Publish, Perish, or Pivot:
Project Completion and Alignment Under an Uncertain Bayesian Signalling Framework

Approximate Word Count: 18,720

Victor Ehrnrooth *

University of Oxford: Department of Economics

Last updated: May 6, 2025

Abstract

We model settings where projects are terminated prematurely or pursued unwisely to the detriment of all stakeholders. A decision maker, with sole decision making power, must decide whether to continue a project or abort it while repeatedly interacting with an agent of uncertain type that is able to send the decision maker fuzzy signals that simultaneously directly affect project success probabilities. We characterise an equilibrium showing that with reasonable parameter choices we observe some agent types seemingly behaving against their own best interest, attempting to maximise the failure probability of a project they view favourably; equally, one can characterise an equilibrium where agents wish to minimise the failure probability of a project they would prefer to fail. Further, there is a positive probability of decision makers terminating a project that both they and the agent would have wanted to complete, even when the true probability of success is high. We then discuss key mechanisms driving the model and implications thereof, demonstrating that many observed real-world counter-intuitive behaviours can be sustained in this model of ambiguous two-agent interaction. Finally, we discuss extensions of the model to demonstrate that various adjacent scenarios can be accommodated within the framework of this model.

*I would like to thank Simon Cowan, Ian Jewitt, my MPhil colleagues, and of course, my wife Maria. The direction, technical expertise, and patience they all provided have been invaluable.

Contents

1 Introduction and Related Literature

1.1 Introduction

We often observe public infrastructure projects being seemingly inefficiently cancelled or initiated, such as HS2 being partially cancelled after a large investment had been made, or the MTA in NYC failing to move forward with productive projects. In some cases, it appears that a majority of stakeholders approve of a project, but that outcomes do not reflect this. Decision makers appear to worry that their pet projects, projects to which they attach their names and reputation, could fail regardless of majority approval. In some cases, it seems decision makers may be catering to a particular, smaller subset of the electorate who hold more direct power over the decision maker. One might, for instance, speculate that Rishi Sunak was not catering to all Britons in cancelling HS2, but rather to those Britons most likely to sway the election in his favour.

This type of dynamic has been explored largely from an empirical or macroeconomic perspective, notably in the specific domain of public infrastructure. We approach the problem more foundationally, building up from a simple microeconomic model. In particular, we model the interaction between decision makers and electors as a simple two agent game between a decision maker and some pivotal agent. We then simulate the uncertainty over the agents' preferences by drawing the preferences of the agent from some distribution, without revealing the draw to the decision maker, who subsequently weighs whether they should proceed with or cancel a project. This model captures the dynamics of interest, where some collective influence exists but it is made fuzzy by the preferences of specific groups, which the decision maker must not only identify, but cater to.

More broadly, the model applies to any scenario with one decision maker and several agents with limited influence on outcomes. Key to the model is that some subset of agents have greater influence than others, represented by the single agent with which the decision maker interacts. In public infrastructure, we have some decision maker and groups of voters which do not have the same power over the decision maker; think of swing states in American elections with varying levels of influence. In a corporate setting, a CEO might require board-approval to advance their goals; while each board-member has some say, there may be a particularly vocal board member with strong influence over the rest of the board. The editorial board of a publisher may have several editors, and while none of them alone determines outcomes, some editors are more vocal than others, aiming to sway the rest of the board to their view-point. All of these examples are captured in the model, allowing it to flexibly describe many asymmetric, multi-agent scenarios.

The discourse surrounding public infrastructure projects seems to justify a foundational microeconomic approach. Often, we do not wonder if a project is good or bad, but whether it will be completed, if it is worth the risk, or if we would be better off cutting our losses and cancelling it. Cases of projects mired in problems eventually yielding positive results (eg: nuclear plant development) do not seem to alter the conversation, nor have contracting processes been substantially altered to avoid starting lousy projects in the first place. Were these misalignments a simple case of bad systems or perverse incentives, they should vanish with better structures in place. Yet in Finland, for instance, where structures are reasonably good, we still struggle constantly with these same problems. This helps affirm that perhaps the reasons for these misalignments are largely an underlying information structure liable to cause potentially unwanted outcomes.

1.2 Related Literature

Several empirical studies analysing project termination and completion exist, with varying conclusions. Beetsma & van der Ploeg (2007) observe that outgoing parties raise spending to "bind the hands of their successors" and thus affect policy in their favour. Conversely, Fiva & Natvik (2013) explore Norwegian public investment to conclude that high re-election odds stimulate investment, due to potential complementaries between spending today and spending in the future. They recognise Beetsma & van der Ploeg's findings, but argue that the direction of the effect varies due to different motivations dominating in different cases. These two are especially relevant, but there is a large body of empirical work tackling this issue and connected problems. These empirical studies focus on macroeconomic and political economy factors as well as ideological motivations, but they validate that the posited effect exists and is worth exploring.

To explain tensions observed empirically using a theoretical microeconomic foundation, we have independently produced a model that bears similarities to a many sub-fields. Foundationally, this thesis follows in the footsteps of Crawford & Sobel (1982) in exploring the implications of fuzzy signalling, as well as developing ideas of jointly governed state transitions as discussed by, for instance, Shapley (1953). Some model mechanisms resemble "garbling" as in Blackwell (1951, 1953), reducing the informativeness of signals relative to other signalling structures (although the connection is high-level, as choosing between information structures does not feature in this thesis).

The literature has developed since, and there are similarities in this model to the literature of Bayesian Persuasion (BP) introduced by Kamenica & Gentzkow (2011) and, to a lesser extent, Optimal Delegation (OD) introduced by Alonso (2008). This thesis does not neatly fit into either field, yet the notion of manipulating the game environment to affect outcomes via Bayesian updating as in BP is central to our model. The set-up introduced in Kamenica & Gentzkow (2016) as well as Dworzak & Martini (2019) is closer to the one used in our model, where optimal behaviour depends only on expected state rather than the state; results follow more easily in such a set-up. Further, as in OD, the dynamic of a single agent making core decisions affecting all agents in the strategic environment plays a key role (unlike in OD, however, the delegation of decision-making power is not voluntary). BP is a closer match for our methods, but Bayes Correlated Equilibria (Bergemann & Morris, 2016) offer an alternative for many-player situations.

Alonso & Câmara (2016) (AC) provide an especially close example, proposing a BP model where politicians persuade voters to vote for their platform. AC conclude that, when no conflict of interest is present, voters benefit from politicians designing experiments to provide information. Their mechanism is effectively the inverse of our model, where instead voters persuade politicians using their influence over politician success as leverage. Our model builds on their results by allowing for even aligned politicians and voters to fail to coordinate, an important addition.

Gill & Sgroi (2008) develop a model where a sequence of tests, chosen by an agent, reveal information about said agent to orient decision making. Their model foreshadows some of the ideas used in BP, but especially relevant for this thesis is the repeated structure of their model. A similar idea forms the core of strategic interaction in this thesis, where information revelation based on the choices of an agent reveal fuzzy information about their type over time.

D'Antoni et al. (2024, working paper) propose a BP model where agents decide whether or not to act, cast as a one-armed bandit problem. Their method and application differ from ours, but they also consider the problem of whether or not to act. We keep their work in mind as an alternative approach to the problem.

Finally, this thesis leans on foundational principles, such as Sequential Equilibria (Kreps & Wilson, 1982), First-Order Stochastic Dominance (Hadar & Russell, 1969), and others. This is however mechanic rather than a consequence of deep connections between our models.

2 Preliminaries and Model Walk-Through

Here, we specify the model and provide a thorough graphical walk-through. A glossary is included at the end of the thesis with key terms that recur frequently. The first occurrence of terms included in the glossary will be highlighted and will re-direct to the glossary, making clear what terms are included therein.

2.1 The Model

As a motivating model, we consider a book Author who hopes to publish a book, but must get it through two rounds of review first. The Author can choose to quit trying to publish the book and instead publish a collection of short stories with certainty, which they find less rewarding. If they attempt to publish the book, but are rejected by the review board, their dejection offsets any benefit gained from publishing their short stories.

The review board has many members, each liking the book in different measure. Each reviewer thus gets some payoff from seeing the book published, depending on their own preference for the book. If the book is not published, the Author's short stories are published instead, which all reviewers view equally favourably. One Reviewer on the board is particularly influential in the reviews, swaying other reviewers in one direction or another, but never single-handedly determining outcomes.

If the book passes the first stage of reviews, the Author is given a report on what the reviewers thought of the book; it might be damning, suggesting the book barely passed review, it might be glowing, suggesting the book was very well received, or it can be anything in between. This report helps the Author know if it is worth continuing to try to get the book published, or simply to cut their losses.

Once a book has passed the first review stage, all reviewers prefer publishing the short stories a little bit less due to the sunk cost of the Author having worked on the book for a prolonged period. The "sunk cost" more reflects the energy already expended on the book than it does a decreasing preference for the short stories; the short stories could have been published from day one, and so expending energy on the book only to abandon it makes the overall value of publishing the short stories a little bit lower. In the context of infrastructure projects, if a mayor is building a bridge, but the city council cancels it when it is half built, then only half the funds earmarked for the bridge remain available to distribute to other ventures.

The model aims to replicate a scenario where a single agent has the power to choose to advance a project or not (the "Author"), but is beholden to a group or electorate that is able to influence the outcomes of the project. Importantly, the electorate should be of such a structure that one segment has greater influence than others (the "Reviewer"); think about a company board with a particularly vocal and thus influential board-member, or American elections where specific swing-states have higher influence, as mentioned in the introduction.

The key innovation of the model is to avoid directly modelling election behaviour and dynamics by instead treating the Author and the Reviewer as the only agents of interest. To accomplish this, we make the Reviewer's power to influence outcomes limited and random. Despite the Reviewer having more influence than their peers, there is still a larger electorate collectively making decisions in the background, and so their influence remains limited. Since we do not model "voter" behaviour, such an abstraction is most appropriate when we are more interested in the behaviour of the decision maker.

Importantly, the Author must be able to respond to Reviewer behaviour. Thus, the Author must get some indication of what type of Reviewer they are facing. In elections, polling offers an indicator that tells candidates if they are doing the right thing. A CEO communicating with their board may receive feedback or attend meetings, which partially reveal the board's sentiment.

2.1.1 Preliminaries

The basic structure of the model and situations to which it applies follows the below logic:

1. A core decision maker must decide whether to pursue an objective or abandon it.
2. The decision maker faces an uncertain evaluation if they choose to pursue their objective.
3. The uncertain evaluation is strongly influenced by some pivotal evaluator, but it is unknown to the decision maker who this evaluator is.
4. While the evaluator can significantly sway outcomes, they do not determine them entirely.
5. Before the decision maker decides whether to abandon or pursue their objective at a second stage, they are made aware of the conclusions of the evaluation. This reveals to them some information about the type of evaluator they are likely facing.

2.1.2 Model Definition

Using the language of the publication story, the initial set-up for the model is as follows: There are 2 agents, a Reviewer and an Author, and two periods, $t = 1$ and $t = 2$. At $t = 1$, the Author can choose whether to attempt to continue to $t = 1$ or to quit. At $t = 2$, the Author can choose whether to attempt to get their book published, or to quit.

The probability of an Author's book successfully passing from one period to the next is given by $\alpha_t \sim F(x|m_t)$, where $F(\cdot)$ is some distribution between 0 and 1 from which each α_t is drawn, and m_t is a parameter that governs the shape of the distribution which the Reviewer has the power to choose. In essence, the limited influence of the Reviewer originates from being able to affect the distribution, not the outcome. Importantly, this means that success, even accounting for the Reviewer's behaviour, is non-deterministic, since $\alpha_t \in [0, 1]$.

At each period, the Reviewer (at the same time as the Author chooses to continue or not) selects $m_t \in M$, where M is a restricted set of choices for m_t. This restriction ensures that the Reviewer has limited decision making power and cannot generate a deterministic distribution.

The Author gets payoff a from quitting, b from publishing, and 0 from being rejected after continuing. a and b are known to all. The Author gets 0 not a from failure to reflect failure being costly following the notion mentioned above of "dejection from failure" offsetting publication benefits. Failure payoff need not be zero, but changing it does not alter results much.

The Reviewer gets payoff s if the Author quits or is rejected at $t = 1$, δs if the Author quits or is rejected at $t = 2$, and r if the paper is published, where $\delta \in (0, 1)$ and $r \sim U[0, R]$. δ represents the reduced outside-option value to reviewers as more time is invested (sunk cost). The Reviewer knows their own, realised, value of r, and the Author knows s, δ, and the distribution of r, but not what type of Reviewer they are facing (ie: their r value).

Crucially, at $t = 2$, before choosing whether to continue or not, the Author observes the realised value of α_1. Consequently, the Author is able to update their beliefs over what type of Reviewer they are facing. The question of how one is able to observe a probability is addressed in detail later, but the broad idea is that in reality a probability is not directly observed, but is inferred based on some observable data. For example, polls provide information about one's likelihood of winning, but do not perfectly reveal the outcome of a subsequent election.

2.1.3 Extensive Form Walk-Through

Full Extensive Form

To make the above clearer, we draw a complete game tree, which we will then walk through step by step. While the graphical depiction below is essentially an accurate representation of the game, it does not capture 100% of the relevant dynamics. Still, it well illustrates broadly what is happening, and wherever details are not shown graphically, they are developed in the specific section discussing that portion of the tree.

