1175, 1988.

[39] Laffont, Jean-Jacques, and Jean Tirole. *Adverse Selection and Renegotiation in Procurement*. The Review of Economic Studies 57 (4): 597–625, 1990.

[40] Lomys, Niccolò, and Yamashita, Takuro. *A Mediator Approach to Mechanism Design with Limited Commitment*, Working Paper, 2021.

[41] Maestri, Lucas. *Dynamic Contracting Under Adverse Selection and Renegotiation*. Journal of Economic Theory 171 (C): 136–173, 2017.

[42] McAdams, David, and Michael Schwarz. *Credible Sales Mechanisms and Intermediaries*. American Economic Review 97 (1): 260–276, 2007.

[43] Mcafee, R. Preston, and Wiseman, Thomas. *Capacity Choice Counters the Coase Conjecture*. The Review of Economic Studies 75 (1): 317–331, 2008.

[44] Montez, João. *Inefficient Sales Delays by a Durable Good Monopoly Facing a Finite Number of Buyers*. The RAND Journal of Economics 44 (3): 425–437, 2013.

[45] Myerson, Roger B., and Reny, Philip J. *Perfect Conditional ε-Equilibria of Multi-Stage Games with Infinite Sets of Signals and Actions*. Econometrica: 88 (2), 495–531, 2020.

[46] Nava, Francesco, and Schiraldi, Pasquale. *Differentiated Durable Goods Monopoly: A Robust Coase Conjecture*. American Economic Review: 109 (5), 1930–1968, 2019.

[47] Ortner, Juan. *Durable Goods Monopoly with Stochastic Costs*. Theoretical Economics 12 (2): 817–861, 2017.

[48] Perea, Andrés. *A Note on the One-Deviation Property in Extensive Form Games*. Games and Economic Behavior 40 (2), 322-338, 2002.

[49] Simon, Leo K. and Zame, William R. *Discontinuous Games and Endogenous Sharing Rules.* Econometrica 58 (4), 861-872, 1990.

[50] Sobel, Joel. *Durable Goods Monopoly with Entry of New Consumers.* Econometrica 59 (5): 1455-1485, 1991.

[51] Steffen, Samuel, Bichsel, Benjamin, Gersbach, Mario, Melchior, Noa, Tsankov, Petar and Vechev, Martin. *Zkay: Specifying and Enforcing Data Privacy in Smart Contracts.* In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19). Association for Computing Machinery, New York, NY, USA, 1759–1776, 2019.

[52] Stokey, Nancy L. *Rational Expectations and Durable Goods Pricing.* Bell Journal of Economics 12: 112–128, 1982.

[53] Strulovici, Bruno *Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts.* Econometrica 85 (2): 585–616, 2017.

[54] Takeyama, Lisa N. *Strategic Vertical Differentiation and Durable Goods Monopoly.* The Journal of Industrial Economics 50 (1): 43–56, 2002.

[55] Tirole, Jean. *From Bottom of the Barrel to Cream of the Crop: Sequential Screening With Positive Selection.* Econometrica 84 (4): 1291–1343, 2016.

[56] Tinn, Katrin. *'Smart' Contracts and External Financing.* Working Paper, Available at SSRN: https://ssrn.com/abstract=3072854, 2019.

[57] von der Fehr, Nils-Henrik, and Kuhn, Kai-Uwe. *Coase versus Pacman: Who Eats Whom in the Durable-Goods Monopoly?.* Journal of Political Economy, 103(4): 785–812, 1995.

[58] Waldman, Michael. *A New Perspective on Planned Obsolescence.* The Quarterly Journal of Economics 108 (1): 273–283, 1993.

[59] Wang, Gyu Ho. *Bargaining Over a Menu of Wage Contracts.* The Review of Economic Studies 65 (2): 295“305, 1998.

2.5 Appendix

This appendix is concerned with constructing an equilibrium when the seller's contract space is restricted to include only *simple and direct* contracts. In section 2.5.1, we use a fixed-point argument to define a set of contracts $D^*(\mu)$ parameterised by $\mu \in [0,1]$, such that the seller deploys a contract in $D^*(\mu)$ when her belief about the buyer having high valuation is μ. In section 2.5.2, we describe and analyse a class of auxiliary games whose sequential equilibria map to off-path parts of our equilibrium assessment. Finally in section 2.5.3, we use the results from sections 2.5.1 and 2.5.2 to build a complete assessment and prove that it constitutes an equilibrium.

2.5.1 Deployed Contracts

Let $v_h > v_l > 0$ and $\delta \in (0,1)$. For $\mu \in [0,1]$, let $\bar{J}(\mu) = \max\{v_l, \mu v_h\}$ and

$$\underline{J}(\mu) = \begin{cases} v_l & \text{if } \mu < 1, \\ v_h & \text{if } \mu = 1. \end{cases}$$

A function $J : [0,1] \to \mathbb{R}$ is piece-wise linear if $[0,1]$ can be partitioned into countably many intervals such that J is affine on each of these intervals. We denote by $\mathcal{J}$ the set of non-decreasing, piece-wise linear and convex functions $J : [0,1] \to [v_l, v_h]$ such that, for all $\mu \in [0,1]$, $\bar{J}(\mu) \geq J(\mu) \geq \underline{J}(\mu)$.

Given $J \in \mathcal{J}$, we construct a function $\mathcal{A}J \in \mathcal{J}$. We set $\mathcal{A}J(0) = v_l$ and $\mathcal{A}J(1) = v_h$. For $\mu \in (0,1)$, $\mathcal{A}J(\mu)$ is the value to the maximisation problem described below.

Fix $\mu \in (0,1)$. A trading time s is a random time whose distribution $\langle s \rangle$ depends on the buyer's valuation $v \in \{v_l, v_h\}$. We identify $\langle s \rangle = (q^h, q^l) \in \mathcal{Q} \times \mathcal{Q}$, where $\mathcal{Q} = \{q \in [0,1]^{\mathbb{N}} : \sum_{t\geq 0} q_t \leq 1\}$. For $t \geq 0$ and $i \in \{h, l\}$, q_t^i is interpreted as the probability that trade occurs in period t if the buyer's valuation in v_i. Given $\langle s \rangle \in \mathcal{Q} \times \mathcal{Q}$ and $t \geq 0$, if $\mathbb{P}(s \geq t) = 1 - \sum_{k=0}^{t-1} \left(\mu q_k^h + (1-\mu) q_k^l\right) > 0$, we define:

$$\mu_t = \mathbb{P}(v = v_h | s \geq t) = \frac{\mu\Big(1 - \sum_{k=0}^{t-1} q_k^h\Big)}{\mu\Big(1 - \sum_{k=0}^{t-1} q_k^h\Big) + (1-\mu)\Big(1 - \sum_{k=0}^{t-1} q_k^l\Big)}.$$

Define, for $\langle s \rangle \in \mathcal{Q} \times \mathcal{Q}$, the objective function:

$$\Omega\big(\langle s \rangle | \mu\big) = \sum_{k=0}^{\infty} \delta^k \Big(q_k^h \mu v_h + q_k^l (v_l - \mu v_h)\Big),$$

and for $t \geq 1$, the constraint mapping:

$$G_t\big(\langle s \rangle | \mu, J\big) = \begin{cases} \mathbb{P}\big(s \geq t\big) J(\mu_t) - \sum_{k=t}^{\infty} \delta^{k-t} \Big(q_k^h \mu v_h + q_k^l (1-\mu) v_l\Big) & \text{if } \mathbb{P}\big(s \geq t\big) > 0, \\ 0 & \text{if } \mathbb{P}\big(s \geq t\big) = 0. \end{cases}$$

Note that $\Omega(\cdot|\mu)$ is linear and, for each $t \geq 1$, $G_t(\cdot|\mu, J)$ is convex on the convex set $\mathcal{Q} \times \mathcal{Q}$. Denoting

$G(\langle s\rangle|\mu, J) = \Big(G_t(\langle s\rangle|\mu, J)\Big)_{t\geq 1}$, the maximisation problem is given by:

$$\begin{aligned} \mathcal{A}J(\mu) = \max_{\langle s\rangle\in\mathcal{Q}\times\mathcal{Q}} & \ \Omega(\langle s\rangle|\mu) \\ \text{s.t. } & G(\langle s\rangle|\mu, J) \leq 0. \end{aligned} \tag{2.12}$$

If $\mu \leq \frac{v_l}{v_h}$, an obvious solution is given by $s = 0$, that is $q_0^h = q_0^l = 1$. Thus, in this case, $\mathcal{A}J(\mu) = v_l$. Therefore, we focus on the case $\mu > \frac{v_l}{v_h}$. We first prove preliminary results which help describe a candidate solution $\langle s^*\rangle(\mu, J)$ to (2.12). Then, we prove that $\langle s^*\rangle(\mu, J)$ indeed solves the maximisation problem, and we verify that $\mathcal{A}J$ indeed belongs to $\mathcal{J}$.

Lemma 2.5.1. *For $\tilde{\mu} \in (0,1)$, there exists $\tilde{\mu}' \in [0,\mu)$ such that:*

$$\delta J(\tilde{\mu}') = v_h - \frac{1-\tilde{\mu}'}{1-\tilde{\mu}}\big(v_h - J(\tilde{\mu})\big),$$

if and only if $\tilde{\mu} \geq \frac{(1-\delta)v_l}{v_h-\delta v_l} \equiv \bar{\mu} \in (0, \frac{v_l}{v_h})$. In this case, $\tilde{\mu}'$ is unique.

Proof. Uniqueness is guaranteed by the fact that J is convex, and $\delta J(\tilde{\mu}) < J(\tilde{\mu})$. If $\delta J(0) \geq v_h - \frac{v_h - J(\tilde{\mu})}{1-\tilde{\mu}}$, then a solution exists by the intermediate value theorem. Otherwise, no solution exists, by convexity. Therefore, a solution exists if and only if:

$$J(\tilde{\mu}) \leq \tilde{\mu} v_h + (1-\tilde{\mu})\delta v_l.$$

This inequality is satisfied if $\tilde{\mu} \geq \frac{v_l}{v_h}$ since $J(\tilde{\mu}) \leq \max\{\tilde{\mu}v_h, v_l\}$. For $\tilde{\mu} < \frac{v_l}{v_h}$, $J(\tilde{\mu}) = v_l$ and the inequality is equivalent to $\tilde{\mu} \geq \bar{\mu}$. □

In view of lemma 2.5.1, define on $[0,1)$ the function $\tilde{\mu}'$ such that $\tilde{\mu}'(\tilde{\mu}) = 0$ if $\tilde{\mu} \leq \bar{\mu}$, and:

$$\delta J(\tilde{\mu}') = v_h - \frac{1-\tilde{\mu}'}{1-\tilde{\mu}}\big(v_h - J(\tilde{\mu})\big),$$

otherwise. Now, given $\mu_1 \in (0,1)$, we construct the non-increasing sequence $(\mu_k)_{k\geq 1}$ such that, for $k \geq 1$, $\mu_{k+1} = \tilde{\mu}'(\mu_k)$.

Lemma 2.5.2. *There exists $T \geq 1$ such that $\mu_T < \bar{\mu}$.*

Proof. Otherwise, $(\mu_k)_{k\geq 1}$ has a limit $\mu_\infty \in [\bar{\mu}, 1)$. Since J is continuous on $[0,1)$, the limit must satisfy:

$$\delta J(\mu_\infty) = J(\mu_\infty),$$

which is impossible since $J \geq v_l > 0$ and $\delta < 1$. □

We define $T = \min\{t \geq 1 : \mu_t < \bar{\mu}\}$. Now, let:

$$\rho^J(\mu_1) = \frac{v_l}{v_h - v_l}\prod_{t=1}^{T}\Big(1 + \frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t) - (1-\mu_t)J'_+(\mu_t)}\Big) \in (0,\infty],$$

where J'_+ is the right-derivative of J. By convexity and piece-wise linearity of J, ρ^J is a non-decreasing and right-continuous step function, which may take infinite values if $v_h = J(\mu_1) + (1 - \mu_1)J'_+(\mu_1)$.

Lemma 2.5.3. *For any* $\mu_1 \in (0,1)$, $\rho^J(\mu_1) > \frac{\mu_1}{1-\mu_1}$.

Proof. If $T = 1$, then:

$$\frac{\mu_1}{1-\mu_1} < \frac{\bar{\mu}}{1-\bar{\mu}} = \frac{(1-\delta)v_l}{v_h - v_l} = \frac{\rho^J(\mu_1)}{\frac{1}{1-\delta} + \frac{1}{\delta}\frac{v_h}{v_h - v_l}} < \rho^J(\mu_1).$$

If $T > 1$, note that for $t < T$, since $\mu_{t+1} < \mu_t$ and J is convex:

$$J'_+(\mu_t) \geq \frac{J(\mu_t) - J(\mu_{t+1})}{\mu_t - \mu_{t+1}}.$$

As a result:

$$1 + \frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t) - (1-\mu_t)J'_+(\mu_t)} \geq \frac{(\mu_t - \mu_{t+1})v_h - \delta(1-\mu_{t+1})J(\mu_t) + \delta(1-\mu_t)J(\mu_{t+1})}{\delta\big[(\mu_t - \mu_{t+1})v_h - (1-\mu_{t+1})J(\mu_t) + (1-\mu_t)J(\mu_{t+1})\big]}.$$

Substitute in the numerator $\delta J(\mu_{t+1}) = v_h - \frac{1-\mu_{t+1}}{1-\mu_t}\big(v_h - J(\mu_t)\big)$, and in the denominator $J(\mu_t) = v_h - \frac{1-\mu_t}{1-\mu_{t+1}}\big(v_h - \delta J(\mu_{t+1})\big)$, to obtain:

$$1 + \frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t) - (1-\mu_t)J'_+(\mu_t)} \geq \frac{(1-\mu_{t+1})J(\mu_t)}{\delta(1-\mu_t)J(\mu_{t+1})}.$$

As a result:

$$\rho^J(\mu_1) \geq \frac{1-\mu_T}{\delta^T}\frac{v_h - \delta v_l}{(v_h - v_l)^2}\frac{J(\mu_1)}{1-\mu_1}.$$

Since $J(\mu_1) = v_h - \frac{1-\mu_1}{1-\mu_2}\big(v_h - \delta J(\mu_2)\big)$, we have:

$$\frac{1-\mu_T}{\delta^T}\frac{v_h - \delta v_l}{(v_h - v_l)^2}\frac{J(\mu_1)}{1-\mu_1} - \frac{\mu_1}{1-\mu_1} = \frac{1-\mu_T}{\delta^{T-1}}\frac{v_h - \delta v_l}{(v_h - v_l)^2}\frac{J(\mu_2)}{1-\mu_2} - \frac{\mu_2}{1-\mu_2} + \left(\frac{1-\mu_T}{\delta^T}\frac{v_h - \delta v_l}{(v_h - v_l)^2}v_h - 1\right)\left(\frac{1}{1-\mu_1} - \frac{1}{1-\mu_2}\right).$$

Now, since $\mu_T < \bar{\mu} = \frac{(1-\delta)v_l}{v_h - \delta v_l}$:

$$\frac{1-\mu_T}{\delta^T}\frac{v_h - \delta v_l}{(v_h - v_l)^2}v_h - 1 > \frac{1}{\delta^T}\frac{v_h}{v_h - v_l} - 1 > 0,$$

and since $\mu_1 > \mu_2$, we obtain:

$$\frac{1-\mu_T}{\delta^T}\frac{v_h - \delta v_l}{(v_h - v_l)^2}\frac{J(\mu_1)}{1-\mu_1} - \frac{\mu_1}{1-\mu_1} > \frac{1-\mu_T}{\delta^{T-1}}\frac{v_h - \delta v_l}{(v_h - v_l)^2}\frac{J(\mu_2)}{1-\mu_2} - \frac{\mu_2}{1-\mu_2}.$$

The same argument applies by induction to establish:

$$\frac{1-\mu_T}{\delta^T}\frac{v_h - \delta v_l}{(v_h - v_l)^2}\frac{J(\mu_1)}{1-\mu_1} - \frac{\mu_1}{1-\mu_1} > \frac{1}{\delta}\frac{v_h - \delta v_l}{(v_h - v_l)^2}v_l - \frac{\mu_T}{1-\mu_T} > \frac{v_l}{v_h - v_l}\frac{(1-\delta)v_h + \delta^2(v_h - v_l)}{\delta(v_h - v_l)} > 0.$$

To summarise:

$$\rho^J(\mu_1) > \frac{1-\mu_T}{\delta^T}\frac{v_h-\delta v_l}{(v_h-v_l)^2}\frac{J(\mu_1)}{1-\mu_1} > \frac{\mu_1}{1-\mu_1},$$

which proves the claim. □

Now, given the prior $\mu \in (\frac{v_l}{v_h}, 1)$, let:

$$\mu_1^* = \min\Big\{\mu_1 \in [0,1) : \rho^J(\mu_1) > \frac{\mu}{1-\mu}\Big\}.$$

As above, we iterate on $\tilde{\mu}'$ to construct the path $(\mu_1^*, \mu_2^*, ..., \mu_{T^*}^*)$, where $\mu_{T^*}^* \in [0, \bar{\mu})$. Defining $\mu_0^* = \mu$ and $\mu_{T^*+1}^* = 0$, the candidate solution $\langle s^*\rangle(\mu, J)$ is characterised by:

$$\forall t \in \{0, ..., T^*\}, \quad q_t^{*h} = \frac{1-\mu}{\mu}\Big(\frac{1}{1-\mu_t^*} - \frac{1}{1-\mu_{t+1}^*}\Big),$$

$$q_{T^*}^{*l} = 1 - q_{T^*+1}^{*l} = \frac{(1-\delta)v_l - \mu_{T^*}^*(v_h - \delta v_l)}{(1-\delta)(1-\mu_{T^*}^*)v_l}.$$

$\langle s^*\rangle(\mu, J)$ solves problem (2.12).

Proof. To simplify notations, we omit the dependence on μ and J of $\langle s^*\rangle$, Ω and G. We first define a sequence of non-negative Lagrange multipliers $(\lambda_t)_{t\geq 1}$ as follows.

If:

$$\frac{v_l}{v_h-v_l}\prod_{t=1}^{T^*}\Big(1+\frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t^*) - (1-\mu_t^*)J'_-(\mu_t^*)}\Big) \leq \frac{\mu}{1-\mu} < \frac{v_l}{v_h-v_l}\prod_{t=1}^{T^*}\Big(1+\frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t^*) - (1-\mu_t^*)J'_+(\mu_t^*)}\Big),$$

where $J'_-(0) = 0$, then there exists $\big(J'_*(\mu_t^*)\big)_{1\leq t\leq T^*} \in \prod_{1\leq t\leq T^*}\big[J'_-(\mu_t^*), J'_+(\mu_t^*)\big]$ such that:

$$\frac{\mu}{1-\mu} = \frac{v_l}{v_h-v_l}\prod_{t=1}^{T^*}\Big(1+\frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t^*) - (1-\mu_t^*)J'_*(\mu_t^*)}\Big).$$

Otherwise, it must be that $\mu_{T^*}^* = 0$, and:

$$\frac{v_l}{v_h-v_l}\prod_{t=1}^{T^*-1}\Big(1+\frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t^*) - (1-\mu_t^*)J'_-(\mu_t^*)}\Big) \leq \frac{\mu}{1-\mu} < \frac{v_l}{v_h-v_l}\prod_{t=1}^{T^*}\Big(1+\frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t^*) - (1-\mu_t^*)J'_-(\mu_t^*)}\Big).$$

Then, we define $J'_*(\mu_t^*) = J'_-(\mu_t^*)$ for all $t \in \{1, ..., T^*\}$.

In both cases, for $t \in \{1, ..., T^*\}$, let:

$$\lambda_t = \frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_t^*) - (1-\mu_t^*)J'_*(\mu_t^*)}\prod_{k=1}^{t-1}\Big(1+\frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_k^*) - (1-\mu_k^*)J'_*(\mu_k^*)}\Big),$$

and for $t > T^*$, let $\lambda_t = 0$. We also introduce the notation, for $t \geq 0$:

$$\Lambda_t = \sum_{k=1}^{t}\lambda_k = -1 + \prod_{k=1}^{\min\{T^*,t\}}\Big(1+\frac{1-\delta}{\delta}\frac{v_h}{v_h - J(\mu_k^*) - (1-\mu_k^*)J'_*(\mu_k^*)}\Big).$$

With these definitions, we establish below that $\langle s^* \rangle$ maximises on $\mathcal{Q} \times \mathcal{Q}$ the Lagrangian:

$$\Omega\big(\langle s \rangle\big) - \sum_{t=1}^{\infty} \delta^t \lambda_t G_t\big(\langle s \rangle\big).$$

It follows that $\langle s^* \rangle$ solves problem (2.12), since for any $\langle s \rangle$ feasible, $G\big(\langle s \rangle\big) \leq 0$, so:

$$\Omega\big(\langle s \rangle\big) \leq \Omega\big(\langle s \rangle\big) - \sum_{t=1}^{\infty} \delta^t \lambda_t G_t\big(\langle s \rangle\big) \leq \Omega\big(\langle s^* \rangle\big) - \sum_{t=1}^{\infty} \delta^t \lambda_t G_t\big(\langle s^* \rangle\big) = \Omega\big(\langle s^* \rangle\big).$$

For $\langle s \rangle \in \mathcal{Q} \times \mathcal{Q}$, let:

$$\begin{aligned} f_{\langle s \rangle} &: [0,1] \to \mathbb{R} \\ &\alpha \mapsto \Omega\big(\alpha\langle s \rangle + (1-\alpha)\langle s^* \rangle\big) - \sum_{t=1}^{\infty} \delta^t \lambda_t G_t\big(\alpha\langle s \rangle + (1-\alpha)\langle s^* \rangle\big). \end{aligned}$$

The desired result is implied if, for any $\langle s \rangle \in \mathcal{Q} \times \mathcal{Q}$, $f_{\langle s \rangle}$ is maximised at $\alpha = 0$. Since $f_{\langle s \rangle}$ is concave, it is sufficient to show that $f_{\langle s \rangle}$ is differentiable at $\alpha = 0$, with $f'_{\langle s \rangle}(0) \leq 0$.

Fix $\langle s \rangle \in \mathcal{Q} \times \mathcal{Q}$ and denote $\langle s_\alpha \rangle = \alpha\langle s \rangle + (1-\alpha)\langle s^* \rangle$. Ω is linear and $\lambda_t = 0$ when $t > T^*$, thus it is sufficient to show differentiability of the term $\alpha \mapsto G_t\big(\langle s_\alpha \rangle\big)$, when $t \in \{1, ..., T^*\}$. In this case, if $\mathbb{P}(s \geq t) > 0$, then:

$$\begin{aligned} G_t\big(\langle s_\alpha \rangle\big) =& \mathbb{P}(s_\alpha \geq t) J\left(\alpha \frac{\mathbb{P}(s \geq t)}{\mathbb{P}(s_\alpha \geq t)} \mu_t + (1-\alpha) \frac{\mathbb{P}(s^* \geq t)}{\mathbb{P}(s_\alpha \geq t)} \mu_t^*\right) - \alpha\Big(\mathbb{P}(s \geq t) J(\mu_t) - G_t\big(\langle s \rangle\big)\Big) \\ & - (1-\alpha)\Big(\mathbb{P}(s^* \geq t) J(\mu_t^*) - G_t\big(\langle s^* \rangle\big)\Big), \end{aligned}$$

where $\mathbb{P}(s_\alpha \geq t) = \alpha\mathbb{P}(s \geq t) + (1-\alpha)\mathbb{P}(s^* \geq t)$. This expression has a right- and left-derivative at any α, which coincide if $\alpha \frac{\mathbb{P}(s \geq t)}{\mathbb{P}(s_\alpha \geq t)} \mu_t + (1-\alpha) \frac{\mathbb{P}(s^* \geq t)}{\mathbb{P}(s_\alpha \geq t)} \mu_t^*$ is not at a kink of J and write:

$$\begin{aligned} \frac{\partial G_t\big(\langle s_\alpha \rangle\big)}{\partial \alpha} =& \frac{\mathbb{P}(s \geq t)\mathbb{P}(s^* \geq t)}{\mathbb{P}(s_\alpha \geq t)} (\mu_t - \mu_t^*) J'\left(\alpha \frac{\mathbb{P}(s \geq t)}{\mathbb{P}(s_\alpha \geq t)} \mu_t + (1-\alpha) \frac{\mathbb{P}(s^* \geq t)}{\mathbb{P}(s_\alpha \geq t)} \mu_t^*\right) \\ & + \Big(\mathbb{P}(s \geq t) - \mathbb{P}(s^* \geq t)\Big) J\left(\alpha \frac{\mathbb{P}(s \geq t)}{\mathbb{P}(s_\alpha \geq t)} \mu_t + (1-\alpha) \frac{\mathbb{P}(s^* \geq t)}{\mathbb{P}(s_\alpha \geq t)} \mu_t^*\right) \\ & - \left(\Big(\mathbb{P}(s \geq t) J(\mu_t) - G_t\big(\langle s \rangle\big)\Big) - \Big(\mathbb{P}(s^* \geq t) J(\mu_t^*) - G_t\big(\langle s^* \rangle\big)\Big)\right), \end{aligned}$$

Now, using the fact that $G_t\big(\langle s \rangle\big) = \mathbb{P}\big(s \geq t\big) J(\mu_t) - \sum_{k=t}^{\infty} \delta^{k-t}\Big(q_k^h \mu v_h + q_k^l (1-\mu) v_l\Big)$ and $G_t\big(\langle s^* \rangle\big) = 0$, we obtain:

$$\left.\frac{\partial G_t\big(\langle s_\alpha \rangle\big)}{\partial \alpha}\right|_{\alpha=0} = \mathbb{P}(s \geq t)\Big(J(\mu_t^*) + (\mu_t - \mu_t^*) J'_{\to \mu_t}(\mu_t^*)\Big) - \sum_{k=t}^{\infty} \delta^{k-t}\Big(q_k^h \mu v_h + q_k^l (1-\mu) v_l\Big), \quad (2.13)$$

or equivalently:

$$\frac{\partial G_t\big(\langle s_\alpha\rangle\big)}{\partial \alpha}\bigg|_{\alpha=0} = \mu\Big(1-\sum_{k=0}^{t-1} q_k^h\Big)\Big(J(\mu_t^*)+(1-\mu_t^*)J'_{\to\mu_t}(\mu_t^*)\Big)+(1-\mu)\Big(1-\sum_{k=0}^{t-1} q_k^l\Big)\Big(J(\mu_t^*)-\mu_t^* J'_{\to\mu_t}(\mu_t^*)\Big)$$
$$-\sum_{k=t}^{\infty}\delta^{k-t}\Big(q_k^h\mu v_h+q_k^l(1-\mu)v_l\Big),$$

where:

$$J'_{\to\mu_t}(\mu_t^*)=\begin{cases} J'_-(\mu_t^*) & \text{if } \mu_t<\mu_t^*,\\ J'_+(\mu_t^*) & \text{if } \mu_t>\mu_t^*,\\ J'_*(\mu_t^*) & \text{if } \mu_t=\mu_t^*.\end{cases}$$

This expression is valid if $\mathbb{P}(s\geq t)=0$, in which case $G_t\big(\langle s_\alpha\rangle\big)=(1-\alpha)G_t\big(\langle s^*\rangle\big)=0$, if we extend the definition $J'_{\to\mu_t}(\mu_t^*)=J'_*(\mu_t^*)$ when $\mathbb{P}(s\geq t)=0$. Thus, the derivative at $\alpha=0$ of $f_{\langle s\rangle}$ writes:

$$f'_{\langle s\rangle}(0)=-\Omega\big(\langle s^*\rangle\big)+\sum_{k=0}^{\infty}\delta^k\Big(q_k^h\mu v_h+q_k^l(v_l-\mu v_h)\Big)-\sum_{t=1}^{T^*}\delta^t\lambda_t\Bigg(\mu\Big(1-\sum_{k=0}^{t-1}q_k^h\Big)\Big(J(\mu_t^*)+(1-\mu_t^*)J'_{\to\mu_t}(\mu_t^*)\Big)$$
$$+(1-\mu)\Big(1-\sum_{k=0}^{t-1}q_k^l\Big)\Big(J(\mu_t^*)-\mu_t^*J'_{\to\mu_t}(\mu_t^*)\Big)-\sum_{k=t}^{\infty}\delta^{k-t}\Big(q_k^h\mu v_h+q_k^l(1-\mu)v_l\Big)\Bigg).$$