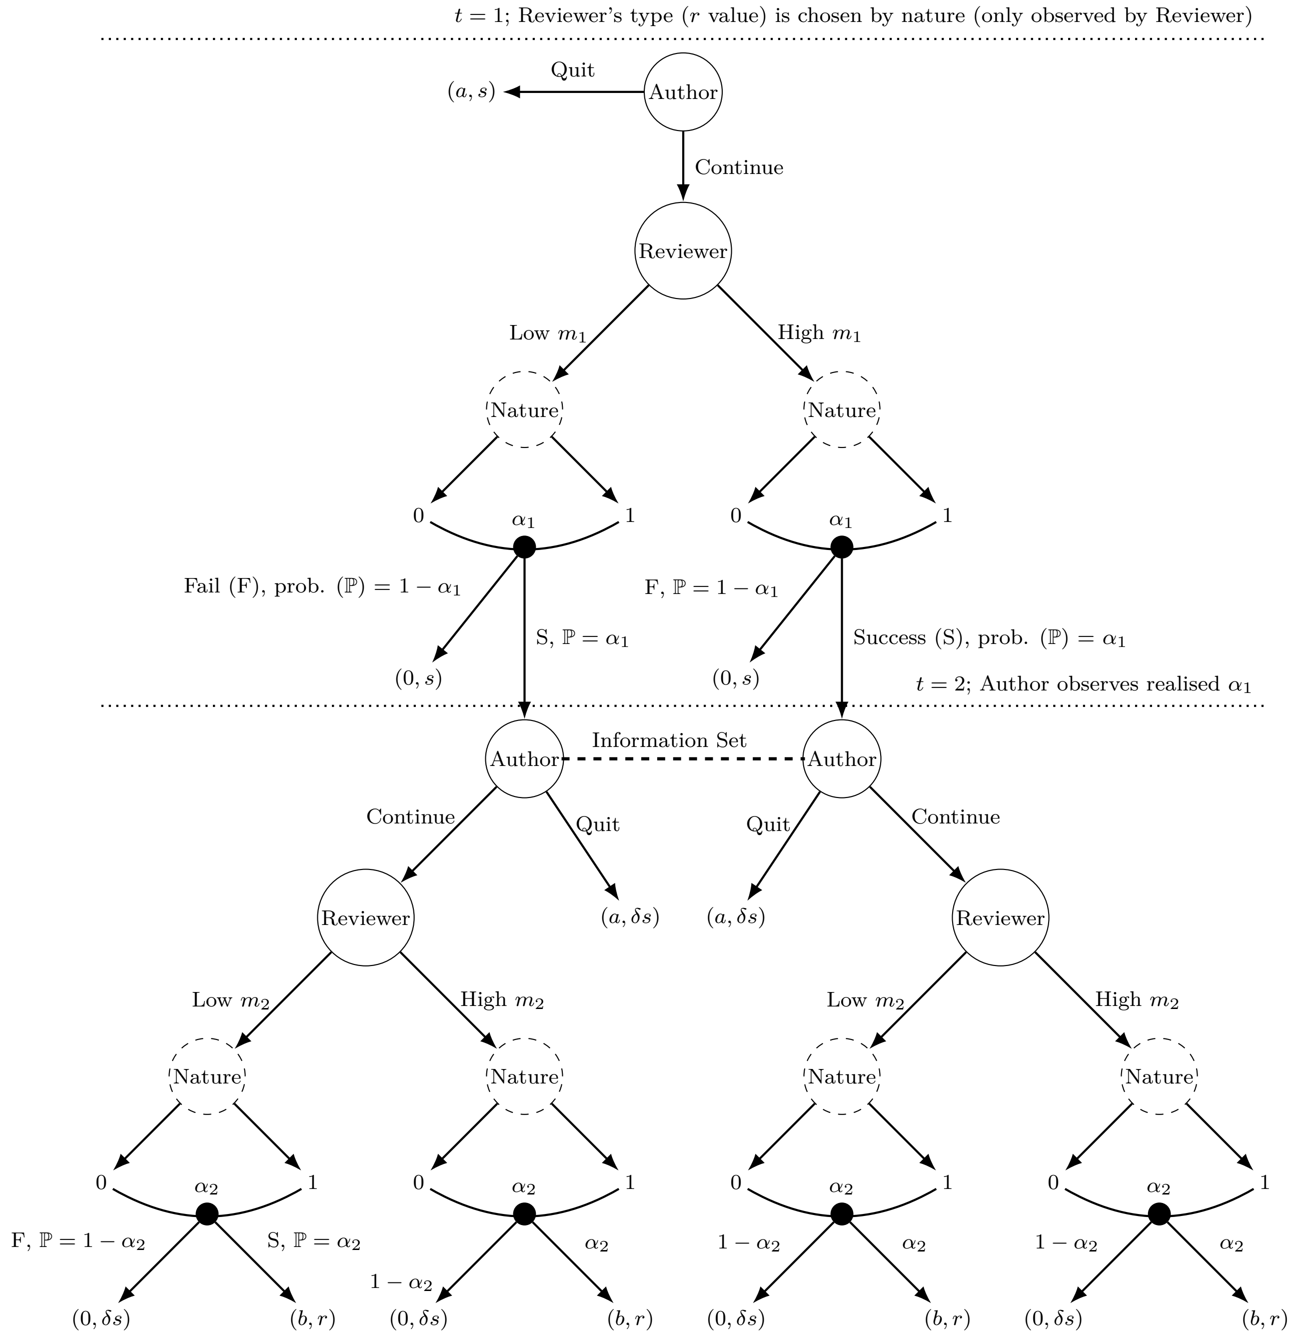

Figure 1: Full Extensive Form.

Walk-Through

Pre-Action Events

The dotted line details an event that occurs before action is taken. Before the Author plays, the type of Reviewer they face is chosen by nature according to $r \sim U[0, R]$. The Author does not observe this.

Author's First Move

The Author's first move is pictured below:

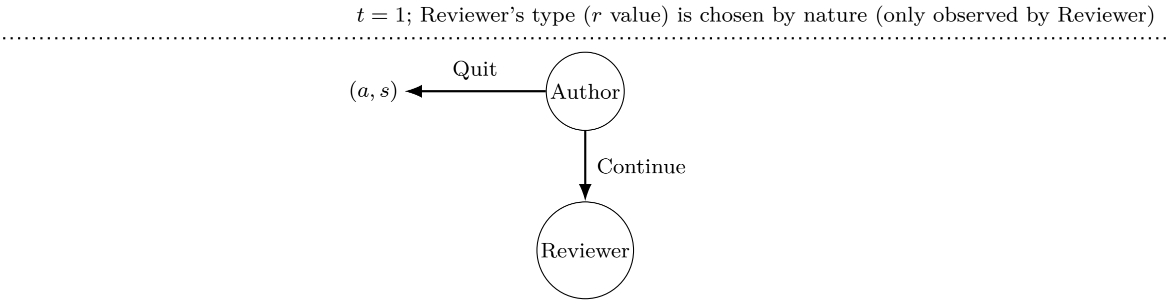

Figure 2: Author's First Move.

The Author can choose to end the game immediately, and receive a guaranteed payoff of a. Otherwise, they can choose to continue, and pass the initiative to the Reviewer. Which action they will choose will depend on their ex-ante probability of reaching an end-note with payoff b, as well as how much higher b is than a. If $b < a$, continuing is strictly worse than quitting, as there is a non-zero chance of receiving payoff 0 after continuing. Even if $b > a$, the probability of reaching end-nodes with payoff b is never equal to 1, so the Author must weigh the odds of success at this stage in making their decision.

Reviewer's First Move

Now, we add in the Reviewer's first move:

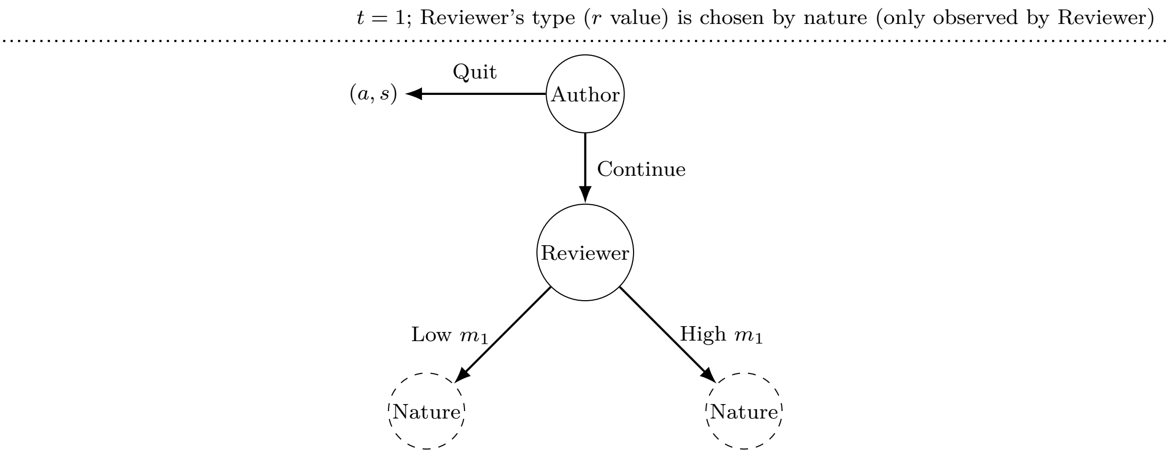

Figure 3: Reviewer's First Move.

Once the Author has chosen to continue, the Reviewer must choose what level of m_1 to set. In this example, we consider a scenario where $m_1 \in \{\text{low, high}\}$. The Reviewer's choice of m_1 affects outcomes via α_1, since $\alpha_1 \sim F(x|m_1)$; the Reviewer can effectively make certain values of α_1 more or less likely via their choice of m_1. However, since the Reviewer does not deterministically select the value of α_1, Nature plays after the Reviewer, since based on the choice of m_1, Nature picks a random α_1 governed by the distribution $F(x|m_1)$.

Nature Chooses α_1

After the Reviewer's move, Nature plays as follows:

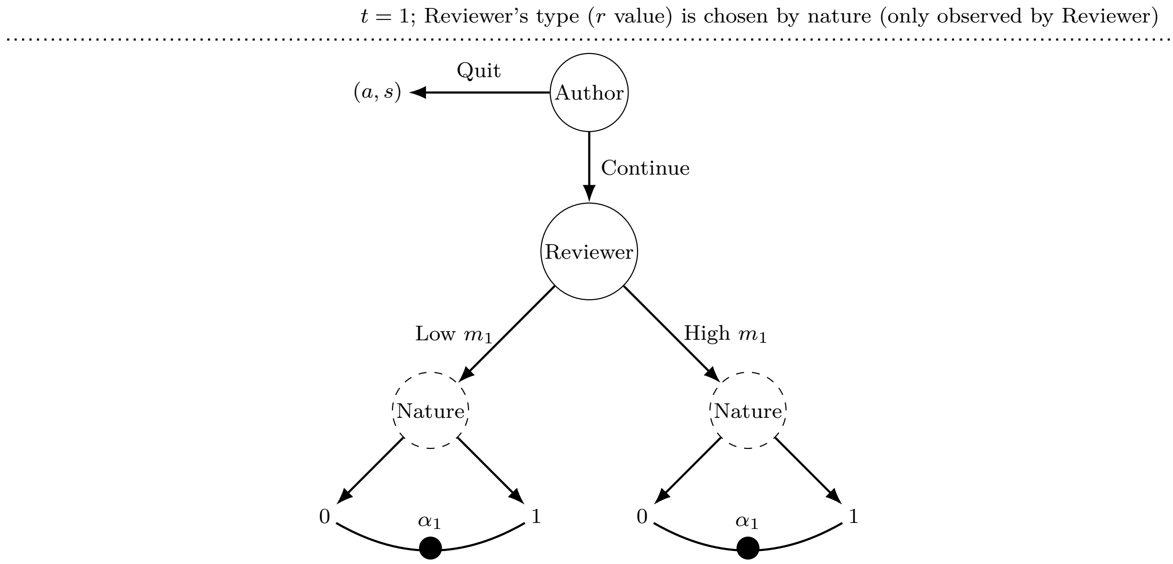

Figure 4: Nature's First Move.

Depending on the Reviewer's choice of m_1, Nature will select α_1 from either $F(x|m_1 = \text{low})$ or $F(x|m_1 = \text{high})$. In either case, $\alpha_1 \in [0, 1]$, hence the arcs used to represent Nature's continuous choice of α_1. The specific point here is merely illustrative, any point along the arc between 0 and 1 is a possible choice by Nature.

Transition to $t = 2$

After Nature's move, we (possibly) transition to $t = 2$:

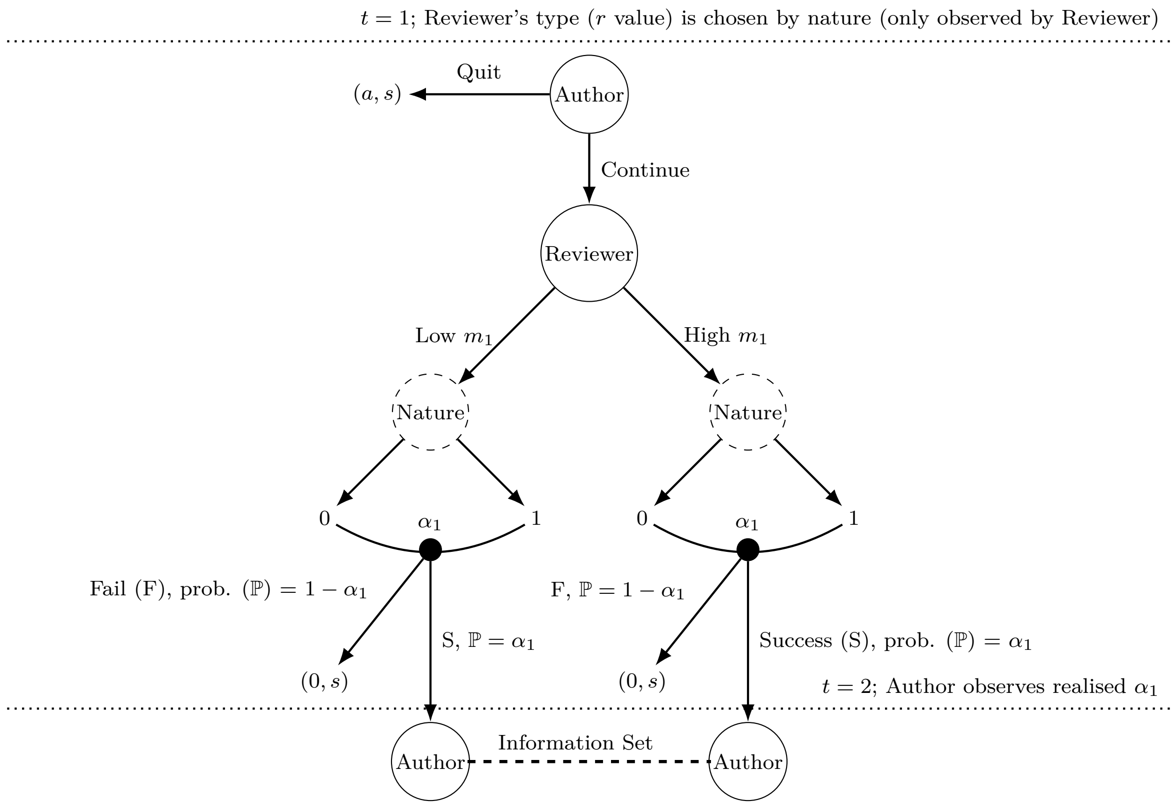

Figure 5: Transition to Second Period.

Once α_1 has been determined, whether the game continues to $t = 2$ depends on chance. With probability α_1, the game continues; with probability $1 - \alpha_1$, it ends immediately, giving payoffs $(0, s)$. This is true in both the low and high m_1 cases, but naturally there is nothing dictating that α_1 on the left or right side of the tree should be the same; as has been stated, α_1 is drawn from a probability distribution, which depends on m_1.

If the game continues to $t = 2$, the Author observes the realised value of α_1, from which they are able to update their beliefs over what type of Reviewer they face. The Author *does not* observe the actual value of r. This is why we include an "Information Set"; the Author cannot actually tell apart the low and high m_1 branches. What this information set is *not* claiming is that observed α_1 is always going to be the same; observed α_1 is ultimately random, as it is drawn from $F(x)$. The use of "Information Set" is imperfect, as there is information that the Author can use to try to determine which side of the tree they are on. Nevertheless, the Author cannot discern between the two sides of the tree with certainty.

Author's Second Move

From now on, we suppress everything happening before $t = 2$ to avoid repeating the entire game-tree. If the game continues to $t = 2$, the Author is given a second opportunity to choose to quit or continue:

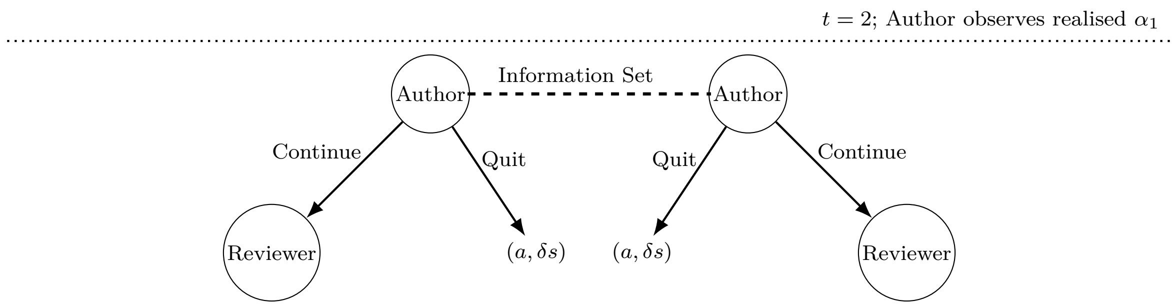

Figure 6: Author's Second Move.

The Author effectively faces the same decision as at $t = 1$, however they now have the added information gained from observing α_1 to help determine if they should quit or continue.

Now, if the Author chooses to quit, payoffs are $(a, \delta s)$, rather than (a, s), reflecting the "sunk-cost" losses to the Reviewer. If the Author continues, the Reviewer once again gets to move.