Rearranging, we get:

$$f'_{\langle s\rangle}(0)=-\Omega\big(\langle s^*\rangle\big)+\sum_{k=0}^{\infty}\delta^k\Big[q_k^h\mu v_h(1+\Lambda_k)+q_k^l\big(v_l-\mu v_h+(1-\mu)v_l\Lambda_k\big)\Big]$$
$$+\sum_{k=0}^{T^*-1}\Bigg(\mu q_k^h\sum_{t=k+1}^{T^*}\delta^t\lambda_t\Big(J(\mu_t^*)+(1-\mu_t^*)J'_{\to\mu_t}(\mu_t^*)\Big)+(1-\mu)q_k^l\sum_{t=k+1}^{T^*}\delta^t\lambda_t\Big(J(\mu_t^*)-\mu_t^*J'_{\to\mu_t}(\mu_t^*)\Big)\Bigg)$$
$$-\sum_{t=1}^{T^*}\delta^t\lambda_t\Big(J(\mu_t^*)+(\mu-\mu_t^*)J'_{\to\mu_t}(\mu_t^*)\Big). \tag{2.14}$$

Using equation (2.13), note that: $f'_{\langle s\rangle}(0)\leq H\big(\langle s\rangle\big)$, where:

$$H\big(\langle s\rangle\big)=-\Omega\big(\langle s^*\rangle\big)+\sum_{k=0}^{\infty}\delta^k\Big[q_k^h\mu v_h(1+\Lambda_k)+q_k^l\big(v_l-\mu v_h+(1-\mu)v_l\Lambda_k\big)\Big]$$
$$+\sum_{k=0}^{T^*-1}\Bigg(\mu q_k^h\sum_{t=k+1}^{T^*}\delta^t\lambda_t\Big(J(\mu_t^*)+(1-\mu_t^*)J'_*(\mu_t^*)\Big)+(1-\mu)q_k^l\sum_{t=k+1}^{T^*}\delta^t\lambda_t\Big(J(\mu_t^*)-\mu_t^*J'_*(\mu_t^*)\Big)\Bigg)$$
$$-\sum_{t=1}^{T^*}\delta^t\lambda_t\Big(J(\mu_t^*)+(\mu-\mu_t^*)J'_*(\mu_t^*)\Big). \tag{2.15}$$

H is linear on $\mathcal{Q} \times \mathcal{Q}$. Denote, for all $t \geq 0$, γ_t^h and γ_t^l the terms multiplying q_t^h and q_t^l respectively. If $t \geq T^*$, $\gamma_t^h = \mu v_h \delta^t (1 + \Lambda_{T^*})$ is positive and decreasing in t. If $t \in \{1, ..., T^*\}$:

$$\begin{aligned}
\gamma_t^h &= \mu v_h \delta^t (1 + \Lambda_t) + \mu \sum_{k=t+1}^{T^*} \delta^k \lambda_k \Big(J(\mu_k^*) + (1 - \mu_k^*) J_*'(\mu_k^*) \Big) \\
&= \gamma_{t-1}^h - (1 - \delta) \mu v_h \delta^{t-1} (1 + \Lambda_{t-1}) + \mu \delta^t \Big(v_h - J(\mu_t^*) - (1 - \mu_t^*) J_*'(\mu_t^*) \Big) \lambda_t \\
&= \gamma_{t-1}^h,
\end{aligned}$$

where we have used the fact that:

$$\lambda_t = 1 + \Lambda_t - (1 + \Lambda_{t-1}) = \frac{1 - \delta}{\delta} \frac{v_h}{v_h - J(\mu_t^*) - (1 - \mu_t^*) J'(\mu_t^*)} (1 + \Lambda_{t-1}).$$

Now, for $t \in \{1, ..., T^*\}$:

$$\begin{aligned}
\gamma_t^l &= \delta^t \big(v_l - \mu v_h + (1 - \mu) v_l \Lambda_t \big) + (1 - \mu) \sum_{k=t+1}^{T^*} \delta^k \lambda_k \Big(J(\mu_k^*) - \mu_k^* J_*'(\mu_k^*) \Big) \\
&= \gamma_{t-1}^l + (1 - \mu) \delta^t \lambda_t \Big(v_l - J(\mu_t^*) + \mu_t^* J_*'(\mu_t^*) \Big) + (1 - \delta) \delta^{t-1} \Big(\mu (v_h - v_l) - (1 - \mu) v_l (1 + \Lambda_{t-1}) \Big).
\end{aligned}$$

By convexity, $v_l = J(0) \geq J(\mu_t^*) - \mu_t^* J_*'(\mu_t^*)$. In addition:

$$1 + \Lambda_{t-1} \leq 1 + \Lambda_{T^*-1} \leq \frac{v_h - v_l}{v_l} \frac{\mu}{1 - \mu}.$$

It follows that $\gamma_t^l \geq \gamma_{t-1}^l$. If $t \geq T^*$, $\gamma_t^l = \delta^t \big(v_l - \mu v_h + (1 - \mu) v_l \Lambda_{T^*} \big)$. When $\mu_{T^*}^* > 0$, $1 + \Lambda_{T^*} = \frac{v_h - v_l}{v_l} \frac{\mu}{1 - \mu}$, so $\gamma_t^l = 0$ for all $t \geq T^*$. In any case, $1 + \Lambda_{T^*} \geq \frac{v_h - v_l}{v_l} \frac{\mu}{1 - \mu}$, and γ_t^l is non-negative and non-increasing in t for $t \geq T^*$.

It follows that H is maximised at $\langle s^* \rangle$, that is, for any $\langle s \rangle \in \mathcal{Q} \times \mathcal{Q}$:

$$f_{\langle s \rangle}'(0) \leq H\big(\langle s \rangle\big) \leq H\big(\langle s^* \rangle\big) = f_{\langle s^* \rangle}'(0) = 0,$$

which proves the desired result. □

Remark 2.5.1. *For each μ and J, we have described a solution $\langle s^* \rangle(\mu, J)$. If there exists $\mu_1 \in [0, 1)$ such that $\rho^J(\mu_1) = \frac{\mu}{1 - \mu}$, we can describe an alternative solution $\langle \hat{s}^* \rangle(\mu, J)$ by setting:*

$$\hat{\mu}_1^* = \min \Big\{ \mu_1 \in [0, 1) : \rho^J(\mu_1) = \frac{\mu}{1 - \mu} \Big\},$$

and the continuing path of beliefs as before. All the arguments in the proof of proposition (2.5.1) directly apply. Since ρ^J is a step function, $\langle \hat{s}^* \rangle(\mu, J)$ is only defined for a discrete set of priors μ.

Remark 2.5.2. *For each μ and J, there exists a direct and simple contract that implements $\langle s^* \rangle(\mu, J)$. The probabilities of trade for each type of the buyer are given by q^{*h} and q^{*l}. For every $t \geq 1$, the trading price if the buyer's report was v_i is $p_t^i = v_i$. In the initial period of deployment, the low-valuation buyer trades at price $p_0^l = v_l$, while the price for the high-valuation buyer is* *such that:*

$$q_0^{*h}(v_h - p_0^h) = \delta^{T^*}\big(q_{T^*}^{*l} + \delta(1 - q_{T^*}^{*l})\big)(v_h - v_l).$$

The same applies to $\langle \hat{s}^* \rangle(\mu, J)$. We denote $D(\mu, J)$ the set of those contracts (which contains either one or two elements).

$\mathcal{A}J$ belongs to $\mathcal{J}$.

Proof. For $\mu \leq \frac{v_l}{v_h}$, $\mathcal{A}J(\mu) = v_l$. For $\mu \in \left(\frac{v_l}{v_h}, \frac{\rho^J(0)}{1+\rho^J(0)}\right]$, $\mathcal{A}J(\mu) = (1-\delta)\mu v_h + \delta v_l$. Note that $\mathcal{A}J$ is thus continuous at $\frac{v_l}{v_h}$. Similarly, if μ_1^* is a point of discontinuity of ρ^J, and $(\mu_1^*, ..., \mu_{T^*}^*)$ the corresponding path of beliefs obtained by iteration on $\tilde{\mu}'$, then for every $\mu \in \left[\frac{\rho_-^J(\mu_1^*)}{1+\rho_-^J(\mu_1^*)}, \frac{\rho^J(\mu_1^*)}{1+\rho^J(\mu_1^*)}\right]$, where $\rho_-^J(\mu_1^*)$ denotes the left limit of ρ^J at μ_1^*, we have:

$$\mathcal{A}J(\mu) = \left((1-\delta)\sum_{t=1}^{T^*} \frac{\delta^{t-1}}{1-\mu_t^*} + \delta^{T^*}\frac{\mu_{T^*}^*}{1-\mu_{T^*}^*}\frac{v_h - v_l}{v_l}\right)\mu v_h + \left(1 - (1-\delta)\sum_{t=1}^{T^*}\frac{\delta^{t-1}}{1-\mu_t^*} - \frac{\delta^{T^*}}{1-\mu_{T^*}^*}\frac{v_h - v_l}{v_h}\right)v_h.$$

Thus $\mathcal{A}J$ is piece-wise linear. In addition, by continuity at the boundaries, $\mathcal{A}J$ is non-decreasing. The term $\left((1-\delta)\sum_{t=1}^{T^*}\frac{\delta^{t-1}}{1-\mu_t^*} + \delta^{T^*}\frac{\mu_{T^*}^*}{1-\mu_{T^*}^*}\frac{v_h - v_l}{v_l}\right)$ is increasing in μ_1^*, therefore $\mathcal{A}J$ is convex. It is clear that $\mathcal{A}J \geq v_l$. Now, for $\mu > \frac{v_l}{v_h}$ and $\langle s \rangle \in \mathcal{Q} \times \mathcal{Q}$:

$$\mu v_h - \Omega\big(\langle s \rangle \big| \mu\big) = \left(1 - \sum_{k=0}^{\infty}\delta^k q_k^h\right)\mu v_h + \sum_{k=0}^{\infty}\delta^k q_k^l(\mu v_h - v_l) \geq 0,$$

so $\mathcal{A}J \leq \bar{J}$. □

Lemma 2.5.4. *For all $J, \hat{J} \in \mathcal{J}$, if for all $\mu \in [0,1]$, $J(\mu) \geq \hat{J}(\mu)$, then for all $\mu \in [0,1]$, $\mathcal{A}J(\mu) \leq \mathcal{A}\hat{J}(\mu)$.*

Proof. Under the assumption of the lemma, for any $\langle s \rangle \in \mathcal{Q} \times \mathcal{Q}$ and $\mu \in [0,1]$, $G\big(\langle s \rangle \big| \mu, J\big) \geq G\big(\langle s \rangle \big| \mu, \hat{J}\big)$. Thus $\langle s^* \rangle(\mu, J)$ is feasible in problem (2.12) for $\hat{J}$, from which the conclusion follows. □

The operator $\mathcal{A} : \mathcal{J} \to \mathcal{J}$ has a unique fixed point J^*.

Proof. By lemma 2.5.4, and since $\bar{J}$ is an upper bound on any element of $\mathcal{J}$, any fixed point J^* of $\mathcal{A}$ must satisfy:

$$\mathcal{A}\bar{J} \leq J^* \leq \bar{J},$$

and by immediate induction:

$$\forall n \geq 0, \quad \mathcal{A}^{2n+1}\bar{J} \leq J^* \leq \mathcal{A}^{2n}\bar{J}.$$

Therefore, existence and uniqueness of a fixed point are guaranteed if $(\mathcal{A}^n\bar{J})_{n\geq 0}$ converges in $\mathcal{J}$. Note that for all $n, m \geq 0$, $\mathcal{A}^n\bar{J}$ and $\mathcal{A}^m\bar{J}$ coincide on $[0, \hat{\mu}_0)$, where $\hat{\mu}_0 = \frac{v_l}{v_h}$. Therefore, $\rho^{\mathcal{A}^n\bar{J}}$ and $\rho^{\mathcal{A}^m\bar{J}}$ also coincide on $[0, \hat{\mu}_0)$, thus $\mathcal{A}^{n+1}\bar{J}$ and $\mathcal{A}^{m+1}\bar{J}$ coincide on $[0, \hat{\mu}_1)$, where $\hat{\mu}_1 = \frac{\rho_-^{\bar{J}}(\hat{\mu}_0)}{1+\rho_-^{\bar{J}}(\hat{\mu}_0)} > \hat{\mu}_0$. By immediate induction, for every $k \geq 0$ and $n, m \geq k$, $\mathcal{A}^n\bar{J}$ and $\mathcal{A}^m\bar{J}$ coincide on $[0, \hat{\mu}_k)$, where the sequence $(\hat{\mu}_k)_{k\geq 0}$ is constructed such that $\hat{\mu}_{k+1} = \frac{\rho_-^{\mathcal{A}^k \bar{J}}(\hat{\mu}_k)}{1+\rho_-^{\mathcal{A}^k \bar{J}}(\hat{\mu}_k)}$. $(\hat{\mu}_k)_{k\geq 0}$ is an increasing sequence. By lemma 2.5.3, its limit must be $\hat{\mu}_\infty = 1$. It follows that for every $\mu \in [0,1]$, $\mathcal{A}^n \bar{J}(\mu)$ remains constant for n sufficiently large (recall that $\mathcal{A}^n \bar{J}(1) = v_h$ for all $n \geq 0$). Denote $J^*(\mu)$ this constant.

By construction, $J^* \in \mathcal{J}$ and $\mathcal{A}J^* = J^*$. □

For $\mu \in [0,1]$, we denote $D^*(\mu) = D(\mu, J^*)$. If a contract $d \in D^*(\mu)$ is deployed actively and truthfully in every period, the seller's payoff is $J^*(\mu)$ and the low-valuation buyer's payoff is always $U_l(\mu) = 0$. The high-valuation buyer's payoff depends on the contract, and we denote $U_h(\mu)$ the convex-hull of those payoffs. Finally, denote $d^*(\mu) \in D^*(\mu)$ the contract such that $\max U_h(\mu)$ is achieved.