Reviewer's Second Move

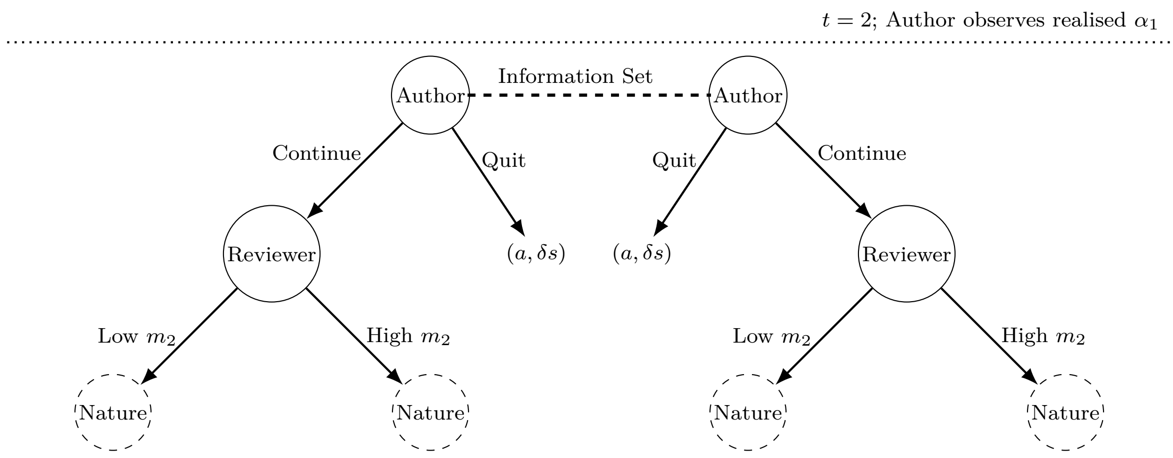

Figure 7: Reviewer's Second Move.

As at $t = 1$, the Reviewer chooses some value for m_2. Note that m_2 need not correlate with m_1.

Nature Chooses α_2

After the Reviewer's choice, Nature makes a similar "choice" as at $t = 1$:

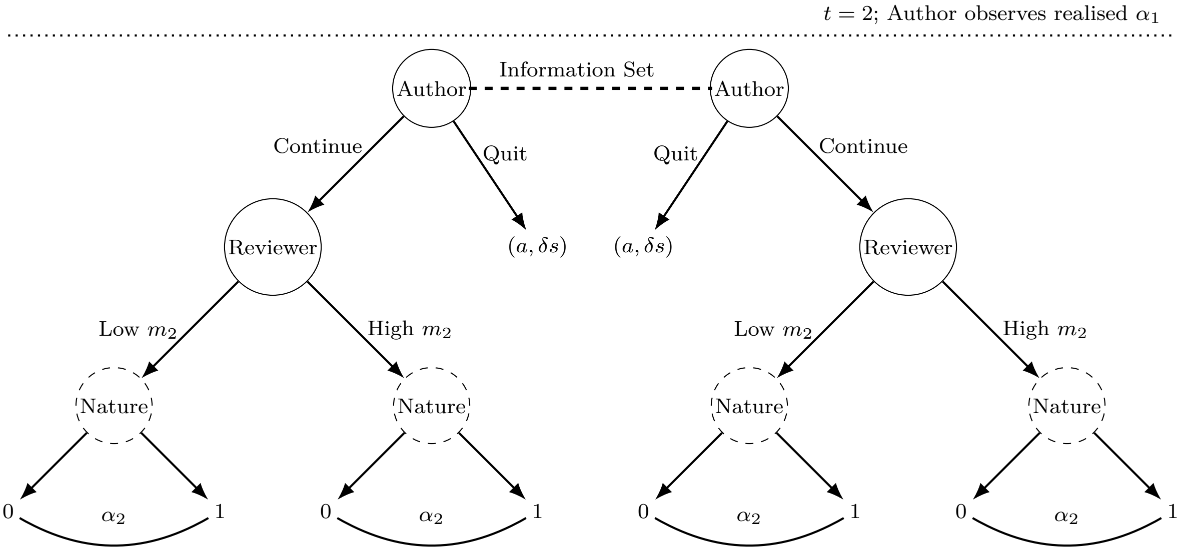

Figure 8: Nature's Second Move.

Nature chooses a value for α_2 between 0 and 1, in accordance with the distribution $F(x|m_2)$. Just as m_2 did not need to depend on m_1, $F(x|m_2)$ does not depend on $F(x|m_1)$; that is, the choice of α_2 is independent of the choice of α_1. That said, since the Reviewer has not changed between periods, it is reasonable to suggest that α_1 and α_2 will correlate; after all, a Reviewer with extremely low r is likely to pick low m_t in both periods!

Game Termination

Once Nature has chosen α_2, the game proceeds much like it did at $t = 2$, but now to its terminal point:

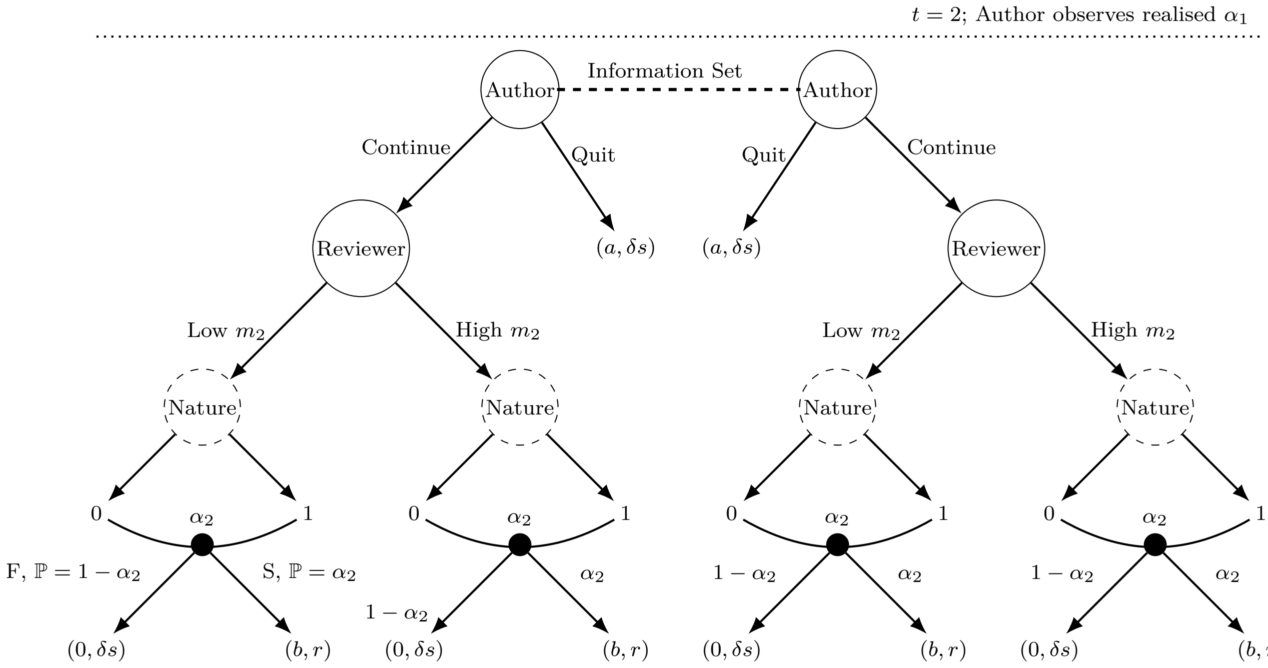

Figure 9: Game Termination.

At each node, with probability α_2, the project is a success at $t = 2$, yielding payoffs (b, r). With probability $1 - \alpha_2$, the project is instead a failure, yielding $(0, \delta s)$, where again δs rather than s

appears to capture the sunk cost dynamic that the Reviewer faces.

3 Model Solution

In this section, we solve the model for a particular equilibrium of interest. We then discuss some implications of the particular equilibrium we compute.

3.1 Set-Up

The set-up is as follows:

- $r \sim U[0, R]$, and $a,\ b,\ s,\ \delta$ are all parameters.
- $\alpha_t \sim F(x|m_t)$, where $F(\cdot)$ is the power distribution, defined as $F(x|m_t) = x^{m_t}$ for $x \in [0, 1]$, $F(x) = 0\ \forall x < 0$ and $F(x) = 1\ \forall x > 1$.
- $m_t \in \{\frac{1}{3}, 3\}$; m_t is effectively the "review" chosen by the Reviewer, where high m_t is a favourable review.
- This yields the following CDF:

$$F(x|m_t) = \begin{cases} 0 & x < 0 \\ x^{m_t} & x \in [0,1] \\ 1 & x > 1 \end{cases}$$

  We note that at each t, $m_t = 3$ First-Order Stochastically Dominates (FOSD) (Hadar & Russell, 1969) $m_t = \frac{1}{3}$ (since $F(x|3) \leqslant F(x|\frac{1}{3})$). This forms the basis of the solution later on, as any Reviewer seeking to maximize success outcomes will pick FOSD m_t.

- It also yields the following PDF:

$$f(x|m_t) = \begin{cases} 0 & x \notin [0,1] \\ m_t x^{m_t - 1} & x \in [0,1] \end{cases}$$

- Finally, the above distributions also yield $E[x|m_t] = \frac{m_t}{m_t+1}$

Solution Concept

The solution concept to be used throughout will be Sequential Equilibrium (Kreps & Wilson, 1982). To evaluate equilibria, we will walk through every information set in the game and ensure rational play given that information set.

Target Equilibrium

Here we look for one specific equilibrium which can be considered the most interesting one. However, several equilibria exist, given the correct choice of parameters. We consider an equilibrium as follows: the Author continues at $t = 1$, and they continue for some but not all realised α_1 at $t = 2$. Some Reviewers choose high m_1 and high m_2, others choose low m_1 but high m_2, and some choose low m_1 and low m_2. This equilibrium is interesting as it places positive probability on a large set of outcomes, allowing project completion while forcing the Author to take Reviewer actions into account (it is not the case that "always continue" dominates regardless of Reviewer type, for instance).

We first consider an intuitive motivation of what dynamics are needed to sustain such an equilibrium. the Author must ex-ante wish to engage in the project, meaning that the expected utility of continuing at $t = 1$ is high. This requires high success payoff and/or success probability. In addition, the Reviewer must be able to dissuade the Author at $t = 2$; if the Author knows with certainty that they are facing an unfavourable Reviewer at $t = 2$, they must prefer to quit rather than continue. What we will find is that this will require giving the Reviewer sufficient power to influence outcomes both in favour and against the Author and/or making failure payoffs sufficiently low for the Author. Finally, Reviewer payoffs δs, s, and r must differ sufficiently and r must take on sufficiently wide range of values to ensure that optimal Reviewer play can follow the desired path described.

3.2 Target Equilibrium Solution

3.2.1 Preliminary Conjectures and Definitions

To avoid confusion going forward, we pre-emptively make some definitions and conjectures at this early stage. These definitions and conjectures will, during the course of the proofs below, prove necessary and valid. They will all also be re-stated for clarity in proofs as needed.

Conjectures:

Conjecture 3.1 (Author Cut-Off Behaviour)**.** *Given sufficient assumptions (1-4) shown below, Author will play a cut-off strategy at* $t = 2$ *such that for sufficiently large observed values of* α_1*, they continue at* $t = 2$.

Conjecture 3.2 (Reviewer Cut-Off Behaviour)**.** *Reviewers play cut-off strategies at both periods, where for sufficiently high values of* r *(defined differently at each period), they choose* $m_t = 3$*, otherwise choosing* $m_t = \frac{1}{3}$.

Definitions:

Definition 3.3 (Reviewer Cut-Off)**.** γ is defined as the Author's belief of the Reviewer's cut-off such that $\forall r \geqslant \delta s + \gamma$, they play $m_1 = 3$, and $\forall r < \delta s + \gamma$, they play $m_1 = \frac{1}{3}$.

Definition 3.4 (Author Cut-Off)**.** q is defined as the Reviewer's belief of the Author's cut-off such that $\forall \alpha_1 \geqslant q$, they continue at $t = 2$, and $\forall \alpha_1 < q$, they quit at $t = 2$.

3.2.2 Sufficient Assumptions

For the target equilibrium, some assumptions prove to be sufficient. These assumptions arise naturally in the proofs below, but are highlighted here for immediate clarity. Note that other equilibria exist when these conditions fail, but they are needed to reach the target equilibrium.

1. $3b - 4a > 0$; that is, success payoff must be high enough to justify ever continuing.

2. Require that

$$\frac{\delta s(4a - b) - \gamma(3b - 4a)}{9(R - \delta s - \gamma)(3b - 4a)} \geqslant 0$$

   This condition ensures that there is a non-zero probability of the Author quitting at $t = 2$. If the above does not hold, the Author wishes to continue no matter what.

3. $\gamma \in [0, R - \delta s)$. First, $\gamma \geqslant 0$ is required to ensure consistency of beliefs, namely that the Author does not believe Reviewers to be behaving more generously than is consistent with their optimal play. Then, $\gamma < R - \delta s$ is required to ensure that the Author believes that at least some Reviewers choose to behave generously, that is, they believe there to be an interior cut-off point for Reviewers. Effectively, without this condition, we find that either the Author has inconsistently optimistic beliefs, believing Reviewers to be willing to sacrifice payoff to benefit the Author, or the Author believes that no Reviewer, regardless of their payoffs, will ever wish to give the Author a positive review.

4. $q \in (0, 1)$. This condition amounts to requiring Reviewer beliefs to be consistent with the Author playing a cut-off strategy, rather than always quitting (implied by $q = 1$) or always continuing (implied by $q = 0$) at $t = 2$. If this fails to hold, the Author will believe their own behaviour to have no influence on the Reviewer, making the conjectured equilibrium impossible.

3.2.3 Reviewer Behaviour at $t = 2$

Proposition 3.5. *$\forall r \geqslant \delta s$, Reviewer with completion-payoff r sets $m_2 = 3$. For all other values of r, Reviewer sets $m_2 = \frac{1}{3}$.*

Proof. The Reviewer, given continue by the Author, faces

$$E[\text{Payoff}] = E[\alpha_2|m_2]r + E[1-\alpha_2|m_2]\delta s = \left(\frac{m_2}{m_2+1}\right)r + \left(\frac{1}{m_2+1}\right)\delta s = \frac{m_2 r + \delta s}{m_2+1}$$

$$\Rightarrow \frac{dE[\text{Payoff}]}{dm_2} = \frac{(m_2+1)r - (m_2 r + \delta s)}{(m_2+1)^2} = \frac{r - \delta s}{(m_2+1)^2}$$

Thus, if $r - \delta s \geqslant 0$, the Reviewer will set $m_2 = 3$, whereas if $r - \delta s < 0$, they'll set $m_2 = \frac{1}{3}$, since in the former case the derivative is strictly positive, inciting the Reviewer to select the highest possible m_2 (reverse is true in the latter case).

□

We note that effectively, any Reviewer with $r \geqslant \delta s$ wants to maximise success odds, and so has direct incentive to pick FOSD m_2.

3.2.4 Author Behaviour at $t = 2$

Proposition 3.6. *For any equilibrium where the Author continues at $t = 1$, inducing Reviewer play as described in proposition 3.5, given observed realised success-probability α_1, the Author plays a cut-off strategy where:*

$$\forall \alpha_1 \geqslant \left(\frac{\delta s(4a-b) - \gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{3}{8}} \tag{3.1}$$

the Author continues. For all other values of α_1, the Author quits.

Proof. First, we note that $\forall r < \delta s$, clearly $m_1 = m_2 = \frac{1}{3}$ is the only rational choice for Reviewers, since for such Reviewers they will always prefer minimizing success odds for the Author in every period; that is, Reviewers seek to *avoid* FOSD choices of m_t in all periods.

Given this, we surmise that there must be some cut-off type of Reviewer such that $\forall r \geqslant \delta s + \gamma$, such Reviewers choose $m_1 = m_2 = 3$, and $\forall r$ *s.t.* $\delta s < r < \delta s + \gamma$, such Reviewers choose

$m_2 = 3, m_1 = \frac{1}{3}$, where then $\gamma > 0$ is the cut-off value.

We assume such a cut-off strategy is played, and given such a strategy, we compute Author behaviour. Note that we will check that this strategy is optimal later.

To compute the Author's optimal strategy, we will want to find $E[\alpha_2|\alpha_1]$. For this, we will need:

$$\mathbb{P}(r < \delta s|\alpha_1) = \frac{f(\alpha_1|r < \delta s)\mathbb{P}(r < \delta s)}{f(\alpha_1)};$$
$$f(\alpha_1) = f(\alpha_1|r < \delta s)\mathbb{P}(r < \delta s) + f(\alpha_1|r \geqslant \delta s + \gamma)\mathbb{P}(r \geqslant \delta s + \gamma)$$
$$+f(\alpha_1|r \in [\delta s, \delta s + \gamma])\mathbb{P}(r \in [\delta s, \delta s + \gamma])$$
$$= \left(\frac{1}{3}\right)\left(\alpha_1^{-\frac{2}{3}}\right)\left(\frac{\delta s}{R}\right) + 3\left(\alpha_1^2\right)\left(1 - \frac{\delta s + \gamma}{R}\right) + \left(\frac{1}{3}\right)\left(\alpha_1^{-\frac{2}{3}}\right)\left(\frac{\gamma}{R}\right) = \frac{1}{R}\left[\frac{\delta s + \gamma}{3\alpha_1^{\frac{2}{3}}} + 3\alpha_1^2(R - \delta s - \gamma)\right]$$
$$\Rightarrow \mathbb{P}(r < \delta s|\alpha_1) = \frac{\frac{1}{3}(\alpha_1^{-\frac{2}{3}})\frac{\delta s}{R}}{\frac{1}{R}\left[\frac{\delta s+\gamma}{3\alpha_1^{\frac{2}{3}}} + 3\alpha_1^2(R - \delta s - \gamma)\right]} = \frac{\delta s}{\delta s + 9\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + \gamma}$$

With this, can find the expected value we are interested in:

$$E[\alpha_2|\alpha_1] = E[\alpha_2|r < \delta s]\mathbb{P}(d < \delta s|\alpha_1) + E[\alpha_2|r \geqslant \delta s]\mathbb{P}(d \geqslant \delta s|\alpha_1)$$
$$= \frac{1}{4}\left(\frac{\delta s}{\delta s + 9\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + \gamma}\right) + \frac{3}{4}\left(1 - \frac{\delta s}{\delta s + 9\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + \gamma}\right)$$
$$= \frac{\delta s + 27\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + 3\gamma}{4[\delta s + 9\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + \gamma]}$$

Note: For the above to be valid, it must hold that

$$4[\delta s + 9\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + \gamma] \neq 0$$

Given sufficient assumptions (1-4), the above follows:

$$\gamma \in [0, R - \delta s) \Rightarrow R - \delta s - \gamma > 0 \Rightarrow 4[\delta s + 9\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + \gamma] > 0$$

Then, the Author will prefer to continue at $t = 2$ than to quit if $E[\alpha_2|\alpha_1]b \geqslant a$, which allows us to solve for a cut-off such that $\forall \alpha_1 \geqslant$ cut-off, the Author continues:

$$b[\delta s + 27\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + 3\gamma] \geqslant 4a[\delta s + 9\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma) + \gamma]$$
$$\Rightarrow 9\alpha_1^{\frac{8}{3}}(R - \delta s - \gamma)(3b - 4a) \geqslant \delta s(4a - b) - \gamma(3b - 4a)$$

If $3b - 4a \leqslant 0$, then this inequality never holds, and so we require $3b - 4a > 0$ for the Author to ever want to continue (hence sufficient condition 1). Then:

$$\alpha_1^{\frac{8}{3}} \geqslant \frac{\delta s(4a - b) - \gamma(3b - 4a)}{9(R - \delta s - \gamma)(3b - 4a)}$$

To find a cut-off α_1, we set the inequality to hold with equality to get:

$$\alpha_1^{\frac{8}{3}} = \frac{\delta s(4a-b) - \gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)} \Rightarrow \alpha_1 = \left(\frac{\delta s(4a-b) - \gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{3}{8}}$$

Note: If the fraction in the brackets above is negative, we have the even-root of a negative number, which would not be a real number. What is in fact happening is that if b is large enough, or a is small enough, then the Author wishes to continue no matter what, and there is no interior cut-off α_1; in effect, $\forall \alpha_1 \geqslant 0 \equiv \forall \alpha_1$, the Author continues (hence sufficient condition 2).

So, overall, we have that the Author's optimal play, given the conditions above, is a cut-off strategy such that

$$\forall \alpha_1 \geqslant \left(\frac{\delta s(4a-b) - \gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{3}{8}}$$

the Author will continue at $t = 2$. For all other α_1, the Author will instead quit.

□

3.2.5 Reviewer at $t = 1$

Proposition 3.7. *For any equilibrium where the Author continues at $t = 1$ and plays a cut-off strategy as shown in proposition 3.6, the Reviewer plays a cut-off strategy where:*

$$\forall r \geqslant \frac{\delta s(3q^{\frac{1}{3}} - 9q^3 - 2) + 8s}{3(2 + q^{\frac{1}{3}} - 3q^3)} \tag{3.2}$$

$m_1 = 3$ is played, for all other types of Reviewer r, $m_1 = \frac{1}{3}$ is played.

Proof. If the Author quits at $t = 1$, the Reviewer has no substantive choice to make, so we assume the Author chooses to continue at $t = 1$. The Reviewer should assume that the Author plays as we have surmised, namely that $\forall \alpha_1$ above some cut-off value, they continue at $t = 2$. We say that the Reviewer believes that cut-off α_1 is some value q; that is, the Reviewer believes that if $\alpha_1 \geqslant q$, the Author continues at $t = 2$. Then, Reviewer payoff can be broken down as follows:

Pass ⇒ Continue ⇒ Complete Project

$$E[\text{payoff}] = E[\alpha_1|m_1]\mathbb{P}(\alpha_1 \geqslant q|m_1)E[\alpha_2|m_2]r = \left(\frac{m_1}{m_1+1}\right)(1-q^{m_1})\left(\frac{m_2}{m_2+1}\right) r$$

Pass ⇒ Continue ⇒ Fail

$$E[\text{payoff}] = E[\alpha_1|m_1]\mathbb{P}(\alpha_1 \geqslant q|m_1)E[1-\alpha_2|m_2]\delta s = \left(\frac{m_1}{m_1+1}\right)(1-q^{m_1})\left(1-\frac{m_2}{m_2+1}\right)\delta s$$

Pass ⇒ Quit

$$E[\text{payoff}] = E[\alpha_1|m_1]\mathbb{P}(\alpha_1 < q|m_1)\delta s = \left(\frac{m_1}{m_1+1}\right)(q^{m_1})\delta s$$

Fail

$$E[\text{payoff}] = E[1-\alpha_1|m_1]s = \left(1 - \frac{m_1}{m_1+1}\right)s$$

Together, we get:

$$E[\text{payoff}] = \left(\frac{m_1}{m_1+1}\right)(1-q^{m_1})\left(\frac{m_2}{m_2+1}\right)r + \left(\frac{m_1}{m_1+1}\right)(1-q^{m_1})\left(1-\frac{m_2}{m_2+1}\right)\delta s$$
$$+\left(\frac{m_1}{m_1+1}\right)(q^{m_1})\delta s + \left(1-\frac{m_1}{m_1+1}\right)s$$

We know that if $r < \delta s,\ m_1 = m_2 = \frac{1}{3}$, so we are only interested in the case where $r \geqslant \delta s$, in which case $m_2 = 3$, but m_1 is still undetermined, giving:

$$E[\text{payoff}] = \frac{3r}{4}\left(\frac{m_1(1-q^{m_1})}{m_1+1}\right) + \frac{\delta s}{4}\left(\frac{m_1(1-q^{m_1})}{m_1+1}\right) + \delta s\left(\frac{m_1 q^{m_1}}{m_1+1}\right) + s\left(1-\frac{m_1}{m_1+1}\right)$$

Since the Reviewer's choice is discrete, we can compare payoffs for the two possible m_1 quite easily to verify if a cut-off strategy is sensible:

$$E\left[\text{payoff}|m_1 = \frac{1}{3}\right] = \frac{3r}{4}\left(\frac{(1-q^{\frac{1}{3}})}{4}\right) + \frac{\delta s}{4}\left(\frac{(1-q^{\frac{1}{3}})}{4}\right) + \delta s\left(\frac{(q^{\frac{1}{3}})}{4}\right) + \frac{3s}{4}$$
$$E[\text{payoff}|m_1 = 3] = \frac{3r}{4}\left(\frac{3(1-q^3)}{4}\right)r + \frac{\delta s}{4}\left(\frac{3(1-q^3)}{4}\right) + \delta s\left(\frac{3q^3}{4}\right) + \frac{s}{4}$$

$m_1 = 3$ is preferred if:

$$E\left[\text{payoff}|m_1 = \frac{1}{3}\right] \leqslant E[\text{payoff}|m_1 = 3]$$
$$\Rightarrow \left(\frac{1}{4}\right)\left(\frac{3}{4}\right)r[3(1-q^3) - 1 + q^{\frac{1}{3}}] \geqslant \left(\frac{1}{4}\right)^2 \delta s[1 - q^{\frac{1}{3}} + 4q^{\frac{1}{3}} - 3 + 3q^3 - 12q^3] + \frac{1}{2}s$$
$$\Rightarrow 3r[2 + q^{\frac{1}{3}} - 3q^3] \geqslant \delta s[3q^{\frac{1}{3}} - 9q^3 - 2] + 8s$$

$\forall q \in [0,1),\ 2 + q^{\frac{1}{3}} - 3q^3 > 0$, and $q = 1 \Rightarrow 2 + q^{\frac{1}{3}} - 3q^3 = 0$, so we assume $q < 1$ to avoid division by zero. Then, we get that:

$$\forall r \geqslant \frac{\delta s(3q^{\frac{1}{3}} - 9q^3 - 2) + 8s}{3(2 + q^{\frac{1}{3}} - 3q^3)}$$

$m_1 = 3$ is preferred to $m_1 = \frac{1}{3}$. This confirms a cut-off type strategy for Reviewers at $t = 1$, where $\exists \gamma > 0$ such that if $r > \delta s + \gamma$, Reviewers prefer $m_1 = 3$.

Importantly, the above cut-off is greater than δs, which can be verified using Wolfram-Alpha or Desmos (once we pick specific values for different variables, it is quite easy to verify analytically as well).

□

3.2.6 Author behaviour at $t = 1$

Proposition 3.8. *Given behaviour in accordance with propositions 3.5, 3.6, and 3.7, Author continues at $t = 1$ under predictable and attainable conditions, where Author continues whenever*

$$b\{E[\alpha_1]\mathbb{P}(\alpha_1 \geqslant \tau)E[\alpha_2|\alpha_1 \geqslant \tau]\} + a\{E[\alpha_1]\mathbb{P}(\alpha_1 < \tau)\} \geqslant a \tag{3.3}$$

and objects $E[\alpha_1]$, $\mathbb{P}(\alpha_1 \geqslant \tau)$, $E[\alpha_2|\alpha_1 \geqslant \tau]$ are all well-defined and unique objects (defined in the proof).

Proof. First, we make the following definition:

Definition 3.9 (Tau)**.**

$$\tau := \left(\frac{\delta s(4a-b) - \gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{3}{8}} \tag{3.4}$$

That is, τ is the cut-off α_1 we computed earlier, such that $\forall \alpha_1 \geqslant \tau$, the Author will choose to continue at $t = 2$. $\tau \neq q$ as q is the Reviewer's *belief* of the Author's cut-off.

Then, the Author's payoff from choosing to continue at $t = 1$ can be broken down as follows:

Pass ⇒ Continue ⇒ Complete Project

$$\text{Probability} = E[\alpha_1]\mathbb{P}\left(\alpha_1 \geqslant \tau\right) E\left[\alpha_2 \left|\alpha_1 \geqslant \tau\right.\right]$$
$$\text{Payoff} = b$$

Where we have the following:

- $E[\alpha_1]$ is the ex-ante probability of success
- $\mathbb{P}\left(\alpha_1 \geqslant \tau\right) \equiv \mathbb{P}\left(\alpha_1 \geqslant \left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{3}{8}}\right)$ is the probability that the Author chooses to continue in the next period
- $E\left[\alpha_2 \left|\alpha_1 \geqslant \tau\right.\right] \equiv E\left[\alpha_2 \middle| \alpha_1 \geqslant \left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{3}{8}}\right]$ is the probability of success at $t = 2$ given that the Author chose to continue (which required α_1 to be above the given threshold).

Pass ⇒ Continue ⇒ Fail

$$\text{Probability} = E[\alpha_1]\mathbb{P}\left(\alpha_1 \geqslant \tau\right)\left(1 - E\left[\alpha_2 \left|\alpha_1 \geqslant \tau\right.\right]\right), \ \text{Payoff} = 0$$

Pass ⇒ Quit

$$\text{Probability} = E[\alpha_1]\mathbb{P}\left(\alpha_1 < \tau\right), \ \text{Payoff} = a$$

Fail

$$\text{Probability} = 1 - E[\alpha_1], \ \text{Payoff} = 0$$

Then, we compute all of the relevant components separately:

$\underline{E[\alpha_1]}$

Under the hypothesised equilibrium, we have that

$$E[\alpha_1|r<\delta s+\gamma]=\frac{\frac{1}{3}}{1+\frac{1}{3}}=\frac{1}{4}\text{ and }E[\alpha_1|r\geqslant\delta s+\gamma]=\frac{3}{1+3}=\frac{3}{4}\text{ so:}$$
$$E[\alpha_1]=\frac{1}{4}\left(\frac{\delta s+\gamma}{R}\right)+\frac{3}{4}\left(\frac{R-\delta s-\gamma}{R}\right)=\frac{3R-2(\delta s+\gamma)}{4R}$$

$\underline{\mathbb{P}(\alpha_1\geqslant\tau)}$

Here, we effectively want to find $1-F(\tau)$, where $F(\cdot)$ is the unconditional CDF of α_1. Using similar logic to before, we argue:

$$\begin{aligned}F(x)&=F(x|r<\delta s+\gamma)\mathbb{P}(r<\delta s+\gamma)+F(x|r\geqslant\delta s+\gamma)\mathbb{P}(r\geqslant\delta s+\gamma)\\&=(x)^{\frac{1}{3}}\left(\frac{\delta s+\gamma}{R}\right)+(x)^{3}\left(\frac{R-\delta s-\gamma}{R}\right)\end{aligned}$$
$$\Rightarrow F(\tau)=\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{1}{8}}\left[\left(\frac{\delta s+\gamma}{R}\right)+\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)\left(\frac{R-\delta s-\gamma}{R}\right)\right]$$

and then, $\mathbb{P}(\alpha_1\geqslant\tau)=1-F(\tau)$, with $F(\tau)$ from above.

$\underline{E[\alpha_2|\alpha_1\geqslant\tau]}$

$$E\left[\alpha_2\,|\alpha_1\geqslant\tau\right]=E[\alpha_2|r<\delta s]\mathbb{P}\left(r<\delta s\,|\alpha_1\geqslant\tau\right)+E[\alpha_2|r\geqslant\delta s]\mathbb{P}(r\geqslant\delta s|\alpha_1\geqslant\tau)$$

The above stems from our previous argument that $\forall r\geqslant\delta s,\ m_2=3$, whereas $\forall r<\delta s,\ m_2=\frac{1}{3}$.