2.5.2 Auxiliary Games

Given a contract $d = (x_\tau, p_\tau)_{\tau=0}^\infty \in \mathcal{D}$ and belief $\mu \in [0,1]$, let $\Gamma(d,\mu)$ be the discontinuous dynamic psychological game defined as follows. The two players are the seller and the buyer. As in the main text, the buyer has two types v_l or v_h, and μ is the seller's common knowledge belief that the buyer has a high valuation. The game is played in discrete time. In the initial period, the contract d is deployed with $\tau = 0$. That is, the buyer chooses among $\{h, l, r\}$. If r is chosen, it is observed by the seller and the game proceeds to the next period. If instead $i \in \{h, l\}$ is selected, trade occurs with probability $x_0(v_i)$ at price $p_0(v_i)$, in which case the game ends. If trade does not occur, the game proceeds to the next period and the index in d is updated to $\tau = 1$. In any following period, the seller chooses among $\{C, S\}$. If S is selected, the game ends. If C is selected, the contract d is deployed again. That is, the buyer makes a report in $\{h, l, r\}$ if $\tau = 0$ or in $\{a, r\}$ if $\tau \geq 1$. As above, when r is selected, it is observed by the seller and the game proceeds to the next period. Otherwise, trade may occur or not as specified by the contract d.

The game ends either when trade occurs or when the seller chooses S. If trade occurs in period $t \geq 0$ at price p, the seller's payoff is $\delta^t p$ and the buyer's payoff is $\delta^t(v - p)$, where $v \in \{v_h, v_l\}$ is his valuation. If the game ends when the seller chooses S in period t, the payoffs are partially determined and depend on the seller's belief about the buyer's valuation at that information set $\hat{\mu}$. In particular, we use the approach of Simon and Zame (1990) and specify payoffs if the seller chooses S in period t with belief $\hat{\mu}$ as $\left(\delta^t J^*(\hat{\mu}), \delta^t U_l(\hat{\mu}), \delta^t U_h(\hat{\mu})\right)$ for the seller, low-valuation buyer and high-valuation buyer respectively. Since U_h is a correspondence, which element of $U_h(\hat{\mu})$ actually determines the high-valuation buyer's payoff is part of the solution concept.

As a result, we define a *payoff selection* u to be a mapping from terminal nodes following S to the real numbers. An *augmented assessment* is a triple (σ, α, u), where σ is a strategy profile, α is the seller's belief system and u is a payoff selection for the high-valuation buyer. u is said to be *consistent* with α if, at every terminal history $\mathfrak{h}$ following S, if the seller's belief that the buyer has a high valuation is $\mu_\mathfrak{h}$, then $u(\mathfrak{h}) \in U_h(\mu_\mathfrak{h})$. A *sequential equilibrium* of $\Gamma(d,\mu)$ is an augmented assessment (σ, α, u) such that (i) α is consistent with σ and the prior μ in the usual sense, (ii) u is consistent with α and (iii) σ is sequentially rational[27] given α and u.

Suppose that d specifies bounded transfers $(p_\tau)_{\tau\geq 0}$. Then $\Gamma(d,\mu)$ has a sequential equilibrium.

Proof. We first consider the truncated version of the game $\Gamma_T(d,\mu)$ such that the game coincides

[27] The notion of sequential rationality extends naturally to psychological games. The reader is referred to Battigalli and Dufwenberg (2009) for details.

with $\Gamma(d,\mu)$ until period T is reached, and for any period $t \geq T$, the seller's action space is restricted to $\{C\}$ and the buyer's action space is restricted to $\{r\}$.

For any $\varepsilon > 0$, there exists a continuous function $U_h^\varepsilon : [0,1] \to [0, v_h - v_l]$ such that any point $\left(\tilde{\mu}, U_h^\varepsilon(\tilde{\mu})\right)$ in the graph of U_h^ε is at a distance less than ε to the graph of U_h. Let $\Gamma_T^\varepsilon(\mu, d)$ be the psychological game corresponding to $\Gamma_T(\mu, d)$ in which the high-valuation buyer's payoff after the seller chooses S with belief $\hat{\mu}$ is $U_h^\varepsilon(\hat{\mu})$. By Theorem 9 of Battigalli and Dufwenberg (2009), $\Gamma_T^\varepsilon(d,\mu)$ has a sequential equilibrium $(\sigma^\varepsilon, \alpha^\varepsilon)$ (where the first component refers to the strategy profile and the second to the belief system).

Given $(\sigma^\varepsilon, \alpha^\varepsilon)$, let $\vec{\mu}^\varepsilon$ the vector listing the seller's beliefs that the buyer has a high valuation at all her information sets in $\Gamma_T^\varepsilon(d,\mu)$, and $U_h^\varepsilon(\vec{\mu}^\varepsilon)$ the vector of high-valuation payoffs whenever S is chosen. Since $\left(\sigma^\varepsilon, \alpha^\varepsilon, U_h^\varepsilon(\vec{\mu}^\varepsilon)\right)_\varepsilon$ lives in a compact set, it possesses an accumulation point $(\sigma_T, \alpha_T, u_T)$ as $\varepsilon \to 0$.

We claim that $(\sigma_T, \alpha_T, u_T)$ is a sequential equilibrium of $\Gamma_T(d,\mu)$. Indeed, α must be consistent for σ since α^ε is consistent for σ^ε for any $\varepsilon > 0$. Moreover, at every terminal node following S, the seller's belief $\hat{\mu}^\varepsilon$ converges to $\hat{\mu}$, while the limit of $\left(\hat{\mu}^\varepsilon, U_h^\varepsilon(\hat{\mu}^\varepsilon)\right)$ must belong to the graph of U_h. Thus u_T must be a payoff selection for the high-valuation buyer. Finally, since the terminal payoffs of any player converge together with the assessment, for any strategy s_i of any player i, the evaluation of player i's expected payoff at any of her information sets under $(s_i, \sigma_{-i}^\varepsilon)$ also converges to that under (s_i, σ_{-i}). Since no profitable deviation exists under σ^ε, the same is true in the limit under σ.

σ_T specifies a full profile in $\Gamma(d,\mu)$. We can also extend α_T and u_T to construct a full assessment and payoff selection in $\Gamma(d,\mu)$. For the belief system α, take the limit of the belief system induced by a fully mixed buyer strategy after period T such that both types make a mistake with the same small probability at every decision node. Given the completed infinite vector $\vec{\mu}$, whenever the seller chooses S after period T with belief $\hat{\mu}$, select $U_h^*(\hat{\mu}) = \max U_h(\mu)$. Denote $\left(\bar{\sigma}_T, \bar{\alpha}_T, \bar{u}_T\right)$ the completed assessment and payoff selection in the infinite-horizon game $\Gamma(d,\mu)$.

Since transfers are bounded, discounting guarantees that $\left(\bar{\sigma}_T, \bar{\alpha}_T, \bar{u}_T\right)$ is a $\bar{p}\delta^T$-equilibrium in $\Gamma(d,\mu)$. In addition, the set of augmented assessments is compact in the topology of Fudenberg and Levine (1983) (note that, given a belief system, a payoff selection is isomorphic to choosing a mixture between $\max U_h(\hat{\mu})$ and $\min U_h(\hat{\mu})$ at every information set of the seller, where $\hat{\mu}$ is her belief at that information set). Thus $\left(\bar{\sigma}_T, \bar{\alpha}_T, \bar{u}_T\right)$ has a converging subsequence, which must be a sequential equilibrium in $\Gamma(d,\mu)$. □

2.5.3 Equilibrium Assessment

We construct an equilibrium in the version of the game presented in the main text in which the seller's contract space is restricted to include only simple and direct contracts, with the additional assumption that all transfers are bounded. Specifically, we assume that any transfer must be in the set $[\underline{p}, \bar{p}]$, where $\underline{p} < 0$ and $\bar{p} > v_h$.

The seller's prior is μ. Throughout we maintain that the seller's belief does not update following her own deviations. At the beginning of the game, the seller deploys the contract $d^*(\mu)$. Consider a history $\mathfrak{h}$ at which the seller's belief is $\mu_\mathfrak{h}$ and the seller offers a new contract d. If $d \in D^*(\mu_\mathfrak{h})$, the buyer accepts the contract and reports his valuation truthfully. In every following period, as long as trade does not occur, the seller's belief updates according to Bayes' rule, she deploys again d and the buyer accepts. If the buyer ever rejects d, the seller's belief does not update and she deploys d

again. If $d \notin D^*(\mu_{\mathfrak{h}})$, the continuing assessment until the game ends or d is replaced is constructed as follows.

First note that, by the axiom of choice and proposition 2.5.2, we can select a sequential equilibrium of the auxiliary game $\Gamma(d, \mu_{\mathfrak{h}})$ for any $(d, \mu_{\mathfrak{h}})$. Let (σ, α, u) the chosen equilibrium. Note that, as long as d is deployed, all the decision nodes of the buyer are identical in $\Gamma(d, \mu_{\mathfrak{h}})$ and in the original game. Thus, the buyer's strategy at those nodes can be taken directly from the profile σ. Now, we translate the seller's strategy in σ to a strategy in the original game. We interpret the action C as deploying the contract d again. Note that, as long as d has been deployed, the seller's belief at every information set can be taken directly from α. Finally, when the seller chooses S in the auxiliary game, with belief $\hat{\mu}$, u specifies a payoff selection for the high-valuation buyer $\hat{u} \in U_h(\hat{\mu})$. We specify that the contract d is replaced by a contract $\hat{d} \in D^*(\hat{\mu})$, which results from the unique mixture across the contracts in $D^*(\hat{\mu})$ that delivers expected payoff $\hat{u}$ to the high-valuation buyer (recall that $D^*(\hat{\mu})$ contains at most two contracts), if $\hat{d}$ is to be deployed actively and truthfully forever after. Once d is replaced, the continuing assessment is described in the previous paragraph.

Next, we prove that the above assessment is indeed an equilibrium. By construction, the assessment satisfies updating consistency in the sense of Perea (2002). Therefore the one-shot deviation principle applies. It is clear that the buyer has no incentive to deviate. Next, we establish that the seller has no profitable one-shot deviation at any information set.

It is sufficient to show that, at any history $\mathfrak{h}$ with belief $\mu_{\mathfrak{h}}$, the seller's payoff from deploying a new contract cannot exceed $J^*(\mu_{\mathfrak{h}})$ given the buyer's strategy. Then, by construction, it is clear that the seller's behaviour is sequentially rational.

Suppose that the seller's belief is $\mu_{\mathfrak{h}}$ and a new contract is offered. This is a one-shot deviation, thus the continuation path is that specified in the above assessment. In particular, at every subsequent history $\hat{\mathfrak{h}}$, if the seller's belief is $\mu_{\hat{\mathfrak{h}}}$, her continuation payoff is at least $J^*(\mu_{\hat{\mathfrak{h}}})$. Let s the random time at which trade occurs induced by the continuation path. Denoting by V, θ_h and θ_l the continuation payoffs for the seller, the high-valuation buyer and the low-valuation buyer respectively, we can express the total continuation surplus as:

$$V + \mu_{\mathfrak{h}}\theta_h + (1 - \mu_{\mathfrak{h}})\theta_l = \mathbb{E}_0[\delta^s|h]\mu_{\mathfrak{h}} v_h + \mathbb{E}_0[\delta^s|l](1 - \mu_{\mathfrak{h}})v_l.$$

Since a new contract is deployed, the high-valuation buyer could mimic the low-valuation buyer's strategy, thus:

$$\theta_h \geq \theta_l + \mathbb{E}_0[\delta^s|l](v_h - v_l).$$

In addition, the buyer can guarantee a non-negative payoff by rejecting the contract in every period, thus: $\theta_l \geq 0$. Together, these inequalities imply that the seller's continuation payoff satisfies:

$$V \leq \mathbb{E}_0[\delta^s|h]\mu_{\mathfrak{h}} v_h + \mathbb{E}_0[\delta^s|l](v_l - \mu_{\mathfrak{h}} v_h).$$

Finally, since J^* is convex, in any period, the total continuation surplus must exceed $J^*(\mu_t)$, where μ_t is the average belief of the seller across all histories reaching period t. That is:

$$\mathbb{E}_t\big[\delta^s|h\big]\mu_t v_h + \mathbb{E}_t\big[\delta^s|l\big](1 - \mu_t)v_l \geq J^*(\mu_t).$$

In other words, it must be the case that $V \leq \mathcal{A}J^*(\mu_{\mathfrak{h}}) = J^*(\mu_{\mathfrak{h}})$. Thus the seller has no profitable one-shot deviation, which concludes the proof.

Chapter 3

On the Equilibrium Maturity of Bail-in Debt

3.1 Introduction

During the Great Financial Crisis (GFC), many banks experienced runs on their short-term debt liabilities due to an investor fear of losses on the asset side. As a consequence, governments were forced to conduct bailouts to prevent the largest banks of the economy from failing. Bailouts however are not only very costly to taxpayers, they also disincentivise banks to internalise asset losses during crises. As a response, regulators around the globe have introduced bail-ins as part of the post-crisis reforms package.[1] During a bail-in, the government enforces equity write-downs and debt-to-equity swaps of unsecured and uninsured bank debt in order to recapitalise banks. While not intended as the silver bullet to end all bank-related financial crises, bail-ins are an essential part of the reforms aimed at stabilising the financial system while committing not to conduct bailouts.

Regulation seems to have leapfrogged ahead of theory, as bail-ins have not yet received considerable attention in academic writing. Many papers in the literature employ frameworks which presuppose the type of bail-in debt that banks issue (see i.e. [11]). Other papers analyse bail-in policies for a given balance sheet composed of different types of bail-in debt (see i.e. [25]).[2] Yet the composition of bank liabilities is endogenous to the government's intended recapitalisation policy. In this context, we endogenise the maturity of bank debt and explore how it responds to the introduction of bail-ins. Given this response, can bail-ins successfully recapitalise failing banks and induce socially optimal behaviour as intended by the reforms?

Model outline. We develop a model of asymmetric information on banks' asset returns. There are two types of bankers: some with safe and high returns, others with risky and low returns in expectation. However, banks need to issue debt to finance the asset. Part of the financing can be raised via insured deposits, but we assume that banks are required to sell bail-in-able debt to investors. The paper develops a benchmark in which the maturity of all bail-in debt—determined endogenously—is either short-term or long-term, and assets can be liquidated pre-maturely without cost. Importantly, a non-contractible signal reveals the type of the banker; short-term debt investors

[1] In the US, bail-ins have been introduced under Title II of the Dodd-Franck Act, whereas the EU has passed the Bank Recovery and Resolution Directive (BRRD).