To find $\mathbb{P}(r<\delta s|\alpha_1\geqslant\tau)$, we use a similar Bayes rule formulation as before:

$$\begin{aligned}\mathbb{P}(r<\delta s|\alpha_1\geqslant\tau)&=\frac{\mathbb{P}(\alpha_1\geqslant\tau|r<\delta s)\mathbb{P}(r<\delta s)}{\mathbb{P}(\alpha_1\geqslant\tau)}\\&=\frac{\left(1-\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{1}{8}}\right)\left(\frac{\delta s}{R}\right)}{1-\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{1}{8}}\left[\left(\frac{\delta s+\gamma}{R}\right)+\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)\left(\frac{R-\delta s-\gamma}{R}\right)\right]}\end{aligned}$$
$$\begin{aligned}\Rightarrow E[\alpha_2|\alpha_1\geqslant\tau]=\frac{1}{4}&\left[\frac{\delta s\left(1-\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{1}{8}}\right)}{R-\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{1}{8}}\left[(\delta s+\gamma)+\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)(R-\delta s-\gamma)\right]}\right]\\+\frac{3}{4}&\left[1-\frac{\delta s\left(1-\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{1}{8}}\right)}{R-\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{1}{8}}\left[(\delta s+\gamma)+\left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)(R-\delta s-\gamma)\right]}\right]\end{aligned}$$

So, given all of the above, we can solve for the Author's $t=1$ action by computing

$$b\{E[\alpha_1]\mathbb{P}(\alpha_1 \geqslant \tau)E[\alpha_2|\alpha_1 \geqslant \tau]\} + a\{E[\alpha_1]\mathbb{P}(\alpha_1 < \tau)\}$$

which is the payoff from continuing at $t = 1$, which we can then compare to the payoff from quitting at $t = 1$ (which is just a payoff of a), to determine the Author's course of action.

□

3.2.7 Necessary and Sufficient Conditions for Target Equilibrium Existence

Finally, to determine the existence of equilibria, we have to solve for the two endogenous parameters q and γ such that they are mutually compatible and generate an actual equilibrium outcome. Here, we generally define necessary equilibrium conditions, which we will then use to showcase a specific equilibrium outcome using a particular choice of parameters.

Theorem 3.10 (Necessary and Sufficient Conditions for Target Equilibrium Existence)**.** *Target equilibrium exists whenever condition 3.3 holds, ensuring Author continuation at* $t = 1$*, and if* $\exists q \in (0, 1),\ \gamma\ \in [0, R - \delta s)$ *such that:*

$$q = \left(\frac{\delta s(4a - b) - \gamma(3b - 4a)}{9(R - \delta s - \gamma)(3b - 4a)}\right)^{\frac{3}{8}} \tag{3.5}$$

$$\delta s + \gamma = \frac{\delta s(3q^{\frac{1}{3}} - 9q^3 - 2) + 8s}{3(2 + q^{\frac{1}{3}} - 3q^3)} \tag{3.6}$$

Proof. First, recall that by proposition 3.6, $\forall \alpha_1 \geqslant \left(\frac{\delta s(4a-b)-\gamma(3b-4a)}{9(R-\delta s-\gamma)(3b-4a)}\right)^{\frac{3}{8}}$, the Author chooses to continue at $t = 2$. Further, by proposition 3.7, $\forall r \geqslant \frac{\delta s(3q^{\frac{1}{3}}-9q^3-2)+8s}{3(2+q^{\frac{1}{3}}-3q^3)}$, Reviewers pick $m_1 = 3$.

γ is defined as the cutoff for which all Reviewers with $r \geqslant \delta s + \gamma$ play $m_1 = 3$ within the Author's belief set.

q is defined as the cutoff for which an Author observing $\alpha_1 \geqslant q$ chooses to continue at $t = 2$ within the Reviewer's belief set.

Sequential Equilibrium will require than each agents' beliefs match the outcome of equilibrium play (consistency of beliefs), and so it will require the following to hold:

$$q = \left(\frac{\delta s(4a - b) - \gamma(3b - 4a)}{9(R - \delta s - \gamma)(3b - 4a)}\right)^{\frac{3}{8}}$$
$$\delta s + \gamma = \frac{\delta s(3q^{\frac{1}{3}} - 9q^3 - 2) + 8s}{3(2 + q^{\frac{1}{3}} - 3q^3)}$$

The above yields two equations with two unknown variables, thus meaning that we can solve explicitly for q and γ.

□

3.2.8 Particular Solution Demonstration

Now, we test a specific choice of parameter values to determine whether an equilibrium exists, and what said equilibrium looks like. Suppose that:

$$a = 1,\ b = 3,\ s = 1,\ \delta = 0.9,\ R = 3$$

Plugging in these values and solving the system of equations given above (in this case done using an online Graphing calculator, namely "Desmos"), we find the following solution:

$$q \approx 0.127,\ \gamma \approx 0.107$$

Note: Another solution does exist, but it is out-of-bounds, that is, produces $q \notin [0, 1]$, and so we take only one of the solutions computed.

Using these values, in addition to all of the parameters given above, we can determine the Author's expected payoff from continuing at $t = 1$ and thus determine their behaviour at $t = 1$. Plugging all of these values into the formula for expected $t = 1$ payoff from continuing yields

$$E[\text{Payoff from continuing}] \approx 1.408 > 1 = a$$

Consequently, the Author prefers to continue at $t = 1$.

Overall, we have found an equilibrium for the specific parameters described, which can be summarised as follows: The Author chooses to continue at $t = 1$. If the Reviewer is such that $r \gtrapprox 1.007$, they set $m_1 = 3$. Otherwise, they set $m_1 = \frac{1}{3}$. If the game then continues to $t = 2$, the Author observes the realised value of α_1, and if $\alpha_1 \gtrapprox 0.127$, the Author continues again. Otherwise, if $\alpha_1 \lessapprox 0.127$, the Author quits. If the Reviewer is such that $r \geqslant 0.9$, they set $m_2 = 3$. Otherwise, they set $m_2 = \frac{1}{3}$.

3.2.9 Implications and Take-Aways

Properties of this Equilibrium

Key points are underlined where they are relevant to highlight main take-aways.

We can begin by discussing some key concepts relating to the specific solution detailed above. First, we note that the Author does not require high observed α_1 to be willing to continue at $t = 2$. The payoff from completing the project is three times higher than that of quitting, so even low success odds are not a deterrent. Notably, once $t = 2$ has already been reached, the Author only needs to survive one more period of review, meaning that even the worst case scenario of $E[\alpha_2 | m_2 = \frac{1}{3}] = \frac{1}{4}$ is a weak deterrent. Further, the share of high r Reviewers relative to low r Reviewers is ex-ante high; this means that a low realised α_1 has a high likelihood of being due to a tail-end realisation of $F(x|3)$, making continuation worthwhile. Given that γ is high in equilibrium, a large share of Reviewers would ex-ante be the type that changes behaviour between periods. Thus, even if low realised α_1 was due to facing a Reviewer who chose $m_1 = \frac{1}{3}$, there is a high probability that m_2 would nevertheless be set to 3.

We also observe that γ is high, so much so that some Reviewers who would prefer project completion to any other outcome still choose to give the Author a negative review. The latter phenomenon is due to the possibility of Author failure or quitting at $t = 2$. If $r > s$, then project completion is better than any other outcome. However, if the odds of the Author quitting or failing at $t = 2$ are high, then Reviewers with small r, but nevertheless with $r > s$, may wish to induce failure early to secure s payoff, rather than risk being stuck with δs payoff. Since $m_t \in \{\frac{1}{3}, 3\}$, Reviewers can induce no better than $E[\alpha_t] = \frac{3}{4}$; thus, at best, success odds ex-ante are $\frac{9}{16} \approx 0.56$. As δs and s are quite close in this example, for small r the approximately 0.44

probability of foregoing 0.1 payoff to gain r with probability 0.56 is not worthwhile.

We also note that in this example, the risk of the Author quitting at $t = 2$ is likely to act to reduce γ. With q being so low, the probability that the Author quits given $m_1 = 3$ is only $F(0.127|3) \approx 0.002$. On the other hand, the risk of the Author quitting if $m_1 = \frac{1}{3}$ is ≈ 0.5, which is high. Thus, choosing $m_1 = \frac{1}{3}$ to try to induce early failure to guarantee payoff s rather than risking δs, significantly raises the odds of quitting at $t = 2$. So, in the event of success at $t = 1$ despite the choice of $m_1 = \frac{1}{3}$, the Reviewer would give themselves poor odds of getting r rather than δs. This reduces the share of Reviewers that are willing to choose $m_1 = \frac{1}{3}$, lowering γ. Therefore, it is not a particularly extreme choice of parameters that has resulted in this perhaps unintuitive result, as it is possible to construct more extreme scenarios with even higher γ.

Finally, the "degree of agency" of the Reviewer is very important. The specifics of this idea will be discussed more in a later section, but it is roughly defined as the Reviewer's ability to influence outcomes; we consider the idea intuitively in this section. Suppose we changed the options available to the Reviewer, such that they could choose actions in $\{\frac{1}{2}, 2\}$, rather than $\{\frac{1}{3}, 3\}$. In effect, we reduce how "strong" the Reviewer's actions are, as the amount by which the Reviewer can shift expected α is now lower. If the Author knew with 100% certainty that they faced an unfavourable Reviewer at $t = 2$ in this new scenario, their expected payoff from continuing would be:

$$E\left[\alpha_2 \middle| m_2 = \frac{1}{2}\right](3) = 1$$

whereas their expected payoff from quitting would be $a = 1$. Therefore, conditional on reaching $t = 2$, the Author would never quit, as even if they were absolutely certain that the Reviewer was unfavourable, they would have no incentive to do so. Thus, the ability of the Reviewer to actually influence expected outcomes is clearly very important. This influence is what will later be defined to be the "degree of agency", but it is an intuitive concept. In effect, the Reviewer must have sufficient "degree of agency" for their threat of a negative review, or promise of a positive review, to affect Author behaviour.

The Story of the Equilibrium

To summarise the idea that this particular equilibrium outcome is capturing, we consider the following story: A mayor is deciding whether to initiate a construction project or not. If they choose not to, or if the city council rejects the project, any unused funds earmarked for the project are distributed towards other, smaller investments. A majority of the city council prefers the completion of the project to many smaller investments. Some members prefer the completion of the project to many smaller investments, but not by much. Given that the project has to survive several council voting rounds, and the mayor might withdraw their own support for the project, these members would be reluctant to offer the project support.

In such a situation, voting dynamics may be complex and varied, but the model describes the outcomes well. The mayor takes the role of the Author, and the mix of council-member preferences are captured by the range and distribution of r. The project's survival is then governed not only by the mayor's choice, obvious in the model, but also voting, which in our model is abstracted by making the Reviewer's influence limited. Someone is inevitably going to be the loudest voice on the council, and they may or may not succeed in swaying a majority to vote in favour or against. This loudest voice is then the Reviewer, and the uncertainty of voting behaviours among the rest of the council we capture by randomising project success, even after the intervention of the Reviewer. Which way they push the council will depend on their own preferences for the project (ie: r), and dynamics then follow the process described in the model.

3.3 Alternative Equilibria and Non-Equilibrium Outcomes

The equilibrium detailed above is just one of many possible equilibria. Which equilibrium ultimately manifests depends on what parameters are chosen, as well as the equilibrium strategies prescribed to each agent. Nevertheless, there are some outcomes that we can rule out definitively. Broad classes of both attainable equilibria and non-equilibrium outcomes are detailed in this section.

Other Plausible Equilibrium Outcomes

First, we consider some plausible equilibrium outcomes, and detail different varieties of some of these equilibrium outcomes.

The Author quits immediately.

The first equilibrium to consider is one where the Author quits immediately. The Author's strategy can prescribe one of the following actions: always quitting at $t = 2$, always continuing at $t = 2$, or sometimes quitting and sometimes continuing at $t = 2$. Further, the Reviewer may: play in the same way as in the target equilibrium, they may always play low m_t, or they may always play high m_t.

Immediately, we notice that this particular equilibrium outcome is very flexible, with several strategies being sustainable under it. There are essentially four reasons for this. First, ex-ante success probability and/or project completion payoff being low are all that is needed to enforce this outcome; we do not need complex interactions to sustain it. Second, there is nothing requiring the Author to prefer success over the Reviewer, meaning it is entirely reasonable for Reviewers to want success at all costs while Authors prefer to quit. Third, many actions remain to be decided, for instance all behaviour at $t = 2$, so as long as we do not make continuing preferable to quitting at $t = 1$ for the Author, we can be "creative" with every agents' play thereafter. Finally, because quitting immediately occurs so early, there are many processes that can lead to this outcome. Failure at $t = 1$, failure at $t = 2$, being incentivised to quit at $t = 2$, or low success payoff could all dissuade the Author from continuing; that is, because there are so many different ways to dissuade the Author from continuing at $t = 1$, there are many equilibria compatible with this outcome.

The Author always continues.

On the other hand, we can define equilibria where the Author always continues. In such an equilibrium the Reviewer may: play in the same way as in the target equilibrium, always play low m_t, or always play high m_t. This is still a fairly flexible scenario, for many of the same reasons as described above. However, it is less flexible for the simple reason that we cannot choose the Author's play at $t = 2$, it being defined already.

Non-Equilibrium Outcomes

Some outcomes cannot be reached in equilibrium, which can be seen quite easily by detailing each outcome in turn and heuristically considering why it fails.

The Author continues at $t = 1$, and always quits at $t = 2$.

First, we consider an outcome where the Author continues at $t = 1$, and always quits at $t = 2$. The obvious problem with this outcome is that the Author would only continue at $t = 1$ if there is a non-zero prospect of success at $t = 2$, and if success gives greater payoff than quitting. By always quitting at $t = 2$, the Author ensures zero probability of success at $t = 2$. Then, by continuing at $t = 1$, the Author introduces non-zero probability of failure at $t = 1$. Clearly, there is no reason for the Author to prefer such an outcome to quitting immediately, and so it cannot be sustained in an equilibrium.

Some Reviewers select high m_1 and low m_2.

Next, we consider an outcome where some Reviewers select high m_1 and low m_2. We observe that failure and quitting payoff is strictly lower at $t = 2$ than at $t = 1$ for Reviewers. We also observe that at $t = 2$, the choice of m_2 no longer has any influence on Author behaviour, so it must be that the Reviewer strictly prefers failure to success at $t = 2$ ceteris paribus. Given that the Reviewer prefers failure to success, high m_1 can at best have no effect on the Author's probability of quitting or continuing, or at worst raise the probability of the Author continuing, which is not desirable to a Reviewer that prefers failure or quitting to success. Thus, any Reviewer that is willing to induce high failure probability at $t = 2$ must be willing to do so at $t = 1$, since payoff of failure is strictly higher. So, it is not possible to sustain this type of Reviewer behaviour in an equilibrium.

Note: If an Author quits immediately, then the letter scenario becomes off-path, making the use of a Sequential Equilibrium solution concept important in ruling it out.

4 Key Model Components and Features to Consider

Having solved for an equilibrium of the model, we might now ask what the actual objective of the model is? What does the model do that others do not? What ideas do we wish to highlight by using it? To better understand some of these questions, we analyse the mechanisms at play in the model. While we primarily consider a single, simplified model that explores an interesting real-life phenomenon theoretically, it is in fact composed of several interacting mechanisms. So, we discuss in detail the main mechanisms at play in the model, from which we can discern what they imply. Further, while none of the mechanisms are particularly useful when deployed alone, they can be combined in different ways than is done in this model; thus, it is also worth understanding the functioning of each mechanism separately to allow this model to be extendible and applicable outside of the context proposed here. Three of these mechanisms are separately developed further later in this section, as they are more complex and merit deeper discussion.

4.1 Model Mechanism Overview and Key Interpretation

Main Mechanisms of the Model:

1. The Reviewer's agency is "twice-limited", and they cannot pick outcomes deterministically.
2. The Author observes realised probability of success.
3. Entanglement of the Reviewer's signal and action.
4. Effectively unilateral decision-making by the Author.
5. Repeated interaction.

Direct Implications of the Mechanisms:

Key points are underlined where they are relevant to highlight main take-aways.

Mutually-beneficial coordination between Author and Reviewer may fail. We can view this as an inability of the electorate to fully communicate its preferences to a decision-maker, leading to policy failure. That is, even when the Author and Reviewer are perfectly aligned, the Author may (depending on model parameters chosen) act against both of their interests. Mechanism (1) prevents the Reviewer from being able to fully reveal their type, even when it would induce the Author to behave desirably. It also enables very low success-probabilities from materialising, even when the Reviewer tries to maximise success odds. Then, mechanism (2) allows the Author to observe low success-probability generated from an unlucky favourable Reviewer, which then induces quitting, despite the review having been favourable. Finally, mechanism (4) means that whatever action the Author chooses determines everyone's outcomes, and so if the Author observes a negative signal, they can unilaterally induce a bad outcome for both agents, even when both would have preferred a different action.

A Reviewer with $r \geqslant s$ may give a negative review at $t = 1$, despite preferring success to any other outcome. We can view this as a voter voting against a policy in their best interest, and this is something we observed in the equilibrium solution detailed earlier. Mechanism (5) dictates that whatever happens at $t = 1$, success is needed at $t = 2$ as well as at $t = 1$ to ensure full project success. Mechanisms (2) and (4) mean that the Author can unilaterally induce a payoff of δs at $t = 2$, even if the Reviewer gave a favourable review, simply due to bad luck combined with the Author's unilateral power. If the above effects make continuing to $t = 2$ too risky, even if success were the Reviewer's preferred outcome, they may choose to minimise success odds at $t = 1$ to "take a safe bet", rather than risking proceeding into $t = 2$ and facing the possibility of the Author quitting or failing (granting reduced payoff δs). Whether continuation at $t = 2$ can be said to be "too risky" depends entirely on parameter choices and equilibrium effects, but as we have already seen, there are reasonable parameter choices that induce precisely this type of behaviour.

A Reviewer with $\delta s < r < s$ may give a positive review at $t = 1$, despite preferring failure to success. We can view this as a voter voting in favour of an unwanted policy. This is, in effect, the opposite of the above phenomenon, and using the language of the equilibrium solution given above, would require $\gamma < (1 - \delta)s$. Mechanism (1) means that even when the Reviewer tries to minimise success odds, the Author may nevertheless succeed, since success is a stochastic process influenced by the Reviewer (but not deterministically so). Mechanism (3) says that whatever decision the Reviewer makes at $t = 1$ acts as both a signal and a direct action. Then, a combination of mechanisms (2), (4), and (5) allow the Author to observe the realised outcome of the Reviewer's decision and then act upon that information unilaterally at the start of $t = 2$, potentially choosing to quit (if α_1 is below a cut-off, for instance, as in the equilibrium previously detailed). Since mechanism (3) means that minimising success odds at $t = 1$ also minimises positive signal odds at $t = 2$, the Reviewer knows that a negative review early on raises the odds of the Author quitting at $t = 2$ due to an unfavourable realisation of α_1. If the odds of reaching $t = 2$ only for the Author to then quit are high enough (again, this is parameter dependent), given a negative signal, the Reviewer prefers to send a positive signal early on to raise the probability that the Author continues at $t = 2$. Thus, they are able to secure $r > \delta s$, while lowering their odds of getting $s > r$; if the risk is high and the gap between r and s is relatively small compared to the gap between r and δs, the Reviewer would prefer this course of action.

These outcomes, which are quite interesting and counter-intuitive, all require each of the 5 mechanisms described. This suggest that we can explain the story of legislators cancelling desirable projects or pushing through undesirable ones via this model, as a combination of coordination failure and electorates behaving counter-intuitively so as to incentivise this type of behaviour.

Interpretation: Abstracting Multi-Agent Scenarios

The above implications are especially interesting when we consider the model as capturing complex group-behaviour dynamics despite modelling only a two-agent interaction scenario.

Take the election motivation. In principle, there are several groups of voters, all with their own preferences, all acting in varying ways, to produce a collective election outcome. By isolating a pivotal voter and allowing all other voters to be abstracted away as random noise, we turn the voting system into a bilateral interaction environment. We still capture the influence and behaviour of the pivotal voter, and we capture the voting behaviour of all other agents by correlating their behaviour to that of the pivotal voter.

In scenarios where we are most interested in the behaviours not of the voters but of the decision maker in question, this simplification can be highly informative. Because the decision maker does not know ex-ante who the pivotal voter actually is, they have to behave *as if* they faced the entire electorate in the first period. They must take the population distribution of preferences over projects into account, adjusting for the likely level of influence of different voting blocs. This is a reasonable approximation of many multi-lateral scenarios with one key decision maker and several other stake-holders with varying levels of influence.

Taken in this light, the implications highlighted which arise naturally from this two-agent model, can indeed be viewed as a prediction regarding multi-agent electoral environments.

4.2 First Mechanism: Twice-Limited Agency

The first theoretical mechanism in the model we call "Twice-Limited Agency". In a sense, it formalises a common principle in many economic models, while introducing an additional complication that proves useful in some contexts. We touched on this notion in the discussion regarding the implications of the equilibrium computed above, but now we formalise the idea in question.

First, we define "Once-Limited Agency":

Definition 4.1 (Once-Limited Agency)**.** Agency is said to be "once-limited" if the agent to which the agency applies cannot directly select an outcome, but can affect the probability of an outcome.

From this follows the definition of twice-limited agency:

Definition 4.2 (Twice-Limited Agency)**.** Agency is said to be "twice-limited" if the agent to which the agency applies can neither directly select an outcome, nor directly select the probability of an outcome, but can affect the distribution governing the probability of an outcome.

4.2.1 Once-Limited Agency

To understand this in practice, we consider what, in this framework, would be an example of a once-limited agency scenario:

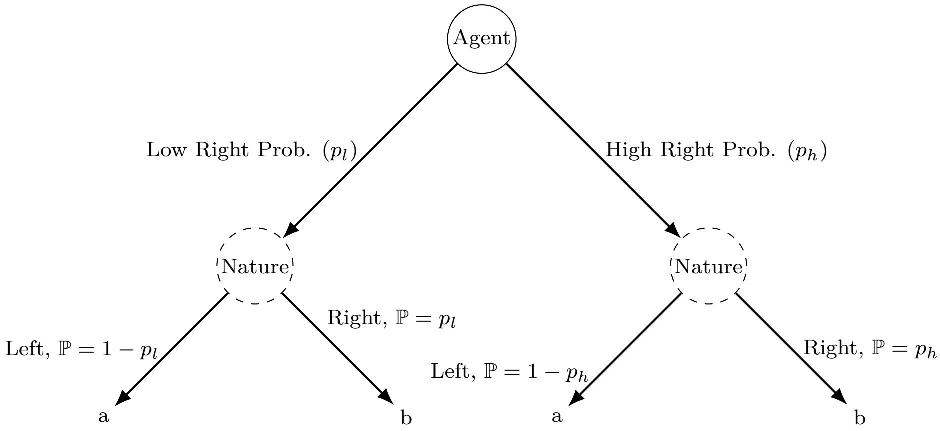

Figure 10: Simple Formulation of Once-Limited Agency

Here, the Agent has once-limited agency; they cannot choose to go left or right, but can only choose the probabilities of going left or right. In general, any scenario where the agent can directly select the probability of an outcome in a deterministic sense is once-limited. While the language "once-limited agency" appears to be unique to this paper, the idea is by no means unique in economics; the need of this language only arises to enable contrasting this more standard principle with "twice-limited agency".

4.2.2 Twice-Limited Agency

Now, we show the simplest possible twice-limited agency setting:

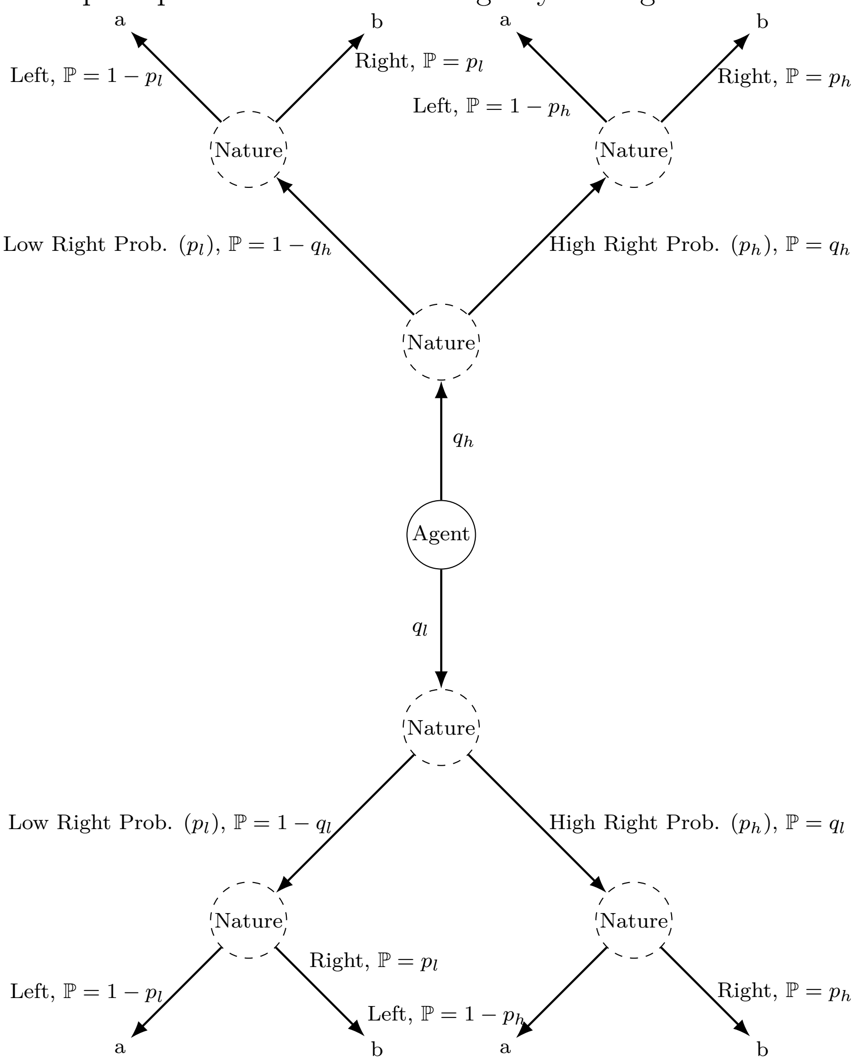

Figure 11: Simple Formulation of Twice-Limited Agency

What we observe is that the agent not only cannot choose the outcome in a deterministic way, but they cannot even choose the probability of the outcome in a deterministic sense. In this way, their agency is twice-limited.

4.2.3 Outcome-Equivalence Proposition

The first obvious question is if this matters at all? What we will find is that if we consider twice-limited agency in isolation, it does not, but when combined with other features of the model, it proves to be a useful tool. First, we show that it is not particularly valuable if not combined with some other mechanism.

Note: Below, the variable $\mathbf{m}$ is used, as opposed to scalar m introduced earlier. This is a generalisation: decision makers can select a vector of m rather than a single m to affect distributions, for example if a distribution has several governing parameters.

Proposition 4.3. *For any expected-utility maximising agent, twice-limited agency can be reduced to an outcome-equivalent once-limited agency problem, as long as expectations are defined throughout, if the distribution of probabilities generated by a twice-limited agent affects outcomes only through the outcome-probabilities it implies.*

Proof. Twice-limited agent chooses some action $\mathbf{m}_k \in M$ which governs the joint-distribution of subsequent outcome-probability distribution: $\mathbf{p} = (p_1, ..., p_n) \sim F(\cdot|\mathbf{m}_k)$, where p_i is the probability that outcome i is realised.

Note: We implicitly restrict F to ensure that probabilities sum to one.

By assumption, expectations are defined throughout, and so any choice of $\mathbf{m}_k$ generates a corresponding $E[\mathbf{p}|\mathbf{m}_k] = (E[p_1|\mathbf{m}_k], ..., E[p_n|\mathbf{m}_k])$.

Let $\mathbf{u} = (u_1, ..., u_n)$ be the payoff vector to the agent corresponding to each outcome i. Since the agent is an expected-utility maximiser, they will choose $\mathbf{m}_k$ to maximise $< \mathbf{u}, E[\mathbf{p}|\mathbf{m}_k] >$ (inner-product).

Now, we can define an outcome-equivalent once-limited agency problem. For every $\mathbf{m}_k \in M$, define $\mathbf{r}_j \in R$ such that for each $\mathbf{m}_k, \exists \mathbf{r}_j = E[\mathbf{p}|\mathbf{m}_k]$. Then, for every choice of $\mathbf{m}_k$, there is a $\mathbf{r}_j$ producing the exact same outcome (in the eyes of the agent) in a once-limited agency setting.

Note: We say "outcome-equivalent" and not "equivalent" because it is possible (in this very general setting) that two $\mathbf{m}_j \neq \mathbf{m}_k$ exist that produce exactly the same expected outcome. When reducing these to corresponding $\mathbf{r}$, we have that the same $\mathbf{r} \in R$ maps onto two different $\mathbf{m}$, meaning the relationship is not equivalent. Further, this proposition requires $\mathbf{m}$ to only affect outcomes by how it affects the draw of $\mathbf{p}$; if, for example, $\mathbf{m}$ had some secondary effect before it generates a vector of probabilities, then there is no once-limited agency set-up that achieves the same result. □

What we understand from this is that it is not unexpected that this concept has not been used frequently, nor has it really been formalised in this particular way; it does not change results in most set-ups. Where it does become very important, however, is when combined with other mechanisms that invalidate the proposition, as is the case in our model; here, the choice of m_1 not only affects the success-probability, but also the signal that the Author ultimately receives.

4.2.4 Value in Strategic-Interaction Environments and Blackwell Parallels

Twice-limited agency is valuable in strategic-interaction environments by allowing a different, more complex, form of imperfect communication. In a once-limited agency environment, one can introduce imperfect communication and coordination by, for example, allowing a second agent to view the outcome of the first agent's choice, but not the actual probability they chose. This already goes a long way, but it can be limited. Scenarios where we want there to exist substantial ambiguity between what action the first agent chose to create a communication failure can be easier to accomplish with twice limited agency.

In particular, for a once-limited agency scenario to allow this type of uncertain communication, we require a repeated game where an intermediate outcome is recorded. Suppose the first agent can choose probability of success $\{\frac{1}{4}, \frac{3}{4}\}$. If the second agent can observe if success or failure was realised, then we get an interesting uncertain interaction environment. If we wanted failure to immediately end the game, while we preserve some of that uncertain environment since success could have occurred with a $\frac{1}{4}$ or $\frac{3}{4}$ probability, it is a less rich decision environment due to the discrete nature of possible outcomes.

With twice-limited agency, we can allow for a complete range of success probability in the set $[0, 1]$ without having to give the first agent the ability to choose from a continuous set of actions, which is easier to handle computationally. We still allow a continuous set of outcome probabilities, preserving complexity, while preserving discrete action-sets for the twice-limited agent. This is exactly what the equilibrium we solved achieves; despite $M = \{\frac{1}{3}, 3\}$, there is a positive probability of observing any $\alpha \in [0, 1]$, complicating decision-making and enabling the counter-intuitive results we observed in the solution given.

In addition, we can introduce the possibility that an agent who selected the higher probability nevertheless induces a lower probability of success; we can reveal that the success probability was low to create a further strategic interaction, where now the agent that observes the realised probability must determine whether this was due to misfortune or deliberate choice.