[2] A notable exception is [4].

then decide whether they want to rollover their debt.

The model further contains a scope for government intervention: at a later stage, banks have an investment opportunity but invest inefficiently when they face a debt overhang. In the absence of a bail-in policy, the model thus features the usual time inconsistency problem on the government's side which gives rise to bailouts.

When the government's preferred form of bank recapitalisation is bail-ins, we show that banks issue short-term debt in equilibrium. Banks with high quality assets attempt to distinguish themselves by shortening the maturity of their debt, leading to an overall contraction in the maturity for all banks. Then in the run up to government intervention for banks with low quality assets, investors decide not to rollover their debt. This forces the banker to liquidate their assets and, eventually, to default. Thus, the government again requires bailouts to prevent the bank from failing and to induce efficient investment.

We then develop a notion of constrained efficiency: only banks with positive asset and investment NPV should enter financing markets; all banks invest efficiently post-intervention; and the government is not required to use public funds to achieve the efficient investment. Short-term debt leads to inefficiencies for multiple reasons. Absent of bailouts, the bank defaults and therefore does not invest. Hence the government is required to use public funds on top of conducting bail-ins. Due to these partial bailouts, negative NPV bankers find it profitable to enter financing markets ex-ante. Bail-ins do not fully achieve their desired disciplining effects.

We show that long-term bail-in debt achieves efficiency. Bail-in debt holders cannot force the bank into liquidating its assets in the run up to an intervention. Hence bail-ins can successfully induce efficient levels of investment. Since the government is not forced to conduct partial bailouts, only positive NPV projects enter the market. In this sense long-term debt has stronger disciplining effects on financial markets than short-term debt. This is in contrast with a large body of the literature (see i.e. [3]), in which the threat of liquidation associated with short-term debt disciplines banks. From a policy perspective, these findings give rise to a long-term bail-in debt requirement.

We then extend the model by allowing banks to issue an endogenous mixture of short-term and long-term debt. We show that any equilibrium again features runs on bail-in debt. Long-term debt requirements are necessary to ensure bail-in effectiveness. However, if short-term debt is treated preferentially during a bail-in, the government needs to impose weaker long-term debt requirements.

Finally, we also discuss costly liquidation of assets and show that it leads to an increase in market discipline through short-term debt. However, this ex-ante disciplining effect only arises because this liquidation is inefficient ex-post.

Literature review. First, the paper relates to the literature on bank recapitalisation using bail-ins. [20] and [23] are concerned with optimal capital regulation and conduct numerical exercises for bail-in debt and equity requirements. [6] discuss liquidity support in a bail-in environment. In [13], bail-ins act as a commitment device not to conduct bailouts because they reduce the cost of liquidating assets. In [19], banks do not impose bail-ins if they anticipate bailouts, leading to runs and even larger required bailouts. In our paper, anticipated future bail-ins imposed by the government lead to an equilibrium contraction in the bail-in debt maturity structure. Runs ensue, leaving the government with no option but to bailout. [25], starting from a fixed balance sheet, find that the government may conduct bail-ins that are too small when they are read as a signal indicating weak bank fundamentals. Runs ensue after the government's intervention. Our paper endogenises the bank balance sheet, giving investors the option to act upon the realisation of a signal but before the government. We find that bail-ins become ineffective unless the government restricts the maturity of bail-in debt. [4] discuss bail-ins and bailouts in the presence of a bank monitoring problem. Banks issue too little (long-term) bail-in debt—which can be written down and helps to avoid costly liquidation—because they do not internalise a firesale externality. If the amount of bail-in debt is sufficient, bailouts are fully replaced. Our paper points out that the nature of bail-in debt contracts responds to policy: unless the government requires bail-in debt to have long maturity, bail-ins are rendered ineffective and do not fully replace bailouts.[3][4]

Second, the paper relates to the literature on short-term debt and fragility. In [8], banks provide liquidity insurance when issuing deposits while generating high returns from illiquid investments. This liquidity mismatch creates fragility.[5] In our paper, banks create a maturity mismatch (leading to a liquidity mismatch) in order to distinguish themselves from low quality projects—it does not provide liquidity services but shields investors from losses imposed by the government. In [3], short-term debt disciplines banks through its threat of liquidation. Our paper points out that short-term bail-in debt leads to pre-mature liquidations followed by government bailouts while failing to fully align incentives of banks and society. Long-term debt avoids such liquidations, ensures bail-in effectiveness and thus corrects ex-ante incentives. [2] and [16] feature a shortening of debt maturity since long-maturity creditors fear dilution by future short-maturity creditors. Our paper finds a shortening of the bail-in debt maturity structure as response to the government's proposed recapitalisation mechanism.[6]

Third, the paper relates to the literature on debt in the presence of asymmetric information. In [12], firms signal quality of their projects through their maturity choice. In [7], borrowers decide between issuing long-term and short-term debt, trading off liquidity risk against future improvements in their credit ratings associated with lower financing costs. [21] derive conditions under which issuing debt is the unique equilibrium in an environment of ex-ante information asymmetries about return distributions. [24] discusses how to optimally intervene in financing markets ex-post to overcome asymmetric information. Our paper is concerned with the anticipated response to a shift in ex-post government intervention from bailouts to bail-ins, also in the presence of asymmetric information on returns. We show that high return banks render bail-ins of low return banks ineffective in a futile attempt to distinguish themselves.

Organisation of paper. Section 3.2 introduces the model environment. Section 3.3 presents a simple notion of bailouts. Turning to bail-ins, the equilibrium debt contract and its consequences for bail-in effectiveness are derived in Section 3.4. Section 3.5 discusses efficiency properties. Mixed maturity bail-in debt is considered in Section 3.6, while Section 3.7 revisits the question on the market disciplining effect of short-term debt. Section 3.8 discusses extensions to the baseline model and Section 3.9 concludes.

[3] See [1] and [5] for bail-ins of investors of multinational banks.

[4] See also [22] where bail-ins are ex-post efficient but lead to a breakdown of financing markets ex-ante. A combination of bail-ins and bailouts are the optimal policy. In our paper, bail-ins do not lead to a breakdown of financing markets, but the maturity structure of debt shortens to render bail-ins ineffective, leaving the government with no option but to conduct bailouts.

[5] See further canonical papers on liquidity creation and fragility by [15], [10] and [14].

[6] See also [17] and [9].

3.2 Environment

The model features four different types of agents: bankers who have access to an asset but require initial financing; investors of bail-in debt; depositors who supply funding but cannot be bailed in; and the government which recapitalises banks.

The model contains four time periods. Banks decide whether to enter financing markets at time-0 before learning their type. Having learnt their type, they then seek financing at time-1. At time-2, a signal reveals the type of the banker to financial markets. Depending on the maturity of bail-in debt, investors have the option not to roll over their debt. The government then decides whether to recapitalise banks for reasons introduced below. The game ends at time-3 when all uncertainty is resolved, and funds are distributed among bankers, investors and depositors.

3.2.1 Bankers

Asset

At time-1, a mass of bankers have access to an asset which generates returns at time-3.[7] There are two types of bankers, $\theta \in \{\theta^L, \theta^H\}$. A bank's type is unobservable to investors. Let the fraction of high types, θ^H, in the economy be given by μ_0. Their asset return, denoted by X^H, is deterministic. The asset return of the low type, θ^L, is drawn from a distribution with CDF $F(x)$. Let X^L denote the mean of the distribution and let $\underline{x} > 0$ denote its lower bound. We refer to this lower bound as the asset's safe return.

At time-2, the period after initial financing, a signal perfectly reveals the type of the banker (but not the return of the low type). Importantly, we assume that the signal is not contractible: agents cannot write contracts contingent on its realisation at time-2.[8]

Lastly, assets can be liquidated pre-maturely without cost. The liquidation value of the low type's asset is thus given by its expected return.[9]

Scope for intervention: Ex-post continuation investment

In the model, banks face an investment opportunity at time-2. This ex-post investment generates a scope for intervention: banks underinvest when they face a debt overhang. The return of the continuation investment i is increasing in effort e in a first-order stochastic dominance sense:

$$i(e) = \int_0^{\bar{\iota}} y \, dG(y|e) \tag{3.1}$$

where $G(y|e)$ denotes the CDF of investment returns, satisfying $G(y|e) > G(y|e')$ if $e' > e$ for all $y \in \big(0, \bar{\iota}\big)$. However, effort incurs a private cost $c(e)$ to the banker who is initially running the bank.[10] Let z denote the expected continuation investment returns net of effort cost: $z(e) = i(e) - c(e)$. We assume that $z(e)$ is strictly concave and twice continuously differentiable. It follows that the

[7] We endogenise the size of the mass of bankers in Section 3.5.

[8] We discuss a contractible signal in Section 3.8.2, and show that it leads to multiplicity: banks either finance using short-term debt, or using CoCo bonds that fully dilute the banker upon conversion.

[9] We discuss costly liquidation in Section 3.7.

[10] Looking ahead, bail-ins involve equity write-downs and debt-to-equity swaps, and hence affect who is running the bank and thus facing the private effort cost.

socially optimal effort level, denoted by e^*, satisfies

$$z'(e^*) = 0 \tag{3.2}$$

Going forward, let z^* denote the efficient expected net investment return: $z^* = z(e^*)$.

When choosing effort levels, banks maximise their expected payoffs. Importantly, they are protected by limited liability. The optimal effort choice by the low type is characterised by

$$e(\hat{B}) \;=\; arg\,\underset{e}{max} \int_0^{\bar{\iota}} \int_{\underline{x}}^{\infty} max\{x + y - \hat{B}, 0\}\, dF(x)\, dG(y|e) \;-\; c(e) \tag{3.3}$$

where $\hat{B}$ denotes the level of debt outstanding. Effort is inefficiently low if—with positive probability—some of the investment proceeds accrue to outside creditors and not the banker who is bearing the effort cost. It follows that banks exert too little effort whenever the asset return does not fully cover all debt outstanding with probability one. Turning to the high type, they exert efficient levels of effort if their deterministic return covers all debt outstanding: $X^H \geq \hat{B}$.

Assumption 3.2.1. *We make the following assumptions on asset and investment returns:*

1. *The total expected returns generated by bankers of the low type, even at the efficient level of expected investment gains, do not cover the initial required financing cost:* $X^L + z^* \;<\; 1$.
2. *The asset returns generated by bankers of the high type are sufficiently high such that financing is always achieved in equilibrium:* $X^H \;>\; \frac{1}{\mu_0}$.

The first assumption implies that the asymmetric information on asset returns is sufficiently severe to create adverse selection: low type bankers cannot achieve financing if they are identified as such. Hence, runs on identified low types occur for fundamental reasons in equilibrium. The second assumption implies that the high type's asset return is sufficiently high such that—even if the low type does not repay anything to bail-in investors who provide unit financing—the high type can repay all debt obligations with probability one. As a consequence, the financing market does not collapse. This assumption also implies that high types always invest efficiently.

Financing

Financing of the asset is obtained from investors of bail-in debt and depositors. We assume that bankers need to raise at least δ units of bail-in debt. Without loss of generality, we assume that this constraint binds. As a benchmark, we further assume that all bail-in debt is either short-term or long-term debt. Banks thus choose a financing strategy $\sigma \in \{d, D\}$, where d denotes short-term debt financing and D denotes long-term debt financing.[11] Short-term debt issued at time-1 matures at time-2. Importantly, short-term debt matures after the signal realisation but before the government can intervene.[12] Long-term debt does not need to be rolled over but instead matures at time-3.

Assumption 3.2.2. *We make the following assumptions on deposit and bail-in debt financing relative to asset returns:*

[11] We endogenise the breakdown of short and long maturity bail-in debt in Section 3.6, and allow banks to issue equity in Section 3.8.1.

[12] Any short-term debt maturing after the government intervention corresponds to long-term debt in our model.

1. *Bankers need to raise $\delta > X^L$ units in bail-in debt.*
2. *In principle, bail-ins are an effective tool to recapitalise banks since the low type's safe asset return covers deposits: $\underline{x} \geq (1 - \delta)$.*

The first assumption implies that the face value of bail-in debt exceeds the liquidation value of the low type's assets. In any equilibrium featuring short-term bail-in debt, such debt is risky since some investors cannot be repaid in full when choosing not to roll over.[13] The second assumption implies that bail-ins are effective absent of runs. This is the most interesting case economically. The point of this paper is to show how bail-ins create fragilities which render bail-ins ineffective and lead to an erosion of market discipline. Without this assumption, banks can never be recapitalised using bail-ins, no matter the debt maturity structure. More interestingly, we ask whether agents write debt contracts that render bail-ins ineffective in an environment in which they generally can be.

3.2.2 Investors

A unit mass of investors purchase the bail-in debt issued by bankers. They are risk-neutral, are endowed with a unit of capital at time-1, and do not discount the future. Importantly, banks have complete market power: they post a financing contract which is then priced competitively by the market, subject to the investors' participation constraint. Investors form beliefs about the type of a bank given their financing decision; upon the realisation of a given signal, investors update their beliefs according to Bayes' rule. These prior and posterior beliefs determine the required interest rate such that investors are indeed willing to supply funding and rollover their maturing claims.

We assume that all investors of the same class of debt face the same losses, either during a run, or when they are imposed on them by the government. Importantly, runs in the paper arise for fundamental reasons, not due to investor miscoordination.

3.2.3 Depositors

Depositors supply funds to the bank. Importantly, depositors in this paper are the simplest way of introducing bank liabilities that cannot be bailed in. They have no features typical of deposits, i.e. liquidity shocks the spirit of [8]. Instead, they are fully insured by government. Given this insurance, they demand an interest rate of one. Depositors never withdraw their funds at time-2.

3.2.4 Government

The government chooses its desired recapitalisation strategy at the beginning of the game and implements it at time-2. In Section 3.3, recapitalisation is achieved through bailouts. In Section 3.4, the government imposes bail-ins; if they do not induce efficient investment levels, bail-ins are followed by bailouts.

3.3 A simple notion of bailouts

We introduce a very simple notion of *bailouts* which fully captures the government's time inconsistency problem at the heart of the analysis on banks that are considered 'too big to fail'.