In most cases once-limited agency will be sufficient. However, in specific decision-making environments where coordination is very hard, twice-limited agency may prove a better fit, both by adding a layer of "confusion" in that it is unclear if an outcome was due to misfortune or deliberate choice, and by more easily enabling continuous outcome-probabilities.

As a side-note, it is reasonable to view twice-limited agency as a very specific form of *garbling* in the **Blackwell Experiment** sense (Blackwell 1951, 1953). Namely, if the strategic interaction environment involves an agent with twice-limited agency sending a signal that depends on their twice-limited agency action, then it is a garbled signal relative to once-limited agency. The garbling function is then the intermediate probability that is selected, which transforms the observed signal. Twice-limited agency is then a Blackwell-dominated experiment, relative to once-limited agency (itself dominated by perfect signalling).

4.2.5 Relevant Concepts in Twice-Limited Agency

To supplement the discussion, we consider some relevant concepts that are influential in shaping outcomes in twice-limited agency strategic-interaction scenarios. These are in effect formal definitions of concepts that intuitively arise when considering twice-limited agency, but it is useful to simplify discussion.

The below definitions all apply to a relatively general twice-limited agency formulation, where an agent chooses some $\mathbf{m} \in M$ which governs the joint-distribution of subsequent outcome-probability distribution, ie: $\mathbf{p} = (p_1, ..., p_n) \sim F(\cdot|\mathbf{m})$.

Definition 4.4 (Limited Agency Action Set)**.** An agent's limited agency action set is defined as the set M. This can be equivalent to their action set, if the only actions they can take exhibit limited agency.

Definition 4.5 (Degree of Agency)**.** An agent's **degree of agency** is defined as:

$$\text{Degree of Agency} = \max_{\mathbf{m}_1, \mathbf{m}_2} \left(\left|\left| \frac{1}{2} (E[\mathbf{p}|\mathbf{m}_1] - E[\mathbf{p}|\mathbf{m}_2]) \right|\right| \right)$$

In words, the degree of agency is the maximum change in expected outcome-probabilities that an agent can induce.

For example, if an agent chooses $m \in \{m_1, m_2\}$ to distribute a single p (and corresponding $(1-p)$), such that $E[p|m_1] = \frac{1}{4}$ and $E[p|m_2] = \frac{3}{4}$, then their degree of agency is $\sqrt{\frac{1}{2}[(\frac{1}{4} - \frac{3}{4})^2 + (\frac{3}{4} - \frac{1}{4})^2]} = \frac{1}{4}$; this is exactly the situation in the equilibrium we solved for.

Note: The $\frac{1}{2}$ in the degree of agency is just a normalisation constant to ensure that Degree of Agency $\in [0, 1]$, such that if Degree of Agency $= 1$, we can interpret that as being "full agency".

Degree of Agency formalises how much power we are giving the agent to determine outcomes. An agent with very limited degree of agency is essentially powerless, acting no different than random noise, whereas at the extreme an agent with high degree of agency can induce near-determinism (at least insofar that to an expected-utility maximiser, they can functionally pick outcomes "deterministically").

Definition 4.6 (Perfect Agency)**.** An agent has perfect agency in a twice-limited agency setting if their degree of agency is equal to 1.

Definition 4.7 (Complete Agency)**.** An agent has complete agency if the set $\{E[\mathbf{p}|\mathbf{m}]\}_{\mathbf{m} \in M}$ is convex.

Complete agency means that if an agency can induce two expected outcome-probabilities $E[\mathbf{p}]_1$ and $E[\mathbf{p}]_2$, they can induce any expected outcome-probabilities "in-between" those two. Complete agency can be difficult to work with in practice.

4.2.6 Limited Agency in the Context of the Target Equilibrium

Now, we can make practical observations regarding the target equilibrium computed before in the context of a limited agency environment. In particular, we observe a few properties of the set-up of the target equilibrium. First, the Reviewer's Limited Agency Action Set M is their entire action set; the have no deterministic actions they can take. Further, the Reviewer's Degree of Agency at any period t is $\frac{1}{2}\left(\frac{3}{4} - \frac{1}{4}\right) = \frac{1}{4}$ which is a moderate degree of agency, but the Reviewer does therefore not have perfect agency. As was discussed in the implications section of the equilibrium solution, this proves quite critical; if we change M to instead be $\{\frac{1}{2}, 2\}$, reducing the degree of agency to $\frac{1}{6}$, the equilibrium detailed ceases to hold. As we discussed, the Reviewer must have a sufficient degree of agency to influence the behaviour of the Author to preserve the target equilibrium. Finally, we note that the Reviewer does not have Complete Agency. In effect, Reviewers cannot give "intermediate" reviews. This is natural in voting scenarios, as you can either vote in favour or against, never "in-between" (abstentions being a somewhat different idea).

It is natural to ask what would happen if complete agency was allowed; for example one might imagine a book reviewer could give a tepid review. Ultimately, it has a small impact on the equilibrium outcome, but makes the problem impossible to solve analytically. This is explored further in the extensions section.

The first two observations have a material impact on outcomes, while the third provides an extension we explore later. That the Reviewer exclusively has twice-limited agency actions available to them makes their chosen action as ambiguous and weak as possible, as they can neither signal nor act in any deterministic way. However, the Reviewer ultimately has sufficient degree of agency to sway the Author, as if their degree of agency was too low, then the Author would behave as-if the Reviewer had no decision making power.

4.2.7 Select Twice-Limited Agency Specifications

To complete discussion of twice-limited agency, we highlight some ways it might be implemented. We have already observed one application of twice-limited agency in the equilibrium we computed, but it is good to highlight more variants to better understand the concept. Notably, all variants given here were at some point considered as a candidate mechanism, before the preferred specification was ultimately chosen.

The agent directly chooses the probability of a second probability.

The first and indeed simplest specification is one where the agent directly chooses the probability of a second probability (as shown in the simple example above). This does not add much compared to once-limited agency, but it still allows further abstraction. A key advantage is that it is relatively easy to do Bayesian updating, as the specification is so simple.

The agent chooses some p, and then outcome-probability is $p + \varepsilon$, where $\varepsilon \sim F(\cdot)$.

Alternatively, the agent may choose some p, and then outcome-probability is $p + \varepsilon$, where $\varepsilon \sim F(\cdot)$. This is already a far more complex specification, creating a quite rich outcome-probability space. It does require limiting what p can be chosen as well as picking $F(\cdot)$ to prevent outcome-probabilities above 1 or below 0, but this is not overly problematic. What is indeed more problematic is that Bayesian updating is hard, as in effect every choice of p changes the overall distribution's support, rather than only changing the distribution itself. An additional limitation is that limiting the size of p and the range of ε limits overlap in outcome-probabilities, slightly weakening the value of twice-limited agency. For instance, if $p \in \{0.3, 0.7\}$, $\varepsilon \in [-0.3, 0.3]$, and we allow the second agent to observe outcome-probabilities, then any outcome probability outside of the set $[0.4, 0.6]$ deterministically reveals what p was chosen. This is not necessarily a problem, but depending on the use case it may or may not be desirable. The first variant of the model given here actually used this specification with F being a uniform distribution, and while it was perfectly solvable, it was far less interesting.

The agent chooses some $m \in M$, and outcome probability is then $p \sim F(\cdot|m)$.

Finally, we consider the preferred specification actually used in the model previously shown, where the agent chooses some $m \in M$, and outcome probability is then $p \sim F(\cdot|m)$. Such a specification allows a modeller essentially unlimited choice of M, which can be chosen flexibly to enable the types of behaviours we are interested in. We do require the support of F to be contained in $[0, 1]$, so that F can only return a probability. Broadly, the Beta class of distributions work well here. What we find is that even a simple discrete choice of m can yield every $p \in [0, 1]$

with positive probability, as we observed in the equilibrium solved before. Further, ambiguity is quite high; any p can be realised with any choice of m, creating an environment that is hard to coordinate in. This is precisely what allowed things like γ being surprisingly large, and this is one of the most desirable features of this particular specification.

4.3 Second Mechanism: Realised-Probability Revelation

The second mechanism is the realised outcome-probability revelation at an intermediate point, from which the agent observing this realised probability can make inferences about the sender.

Realised-probability revelation serves as a way to enable interaction between agents in a more flexible manner than simply revealing actions taken. The observing agent is given *partial* information about the state of the world which allows them to conclude broader information about the state of the world, essentially a partially informative signal.

The fact that indirect actions taken by a counter-party have generated this realised-probability are what create an interaction environment in the model. In once-limited agency, revealing the probability of an outcome is equivalent to revealing the action taken by the counter-party, effectively fully revealing the state of the world. Combined with twice-limited agency, however, it means that the observing agent knows only the relative probabilities of facing different agent types.

The objective is to allow misalignment. In order to explain situations where agents appear to act against their best interest, or act against the will of a majority, it is necessary to introduce some coordination failure. The partial understanding of the world due to only observing the probabilities of an outcome generated by a twice-limited agent creates the desired coordination failure.

4.3.1 Realised-Probability Revelation as Incomplete Feedback

A good way to think about realised-probability revelations is as incomplete feedback that can only be partially trusted or adhered to. Pertinent examples of incomplete feedback include website cookies, which paint a good but imperfect image of a user that should be taken with a grain of salt, or student-made reviews of professors, which are coloured by the reviewers' biases.

Any agent wishing to make decision based on incomplete feedback that may have been sincerely and honestly given but cannot be fully trusted, needs to take into account the probability of the feedback being erroneous. This is precisely captured by a realised-probability revelation mechanism as described, since any realised-probability could have been generated by a variety of agents (which may or may not be cooperative), and so it embeds the notion that whatever the intended feedback of the sender, the actual signal could well vary wildly.

4.3.2 Justification of Probability Revelation in Real-World Scenarios

This type of mechanism requires some justification, as it is not immediately intuitive. It generates an interesting fuzzy interaction environment, but it is reasonable to ask how one can justify this type of mechanism "in the real world". We consider one particular example first, which helps ground the idea.

A Presidential candidate in country X aims to make decisions regarding which policies to hold on to and which to discard as an election nears. The candidate's party is centre-right, but the candidate is uncertain where the electorate stands. While their traditional voting base is certain to vote for the party, the candidate does not know where the decisive vote will come in. The candidate believes that centrists are lost to the party, and that the make-or-break voters will be far-right. Consequently, the candidate chooses to cater heavily to the far-right, ensuring that the policies they are proposing would appeal to them while driving the further left elements away.

Before the election, the party conducts an internal vote to see who will lead their party in the election cycle. Our candidate wins, but only just; they edge out their opponent, winning just 51% of the vote. For this story, we suppose that the electorate contributes to this selection process.

What does this mean for the candidate? They know they won, but was their choice sensible or were they lucky? As the candidate does not know who exactly voted how, they have to discern what they can from their vote share. If those further to the left had been completely lost to the party, they may not have taken an interest whatsoever in this intermediate election, and so the far-right should have secured the candidate a decisive win; perhaps the left-faction is in fact the dominant one, and losing them means losing the actual election? On the other hand, might it be that securing the far-right is the only thing that secured our candidate the win? Indeed, those voting against the candidate may have been the bulk of "safe" voters, so once faced with the choice of this candidate or a left-party candidate they will surely vote in favour of the right candidate?

This type of electoral ambiguity well captures the more general story that justifies this revelation mechanism. Wherever it is known that success occurred, by what margin this success occurred, but it is not known *why* the success occurred, it is reasonable to use this probability revelation mechanism.

4.4 Third Mechanism: Payoff-Critical Signal

The final mechanism to address is that the only signal providing any information has a direct influence on actual outcomes, beyond just how it informs the behaviour of agents.

The mechanism in question is that the revealed signal, the realised outcome-probability, is not purely a signal. Instead, it has a direct impact on payoffs by governing the probability of reaching different outcomes. This is not an especially unusual specification in economics, but it proves vital here in confounding agent behaviours and making coordination failure possible.

The key influence on outcomes of this mechanism is that any agent wishing to reduce the probability of success at $t = 1$ must contend with the fact that this decision impacts the probability of sending a positive signal at $t = 2$. Agents are unable to disentangle the signal they send about their behaviour at $t = 2$ from the payoff-relevant decision they take at $t = 1$.

Consider an agent that might desire to succeed at $t = 2$, but ex-ante at $t = 1$, would actually prefer failure to success. Maximizing failure odds at $t = 1$ also maximizes the chance of sending a negative signal. If the agent maximizes failure odds at $t = 1$ they cannot, then, distinguish themselves from an agent that wishes to ensure failure at all periods.

As agents must weigh these trade-offs, we enable counter-intuitive behaviour. Namely, agents that desire failure ex-ante may nevertheless elect to send a positive signal (thus minimizing failure odds), if they believe the risk induced by a low signal outweighs the benefit of lowering success odds early on.

Relation to Twice-Limited Agency

While the payoff-critical signal very clearly interacts with the realised-probability revelation mechanism, it is worth elaborating on the interaction with twice-limited agency.

Consider **complete agency** as defined before. Why would any agent in a binary-outcome scenario care about complete agency? After all, if we only have two outcomes A and B, they either prefer A or B and thus should select whatever $\mathbf{m}$ is required to maximize the expected probability of their preferred outcome being selected. However, because such a choice invariably also leads to higher odds of sending a negative signal, this is no longer true. The trade-off between signal and direct payoff value of any choice of $\mathbf{m}$ is what makes complete agency relevant in any binary outcome scenario.

Ultimately, the model explored does not allow complete agency, as the model becomes far more tractable as a consequence (and added benefit of complete agency turns out to be very minor in this case). However, it is highlighted here to show how these two concepts are actually very closely related.

5 Extensions, Common Questions, and Conclusion

In this section, we consider some possible extension to the model presented as well as possible directions of future research. Moreover, we highlight some interesting questions we might ask about the model and briefly address them. Finally, we conclude this thesis, summarising key findings.

5.1 Extensions and Future Research

5.1.1 Complete Agency

Recall that in the solution to the model we considered, the Reviewer's agency was incomplete. While it captures well any scenario where a vote is cast, it might miss situations where an intermediate-quality review is reasonable. The main point of this section is to highlight that an equilibrium as described can still exist, although it will not be fully solved. However, the minimum necessary proofs to demonstrate the claim will be provided.

To model Complete Agency, the set-up is precisely as before, except that now, $m_t \in [\frac{1}{3}, 3]$ (rather than $m_t \in \{\frac{1}{3}, 3\}$). We note first of all that the Reviewer's optimal choice at $t = 2$ is unchanged relative to the solution given above, as it never relied on specific choices of the set M beyond its upper and lower bound. We continue to assume that the Author follows a cut-off strategy, as that is the most sensible action for them to take given that higher observed values of α_1 are likelier the more favourable a Reviewer is.

<u>Reviewer Behaviour at $t = 1$.</u>

Proposition 5.1. *For any equilibrium where the Author continues at $t = 1$, optimal Reviewer choice of m_1 solves the equation*

$$q^{m_1}\left[\ln(q)m_1(1+m_1)+1\right] = \frac{4c-3r-\delta s}{3(\delta s - r)} \tag{5.1}$$

Proof. The Reviewer at $t = 1$ now solves

$$\max_{m_1} \{E[\text{payoff}|m_1]\} =$$

$$\max_{m_1} \left\{ \frac{3r}{4} \left(\frac{m_1(1-q^{m_1})}{m_1+1} \right) + \frac{\delta s}{4} \left(\frac{m_1(1-q^{m_1})}{m_1+1} \right) + \delta s \left(\frac{m_1 q^{m_1}}{m_1+1} \right) + s \left(1 - \frac{m_1}{m_1+1} \right) \right\}$$

which follows from the proof of proposition 3.7, up to the point where the discreteness of M is introduced. We solve the FOC:

$$\begin{aligned} \frac{\partial E[\text{payoff}|m_1]}{\partial m_1} &= \left(\frac{1-q^{m_1}}{(m_1+1)^2} \right) \frac{3r}{4} - \ln(q) q^{m_1} \frac{3r}{4} \left(\frac{m_1}{m_1+1} \right) + \left(\frac{1-q^{m_1}}{(m_1+1)^2} \right) \frac{\delta s}{4} \\ &- \ln(q) \left(\frac{q^{m_1} m_1}{m_1+1} \right) \frac{\delta s}{4} + \frac{q^{m_1} \delta s}{(m_1+1)^2} + \ln(q) \left(\frac{m_1 q^{m_1}}{m_1+1} \right) \delta s - \frac{c}{(m_1+1)^2} = 0 \\ &\Rightarrow q^{m_1} \left[\ln(q) m_1 (1+m_1) + 1 \right] = \frac{4c - 3r - \delta s}{3(\delta s - r)} \end{aligned}$$

Note that SOCs hold to make this a local maximum and not minimum.

□

This presents us with a significant complication, as the equation governing Reviewer action cannot be solved for m_1, forcing us to use inverse-integral techniques to compute Author behaviour.

<u>Author Behaviour at $t = 2$.</u>

We now want to verify that the Author will actually employ a cut-off type behaviour as conjectured, since if that is the case, we can ensure that the posited equilibrium exists. We recall that the Author continues if

$$E[\alpha_2|\alpha_1] r \geqslant s$$

and that to compute $E[\alpha_2|\alpha_1]$, we require

$$\mathbb{P}(r < \delta s|\alpha_1) = \frac{f(\alpha_1|r < \delta s)\mathbb{P}(r < \delta s)}{f(\alpha_1)}$$

This follows from Reviewer behaviour at $t = 2$ being unchanged. Now, however, the computation of $f(\alpha_1)$ is harder, as Reviewer behaviour at $t = 1$ is more complex.

Conjecture 5.2. *For sufficiently small r_l, $m_1 = \frac{1}{3}$ $\forall r \leqslant r_l$. For sufficiently large r_h, $m_1 = 3$ $\forall r \geqslant r_l$.*

This can be informally verified by plotting the expected payoff of the Reviewer with varying values of r and m_1. Moreover, one could let $r_l = \frac{1}{3}$, $r_h = 3$ and the below holds with minor changes, but this is not necessary. Given this, we have the following:

$$\begin{aligned} f(\alpha_1) &= f(\alpha_1|r \leqslant r_l)\mathbb{P}(r \leqslant r_l) + f(\alpha_1|r \geqslant r_h)\mathbb{P}(r \geqslant r_h) \\ &+ \int_{r_l}^{r_h} f(\alpha_1|r_l < r < r_h) f(r|r \sim U[r_l, r_h]) dr \\ &= \frac{1}{3} \left(\alpha_1^{-\frac{2}{3}} \right) \left(\frac{r_l}{R} \right) + 3\alpha_1^2 \left(1 - \frac{r_h}{R} \right) + \int_{r_l}^{r_h} m_1(r) \alpha_1^{m_1(r)-1} \left(\frac{1}{r_h - r_l} \right) dr \end{aligned}$$

If we had to set $r_l = \frac{1}{3},\ r_h = 3$, we could simply remove terms outside the integral, but again this is not necessary. $m_1(r)$ cannot be explicitly defined, as we cannot write Reviewer behaviour in the explicit form $m_1(r)$. We will thus make use of the fact that the inverse function $r(m_1)$ can be written explicitly.

We get the below, where the constant $\frac{1}{r_h - r_l}$ has been omitted [1]:

$$\int_{r_l}^{r_h} m_1(r)\alpha_1^{m_1(r)-1} dr = \int_{r_l}^{r_h} \frac{1}{m_1'(r)} \frac{d}{dr}\alpha_1^{m_1(r)} dr = \int_{r_l}^{r_h} \frac{1}{m_1'(r)} d\alpha_1^{m_1(r)}$$

Then, we employ a change of variables and combine with the above, where we let $y = m_1(r)$:

$$\int_{r_l}^{r_h} \frac{1}{m_1'(r)} d\alpha_1^{m_1(r)} = \int_{m_1(r_l)}^{m_1(r_h)} \frac{1}{m_1'(m_1^{-1}(y))} d\alpha_1^{y} = \int_{m_1(r_l)}^{m_1(r_h)} \frac{dm_1^{-1}(y)}{dy} \ln(\alpha_1)\alpha_1^{y} dy$$

Recall that $m_1(r)$ was defined via the equation

$$q^{m_1}\left[\ln(q)m_1(1+m_1)+1\right] = \frac{4c - 3r - \delta s}{3(\delta s - r)}$$

which then means that we can write $m_1^{-1}(y) = r(m_1)$ as:

$$r(m_1) = \frac{3\delta s q^{m_1}(\ln(q)m_1(1+m_1)+1) - 4s + \delta s}{3\left[q^{m_1}(\ln(q)m_1(1+m_1)+1) - 1\right]}$$

Computing the derivative of $r(m_1)$ with respect to m_1 yields:

$$\frac{\partial r(m_1)}{\partial m_1} = \frac{4(1-\delta)(m_1+1)s q^{m_1}\ln(q)(m_1\ln(q)+2)}{3(q^{m_1} + m_1(m_1+1)q^{m_1}\ln(q) - 1)^2}$$

By the way we defined the boundaries, we have that $m_1(r_l) = \frac{1}{3}$ and $m_1(r_h) = 3$, and so we wish to solve the following definite integral:

$$\int_{\frac{1}{3}}^{3} \left[\frac{4(1-\delta)(m_1+1)s q^{m_1}\ln(q)(m_1\ln(q)+2)}{3(q^{m_1} + m_1(m_1+1)q^{m_1}\ln(q) - 1)^2}\right] \ln(\alpha_1)\alpha_1^{m_1} dm_1$$

This integral does not have an analytical solution, but it can be solved computationally using Python and the scipy function "quad". Code used is available on request.

We find that given, say, $q = 0.15$, the definite integral above ranges from 0.002 when $\alpha_1 = 0.9$ to 0.036 when $\alpha_1 = 0.1$. These are ultimately very small values, which may be inflated once we multiply this number by $\frac{1}{r_h - r_l}$ if $r_h - r_l < 1$, but nevertheless we observe that we *can* compute Author behaviour even in the complete agency case, and importantly, the impact of complete agency is quite small.

Take-Away

While we have not competed a full equilibrium for the complete agency case, we can follow the procedure outlined in the incomplete agency case with discrete choices of m_1 to compute an equilibrium. The point is that this extension is possible, but does not provide many new interesting insights while massively increasing the complexity of the problem; after all, the change in $f(\cdot)$ from adding the integral is very small, as shown. This is why the discrete-choice M set-up was preferred.

[1] Note that the following would not have been possible without Professor Ian Jewitt, who kindly offered his help with these more difficult integral operations.

5.1.2 Direct Model Extensions

Here, we consider some further extensions to the model itself. These are not altogether different research directions, but instead simple modifications that could reasonably be made to the model.

Changing the Distribution of F.

As has already been briefly mentioned, any F in the Beta family of distributions can be used without needing to alter the model much. Indeed, the power distribution is a special case of a Beta distribution. Note that if an alternative F is chosen, the Reviewer's action set may need to be adjusted. One example would be to select F to be the Uniform Distribution (another special case of a Beta distribution). Such a set-up requires changing the support of the Uniform Distribution to produce twice-limited agency, as was discussed in a previous section, but it can be shown that the principles seen in the equilibrium computed here do not change.

Adding Additional Periods.

Adding periods can allow for greater levels of coordination failure. For example, the threat of quitting after s is discounted several times and may encourage Reviewers to provide even more favourable reviews than otherwise, or on the other hand the risk of failure at future periods may prompt Reviewers to be even harsher. As observed in the proofs above, the problem is already very tricky with only two periods due to the extensive updating and it is not clear that the additional insights from adding periods are worth the computational complexity.

Discounting the Author's Quitting Payoff.

It is reasonable to suggest that if the Reviewer faces some sunk-cost loss, so should the Author. The choice not to discount the Author's outside option was largely one of convenience, but it is fine to discount it. Discounting quitting payoff would reduce the range of parameters under which the Author would be willing to continue at $t = 1$ as well as make them less willing to quit at $t = 2$, but the impact is minor.

5.1.3 Future Research

The above are not research directions entirely of their own, being only small modifications to the model as presented here. Now, we consider questions that remain unanswered, and would be worth exploring in the future.

Applying Twice-Limited Agency Beyond this Model.

There is potential for a garbled signalling method such as twice-limited agency to be used in other settings. It would be particularly interesting to find some alternative signal to pair with it, as probability revelation has already been explored quite extensively here.

Develop Computer Programmes to Solve for Equilibria and Corresponding Parameter Combinations.

Developing a computer programme to solve for equilibria and parameter combinations yielding said equilibria is less theoretically involved, but nevertheless important. While classes of equilibria have been defined here, what parameters yield which equilibria would be of interest. Since equilibria can be solved for explicitly as per the propositions given here, it should be quite straight-forward and would be an easy contribution.

Empirically Testing the Model and Calibrating Parameters.

It is crucial to move beyond the theoretical framework proposed here and thus to empirically test the model and calibrate parameters. While the model is unlikely to perfectly model reality, it would be useful to know how well it predicts real outcomes. Notably, it would be interesting to know if parameters can be calibrated and discerned sufficiently accurately to make at least directionally accurate predictions about project outcomes. The randomness inherent in the model means that empirical work is difficult, as prediction errors can be either due to randomness or model failure, but it would be worth exploring.

5.2 Questions of Interest

We summarise key questions of interest that arise naturally from the model. They have all been addressed above, but we consolidate them here so as to highlight these points that might otherwise be difficult to find elsewhere in the document.

"How can probabilities be observed?"

The first of these is "how can probabilities be observed?" In a real-life setting, they would be inferred from some feedback more than directly observed. Think, for example, of a politician trying to glean information from poll results and what that means for their final election odds. For a more detailed discussion, see section 4.3.2 on page 31.

"What do we gain from repeating the game?"

Another important question is "what do we gain from repeating the game?" Repeating the game allows the Author to learn something about the Reviewer, generating a rich updating environment. In terms of outcomes, we enable coordination failure, since only by repeating the game can we observe situations where Reviewers select outcomes against their best interest in the first period. Section 4.1 on page 23 discusses this a bit more, but broadly this effect is implicitly critical throughout.

"Why do we want partial revelation?"

Finally, we might wonder "why do we want partial revelation?" With full revelation, coordination failure would not really be possible. If full revelation was possible, there would be no risk to a favourable Reviewer of the Author quitting at $t = 2$ despite a favourable review. Counter-intuitive results observed in reality, where outcomes against electorate best-interest are upheld, would not be explained in this model. The discussion in section 4.2 on page 25 touches on this indirectly, but discussion in section 4.1 on page 23 covers it well also.

5.3 Conclusion

This thesis explores a theoretical model that allows us to explain the misalignment of group interests in a variety of settings. By abstracting a many-agent environment to a simpler two agent model, leveraging noisy signals and uncertainty to emulate the effect of a large electorate, we are able to tractably solve an otherwise impractical interaction. Depending on specific parameter choices, we find that voters might actively vote against their best interest, voting against desirable projects or in favour of undesirable ones due to fear of them being cancelled in the future. Equally, decision makers (legislators) may reasonably cancel projects both they and the voter are in favour of, or push through unpopular projects, due to misinterpreting signals such as polls or primary elections. Crucially, all of these results can arise in well-designed environments with no need for institutional failure or perverse incentives; poor signalling capacity, combined with uneven power distributions within the electorate, proves sufficient for all of these results to arise. Ultimately, the full range of observed coordination failures in project completion environments can be sustained with a relatively simple two-agent model, providing a strong basis for understanding these phenomena as a case of poor coordination.

6 References and Glossary

Numbers next to entries are the page number of the first occurrence of each term in the glossary.

Glossary

M Set of possible m_t from which Reviewer can choose. 5

R Greatest possible Reviewer type r. 5

α_t Realised probability of continuing to the next period (or completing the project) at time t. 5

δs Reviewer payoff given Author quits or project fails at $t = 2$, where δ is the "sunk-cost" loss of quitting/failure in the second period. 5

γ Author's belief of the Reviewer's cut-off such that $\forall r \geqslant \delta s + \gamma$ they play high values of m_1, otherwise playing low m_1. 12

a Author payoff given Author quits. 5

b Author payoff given project completion. 5

m_t Reviewer's action at time t, which influences distribution of α_t. 5

q Reviewer's belief of the Author's cut-off such that $\forall \alpha_1 \geqslant q$ they continue at $t = 2$, otherwise quitting at $t = 2$. 12

r Reviewer payoff given project completion. 5

s Reviewer payoff given Author quits or project fails at $t = 1$. 5

Author Decision-making agent without private information, comparable to receiver in standard signalling game. 4

Complete Agency Agent has complete agency if the set of expected outcome-probabilities they can induce is convex. That is, if they can induce two extreme expected outcome-probabilities, they can induce all expected outcome-probabilities in-between the extreme ones as well. 29

Degree of Agency Maximum change in expected outcome-probabilities that an agent can induce. For mathematical definition, see definition. 29

Limited Agency Action Set Set M of limited-agency actions an agent with limited agency can take. 29

Once-Limited Agency Agency is once-limited if the agent to which it applies cannot directly select an outcome but can affect the probability of an outcome. 25

Perfect Agency Agent has perfect agency if degree of agency is equal to 1. 29

Reviewer Private information agent with limited decision-making ability, comparable to sender in a standard signalling game. 4

Twice-Limited Agency Agency is twice-limited if the agent to which it applies can neither directly select an outcome, nor directly select the probability of an outcome, but can affect the distribution governing the probability of an outcome. 25

References

Alonso, R., & Câmara, O. (2016). Persuading voters. *American Economic Review*, *106*. https://doi.org/10.1257/aer.20140737

Beetsma, R. M., & van der Ploeg, F. (2007). *Partisan public investment and debt: The case for fiscal restrictions* (Economics Working Papers No. ECO2007/37). European University Institute. https://ideas.repec.org/p/eui/euiwps/eco2007-37.html

Bergemann, D., & Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, *11*. https://doi.org/10.3982/te1808

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, *24*. https://doi.org/10.1214/aoms/1177729032

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, *50*. https://doi.org/10.2307/1913390

DAntoni, M., Lehrer, E., Tabbach, A., & Solan, E. (2024). Bayesian dissuasion with bandit exploration. https://arxiv.org/abs/2401.04974

Dworczak, P., & Martini, G. (2019). The simple economics of optimal persuasion. *Journal of Political Economy*, *127*. https://doi.org/10.1086/701813

Fiva, J. H., & Natvik, G. J. (2013). Do re-election probabilities influence public investment? *Public Choice*, *157*. https://doi.org/10.1007/s11127-012-9946-8

Gentzkow, M., & Kamenica, E. (2016). A rothschild-stiglitz approach to bayesian persuasion. *American Economic Review*, *106*. https://doi.org/10.1257/aer.p20161049

Gill, D., & Sgroi, D. (2008). Sequential decisions with tests. *Games and Economic Behavior*, *63*. https://doi.org/10.1016/j.geb.2006.07.004

Glazer, A. (1989). Politics and the choice of durability. *American Economic Review*, *79*. http://www.jstor.org/stable/1831446

Hadar, J., & Russell, W. R. (1969). Rules for ordering uncertain prospects. *American Economic Review*, *59*. http://www.jstor.org/stable/1811090

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, *101*. https://doi.org/10.1257/aer.101.6.2590

Kreps, D. M., & Wilson, R. (1982). Sequential equilibria. *Econometrica*, *50*. https://doi.org/10.2307/1912767

Mathevet, L., Perego, J., & Taneva, I. (2020). On information design in games. *Journal of Political Economy*, *128*. https://doi.org/10.1086/705332

Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences - PNAS*, *39*, 1095–1100. https://doi.org/10.1073/pnas.39.10.1095