[13] Alternatively, we could assume that only part of the asset can be liquidated as short-term bail-in debt investors withdraw; or we could include costs that reduce the liquidation value of assets.

Assumption 3.3.1. *Public funds are free at the time of intervention, but the government needs to raise distortionary taxes at the beginning of the game. The government is unable to commit not to conduct bailouts by not raising taxes initially.*[14]

As a result, the government fully removes the debt overhang at time-2. This is achieved by using public funds to decrease outstanding debt, $\hat{B}$, to match the minimum asset return. Thus, the low type receives a bailout of $b = \hat{B} - \underline{x}$. Combined with our assumption on the deterministic return of the high type, Assumption 3.3.1 ensures that all banks can always repay all debt. It follows that $\hat{B} = 1$, as all interest rates are equalised to one. Anticipating full repayment due to the government's guarantees, investors never have any incentives to withdraw funds from the bank. Since the government fully removes the debt overhang, the low type's expected payoff corresponds to the risky component of the asset, yielding $X^L - \underline{x}$ in expectation, and the expected net investment gain z^*. Given the the lowest possible debt burden of $\hat{B} = 1$, the high type achieves their maximum payoff:

$$V^{L,bailout} = x^L + z^* - \underline{x} \qquad\qquad V^{H,bailout} = V^{H,max} = x^H + z^* - 1 \tag{3.4}$$

Anticipating the equilibrium section, any maturity structure of debt can be supported in equilibrium as both types achieve the same payoffs for any financing decision. Investors of short-maturity debt are willing to rollover even when they learn to have lent to the low type, as the government ensures full repayment when inducing efficient investment levels.

3.4 Bail-ins

In this section, suppose the government's preferred recapitalisation strategy are *bail-ins.* During this type of intervention, all bail-in debt claims are converted into equity in order to induce efficient investment. Before proceeding to compute equilibrium payoffs for different financing strategies under a bail-in policy, we provide some details on how we model bail-ins.

Assumption 3.2.2 ensures that bail-ins—if no assets are liquidated—fully remove the debt overhang and enable efficient investment. The definition for bail-in ineffectiveness follows naturally:

Definition 3.4.1 (Bail-in ineffectiveness)**.** *Bail-ins are ineffective whenever they are unable to induce efficient levels of the continuation investment:* $z^{bail\text{-}in} < z^*$.

Secondly, we need to specify how the government acts if bail-ins are indeed unable to induce efficient investment behaviour. Continuing the simple notion of bailouts introduced in the previous section, bailouts follow ineffective bail-ins: whenever bail-ins do not induce the socially optimal level of investment, the government uses bailouts of the minimum size required to fully remove the remaining debt overhang.

Thirdly, we need to specify the distribution of resources among investors and the banker during a bail-in. This boils down to determining the post-intervention equity shares. For this purpose, let $\underline{V}^\theta$ denote the virtual value of the banker's equity claim absent of intervention, anticipating future

[14]Alternatively, bailouts could be modelled as in [4] where a proportional administration cost $1 > \kappa > 0$ accrues at the time of bailouts, and a proportional cost of distortionary taxation $\tau > 1$ accrues at time-1 (rather than time-2 when bailouts occur). This generates a cut-off rule for bailouts where the marginal gain of bailouts is traded off against their marginal cost, κ. Importantly, if the cost of bailouts is too large at the time of intervention, then this cost would act as a credible commitment device not to bail out the banker, and the model would never feature bailouts.

insolvency conditional on low investment and asset returns. Let V^θ denote the banker's payoff post-intervention. Motivated by existing bail-in regulation we consider the following condition:

Definition 3.4.2 (No creditor worse off (NCWO) condition)**.** *Any government intervention cannot make any agent—neither the banker nor investors nor depositors—strictly worse off. With regards to the banker, it must be that*

$$V^\theta \geq \underline{V}^\theta \tag{3.5}$$

Assumption 3.4.1. *The government converts bail-in debt to equity at a conversion rate such that the banker's NCWO condition as outlined in Equation (3.5) holds with equality.*

Existing bail-in regulation requires that any bank shareholders or creditors cannot be worse off under resolution than under a no-intervention policy. Since equity has a strictly positive virtual value absent of intervention—even with insolvency on the horizon for low investment and asset returns—the banker must not be fully diluted during a bail-in.[15] In the model, the NCWO condition can be derived from an intervention participation constraint that needs to be satisfied for each agent. Whenever the government intervenes and leaves any agent strictly worse off, this agent has the option of suing the government and forcing them to raise taxes to generate the no-intervention payoff. When honouring the NCWO condition during intervention, the government can avoid any such legal costs.

Before discussing the implications of the NCWO condition for our results, we introduce a notion of default which may be arising as a result of withdrawing creditors with maturing debt claims:

Definition 3.4.3 (Default)**.** *Whenever a debtor is unable to repay a creditor with a maturing debt claim, the debtor is defaulting. In this state, they receive a payoff of zero.*

Importantly for the results of the paper, the banker—having issued long-term debt—must not be fully diluted in all states of the world during a bail-in. At full dilution, the low type always achieves a zero payoff. The crucial difference between short-term and long-term debt is that short-term debt experiences runs in anticipation of losses which may not even realise. These runs lead to default. Thus, avoiding runs through long-term debt financing preserves value to low type bankers.[16]

3.4.1 Long-term debt

Suppose banks issue long-term debt: $\sigma = D$. While markets price bank debt competitively, the investors' participation constraint needs to be satisfied. Given Assumption 3.2.1 on the model's parameters, the low type does not generate returns that cover the initial cost of financing. Hence there exists a threshold belief $\underline{\mu}^D$ such that investors do not supply funding for all $\mu < \underline{\mu}^D$. All banks receive a zero payoff for such beliefs.

Consider an arbitrary belief $\mu \geq \underline{\mu}^D$ of lending to the high type for which banks do achieve financing. By the nature of long-term debt, it only matures at the end of the game and thus investors cannot respond to the signal. Once the government identifies the low type bankers, they

[15] See e.g. the *Bank of England's approach to resolution*, October 2017, based on the European Commission's Bank recovery and resolution directive (BRRD), Article 73.

[16] Equivalently, we could assume that uncertainty over returns is realised before the investment decision takes places and thus before the government intervenes. If the banker has issued long-term debt, avoiding runs and early liquidation, and invests efficiently with some strictly positive probability, avoiding a bail-in, they receive a strictly positive expected payoff.

bail in all long-term debt.[17] The investor post-intervention equity share is denoted by γ. The participation constraint for a given investor is given by

$$\delta \;\leq\; \mu R_1 \delta \;+\; (1-\mu)\gamma\big[X^L + z^* - (1-\delta)\big] \tag{3.6}$$

where R_1 denotes the long-term interest rate. Bankers raise δ units of bail-in debt. Given their belief, investors lend to a high type who is able to fully repay all debt at the promised long-term interest rate with probability μ. Investors believe to be bailed in with probability $(1-\mu)$, and to receive a share γ of equity returns. Clearly, the higher the rate of debt-to-equity conversion—captured by the equity share γ—the lower is the debt burden on the high type who needs to compensate investors for losses arising when lending to the low type.

This rate of conversion is determined according to the banker's NCWO condition. Consider a scenario without any intervention. Having issued long-term debt, and given deposit insurance that prevents depositors from withdrawing, the low type does not experience runs despite being identified as such upon the realisation of the low signal. As long as the cumulative upper bound of the asset and investment return distributions is sufficiently high, they face a strictly positive probability of being able to fully repay all debt. Hence, their expected payoff is strictly positive:

$$\underline{V}^L(D,\mu) \;=\; \max_e \int_0^{\bar{\iota}} \int_{\underline{x}}^{\bar{x}} max\big\{x + y - \hat{B}(\mu), 0\big\}\, dF(x)\, dG(y|e) \;-\; c(e) \;>\; 0 \tag{3.7}$$

where $\hat{B}(\mu)$ denotes the debt burden in the absence of intervention at belief μ. By Assumption 3.4.1, the government dilutes the banker such that their NCWO condition binds:

$$V^L(D,\mu) \;=\; (1-\gamma)\big[X^L + z^* - (1-\delta)\big] \;=\; \underline{V}^L(D,\mu) \tag{3.8}$$

Bankers issue debt at an interest rate R_1 such that Equation (3.6) holds with equality. Combining this with Equation (3.8), the payoff to the high type is then given by:

$$\begin{aligned} V^H(D,\mu) \;&=\; X^H \;+\; z^* \;-\; (1-\delta) \;-\; R_1\delta \\ &=\; X^H \;+\; z^* \;-\; 1 \;-\; \frac{1-\mu}{\mu}\Big[\underline{V}^L(D,\mu) + 1 - X^L - z^*\Big] \end{aligned} \tag{3.9}$$

The high type's payoff is a combination of the asset and net investment returns, minus the unit financing cost and the investor loss of lending to the low type for which each high type banker needs to compensate. Note that $V^H(D,\mu)$ is strictly increasing in μ: the higher are investor beliefs to be lending to a high type, the lower is the required interest compensation for the loss arising due to the presence of the low type. The high type's payoff reaches their maximum value $V^{H,max}$ as $\mu \nearrow 1$.

[17] We could equivalently assume that investors are only bailed in to the degree that the debt overhang is exactly removed. Bail-in creditors receive all of the asset's safe return in excess of deposits, plus a share of equity returns post-intervention that is lower than the one in the main body of text. Given our two assumptions that a) banks issue debt with market power, only honouring the creditors' participation constraint, and that b) bail-ins dilute previous equity holders to the degree that their NCWO constraint binds, both models yield the same payoff to investors and bankers.

3.4.2 Short-term debt

Suppose that banks issue short-term debt: $\sigma = d$. As for long-term debt, there exists a threshold belief $\underline{\mu}^d$ such that financing is not achieved and banks receive a zero payoff for all $\mu < \underline{\mu}^d$. Conditional on achieving initial financing at time-1, we solve the game backwards from time-3 to establish that investors—in anticipation of government bail-ins and being subordinated to depositors—do not want to roll over once they identify the low type at time-2, leading to runs on bail-in debt.

Lemma 5. *When issuing short-term bail-in debt, runs ensue, rendering bail-ins ineffective.*

At time-2, investors update their beliefs about the banker's type according to Bayes' rule. Given the signal structure, the type of the banker is perfectly revealed. Short-term bail-in debt matures, and investors decide whether to withdraw or rollover their debt. Afterwards, the government intervenes to induce efficient investment levels.

Let the short-term interest rate at time-1 be denoted by r_1. The face value of short-term debt is thus given by $r_1\delta$. Suppose investors learn of having lent to a high type banker. Our parametric assumptions ensure that identified high type bankers are always able to rollover their maturing short-term bail-in debt. Given the banks' market power in financing markets, the required interest rate on newly issued short-term debt is determined according to the investor participation constraint. Lending to the identified high type is safe and the rollover interest rate is given by one. Investors can withdraw the face value of their debt today, or rollover and receive the face value tomorrow. Thus, conditional on a high signal, investors receive a payoff of $r_1\delta$.

Suppose investors learn of having lent to a low type banker. Suppose they roll over their maturing debt. Given the debt overhang $\hat{B} > \underline{x}$, bankers invest inefficiently, and the government bails in investors. The bank now invests efficiently, and the total expected net returns to be distributed among depositors, investors and bankers are given by $X^L + z^* < 1$. Depositors hold the most senior claims on bank returns. At the same time, the government does not fully dilute the banker to comply with the NCWO condition, in analogy to Equation (3.8). It follows that the expected payoff to investors at time-3 given the anticipated bail-in must be strictly less than the face value of debt maturing at time-2:

$$X^L + z^* - (1-\delta) - \underline{V}^L(d,\mu) \; < \; \delta \; \leq \; r_1\delta \tag{3.10}$$

Therefore, all investors choose to withdraw at face value upon the realisation of a low signal. The bank is forced to liquidate its asset. Since $X^L < \delta$, the banker cannot repay all investors, even at an interest rate of one. It follows that the banker defaults and receives a zero payoff for all beliefs μ:[18]

$$V^L(d,\mu) \; = \; 0 \tag{3.11}$$

Having established that the bank is forced into default by withdrawing investors, suppose that the government intervenes subsequently by bailing in those investor claims that have not been repaid. The previous banker is fully diluted, given their default payoffs of zero. Since all of the assets have been depleted to serve withdrawals, but the deposits remain in place, the new bank

[18] We cannot compute the low type's payoff when financing with short-term debt for a belief of $\mu = 1$ since the occurrence of a low signal is a zero probability event, and Bayes' rule is not defined. We assume that the low type achieves a payoff of zero for such a belief, as they do for all other beliefs $\mu \in [0, 1)$ where Bayes' rule is defined.

equity holders exert sub-optimal levels of effort:

$$arg\,\underset{e}{max} \int_0^{\bar{\iota}} max\big\{y - (1-\delta), 0\big\}\, dG(y|e) \ - c(e) \ < \ e^* \tag{3.12}$$

It follows that the government conducts bailouts to induce efficient investment levels. In particular, the government fully removes the debt overhang: $b = (1 - \delta)$. As a result, the expected net gains from investing to the new equity holders—the previous investors—is given by z^*, leading to a total expected payoff of low type investors of $X^L + z^*$.

Having characterised payoffs of the low type banker and their investors, we can now turn to the payoffs of high type bankers and their investors. Given belief μ of facing a high type, the investor participation constraint is given by

$$\delta \ \leq \ \mu\, r_1 \delta \ + \ (1-\mu)\,(X^L + z^*) \tag{3.13}$$

Banks issue short-term debt with interest rate r_1 such that this condition holds with equality. The payoff to the high type is then given by

$$\begin{aligned} V^H(d,\mu) \ &= \ X^H + z^* - (1-\delta) - r_1\delta \\ &= \ X^H + z^* \ - \ 1 - \frac{1-\mu}{\mu}\big(\delta - X^L - z^*\big) \end{aligned} \tag{3.14}$$

As for long-term debt, $V^H(d, \mu)$ is strictly increasing in μ, and it reaches its maximum possible value, $V^{H,max}$, at $\mu = 1$. Notably, the payoff now depends on the amount of bail-in debt financing δ. Runs force the government to bailout depositors, but losses are still inflicted on investors as asset return and net investment gain do not cover their initial financing. The higher δ, the larger is the shortfall for investors for which the high type needs to compensate.

3.4.3 Equilibrium

The subgame beginning at time-1 only admits pooling equilibria. If the low type is identified as such when seeking financing, no investor would purchase any of their bail-in debt: even at the efficient investment level, the total returns cannot cover the cost of financing in expectation: $X^L + z^* < 1$. It follows that the low type always mimics the high type, and any equilibrium of this game must be a pooling equilibrium. The share of high types in the economy, μ_0, is then the equilibrium prior belief of investors to be lending to a high type.

With that in mind, we begin by establishing that financing with short-term debt is equilibrium dominated by financing with long-term debt for the low type; and that no financing strategy is equilibrium dominated for the high type.

Definition 3.4.4 (Equilibrium dominance)**.** *An equilibrium strategy σ dominates strategy σ' for a given type θ if the payoff for strategy σ at the equilibrium prior belief μ_0 is strictly larger than maximum payoff for strategy σ' at the best possible belief:*

$$V^\theta(\sigma, \mu_0) \ > \ \underset{\mu}{sup}\, V^\theta(\sigma', \mu) \ = \ V^\theta(\sigma', 1) \tag{3.15}$$

The (weakly) highest expected payoff when financing using a particular class of debt is achieved whenever markets believe that a banker is the high type with probability one. Accordingly, the high type achieves the maximum payoff $V^{H,max}$ for both short-term and long-term debt if $\mu = 1$. This maximum payoff is strictly larger than the payoff, as characterised by Equations (3.9) and (3.14) for long-term or short-term debt, at the equilibrium prior μ_0. Thus, neither strategy is equilibrium dominated. Turning to the low type, they achieve a zero payoff when financing using short-term debt due to the runs that ensue upon the realisation of a low signal. Given the government's bail-in policy, they achieve a strictly positive equilibrium payoff $V^L(D, \mu_0) > 0$ when financing using long-term debt, as by Equation (3.7). It follows that financing using short-term debt is equilibrium dominated for the low type.

Assumption 3.4.2 (Equilibrium refinement: Intuitive criterion)**.** *Consider a financing strategy σ that equilibrium dominates strategy σ' for type θ but not for type θ'. Then if investors observe a deviation from strategy σ to strategy σ', they must believe that it is type θ' who is deviating:*

$$\mu(\theta'|\sigma') = 1 \tag{3.16}$$

We are now ready to define the equilibrium of the sub-game beginning at time-1:

A Perfect Bayesian Equilibrium of the sub-game beginning at time-1 is given by a set of financing strategies $\sigma \in \{d, D\}$ and an investor belief system μ. The financing strategies maximise the banks' expected payoffs for both types $\theta \in \{\theta^L, \theta^H\}$, given investor beliefs. In equilibrium, the investors' prior beliefs for each form of financing must be correct and are updated according to Bayes' Law whenever possible. Off-equilibrium beliefs are refined according to the intuitive criterion.

Proposition 6. *In the unique equilibrium of the sub-game starting at time-1, both firms finance using short-term debt. Runs occur upon the realisation of the low signal, rendering bail-ins ineffective.*

The result of Proposition 6 follows from above discussion. Consider first an equilibrium candidate in which both types pool on financing using long-term debt. According to the refinement of off-equilibrium beliefs, investors believe that any deviation to short-term debt must be coming from the high type. It follows that the high type achieves their maximum payoff when financing using short-term debt given this belief, and thus faces a profitable deviation. This rules out pooling equilibria on long-term debt.

Next, we show that pooling on short-term debt indeed constitutes an equilibrium. The strategy of financing using short-term debt does not equilibrium dominate financing using long-term debt for either type: markets are free to form any off-equilibrium belief for a deviation to long-term debt. Then, for any off-equilibrium belief that assigns sufficiently high probability to a deviation to long-term debt to the low type, neither type has an incentive to deviate to long-term debt as they face highly unfavourable credit conditions, or worse, cannot achieve financing altogether. Thus, pooling on short-term debt is indeed an equilibrium strategy.

Intuitively, the short-term debt equilibrium arises not because it is necessarily associated with higher expected payoffs for the high type. Instead, the high type attempts to distinguish themself from the low type by choosing a financing strategy that is very painful to the low type. Pooling on short-term debt induces frequent runs on the low type and this emergent fragility is exactly what renders bail-ins ineffective.

3.5 Entry into financing markets and efficiency

This section models the bankers' entry decision into financing markets and develops a notion of efficiency. At time-2, there is a unit mass of bankers who only learn their type once they have entered financing markets. Each banker has the same probability μ_0 to be the high type. Applying the law of large numbers, μ_0 is thus the investors' equilibrium prior belief of facing a high type at time-1.

However, entering financing markets is costly. Each banker draws their fixed cost k^i from a continuous distribution over the support $\left[\underline{k}, \bar{k}\right]$. Thus, banker i enters if

$$\mu_0 V^H(\sigma, \mu_0) \; + \; (1 - \mu_0) V^L(\sigma, \mu_0) \; \geq \; k^i \tag{3.17}$$

where V^H and V^L denote equilibrium payoffs for high and low types at the equilibrium financing strategy σ and prior belief μ_0.

Once the bankers' types have been determined, the regulator would like to prevent the low types from financing. We assume that this is not possible ex-post. Instead, we define a notion of ex-ante constrained efficiency:

Definition 3.5.1 (Constrained efficiency)**.** *Three conditions need to be satisfied for efficiency:*

1. *Banks of both types of asset returns invest efficiently post-intervention.*
2. *The fixed cost of the marginal type that is indifferent between entering or not entering into financing markets is given by k^* which satisfies*

$$\underbrace{\mu_0 \, X^H \; + \; (1 - \mu_0) \, X^L \; + \; z^*}_{\text{expected (net) returns}} \; = \; \underbrace{1 \; + \; k^*}_{\text{total cost}} \tag{3.18}$$

   This implies that only (weakly) positive NPV projects enter financing markets.

3. *The government is not forced to use public funds, neither for bailouts, nor for deposit insurance.*

Applying this constrained efficiency condition to the equilibrium described in Proposition 6 yields the paper's second main result:

Proposition 7. *Forcing banks to issue long-term bail-in debt achieves constrained efficiency. The market equilibrium short-term debt contract induces excessive entry into financing markets and does not lead to socially optimal investment unless the government conducts bailouts.*

When the government recapitalises banks using bail-ins and banks are forced to finance themselves using long-term debt, then all criteria for efficiency are met. First, bail-ins achieve the socially efficient investment for both high and low types. Second, combining Equations (3.7) and (3.9), and verifying against Equation (3.18), shows that the marginal type that enters financing markets pays a fixed cost of k^*. Third, the government is not required to use any public funds.

Bail-ins in the presence short-term debt lead to inefficiencies for two reasons. First, the government uses public funds to achieve the efficient continuation investment. Second, entry into financing markets is inefficiently high. To show the latter, combining Equations (3.11) and (3.14) yields the following expression for the average expected payoff from entering financing markets:

$$\begin{aligned} V^{bail-in}(d, \mu_0) &= \mu_0 \, V^H(d, \mu_0) \, + (1 - \mu_0) \, V^L(d, \mu_0) \\ &= V^* \, + \, (1 - \mu_0) \, (1 - \delta) \end{aligned} \tag{3.19}$$

where V^* denotes the expected average payoff from entering associated with efficiency. Entry is inefficiently high whenever $V^{bail-in}(d, \mu_0) > V^*$. The term $(1 - \mu_0)\,(1 - \delta) > 0$ captures the government subsidy to banks due to the bailouts conducted after bail-ins. The government subsidises the new bank equity holders, the previous bail-in debt holders, by removing the debt overhang to induce efficient levels of investment. This in turn reduces the debt burden of the high types, increases their payoffs and leads to excessive incentives to enter financing markets.

Lastly, and for completeness, bailouts lead to inefficiently high entry into financing markets. In particular, the expected payoff from entering financing markets in the presence of bailouts satisfies

$$V^{bailout}(\sigma, \mu_0) \, = \, V^* \, + \, (1 - \mu_0) \, b \tag{3.20}$$

for all financing strategies σ, and where $b = 1 - \underline{x}$. Additionally—while efficient levels of the continuation investment are achieved—bailing out the low type bankers requires the use of public funds.

3.6 Mixed maturity debt and long-term debt requirements

A brief inspection of balance sheets reveals that banks issue a mixture of long-term and short-term debt that is not collateralised or protected by deposit insurance. Additionally, banks face long-term debt requirements. In the case of the UK, bail-in debt under the minimum requirement for own funds and eligible liabilities (MREL) requires a minimum maturity of one year.[19] Nevertheless, losses can be imposed on unsecured debt beyond MREL during a resolution. This section points out how run incentives and consequences in terms of bail-in effectiveness depend on whether short-term and long-term debt are treated differently during intervention. To demonstrate, we consider bail-in debt with mixed maturities for two policies: one in which short-term and long-term debt experience the same loss rates (or equivalently, are ranked *pari passu*); and one in which long-term debt is subordinated to short-term debt.

3.6.1 Pari passu

We extend the model to allow for mixed maturity debt which is ranked pari passu. Banks issue both short-term debt, $d > 0$, and long-term debt, $D > 0$, to finance the required δ units.

Proposition 8. *In equilibrium, banks issue a combination of short-term and long-term debt that induces runs and thus renders bail-ins ineffective.*

The equilibrium logic of Section 3.4 is unchanged. To illustrate, consider two candidate debt combinations (d, D) and (d', D'). Once investors learn that they have lent to the low type upon the realisation of a low signal, short-term investors withdraw their funds. In expectation, the low type is unable to repay its total debt obligations even after bail-ins induce efficient investment since

[19] See the *The Bank of England's approach to setting a minimum requirement for own funds and eligible liabilities (MREL)*, p.7, Article 5.2.

$X^L + z^* < 1$. Since short-term debt is ranked pari passu with long-term debt, investors prefer not to roll over. The two debt combinations differ in terms of their ability to fully serve all withdrawing investors:

$$d' \; > \; X^L \; > d \tag{3.21}$$

It follows that a debt combination with short-term debt exceeding the value of assets conditional on a low signal induces default, yielding a zero payoff to the low type. Whenever banks are not forced into default, part of the asset remains on their balance sheet which, combined with the possibility of high investment outcomes even for low levels of effort, yields a strictly positive expected payoff to equity. Since the government intervenes but honours the banker's NCWO condition, they achieve a strictly positive payoff after a bail-in. It follows that financing using d units of short-term debt equilibrium dominates financing using d' units of short-term debt for the low type:

$$V^L(d, \mu_0) \; > V^L(d', 1) \tag{3.22}$$

Thus, any equilibrium contract with mixed maturity debt and short-term debt ranked pari passu must induce defaults due to runs on short-term bail-in debt, rendering bail-ins ineffective and leading to inefficiently high entry into financing markets.

Long-term debt requirements and bail-in effectiveness. Bail-ins are effective whenever they are able to fully remove the debt overhang. This requires that—after a bail-in—the safe return of any non-liquidated share of the asset exceeds the deposits outstanding. All short-term investors withdraw upon the realisation of a low signal. The bank liquidates assets at price X^L in order to repay d units of debt. It follows that the non-liquidated share of the asset is given by $\left(1 - \frac{d}{X^L}\right)$. Then bail-ins are effective whenever the amount of short-term bail-in debt, d, satisfies

$$\underbrace{\left(1 - \frac{d}{X^L}\right) \underline{x}}_{\text{safe asset return}} \;\; \geq \;\; \underbrace{(1 - \delta)}_{\text{deposits}} \tag{3.23}$$

In this context, let us interpret δ as capital requirement with the goal of achieving bail-in effectiveness. If the regulator values deposits (for reasons of liquidity and maturity transformation along the lines of [8] but unmodelled here), then they would set δ such that bailing in equity and investor debt exactly removes the debt overhang: $(1 - \delta) = \underline{x}$. If short-term and long-term bail-in debt rank pari passu, then all bail-in debt must be long-term debt in order not to render bail-ins ineffective.

3.6.2 Subordination of long-term debt to short-term debt

In this subsection, we discuss an alternative bail-in policy in which long-term debt is subordinated to short-term debt.[20] As in the previous subsection, we show that there exists a level of short-term debt inducing runs. All debt combinations with short-term debt exceeding this level are equilibrium dominated by other debt combinations for the low type, and thus any equilibrium must feature runs.

[20] In principle, there are many ways in which long-term debt could become subordinated to short-term debt. In terms of regulation, subordination of long-term debt boils down to preferential treatment of short-term debt during insolvency.

Suppose all short-term debt holders roll over in anticipation of a bail-in during which they are treated preferentially. On expectation, total returns of $X^L + z^*$ are distributed among banker, investors and depositors after the government has conducted a bail-in. Depositors are still served first, receiving full repayment of $(1 - \delta)$. Although long-term debt and equity are subordinated to short-term debt, both must receive strictly positive expected payoffs by the NCWO policy (Assumption 3.4.1). Let the payoff to long-term debt denoted by $\underline{D}(d, \mu_0)$. As before, the payoff of the banker is denoted by $\underline{V}^L(d, \mu_0)$. It follows that short-term debt holders are willing to rollover into a bail-in if

$$r_1 d \ \leq \ X^L + z^* - (1 - \delta) - \underline{D}(d, \mu_0) - \underline{V}^L(d, \mu_0) \tag{3.24}$$

If this condition holds, then it must be that $r_1 = 1$. Both for the low and the high type, short-term creditors rollover in future; this is only true if they expect to be repaid at the end of the game (in expectation), and face no credit risk at time-1. If Equation (3.24) does not hold, and the level of short-term debt exceeds the liquidation value of the asset ($d > X^L$), then short-term debt commands an interest rate $r_1 > 1$. It follows that short-term creditors a) do not want to rollover, and b) force the banker into default if

$$d \ > \ max\Big\{X^L + z^* - (1 - \delta) - \underline{D}(d, \mu_0) - \underline{V}^L(d, \mu_0), \ X^L\Big\} \tag{3.25}$$

Importantly, such a contract (D, d) exists given the parameter space considered in this paper. To illustrate, consider a combination of debt with $d \nearrow \delta$. By Assumptions 3.2.1 and 3.2.2, there exists a level of short-term debt for which Equation (3.25) is satisfied.

Long-term debt requirements and bail-in effectiveness. This section highlights that if regulators want to achieve bail-in effectiveness while minimizing the required amount of long-term bail-in debt (relative to short-term bail-in debt), then subordinating long-term debt to short-term debt becomes a useful policy tool. As long as Equation (3.24) is satisfied, such a policy helps to avoid any liquidation of assets even though some bail-in debt has short maturity.[21][22] Bail-in effectiveness is preserved. This is in stark contrast to the scenario in which all bail-in debt claims rank pari passu and short-term investors always withdraw.

3.6.3 Discussion of current UK policy

The UK's insolvency creditor hierarchy, to which the resolution authority adheres, is laid out in the *Bank of England's approach to resolution* (2017).[23] Deposits below the insurance threshold of GBP 85,000 are protected by the FSCS deposit insurance scheme. Further deposits by individuals and SMEs in excess of the threshold are ranking in the creditor hierarchy above senior unsecured and thus above designated bail-in debt. However, deposits in excess of the insurance threshold, which have not been deposited by individuals or SMEs, are ranked pari passu with other unsecured bail-in debt. The model predicts that the latter deposits should be prone to runs in anticipation of bail-ins, leading to outflows out of banks in the run-up of intervention. If these outflows are sufficiently large, then bail-ins are rendered ineffective.

[21] Any runs that ensue for pure liquidity reasons as $d > X^L$ but Equation (3.24) is satisfied can be avoided by a lender of last resort which is never forced to extend any emergency loans in equilibrium.

[22] Costly liquidation of assets is considered in Section 3.7.

[23] Source: link, page 18.

3.7 Disciplining effect of short-term debt

Section 3.5 demonstrated that short-term debt financing leads to excessive entry into financing markets. In this version of the model, we show that costly liquidation of assets may prevent positive NPV projects from entering. In other words, disciplining of financial markets may be excessive.

Suppose that the low type, once identified as such, can liquidate their assets at rate $\ell^L = \ell < 1$. Thus, liquidation is costly. However, we assume that there is no liquidation cost for the high type, $\ell^H = 1$. This avoids illiquidity issues for identified high types and rules out the need for intervention.[24][25] Since pre-mature liquidation of the asset now implies a waste of resources, we need to expand the notion of efficiency by one condition:

Definition 3.7.1. *Additional efficiency condition: No liquidation of assets in equilibrium.*

Adjusting the analysis of Sections 3.4 and 3.5 to incorporate the cost of liquidation, yields the paper's next result:

Proposition 9. *Market discipline increases if liquidation is costly. However, this disciplining effect a) is excessive whenever the liquidation loss exceeds the government's bailout, and positive NPV projects do not enter the financing market; and b) comes at the cost of ex-post inefficiencies.*

Suppose banks issue short-term debt. Since runs occurred without liquidation cost, they must also occur whenever liquidation is costly: $\ell X^L + z^* < X^L + z^* < \delta$. The low type again defaults. As before, the high type is always able to rollover short-term debt. The required short-term interest rate is then derived from the investor participation constraint as in Equation (3.13), adapted to incorporate the liquidation cost and again holding with equality:

$$r_1(\mu_0|\ell < 1) \; = \; \frac{1}{\mu_0} - \frac{1-\mu_0}{\mu_0}\frac{\ell X^L + z^*}{\delta} \; > \; r_1(\mu_0|\ell = 1) \tag{3.26}$$

Effectively, the high type now also needs to compensate for the liquidation loss which is further reducing the total repayment by the low type. Using the expression for the short-term interest rate, the payoff to the high type when financing using short-term debt is given by

$$V^H(d, \mu_0|\ell < 1) \; = \; X^H + z^* \; - \; \left[1 + \frac{1-\mu_0}{\mu_0}\left(\delta - \ell X^L - z^*\right)\right] \; < \; V^H(d, \mu_0|\ell = 1) \tag{3.27}$$

The properties of V^H are unchanged but there is a level reduction due to the increase in interest payments on each unit of short-term bail-in debt. This increase in the cost of borrowing also manifests itself in an increase in market discipline. Given Equation (3.27), the expression for the average expected payoff from entering financing markets becomes

$$V^{bail-in}(d, \mu_0|\ell < 1) \; = \; V^* \; + \; (1-\mu_0)\left[(1-\delta) \; - \; (1-\ell)X^L\right] \tag{3.28}$$

[24] The liquidation cost can be microfounded as in [4]: in their model, arbitrageurs purchase liquidated assets but need to hold costly capital to buffer against potential losses. Since the identified high type never generates losses, no costly capital is required, and there is no liquidation cost.

[25] If liquidation of high type assets is costly, a lender of last resort solves the miscoordination problem and prevents any liquidations without ever having to extend any emergency loans.

As before, the term $(1-\mu_0)(1-\delta)$ captures the government subsidy to banks due to bailouts. The higher the subsidy, the higher are incentives to enter financing markets. This subsidy is now traded off against the liquidation loss, captured by the term $(1-\mu_0)(1-\ell)X^L$, which discourages entry.

3.8 Extensions

3.8.1 Equity

In this section, we expand the contract space and allow banks to issue equity: $\sigma \in \{d, D, e\}$. We treat all inside and outside equity equally: both banker and outside equity holders bear the cost of effort. As for long-term debt, equity holders cannot run. Since there is no debt overhang, the bank always invests efficiently.

Suppose bankers issue equity with the share of outside ownership given by γ^e. The investors' participation constraint is given by:

$$\delta \;\le\; \gamma^e\Big\{\mu\big[X^H + z^* - (1-\delta)\big] + (1-\mu)\big[X^L + z^* - (1-\delta)\big]\Big\} \tag{3.29}$$

Conditional on being able to achieve financing without requiring full outside ownership ($\gamma^e < 1$), it follows that the low type achieves strictly positive payoffs:

$$V^L(e,\mu) \;=\; (1-\gamma^e)\big[X^L + z^* - (1-\delta)\big] \;>\; 0 \tag{3.30}$$

In equilibrium, financing is always achieved at $\gamma^e < 1$ given our assumption on the high type's returns. It follows that equity financing again equilibrium dominates short-term financing for the low type:

$$V^L(e,\mu_0) \;>\; V^L(d,1) \;=\; 0 \tag{3.31}$$

Thus, in the unique equilibrium, banks again issue short-term debt.

Efficiency. Equity financing has the same efficiency properties as long-term debt financing: all banks invest efficiently, the government is not forced to use public funds, and only positive NPV projects enter financing markets. Importantly, debt does not feature positive incentive effects in our framework as in [18], or more recently, as in [4] with regards to bail-ins. Those models feature a classic monitoring problem and debt induces bankers to exert costly yet socially desirable effort.

3.8.2 Market completeness

In this section, we assume that the signal is contractible. Bankers are then able to issue contingent convertible debt contracts (CoCos) which are triggered by a low signal realisation. The expanded contract space is thus given by $\sigma \in \{c, d, D\}$.

Proposition 10 (Multiple equilibria for a contractible signal)**.** *If the signal is contractible, the model features multiple equilibria. Additional to short-term debt, contingent convertible debt that fully writes down all equity and converts bail-in debt to equity can also be sustained in equilibrium. The equilibrium in which banks issue CoCos is constrained efficient.*

Consider a debt contract which converts into equity upon the arrival of a low signal. Such a contract is characterised not only by its rate of interest, but also by its rate of debt-to-equity conversion. Suppose that investors hold γ^c equity shares after the contract clause has been triggered. Note that a full dilution of the banker, $\gamma^c = 1$, corresponds to a conversion rate of infinity, as all of the previous equity holder's shares have a relative value of zero. Given the rate of conversion and prior belief μ, the investors' participation constraint is given by:

$$\delta \;\leq\; \mu R_1^c \;+\; (1-\mu)\gamma^c\big[X^L + z^* - (1-\delta)\big] \tag{3.32}$$

The bankers' expected payoffs are then given by:

$$\begin{aligned} V^H(c,\mu) \;&=\; X^H \;+\; z^* \;-\; R_1^c(\mu,\gamma^c)\,\delta \;-\; (1-\delta) \\ V^L(c,\mu) \;&=\; (1-\gamma^c)\,\big[X^L \;+\; z^* \;-\; (1-\delta)\big] \end{aligned} \tag{3.33}$$

where the interest rate $R_1^c(\mu,\gamma^c)$ is determined by the binding investor participation constraint.

Compare two different CoCo contracts: one with partial dilution, $\gamma^c < 1$, and one with full dilution, $\gamma^c = 1$. The low type achieves a zero payoff for the latter CoCo contract. It follows that such a contract is equilibrium dominated by any CoCo contract with a lower rate of conversion as well as the long-term debt contract. Therefore, investors believe that any deviation to a fully dilutive CoCo contract or to a short-term debt contract must be coming from the high type. Since the infinite dilution CoCo contract does not equilibrium dominate any other contract for either type, markets are again free to form off-equilibrium beliefs, and both short-term debt and infinite dilution CoCo financing can be supported as pooling equilibria.

This section highlights an *important similarity* between short-term debt and CoCo financing: both have the potential to fully wipe out of low type bankers, although the mechanism vastly differs. With short-term debt, investors liquidate their assets, forcing the bank into default. With CoCos, all contracts are automatically converted into equity contracts for low signal realisations at a conversion rate of infinity. This closely resembles the government's bail-in policy of long-term debt but without having to honour the NCWO condition: agents cannot sue the government because a contract clause—on which both parties have mutually agreed at the financing stage—has been activated. However, we do not observe banks issuing CoCo bonds with a conversion rate of infinity.

The *key difference* between these two types of contracts—and currently unmodelled here—is that short-term debt allows individual creditors to also act on their private information. In the version of the model above, both short-term debt and CoCos respond only to public news. If short-term debt issued by the low type experiences runs more frequently, then financing using short-term debt is more 'punishing' to low type bankers. This should help to rule out CoCos as equilibrium contracts which we do not observe in reality. This suggests that not only public information, but also private information has a role in creating fragilities in financial markets. We certainly aim to address this question in future work.

3.9 Conclusion

This paper demonstrates that bail-in debt is prone to runs in the presence of asymmetric information on asset returns. Requiring banks to issue long-term bail-in debt leads to an efficiency improvement since it allows governments to avoid costly bailouts and sets the socially appropriate incentives to enter financing markets. The model provides a motive to subordinate long-term debt to short-term debt: it prevents liquidation of assets in the run up to intervention and allows for weaker long-term debt requirements to achieve bail-in effectiveness.

The paper also revisits a well-established claim that short-term debt disciplines markets through its treat of early liquidation. We show how the disciplining effect of short-term debt depends on the efficiency loss associated with liquidation and the subsequent actions of a government that lacks commitment.

Bibliography

[1] Bolton, Patrick and Oehmke, Martin. *Bank Resolution and the Structure of Global Banks*. Review of Financial Studies, 2019.

[2] Brunnermeier, Markus K and Oehmke, Martin. *The Maturity Rat Race*. Journal of Finance, 2013.

[3] Calomiris, Charles W and Kahn, Charles M. *The Role of Demandable Debt in Structuring Optimal Banking Arrangements*. American Economic Review, 1991.

[4] Clayton, Christopher and Schaab, Andreas. *Bail-Ins, Optimal Regulation, and Crisis Resolution*. Working Paper, 2020.

[5] Clayton, Christopher and Schaab, Andreas. *Multinational Banks and Financial Stability*. working paper, 2020.

[6] Dewatripont, Mathias and Tirole, Jean. *Liquidity Regulation, Bail-ins and Bailouts*. mimeo, 2018.

[7] Diamond, Douglas W. *Debt Maturity Structure and Liquidity Risk*. Quarterly Journal of Economics, 1991.

[8] Diamond, Douglas W and Dybvig, Philip H. *Bank Runs, Deposit Insurance, and Liquidity*. Journal of Political Economy, 1983.

[9] Diamond, Douglas W and He, Zhiguo. *A Theory of Debt Maturity: the Long and Short of Debt Overhang*. Journal of Finance, 2014.

[10] Diamond, Douglas W and Rajan, Raghuram G. *Liquidity Risk, Liquidity Creation, and Financial Fragility: A Theory of Banking*. Journal of Political Economy, 2001.

[11] Farhi, Emmanuel and Tirole, Jean. *Shadow Banking and the Four Pillars of Traditional Financial Intermediation*. Review of Economic Studies, 2020.

[12] Flannery, Mark J. *Asymmetric Information and Risky Debt Maturity Choice*. Journal of Finance, 1986.

[13] Gimber, Andrew R. *Bank Resolution, Bailouts and the Time Consistency Problem*. mimeo, 2012.

[14] Goldstein, Itay and Pauzner, Ady. *Demand–deposit Contracts and the Probability of Bank Runs*. Journal of Finance, 2005.

[15] Gorton, Gary and Pennacchi, George. *Financial Intermediaries and Liquidity Creation*. Journal of Finance, 1990.

[16] He, Zhiguo and Milbradt, Konstantin. *Dynamic Debt Maturity*. Review of Financial Studies, 2016.

[17] He, Zhiguo and Xiong, Wei. *Dynamic Debt Runs*. Review of Financial Studies, 2012.

[18] Innes, Robert D. *Limited Liability and Incentive Contracting with Ex-ante Action Choices*. Journal of Economic Theory, 1990.

[19] Keister, Todd and Mitkov, Yuliyan. *Allocating Losses: Bail-ins, Bailouts and Bank Regulation*. mimeo, 2021.

[20] Mendicino, Caterina and Nikolov, Kalin and Suarez, Javier. *Equity versus Bail-in debt in Banking: An Agency Perspective*. CEPR Discussion Paper No. DP12104, 2017.

[21] Nachman, David C and Noe, Thomas H. *Optimal Design of Securities Under Asymmetric Information*. Review of Financial Studies, 1994.

[22] Pandolfi, Lorenzo. *Bail-in and Bailout: Friends or Foes?*. Management Science, forthcoming, 2021.

[23] Tanaka, Misa and Vourdas, John. *Equity, Debt and Moral Hazard: the Optimal Structure of Banks' Loss Absorbing Capacity*. Bank of England Working Paper No. 745, 2018.

[24] Tirole, Jean. *Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning*. American Economic Review, 2012.

[25] Walther, Ansgar and White, Lucy. *Rules versus Discretion in Bank Resolution*. Review of Financial Studies, 2020.